



INTERNATIONAL ATOMIC ENERGY AGENCY
UNITED NATIONS EDUCATIONAL, SCIENTIFIC AND CULTURAL ORGANIZATION
INTERNATIONAL CENTRE FOR THEORETICAL PHYSICS
I.C.T.P., P.O. BOX 586, 34100 TRIESTE, ITALY, CABLE: CENTRATOM TRIESTE



**WINTER COLLEGE ON "MULTILEVEL TECHNIQUES IN
COMPUTATIONAL PHYSICS"**

**Physics and Computations with Multiple Scales of Lengths
(21 January - 1 February 1991)**

H4.SMR 539/9

***Multigrid Techniques: 1984 Guide
With Applications to Fluid Dynamics***

**A. Brandt
The Weizmann Institute of Science
Rehovot, Israel**

**MULTIGRID TECHNIQUES : 1984 GUIDE
WITH APPLICATIONS TO FLUID DYNAMICS**

Achi Brandt

Department of Applied Mathematics
The Weizmann Institute of Science
Rehovot 76100, Israel

This edition of the Guide has been prepared as Lecture Notes for the Computational Fluid Dynamics Lecture Series at the von-Karman Institute for Fluid Dynamics, Rhode-Saint-Genèse, Belgium, March 26-30, 1984. It is reproduced with the kind permission of the Lecture Series Director, Professor J.A. Essers.

The research reported herein is sponsored by the Air Force Wright Aeronautical Laboratories, Air Force Systems Command, United States Air Force, under Grant AFOSR 82-0063.

Abstract

Starting with an elementary exposition of multigrid fast solvers and with insights into their analyses and their most general algebraic applicability, detailed practical guidelines are then given how to obtain, stage by stage, the full multigrid efficiency for general elliptic and non-elliptic problems, linear as well as nonlinear, scalar or vectorial, smooth or strongly discontinuous, with various possible singularities, boundary conditions and supplementary global conditions. Quantitative insights through local mode analyses, combined with gradual algorithm development, are emphasized throughout, and general rules and approaches are explained for the design of relaxation, coarsening and interpolation. Beyond these fast-solver aspects of multigrid, advanced methods are then described, including various applications of the full approximation scheme (FAS), local refinements and local coordinate transformations, error estimation and grid adaptation criteria, small storage algorithms, the double discretization and other techniques for high-order approximations, algebraic multigrid (AMG), multi-level reduction of complexity for integral equations and for chains of problems, and treatment of time-dependent, eigenvalue and optimization problems. Special chapters describe the solution of Cauchy-Riemann, Stokes and incompressible and compressible Navier-Stokes equations, with numerical results for staggered and non-staggered grids.

CONTENTS

	<u>Page</u>
0. Introduction	1
0.1 Where and why multigrid can help	1
0.2 Purpose of this Guide	3
1. Elementary acquaintance with multigrid	7
1.1 Properties of slowly converging errors	7
1.2 Error smoothing and its analysis: Example	9
1.3 Coarse grid correction	12
1.4 Multigrid cycle	13
1.5 Model program and output	14
1.6 Full multigrid (FMG) algorithm	20
1.7 General warnings. Boundary conditions. Nonlinearity	20
PART I. STAGES IN DEVELOPING FAST SOLVERS	23
2. Stable discretization of boundary value problems	23
2.1 Interior stability measures: h-ellipticity	25
2.2 Boundaries, discontinuities	28
3. Interior relaxation and smoothing factors	29
3.1 Local analysis of smoothing	29
3.2 Work, robustness and other considerations	31
3.3 Block relaxation rule. Semi smoothing	33
3.4 Distributive, weighted, collective and box Gauss-Seidel. Principal linearization	35
3.5 Simultaneous displacement (Jacobi) schemes	39
3.6 Relaxation ordering. Vector and parallel processing	40
3.7 Principles of relaxing general PDE systems	42
3.8 ILU smoothers	44
4. Interior two-level cycles	46
4.1 Two-level mode analysis. Switching criteria	48
4.2 Choice of coarse-grid	51
4.2.1 Semi coarsening	51
4.2.2 Modified coarse-grid functions	52
4.3 Orders of interpolations and residual transfers	52
4.4 Variable operators. Full weightings	55
4.5 Coarse-grid operator. Variational and Galerkin coarsening	56
4.6 Strongly discontinuous, strongly asymmetric operators	57

	<u>Page</u>		<u>Page</u>
5. Boundary conditions and two-level cycling	59	10. High-order techniques	103
5.1 Simplifications and debugging	59	10.1 Fine-grid deferred corrections. Pseudospectral methods	103
5.2 Interpolation near boundaries and singularities	61	10.2 Double discretization: High-order residual transfers	105
5.3 Relaxation of boundary conditions	61	10.3 Relaxation with only subprincipal terms	107
5.4 Residual transfers near boundaries	62	11. Coarsening guided by discretization	107
5.5 Transfer of boundary residuals	63	12. Real role of relaxation	110
5.6 Treatment and use of global constraints	64	13. Dealgebraization of multigrid	112
5.7 Structural singularities. Reentrant corners. Local relaxation	66	13.1 Reverse trend: Algebraic multigrid (AMG)	114
6. Many-level cycles	68	14. Practical role of rigorous analysis	116
6.1 Multigrid cycles. Initial and terminal relaxation	68	15. Chains of problems. Frozen τ	120
6.2 Switching criteria. Types of cycles	68	16. Time dependent problems	122
6.3 Coarsest grids. Inhomogeneous and indefinite operators	70	PART III. APPLICATIONS TO FLUID DYNAMICS	125
7. Full multigrid (FMG) algorithms	72	17. Cauchy-Riemann equations	127
7.1 Order of the FMG interpolation	72	17.1 The differential problem	127
7.2 Optimal switching to a new grid	73	17.2 Discrete Cauchy-Riemann equations	128
7.3 Total computational work. Termination criteria	74	17.3 DGS relaxation and its smoothing rate	131
7.4 Two-level FMG mode analysis	76	17.4 Multigrid procedures	133
7.5 Half-space FMG mode analysis. First differential approximations	78	17.5 Numerical results	135
PART II. ADVANCED TECHNIQUES AND INSIGHTS	80	17.6 Remark on non-staggered grids	136
8. Full-Approximation Scheme (FAS) and applications	80	18. Steady-state Stokes equations	137
8.1 From CS to FAS	80	18.1 The differential equations	137
8.2 FAS: Dual point of view	82	18.2 Finite-difference equations	138
8.3 Nonlinear problems	83	18.3 Distributive relaxation	141
8.3.1 Eigenvalue problems	84	18.4 Multigrid procedures	143
8.3.2 Continuation (embedding) techniques	85	18.5 Numerical results	145
8.4 Estimating truncation errors. τ extrapolation	86	18.6 Non-staggered grids: Recent results	145
8.5 FAS interpolations and transfers	88	19. Steady-state incompressible Navier-Stokes equations	149
8.6 Application to integral equations	89	19.1 The differential problem	149
8.7 Small storage algorithms	90	19.2 Staggered finite-difference approximation	150
9. Local refinement and grid adaptation	91	19.3 Distributive relaxation	151
9.1 Non-uniformity organized by uniform grids	91	19.4 Multigrid procedures and numerical results	153
9.2 Anisotropic refinements	94	19.5 Results for non-staggered grids	154
9.3 Local coordinate transformations	94		
9.4 Sets of rotated cartesian grids	97		
9.5 Self-adaptation techniques	97		
9.6 Exchange rate algorithms. λ -FMG	101		

	<u>Page</u>
20. Compressible Navier-Stokes and Euler equations	155
20.1 The differential equations	155
20.1.1 Conservation laws and simplification	155
20.1.2 The principal part	157
20.1.3 Elliptic singular perturbation	158
20.1.4 Inviscid (Euler) and subprincipal operators	159
20.1.5 Incompressible and small Mach limits	161
20.2 Stable staggered discretization	162
20.2.1 Discretization of the subprincipal part	162
20.2.2 The full quasi-linear discretization	164
20.2.3 Boundary conditions	164
20.3 Distributed relaxation for the simplified system	165
20.3.1 General approach to relaxation design	165
20.3.2 Possible relaxation scheme for inviscid flow	167
20.3.3 Distributed collective Gauss-Seidel	168
20.3.4 Relaxation ordering and smoothing rates	169
20.3.5 Summary: relaxation of the full system	172
20.4 Multigrid procedures	174
21. Remarks on solvers for transonic potential equations	175
21.1 Multigrid improvements	175
21.2 Artificial viscosity in quasi-linear formulations	176

0. INTRODUCTION

0.1 Where and why multigrid can help

The starting point of the multigrid method (or, more generally, the Multi-Level Adaptive Technique - MLAT), and indeed also its ultimate upshot, is the following "golden rule":

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.

That is, whenever the computer grinds very hard for very small or slow real physical effect, there must be a better computational way to achieve the same goal. Common examples of such stalling are the usual iterative processes for solving the algebraic equations arising from discretizing partial-differential, or integro-differential, boundary-value (steady-state) problems, in which the error has relatively small changes from one iteration to the next. Another example is the solution of time-dependent problems with time-steps (dictated by stability requirements) much smaller than the real scale of change in the solution. Or, more generally, the use of too-fine discretization grids, where in large parts of the computational domain the meshsize and/or the timestep are much smaller than the real scale of solution changes. Etc.

If you have such a problem, multi-level techniques may help. The trouble is usually related to some "stiffness" in your problem; i.e., to the existence of several solution components with different scales, which conflict with each other. For example, smooth components, which are efficiently approximated on coarse grids but are slow to converge in fine-grid processes, conflict with high-frequency components which must be approximated on fine grids. By employing interactively several scales of discretization, multilevel techniques resolve such conflicts, avoid stalling and do away with the computational waste.

In fully developed MLAT processes the amount of computations should be determined only by the amount of real physical information.

The main development of multilevel techniques has so far been limited to their role as fast solvers of the algebraic equations arising in discretizing boundary-value problems (steady-state problems or implicit steps in evolution problems). The multigrid solution of such problems usually requires just few (four to ten) work units, where a work unit is the amount of computational work involved in expressing the algebraic equations (see Sec. 7.3). This efficiency is obtained for all problems on which sufficient research has been made, from simple model

problems to complicated nonlinear systems on general domains, including diffusion problems with strongly discontinuous coefficients, integral equations, minimization problems with constraints; from regular elliptic to singular-perturbation and non-elliptic boundary-value problems. Due to the iterative nature of the method, nonlinear problems require no more work than the corresponding linearized problems. Linearization is thus neither needed nor usually recommended (see Sec. 8.3). Problems with global constraints are solved as fast as the corresponding unconstrained difference equations, using a technique of enforcing the constraints only at the coarse-grid stages of the algorithm (Sec. 5.6). Few work units are also all the work required in calculating each eigenfunction of discretized eigenproblems (Sec. 8.3.1). Moreover, all multigrid processes can fully be parallelized and vectorized. A model multigrid program on the Cyber 205 solves 3 million equations per second [B25].

Beyond the fast solvers, multilevel techniques can be very useful in other ways related to stiffness. They can provide very efficient grid-adaptation procedures for problems (either boundary-value or evolution problems) in which different scales of discretization are needed in different parts of the domain (see Sec. 9). They can give new dimension of efficiency to stiff evolution problems (Sec. 16). And they can resolve the conflict between higher accuracy and stability in cases of non-elliptic and singular perturbation boundary-value problems (Sec. 10.2). In addition, multi-level techniques can enormously reduce the amount of discrete relations employed in solving chains of similar boundary-value problems (as in processes of continuation, and in optimization problems; see Secs. 15, 13), or in solving integral equations (see Sec. 8.6). They can also be used to vastly cut the required computer storage (Sec. 8.7). All these topics are now under active research.

Multilevel processes can also cut, sometimes by several orders of magnitude, the computer resources needed to solve some large systems which do not originate from partial-differential or integral equations. The common feature in those systems is that they involve many unknowns related in a low-dimensional space; i.e., each unknown u_p is defined at a point $P = (x_1, \dots, x_d)$ of a low-dimensional space (d is usually 2 or 3), and the equations are given in terms of these coordinates x_j . Moreover, the coupling between two values u_p and u_q generally becomes weaker or smoother as the distance between P and Q increases, except perhaps for a small number of particular pairs (P, Q) . Examples are: the equations of multivariate interpolation of scattered data [B33], [M1]; geodetic problems of finding the locations of many stations that best fit a large set of local observations [M2]; problems in transportation, economy

[V0], and queuing theory [K3]; statistical problems on lattices, arising in statistical mechanics (e.g. Ising model) and in "gauge" theories of elementary particles; and various systems of tomography, image processing, picture reconstruction and pattern recognition [N1], [T1], [G1]. There is in fact strong evidence that the human vision processes themselves are multilevelled [C2], [W5], [T4]. In several of these areas multigrid research has just recently started.

0.2 Purpose of this Guide

The opening chapter of this Guide is dedicated to numerical analysts who have no previous acquaintance with multigrid methods. It also gives some references to other introductory material. (Secs. 1.1 and 1.7 in that chapter may interest veteran multigridders, too.)

The main parts of this Guide are chiefly intended for people with some multigrid knowledge, or even experience. In fact, I was mainly motivated by the following situation, so often encountered in the last few years: A good numerical analyst tries a multigrid solver on a new problem. He knows the basics, he has seen it implemented on another problem, so he has no trouble writing the program. He gets results, showing a certain rate of convergence, perhaps improving a former rate obtained with a one-grid program. Now he is confronted with the question: Is this the real multigrid efficiency? or is it many times slower, due to some conceptual error or programming bug? The algorithm has many parts and aspects: relaxation sweeps and coarse-to-fine and fine-to-coarse transfers at interior points and at points near boundaries; relaxation and transfers of the boundary conditions themselves; treatment of boundary and interior singularities and/or discontinuities; choosing the coarse-grid variables and defining its equations; the method of solving on the coarsest grid; the general flow of the algorithm; etc. A single error (a wrong scheme or a bug) in any of these parts may degrade the whole performance very much, but is still likely to give an improvement over a one-grid method, misleading the analyst to believe he has done a good job. How can an error be suspected and detected? How can one distinguish between various possible troubles? And what improved techniques are available?

The key to a fully successful code is to know in advance what efficiency is ideally obtainable, and then to construct the code gradually in a way which ensures approaching that ideal, telling us at each stage which process may be responsible for a slowdown. It is important to work in that spirit: Do not just observe what efficiency is obtained

by a given multigrid algorithm, but ask yourself what is the *ideal* efficiency and *find out* how to obtain it. To guide inexperienced multigridders in that spirit is the main purpose of this Guide.

We believe that any discrete system derived from a continuous problem is solvable "to the level of truncation errors" in just few "work units" (see Sec. 7.3). To obtain this performance, the first crucial step is to construct a relaxation scheme with a high "smoothing rate" (see Sec. 3). Then the interior inter-grid transfers and coarse-grid operator should be designed (Sec. 4), and full numerical experiments can be started with cycling algorithms, aiming at obtaining the interior rate (Secs. 5, 6). Finally, "Full Multi-Grid" (FMG) algorithms can then be implemented, and "solvability in just few work units" can be tested (Sec. 7). These stages of development are outlined in Part I below, pointing out many possibilities and technical points, together with theoretical tools needed for quantitative insights into the main processes.

The quantitative aspect in these theoretical tools is important, since we want to distinguish between the efficiency of several candidate multigrid algorithms, all of which may be "asymptotically optimal" (i.e., solving the problem in a uniformly bounded number of work units), but some of which may still be several orders of magnitude faster than others. Except for some model problems, most present-day rigorous mathematical theories of multigrid algorithms do not give us accurate enough insights (see Sec. 14), hence the present guide will emphasize the role of "local mode analyses". These analyses (see Secs. 2.1, 3.1, 4.1, 7.4, 7.5) neglect some of the less work-consuming processes so as to obtain a clear and precise picture of the efficiency of the more important processes. The predictions so obtained can be made accurate enough to serve in program optimization and debugging. Experience has taught us that careful incorporation of such theoretical studies is essential for producing reliable programs which fully utilize the potential of the method.

Part II of this Guide summarizes more advanced multigrid techniques and insights. Mainly, it is intended to show how to use the multi-level techniques far beyond their more familiar capacity as fast linear algebraic solvers. See the survey in Sec. 0.1 and the list of contents.

In Part III we bring applications to fluid dynamics. Whereas in Part I the information about technique for all problems is ordered according to their common stages of development, in Part III we study specific problems, separately from each other. But the order is again according to a certain line of development, namely, starting from simple problems and gradually learning our way to more complicated ones. The em-

phasis is on *systems* of differential equations; scalar problems are not separately treated.

This Guide can be viewed as an extension of an earlier Guide [B27], with numerous updates, few new sections (Secs. 3.8, 5.7), a new chapter (Chapter 1) and a whole new part (Part III). It is not intended just for teaching, but also for organizing and unifying the material. It is also used as an opportunity to mention some advances which have not appeared in the literature before. In particular, [B27] already included some new relaxation schemes such as "Box Gauss-Seidel" (Sec. 3.4) and relaxation with only sub-principal terms (Sec. 10.3); the general rule of block relaxation (Sec. 3.3); an analysis of the orders of interpolation and residual transfers which should be used in solving *systems* of differential equations (Secs. 4.3, 7.1); the multigrid treatment of global constraints (Sec. 5.6); an application of FAS to obtain much more efficient discretization to *integral* equations, leading sometimes to solutions in $O(n)$ operations, where n is the number of discrete unknowns (Sec. 8.6); some innovations in higher-order techniques (Sec. 10.2); unified switching and adaptation criteria (Sec. 9.6); multi-level approach to optimization (Sec. 13); and the Algebraic Multi-Grid (AMG) method (Sec. 13.1). Results from a recent work, still unpublished, on non-elliptic and singular perturbation problems [B17] are also mentioned, including a summary of stability requirements (Sec. 2.1); the double-discretization scheme (Sec. 10.2); the two-level FMG mode analysis, which tends to replace the usual two-level mode analysis (Secs. 7.4, 7.5); the F cycle (a hybrid of V and W cycles; Sec. 6.2). (The main topic from [B17] hardly mentioned here is the multigrid treatment of discontinuities, in which research is currently underway. But see Secs. 2.2 and 8.5.) The discussions on the real role of relaxation (Sec. 12), on the general approach to coarsening questions (Sec. 11), and on "algebraization" and "dealgebraization" trends in multigrid development (Sec. 13) were added as general new viewpoints, related to each other, somewhat philosophical, but certainly useful.

This new version of the Guide includes in addition some remarks about the "principal linearization" used in relaxation (usually meaning no linearization at all -- see Sec. 3.4); the general algebraic property of slowly converging relaxation schemes (Sec. 1.1); the superfluity of "perfect smoothers" for non-elliptic or slightly elliptic systems (Secs. 3.3, 3.6); the principle of relaxing general PDE operators in terms of the factors of their subprincipal-part determinant (Sec. 3.7); some new debugging devices (e.g. Secs. 4, 5.1); stabilizing coarsening by added global constraints (Sec. 5.6); the treatment of structural singularities

such as reentrant corners (Secs. 5.7, 9.6); and new numerical results for the Stokes and compressible and incompressible Navier-Stokes equations, including results for non-staggered grids (Secs. 18.6, 19.5, 20.5).

1. ELEMENTARY ACQUAINTANCE WITH MULTIGRID

The following few pages would acquaint you with the conceptual basis of all multigrid solvers, with elementary mode-analysis and with an example of a simple algorithm together with its two-page Fortran code and output.

For a more detailed introduction to multigrid techniques, through a comprehensive treatment of some *model* problems by a variety of multigrid algorithms, mode analyses and numerical experiments, see [S4]. The latter appears in a book [H4] which also includes a previous version of the present Guide [B27], a complete (as of 1982) multigrid bibliography, and many other multigrid papers. Additional material, and in fact summaries of all new multigrid papers, appear quarterly in the MULTIGRID NEWSLETTER, obtainable in North and South America from its Editor (Steve McCormick, Department of Mathematics, Colorado State University, Fort Collins, CO 80523, U.S.A.) and in other countries from the Managing Editor (Kurt Brand, GMD/FIT, Postfach 1240, D-5205 St. Augustin 1, Federal Republic of Germany). Periodically, it publishes a complete list of all past papers. Also available is a set of Videod multigrid lectures [B33].

1.1 Properties of slowly converging errors

The origin of multi-level (multigrid) fast solvers is a certain insight concerning the nature of the algebraic errors that become dominant when conventional iterative schemes are slow to converge. Let us first present this insight in its most general algebraic setting, where a matrix equation $Ax = b$ is being solved for any (possibly rectangular) matrix A .

For any approximate solution \tilde{x} , denote by $e = x - \tilde{x}$ the error vector and by $r = Ae = b - A\tilde{x}$ the vector of residuals. The common feature of all iterative schemes is that at each step some corrections to \tilde{x} are calculated based on the magnitude of certain residuals. As a result, convergence must be slow if the individual residuals do not show the true magnitude of the errors, i.e., if r is in some sense small compared with e . The converse is also true: If convergence of a *suitable* relaxation scheme is slow, residuals must in some sense be small compared with e .

To see this more concretely, consider for example Kaczmarz relaxation applied to the above matrix equation and converging to a solution x . Denoting by a_i the i -th row of A , the Kaczmarz step corresponding to that row is to replace \tilde{x} by $\tilde{x} + (r_i/a_i a_i^T) a_i^T$, thereby forcing r_i to zero. A full Kaczmarz sweep is the employment of such a step for

each row of A , in the natural ordering. (We take this scheme as our example because it applies to the most general matrix: It converges whenever a solution exists [T2].) Let

$$E = e^T e = \sum_i e_i^2 \quad \text{and} \quad R = \sum_i r_i^2 / a_i a_i^T$$

be square norms for errors and residuals, respectively, evidently scaled so that they are comparable. One can then prove the following result [B26, Theorem 3.4]:

THEOREM 1.1 A Kaczmarz sweep reduces E at least by $\max(\gamma_0 R_0, \gamma_1 R_1)$ where R_0 and R_1 are the values of R before and after the sweep, respectively, and

$$\gamma_0 = ((1+\gamma_-)(1+\gamma_+))^{-1}, \quad \gamma_1 = (\gamma_- \gamma_+)^{-1}$$

$$\gamma_- = \max_i \sum_{j < i} |a_i a_j^T| / a_i a_i^T, \quad \gamma_+ = \max_i \sum_{j > i} |a_i a_j^T| / a_i a_i^T.$$

The theorem in essence says that slow convergence can occur only when R is small compared with E . (Observe that in case A arises from the discretization of differential equations, γ_i are completely local quantities, independent of the size of A . The above values of γ_0 and γ_1 are close to the best possible ones.) Similar theorems (with suitably modified E , R and γ_i) hold for all familiar relaxation schemes, such as Gauss-Seidel, Jacobi, SOR, and block schemes (line relaxation, etc. See [B26]).

Since E and R were scaled to be comparable, it is generally only for special types of error components that the slow convergence condition

$$R \ll E \tag{1.1}$$

can be satisfied. The deeper (1.1) is satisfied, the more special must be the error, and hence the fewer is the number of parameters needed to approximate it. Thus, broadly speaking, relaxation efficiently reduces the information content of the error and quickly makes it approximable by far fewer variables. This can be shown to be true even for nonlinear systems.

When the matrix equation $Ax = b$ is a discretization of a differential system $Lu = f$ on some grid, condition (1.1), rewritten in the form $\|Ae\| \ll \|A\| \|e\|$, can be interpreted as saying that the error e approximates a continuous function v satisfying $\|Lv\| \ll \|L\| \|v\|$, in

some corresponding norms. In case L is a uniformly elliptic operator, this implies that e is a smooth function. (In case L is not elliptic, a certain smoothness along characteristics is at least implied.) Hence, relaxation efficiently reduces non-smooth error components, thus making the error approximable on a coarser grid (where solution is much cheaper). A precise measure for this efficiency is discussed next.

1.2 Error smoothing and its analysis: Example

For clarity, consider a simple example. Suppose the partial differential equation

$$Lu(x,y) \equiv a \frac{\partial^2 u(x,y)}{\partial x^2} + c \frac{\partial^2 u(x,y)}{\partial y^2} = f(x,y), \quad (a, c > 0) \tag{1.2}$$

is to be solved with some suitable boundary conditions. Denoting by u^h and f^h approximations to u and f , respectively, on a grid with meshsize h , the usual second-order discretization of (1.2) is

$$L^h u_{\alpha,\beta}^h \equiv a \frac{u_{\alpha+1,\beta}^h - 2u_{\alpha,\beta}^h + u_{\alpha-1,\beta}^h}{h^2} + c \frac{u_{\alpha,\beta+1}^h - 2u_{\alpha,\beta}^h + u_{\alpha,\beta-1}^h}{h^2} = f_{\alpha,\beta}^h, \tag{1.3}$$

where

$$u_{\alpha,\beta}^h = u^h(\alpha h, \beta h), \quad f_{\alpha,\beta}^h = f^h(\alpha h, \beta h); \quad \alpha, \beta \text{ integers.}$$

(In the multigrid context it is important to define the difference equations in this divided form, without, for example, multiplying throughout by h^2 , in order to get the proper relative scale at different grids.) Given an approximation \tilde{u} to u^h , a simple example of a relaxation scheme to improve it is the following.

Gauss-Seidel Relaxation. The points (α, β) of the grid are scanned one by one in some prescribed order; e.g., lexicographic order. At each point the value $\tilde{u}_{\alpha,\beta}$ is replaced by a new value, $\bar{u}_{\alpha,\beta}$, such that Eq. (1.3) at that point is satisfied. That is, $\bar{u}_{\alpha,\beta}$ satisfies

$$a \frac{\bar{u}_{\alpha+1,\beta} - 2\bar{u}_{\alpha,\beta} + \bar{u}_{\alpha-1,\beta}}{h^2} + c \frac{\tilde{u}_{\alpha,\beta+1} - 2\bar{u}_{\alpha,\beta} + \tilde{u}_{\alpha,\beta-1}}{h^2} = f_{\alpha,\beta}^h, \tag{1.4}$$

where the new values $\bar{u}_{\alpha-1,\beta}$, $\bar{u}_{\alpha,\beta-1}$ are used since, in the lexicographic order, by the time (α, β) is scanned new values have already replaced old ones at $(\alpha-1, \beta)$ and $(\alpha, \beta-1)$.

A complete pass, scanning in this manner all the gridpoints, is called a (Gauss-Seidel lexicographic) relaxation sweep. The new approximation \bar{u} does not satisfy (1.3), and further relaxation sweeps may

be required to improve it. An important quantity therefore is the convergence factor, μ say, which may be defined by

$$\mu = \|\bar{v}\| / \|v\|, \quad \text{where } v = u^h - \tilde{u}, \quad \bar{v} = u^h - \bar{u}, \quad (1.5)$$

$\|\cdot\|$ being any suitable discrete norm. For the Gauss-Seidel scheme, with the possible exception of its first few sweeps, $\mu = 1 - O(h^2)$. This means that $O(h^{-2})$ relaxation sweeps are needed to reduce the error order of magnitude.

In multigrid methods, however, the role of relaxation is not to reduce the error, but to smooth it out so that it becomes well approximable on a coarser grid. This relaxation can do very effectively. Indeed, subtracting (1.3) from (1.4), the relation

$$a(v_{\alpha+1,\beta} - 2\bar{v}_{\alpha,\beta} + \bar{v}_{\alpha-1,\beta}) + c(v_{\alpha,\beta+1} - 2\bar{v}_{\alpha,\beta} + \bar{v}_{\alpha,\beta-1}) = 0 \quad (1.6)$$

shows that $\bar{v}_{\alpha,\beta}$ is a weighted average of neighboring values (of both v and \bar{v}), so that, if the old error v is not smooth, the new error \bar{v} must be much smoother.

To analyze the smoothing effect of a relaxation sweep quantitatively, we take advantage of its local nature (points several meshsizes apart affecting each other exponentially little). It allows us, for the purpose of studying the smoothing well in the interior, to regard the grid as embedded in a rectangular domain. We can then expand both v and \bar{v} in Fourier series

$$v_{\alpha,\beta} = \sum A_{\underline{\theta}} e^{i(\theta_1\alpha + \theta_2\beta)}, \quad \bar{v}_{\alpha,\beta} = \sum \bar{A}_{\underline{\theta}} e^{i(\theta_1\alpha + \theta_2\beta)}, \quad (1.7)$$

where $\underline{\theta} = (\theta_1, \theta_2)$ and the summations are over a subset of the square $|\underline{\theta}| = \max(|\theta_1|, |\theta_2|) \leq \pi$. Substituting (1.7) into (1.6) yields

$$(ae^{i\theta_1} + ce^{i\theta_2})\bar{A}_{\underline{\theta}} + (ae^{-i\theta_1} + ce^{-i\theta_2} - 2a - 2c)\bar{A}_{\underline{\theta}} = 0. \quad (1.8)$$

Hence, the amplification factor of the $\underline{\theta}$ component due to one relaxation sweep is

$$\mu(\underline{\theta}) = \left| \frac{\bar{A}_{\underline{\theta}}}{A_{\underline{\theta}}} \right| = \left| \frac{ae^{i\theta_1} + ce^{i\theta_2}}{2a + 2c - ae^{-i\theta_1} - ce^{-i\theta_2}} \right|. \quad (1.9)$$

Observe that $\mu(\underline{\theta}) \rightarrow 1$ as $\underline{\theta} \rightarrow (0,0)$. In domains of diameter $O(1)$ the lowest non-trivial Fourier components have $|\underline{\theta}| = O(h)$, for which

$\mu(\underline{\theta}) = 1 - O(h^2)$, showing why convergence factors are that bad. Here, however, we are only interested in the smoothing effect, i.e., in the amplification factors of those components not approximable on a coarser grid. These are the components for which $(h/H)\pi \leq |\underline{\theta}| \leq \pi$, where H is the meshsize of the next coarser grid. We usually assume $H/h = 2$, because it is the most convenient, and as effective as any other mesh-size ratio (cf. Sec. 4.2). The smoothing factor is thus defined to be

$$\bar{\mu} = \max_{\pi/2 \leq |\underline{\theta}| \leq \pi} \mu(\underline{\theta}). \quad (1.10)$$

It essentially gives the relaxation convergence factor for those components which need to converge through relaxation; others will converge through their approximation on the coarser grid.

Consider first the case $a = c$ (Poisson equation). A simple calculation shows that $\bar{\mu} = \mu(\pi/2, \arccos 4/5) = .5$. This is a very satisfactory rate; it implies that three relaxation sweeps reduce the high-frequency error-components by almost an order of magnitude. Similar rates are obtained for general a and c , provided a/c is of moderate size.

The rate of smoothing is less remarkable in the degenerate case $a \ll c$ (or $c \ll a$). For instance

$$\mu\left(\frac{\pi}{2}, 0\right) = \left[\frac{a^2 + c^2}{a^2 + (c+2a)^2} \right]^{1/2}$$

which approaches 1 as $a \rightarrow 0$. Thus, for problems with such a degeneracy, Gauss-Seidel relaxation is not a suitable smoothing scheme. But other schemes exist. For example,

Line Relaxation. Instead of treating each grid point (α, β) separately, one takes simultaneously a vertical line of points at a time, i.e., the set of all points (α, β) with the same α . All the values $\tilde{u}_{\alpha,\beta}$ on such a line are simultaneously replaced by new values $\bar{u}_{\alpha,\beta}$ which simultaneously satisfy all the Eqs. (1.3) on that line. (This is easy and inexpensive to do, since the system of equations to be solved for each line is a tridiagonal, diagonally dominant system.) As a result, we get the same relation as (1.4) above, except that $\tilde{u}_{\alpha,\beta+1}$ is replaced by $\bar{u}_{\alpha,\beta+1}$. Hence, instead of (1.9) we now obtain

$$\mu(\underline{\theta}) = \left| \frac{a}{2(a+c-c \cos \theta_2) - ae^{-i\theta_1}} \right|, \quad (1.11)$$

from which one can derive the smoothing factor

$$\bar{\nu} = \max\left\{5^{-4}, \frac{a}{a+2c}\right\}, \quad (1.12)$$

which is very satisfactory, even in the degenerate case $a \ll c$.

This situation is very general. Namely, for any stable discretization of a well-posed differential boundary-value problem there exists a relaxation scheme which very efficiently reduces non-smooth error components (see Secs. 3 and 5.3). Moreover, the smoothing factor (1.10) for any candidate relaxation scheme is usually easy to calculate (e.g., by a standard Fortran routine called SMORATE [M3]), even for nonlinear equations or equations with non-constant coefficients, by local linearization and coefficients freeze (see Sec. 3.1). This gives us a general tool for optimizing the relaxation scheme and predict its efficiency. It is the first example of the *local mode analysis*, extensively used in multigrid analysis (see Secs. 3.1, 4.1, 7.4 and 7.5).

1.3 Coarse grid correction

We have seen that relaxation sweeps very quickly reduce all high-frequency components of the error. Its smoother part should then be reduced by being approximated on a coarser grid, a grid with meshsize $H = 2h$, say. Generally, for any linear fine-grid equation $L^h u^h = f^h$ (for the nonlinear case see Sec. 8.1), and any approximate solution \tilde{u}^h , the error $v^h = u^h - \tilde{u}^h$ satisfies

$$L^h v^h = r^h, \quad \text{where } r^h = f^h - L^h \tilde{u}^h. \quad (1.13)$$

It can therefore be approximated by the coarse-grid function v^H which satisfies

$$L^H v^H = I_h^H r^h, \quad (1.14)$$

where L^H is some coarse-grid approximation to L^h (e.g., a finite-difference approximation on grid H to the same differential operator approximated by L^h), and I_h^H is a fine-to-coarse transfer operator, called *residual weighting* or *restriction*. That is, $I_h^H r^h$ is a coarse-grid function whose value at each point is a certain weighted average of values of r^h at neighboring fine-grid points. (See much more on this process of "coarsening" in Secs. 4 and 11, and on the treatment of boundary conditions and global conditions in Secs. 5.4, 5.5 and 5.6.)

Having obtained an approximate solution \tilde{v}^H to Eq. (1.14) (in a way to be discussed below), we use it as a correction to the fine-grid

solution. Namely, we replace

$$\tilde{u}^h + \tilde{u}^h + I_H^{h \rightarrow H} \tilde{v}^H \quad (1.15)$$

where $I_H^{h \rightarrow H}$ is a coarse-to-fine *interpolation* (also called *prolongation*). That is, at each fine-grid point the value of $I_H^{h \rightarrow H} \tilde{v}^H$ (designed to approximate the error v^h) is interpolated from values of \tilde{v}^H at neighboring coarse-grid points. Linear interpolation can be used in most cases. (More on interpolation orders, see in Sec. 4.3.) The whole process of calculating $I_h^{H \rightarrow h} \tilde{v}^H$, solving (1.14) and interpolating the correction (1.15) is called a *coarse-grid correction*.

1.4 Multigrid cycle

To efficiently get an approximate solution to the coarse-grid equation (1.14), we employ the above solution process recursively; i.e., (1.14) is itself solved by relaxation sweeps combined with a still-coarser-grid corrections. We thus have a sequence of grids with mesh-sizes $h_1 > h_2 > \dots > h_M$, where usually $h_k = 2h_{k-1}$. The grid- h_k equation is generally written as

$$L^k u^k = f^k, \quad (1.16)$$

where all the L^k approximate each other (e.g., they are all finite-difference approximations to the same differential operator), and unless k is the finest level ($k = M$), equation (1.16) is of the form (1.14); i.e., u^{k-1} is always designed to be the coarse correction to \tilde{u}^k (the current approximation on the next finer grid), and hence

$$\tilde{f}^{k-1} = I_k^{k-1} (f^k - L^k \tilde{u}^k). \quad (1.17)$$

(Superscripts and subscripts k are now used instead of h_k in the notation of Sec. 1.3. Also, u^{k-1} is used instead of v^{2h} for the purpose of uniform expressions at all levels.)

The exact algorithm for improving a given approximate solution \tilde{u}^k to (1.16) is usually the *multigrid cycle* (MGC)

$$\tilde{u}_k + \text{MGC}(k, \tilde{u}^k, f^k), \quad (1.18)$$

defined recursively as follows:

If $k = 1$, solve (1.16) by Gaussian elimination or by several relaxation sweeps (either is usually cheap, since the grid is extremely

coarse. For additional remarks concerning the coarsest-grid solution, see Sec. 6.3). Otherwise do the following 4 steps:

- A. Perform ν_1 relaxation sweeps on (1.16), resulting in a new approximation \tilde{u}^k .
- B. Starting with $\tilde{u}^{k-1} = 0$, make γ successive cycles of the type

$$\tilde{u}_{k-1} + \text{MGC}(k-1, \tilde{u}^{k-1}, I_k^{k-1}(f^k - L^k \tilde{u}^k)) .$$

C. Calculate

$$\bar{u}^k = \tilde{u}^k + I_{k-1}^k \tilde{u}^{k-1} . \tag{1.19}$$

- D. Finally, perform ν_2 additional relaxation sweeps on (1.16), starting with \bar{u}^k and yielding the final \tilde{u}^k of (1.18).

The sweep counts ν_1 and ν_2 are usually either 0, 1 or 2, with $\nu = \nu_1 + \nu_2$ usually being 2 or 3. (See more about this in Secs. 4.1 and 6.1.) The cycle count γ is usually either 1 or 2. The cycle with $\gamma = 1$ is called a *V cycle*, or $V(\nu_1, \nu_2)$, in view of the shape of its flowchart (see Fig. 1.1). For a similar reason the cycle with $\gamma = 2$ is called a *W cycle*, or $W(\nu_1, \nu_2)$. (See more about different cycles and alternative switching criteria in Sec. 6.2.)

The ν sweeps performed in each *V cycle* on any grid h_k are expected to reduce error components with wave-length between $2h_k$ and $4h_k$ at least by the factor $\bar{\mu}^\nu$, where $\bar{\mu}$ is the smoothing factor (1.10). Since all grids are so traversed, the cycle should reduce all error components at least by the factor $\bar{\mu}^\nu$. Experience and more advanced theory show that for regular elliptic problems this is indeed the case, provided the boundary conditions are properly relaxed and correct inter-grid transfers are used. Thus, $\bar{\mu}$ can serve as an excellent predictor of the multigrid performance one should be able to obtain.

1.5 Model program and output

To give a more concrete idea about multigrid algorithms and their performance, and to introduce the reader to standard multigrid programming techniques, a simple Fortran program is presented here in full together with its output. The program solves a Dirichlet problem for the Poisson equation on a rectangle by employing NCYC times the $V(\nu_1, \nu_2)$ cycle shown in Fig. 1.1. Playing with the program, modifying equations etc., can help establish full understanding of various conceptual and technical aspects.

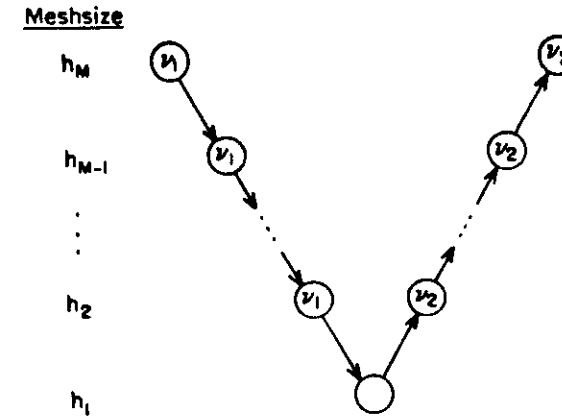


Figure 1.1: Multigrid cycle $V(\nu_1, \nu_2)$

- ν_1 stands for ν_1 relaxation sweeps on the meshsize shown to the left. (On the coarsest grid $\nu_1 + \nu_2$ sweeps are usually performed or the equations are solved directly.)
- \blacktriangleright is the fine-to-coarse ($k+1$ to k) transfer. f^k is defined by (1.17) and u^k is trivially initialized ($u^k \leftarrow 0$).
- \blacktriangleleft is the coarse-to-fine ($k-1$ to k) interpolation of correction (1.19).

The data for the problem are shown in the first two lines of the program, and its two functions: The differential equation is $\Delta u = F(x,y)$, with $u = G(x,y)$ given on the boundary. The function $G(x,y)$ also serves as the first approximation for u^M throughout the finest grid. The coarsest grid has $NX1 \times NY1$ intervals of length $H1$ each. Subsequent grids are defined as straight refinements, with meshsizes $h_k = 2^{1-k} H1$. There are M grids; h_M is thus the meshsize on which the problem is solved, coarser grids being just used for obtaining the fast convergence.

The same 5-point approximation to the Laplace operator ((1.3) with $a = c = 1$) is used on all grids. The residual transfer I_k^{k-1} is the trivial one (injection), the I_{k-1}^k interpolation is linear. For each grid we store both \tilde{u}^k and $h_k^2 f^k$ ($k = 1, \dots, M$). For handling these arrays uniformly, the arrays for \tilde{u}^k and $h_k^2 f^k$ are called U^k and U^{k+M} , respectively, ($k = 1, \dots, M$).

Note the key role of the GRDFN and KEY routines. The first is


```

SUBROUTINE I177DI(KC,FF)
COMMON C (I1C(200),I27C(200),I37F(200))
CALL FEY(KC,I17C,I17C,JJC,FC)
CALL FEY(KF,I17F,I17F,JJF,FF)
DC 1 IC=2,I1C
IF=2*I1C-1
JF=1
I1C=I17C(I17F)
I1F=I17F(I17F-1)
I2C=I17C(I1C)
I2F=I17F(I1C-1)
DC 1 JC=2,JJC
JF=JF+2
A=.5*(C(I1C+JC)+C(I2C+JC-1))
AM=.5*(C(I1F+JC)+C(I2F+JC-1))
C(I1C+JF) = C(I1C+JF)+C(I2C+JC)
C(I1F+JF) = C(I1F+JF)+.5*(C(I1C+JC)+C(I2F+JC))
C(I1C+JF-1)=C(I1C+JF-1)+A
1 C(I1F+JF-1) = C(I1F+JF-1)+.5*(A+AM)
RETURN
END

```

Linear interpolation
and correction
 $\tilde{u}^k + \tilde{u}^k + I_{k-1}^k \tilde{u}^{k-1}$
where $u^k = U^{KF}$
 $u^{k-1} = U^{KC}$

```

SUBROUTINE FESCAL(KF,FFF,KFC)
COMMON C (I1C(200),I1F(200),I1F(200),I1C(200))
CALL FEY(KF,I1F,I1F,JJF,FF)
CALL FEY(KFF,I1F,I1F,JJF,FFF)
CALL FEY(KFC,I1C,I1C,JJC,FC)
I1C1=I1C-1
JJC1=JJC-1
DC 1 IC=2,I1C1
I1F=I1F(I1C)
IF=2*I1C-1
JF=1
I1F=I1F(I1F)
I1C=I1F(I1F)
I1F=I1F(I1F-1)
I1F=I1F(I1F+1)
DC 1 JC=2,JJC1
JF=JF+2
S=C(IFC+JF+1)+C(IFC+JF-1)+C(I1F+JF)+C(I1F+JF)
1 C(I1F+JC)=4.*(C(I1F+JF)-S)+4.*C(IFC+JF)
RETURN
END

```

Residual injection
 $f^{k-1} + I_{k-1}^{k-1} (f^k - \Delta_k \tilde{u}^k)$
where $\tilde{u}^k = U^{KF}$
 $f^k = h_k^{-2} U^{KFF}$
 $f^{k-1} = h_{k-1}^{-2} U^{KFC}$

The output of this program is shown on the next page. It exhibits an asymptotic convergence factor of about .08 per cycle, which is just slightly better than the predicted $\bar{u} = .5^3 = .12$. (The two-level local mode analysis, described in Sec. 4.1, predicts the observed factor exactly.) This means a convergence factor of $.08^4 = .53$ per relaxation work unit. The same algorithm achieves the very same efficiency on general non-rectangular domains -- with a more complicated program, of course. A collection of multigrid programs, with varying degrees of simplicity vs. generality, is available [M3].

CYCLEV Output

LEVEL	RESIDUAL NCRM	WCFK	cycle
LEVEL 6	RESIDUAL NCRM= 1.624E 01	WCFK= 1.000	cycle 1
LEVEL 5	RESIDUAL NCRM= 1.561E 01	WCFK= 1.250	
LEVEL 4	RESIDUAL NCRM= 1.422E 01	WCFK= 1.212	
LEVEL 3	RESIDUAL NCRM= 1.141E 01	WCFK= 1.328	
LEVEL 2	RESIDUAL NCRM= 7.217E 00	WCFK= 1.322	
LEVEL 1	RESIDUAL NCRM= 3.027E 00	WCFK= 1.333	
LEVEL 1	RESIDUAL NCRM= 1.042E-01	WCFK= 1.324	
LEVEL 1	RESIDUAL NCRM= 6.510E-03	WCFK= 1.325	
LEVEL 2	RESIDUAL NCRM= 2.322E 00	WCFK= 1.325	
LEVEL 2	RESIDUAL NCRM= 3.861E-01	WCFK= 1.342	
LEVEL 3	RESIDUAL NCRM= 4.061E 00	WCFK= 1.258	
LEVEL 3	RESIDUAL NCRM= 6.954E-01	WCFK= 1.374	
LEVEL 4	RESIDUAL NCRM= 5.369E 00	WCFK= 1.427	
LEVEL 4	RESIDUAL NCRM= 7.610E-01	WCFK= 1.459	
LEVEL 5	RESIDUAL NCRM= 6.062E 00	WCFK= 1.749	
LEVEL 5	RESIDUAL NCRM= 7.502E-01	WCFK= 1.999	
LEVEL 6	RESIDUAL NCRM= 6.446E 00	WCFK= 2.599	
LEVEL 6	RESIDUAL NCRM= 7.426E-01	WCFK= 3.599	
LEVEL 6	RESIDUAL NCRM= 6.066E-01	WCFK= 4.999	cycle 2
LEVEL 5	RESIDUAL NCRM= 6.174E-01	WCFK= 5.249	
LEVEL 4	RESIDUAL NCRM= 5.180E-01	WCFK= 5.212	
LEVEL 3	RESIDUAL NCRM= 3.551E-01	WCFK= 5.327	
LEVEL 2	RESIDUAL NCRM= 1.759E-01	WCFK= 5.231	
LEVEL 1	RESIDUAL NCRM= 3.797E-02	WCFK= 5.232	
LEVEL 1	RESIDUAL NCRM= 9.513E-03	WCFK= 5.232	
LEVEL 1	RESIDUAL NCRM= 5.946E-04	WCFK= 5.334	
LEVEL 2	RESIDUAL NCRM= 6.029E-02	WCFK= 5.238	
LEVEL 2	RESIDUAL NCRM= 1.875E-02	WCFK= 5.342	
LEVEL 3	RESIDUAL NCRM= 1.246E-01	WCFK= 5.357	
LEVEL 3	RESIDUAL NCRM= 3.120E-02	WCFK= 5.273	
LEVEL 4	RESIDUAL NCRM= 2.052E-01	WCFK= 5.426	
LEVEL 4	RESIDUAL NCRM= 4.014E-02	WCFK= 5.499	
LEVEL 5	RESIDUAL NCRM= 2.528E-01	WCFK= 5.748	
LEVEL 5	RESIDUAL NCRM= 4.117E-02	WCFK= 5.998	
LEVEL 6	RESIDUAL NCRM= 2.794E-01	WCFK= 6.998	
LEVEL 6	RESIDUAL NCRM= 4.071E-02	WCFK= 7.998	
LEVEL 6	RESIDUAL NCRM= 3.480E-02	WCFK= 8.998	cycle 3
LEVEL 5	RESIDUAL NCRM= 3.255E-02	WCFK= 9.248	
LEVEL 4	RESIDUAL NCRM= 2.729E-02	WCFK= 9.211	
LEVEL 3	RESIDUAL NCRM= 1.850E-02	WCFK= 9.226	
LEVEL 2	RESIDUAL NCRM= 9.278E-03	WCFK= 9.320	
LEVEL 1	RESIDUAL NCRM= 3.332E-03	WCFK= 9.231	
LEVEL 1	RESIDUAL NCRM= 4.292E-05	WCFK= 9.322	
LEVEL 1	RESIDUAL NCRM= 2.663E-06	WCFK= 9.233	
LEVEL 2	RESIDUAL NCRM= 2.742E-03	WCFK= 9.237	
LEVEL 2	RESIDUAL NCRM= 5.889E-04	WCFK= 9.241	
LEVEL 3	RESIDUAL NCRM= 6.558E-03	WCFK= 9.256	
LEVEL 3	RESIDUAL NCRM= 1.528E-03	WCFK= 9.272	
LEVEL 4	RESIDUAL NCRM= 1.040E-02	WCFK= 9.425	
LEVEL 4	RESIDUAL NCRM= 2.164E-03	WCFK= 9.497	
LEVEL 5	RESIDUAL NCRM= 1.263E-02	WCFK= 9.747	
LEVEL 5	RESIDUAL NCRM= 2.331E-03	WCFK= 9.997	
LEVEL 6	RESIDUAL NCRM= 1.420E-02	WCFK= 10.997	
LEVEL 6	RESIDUAL NCRM= 3.066E-03	WCFK= 11.997	


1.6 Full multigrid (FMG) algorithm

The multigrid cycles described above can be applied to any first approximation given on the finest grid. In a full multigrid (FMG) algorithm, the first approximation is obtained by interpolation from a solution on the next coarser grid, which has previously been calculated by a similar FMG algorithm. With such a first approximation, and provided that the interpolation correctly corresponds to the error norm used, one multigrid cycle should suffice to solve the fine-grid equations to the level of truncation errors (see Sec. 7). A typical FMG algorithm, with one V cycle per refinement, is shown in Fig. 1.2.

The FMG algorithms are less sensitive than the multigrid cycles. That is, in many irregular cases the asymptotic convergence factor of the cycles is not good, but the FMG algorithm with one cycle per refinement still guarantees solution to the level of truncation errors. In fact, even this guarantee is not necessary: From differences between the final solutions at different meshsizes (e.g., differences between the solutions at the doubly-circled stages in Fig. 1.2), one can directly calculate the rate of convergence to the *differential solution*, which is all that really matters.

1.7 General warnings. Boundary conditions. Nonlinearity

Attempts to extend existing multigrid software often fail. For example, almost everyone who tries to extend the program of Sec. 1.5 from Dirichlet to Neumann boundary conditions, first obtains a much slower solver. Some would then hastily announce that the method is inherently slower for general non-Dirichlet boundary conditions. Others, realizing the conceptual generality of the basic multigrid approach, would ask themselves what caused the slowness and how to correct it. They will eventually discover that various additional processes should be done in case of non-Dirichlet conditions, such as relaxation sweeps over these conditions and transfers of their residuals to coarser grids, concurrently with the corresponding interior processes. It will take some more effort to realize that the best relaxation of boundary conditions in multigrid solvers is often quite markedly different from the scheme one would use in relaxation solvers. Eventually, when everything is done correctly, the algorithm will regain the efficiency it showed in the Dirichlet case. Indeed, with proper treatment, the multigrid efficiency should never depend on boundary conditions, only on the interior equations (and in fact only on the interior smoothing rates, hence only on the *factors* of the subprincipal determinant of the interior operator



global conditions (Kutta condition, total mass, etc.); unbounded domains; complex geometries; three dimensions; not to mention physical instabilities and turbulence.

Each and every one of these features requires a thorough understanding as to what it implies in terms of each multigrid process. Mistreating just one of them may cause the solution time to increase very significantly, sometimes by orders of magnitude. Since convergence will often still be obtained, due to the corrective nature of the fine-grid relaxation, one may be misled to believe that nothing is wrong. Even if he suspects errors, he is unlikely to find them all, because they confusingly interact with each other, an impossible network of conceptual mistakes with programming bugs.

As stated in the Introduction, we believe that every problem should be solvable in just few work units. But only a systematic development is likely to produce this top performance. Each feature should first be separately studied in the framework of as simple a problem as possible. The solver should then be constructed step by step, adding one feature at a time, making sure that full efficiency is always maintained. This Guide may help in the process.

It may be useful to add here a preliminary remark about nonlinearity. The multigrid processes are *not* inherently linear. The basic idea described in Sec. 1.1, namely, the relations between slow convergence, smallness of residuals and error smoothness, has nothing to do with linearity. Multigrid can thus be applied *directly to nonlinear problems*, as efficiently as to the corresponding linearized problems (see Sec. 8). Hence repeated linearizations are neither required nor advised; they are especially wasteful in case the nonlinear problem is autonomous (as is usual in fluid dynamics), because the linearized problems are not autonomous. (See more detailed arguments in Sec. 8.3.) Moreover, the multigrid version developed for nonlinear cases, called FAS, is useful in many other ways. In particular it gives a convenient way to create non-uniform adaptable discretization patterns, based on the interaction between the levels and therefore very flexible, allowing fast local refinements and local coordinate transformations, with equations being still solved in the usual multigrid efficiency (see Sec. 9). Using FAS, one can integrate into a single FMG process various other processes, such as continuation, design and optimization, solution to eigenvalue problems and to inverse problems, grid adaptation, etc.

PART I. STAGES IN DEVELOPING FAST SOLVERS

The intention of this part is to organize existing multigrid approaches in an order which corresponds to actual stages in developing fast multigrid solvers. Each section (Secs. 2 through 7) represents a separate stage. To get an overview of these stages, the reader may first go through the opening remark of all the sections, skipping the subsections. The actual sequence of development may correspond to the actual order of the sections; but Secs. 5, 6, and 7.4 represent three independent stages, which can be taken in any order following Sec. 4. In fact, an increasing current tendency is to replace the usual two-level mode analysis (Sec. 4) by the two-level FMG mode analysis (Sec. 7.4). Generally, one can skip a stage, risking a lesser control of potential mistakes. Even when one does, the information and advice contained in the corresponding subsections are still important.

This part emphasizes the *linear* solver: relaxation of nonlinear equations is described, but the Full Approximation Scheme (FAS) used in inter-grid transfers of nonlinear solvers is deferred to the next part (Sec. 8).

2. STABLE DISCRETIZATION

The formulation of good discretization schemes is of course the first step in any numerical solution of continuous equations. For multigrid solutions some additional considerations enter. First, discrete equations should be written for general meshsizes h , including large ones (to be used on coarse grids). Also, the multigrid processes offer several simplifications of the discretization procedures.

For example, only uniform-grid discretization is often needed. Non-uniform discretization can be effected through multigrid interactions between uniform grids (see Sec. 9). In fact, if the basic grids are non-uniform, various structural and algorithmic complications are introduced, in multigrid as well as, and even more than, in unigrid processes. The uniform discretization can be made either with finite element or finite difference formulations. For finite elements, it is preferable to use either piecewise uniform partitions as in [B2] and [B7, Sec. 7.3], or uniform partitions modified at special parts (e.g., at boundaries) as in [B11, Fig. 2], and produce local refinements by the multigrid process (Sec. 9). When general non-uniform partitions must be used, algebraic

multigrid solvers (see Sec. 13.1) are recommended: Their set-up processing (coarse-levels assembly) is much more costly, but it is still small compared with the very expensive processing required for assembling the (finest-level) equations.

Another simplification is that one has to write low-order (first or second order) discretization only. Higher-order schemes can later be superposed for little extra programming effort and computer time (Sec. 10). The low order makes it easier to write stable equations, easier to devise and analyze relaxation schemes, and cheaper to operate them. Furthermore, designing stable discretization with efficient relaxation is only required in terms of the (sub)principal part of the operator (see Sec. 2.1).

Still another simplification is the easier implementation and more flexible control of *global constraints*, by effecting them only at the coarsest levels of the multigrid processing (with suitable inter-grid transfers. See Sec. 5.6). This makes it possible to free local differencing from complicated forms aimed at precise conservation of global quantities (such as total kinetic energy and square vorticity, as in [A3]), to add suitable controls to ill-posed problems, etc.

Low order finite elements on uniform grids yield kinds of difference equations. The description below will therefore be in terms of finite difference formulations only. It should be emphasized, however, that variational formulations, where appropriate, automatically yield good prescriptions for the main multigrid processes. See [B7, App. A.5], [N3] and [B11]. This is especially useful in some complicated situations, as in [A1]. We return to this issue in Secs. 4.5, 4.6 and 11.

To be solvable by a fast multigrid algorithm, the discretization of a boundary-value problem should be suitably stable. More precisely, the type of stability determines the kind of multigrid algorithm that can be efficient. The simplest differencing of regular elliptic equations yield equations which are stable in every respect, so the reader interested in only such problems can skip the rest of this section. But remember:

Unstable (or very inaccurate) discretization must lead to slow multigrid convergence (unless special techniques are adopted to the case or algebraic multigrid is used). It is indeed an important *advantage* of multigrid solvers that bad discretization cannot be passed unnoticed; it must show up as slow algebraic convergence.

2.1 Interior stability measures: h-ellipticity

Numerical stability is a *local* property, i.e., a property significant only for non-smooth components of the solution (components that change significantly over a meshsize), whereas the smooth-component stability depends, by consistency, on the *differential* system, not on its discretization. Indeed, in multigrid solvers, stability of the discrete operator is needed only in the *local* process of relaxation (cf. Sec. 10.2). Moreover, what counts is not really the stability of the static difference operator itself, but the overall efficiency with which the dynamic process of relaxation smoothes the differential error (cf. Sec. 12); numerical stability of the operator is just a necessary condition for achieving that smoothing.

Because of the local character of the required stability (corresponding to the local task of relaxation), it is easy to get a very good quantitative idea about it, for any given difference operator L^h , by local mode analysis, analogous to the Von-Neumann stability analysis for time-dependent problems. It turns out, however, that for steady-state problems, especially non-elliptic or singular perturbation ones, the distinction between stable and unstable discrete operators is not enough. More important is the *measure* of stability. When that measure, for a given meshsize, is low, the scheme is still formally stable, but its actual behavior can be intolerably bad (see example in [B17, Sec. 3]).

Briefly, the basic relevant measure of stability of an interior (not at boundaries) linear difference operator L^h with constant coefficients is its h-ellipticity measure $E^h(L^h)$, defined for example by

$$E^h(L^h) = \min_{\delta\pi \leq |\underline{\theta}| \leq \pi} |\tilde{L}^h(\underline{\theta})| / |L^h|, \quad (2.1)$$

where the complex function $\tilde{L}^h(\underline{\theta})$ is the "symbol" of L^h , i.e., $L^h e^{i\underline{\theta} \cdot \underline{x}/h} = \tilde{L}^h(\underline{\theta}) e^{i\underline{\theta} \cdot \underline{x}/h}$; $\underline{\theta} = (\theta_1, \dots, \theta_d)$, $\underline{\theta} \cdot \underline{x}/h = \theta_1 x_1/h_1 + \dots + \theta_d x_d/h_d$, $|\underline{\theta}| = \max(|\theta_1|, \dots, |\theta_d|)$; d is the dimension; h_j is the meshsize in direction x_j ; and $|L^h|$ is any measure of the size of L^h , e.g., $|L^h| = \max|\tilde{L}^h(\underline{\theta})|$. The constant $0 < \delta < 1$ is in fact arbitrary, but for convenient multigrid applications a natural choice is the meshsize ratio, hence usually $\delta = \frac{1}{2}$. The range $\delta\pi \leq |\underline{\theta}| \leq \pi$ is then the range of "high frequency" components on grid h , i.e., components $\exp(i\underline{\theta} \cdot \underline{x}/h)$ which on the next coarser grid, with meshsize h/δ , coincide (alias) with lower components.

In case of systems of equations, L^h and $\tilde{L}^h(\underline{\theta})$ are matrices and $|\tilde{L}^h(\underline{\theta})|$ should then be understood as a measure of the non-singularity

of $\tilde{L}^h(\theta)$ (e.g., its smallest eigenvalue, or its determinant). See more details and explanations in [B12, Sec. 3], [B17, Sec. 3.1]. In fact, good discretization schemes can generally be arrived at by requiring $\det L^h$ to be a good discretization for $\det L$, the determinant of the given differential operator: See the examples in Secs. 17.2, 18.2, 19.2 and 20.2.

In case the differential operator L , and hence also L^h , have variable coefficients, L^h is called h-elliptic if $E^h(L^h) = O(1)$ for each combination of coefficients appearing in the domain. If L is nonlinear, L^h is called h-elliptic if its linearizations, around any approximate solution encountered in the calculations, are all h-elliptic.

A major simplification in selecting the discretization scheme for complicated systems is the fact that, being interested in local properties only, we can confine our considerations to those terms which are locally important. In the discretized and linearized operator L^h , the locally important terms, called the h-principal terms are simply those with large coefficients (relative to other coefficients, if any). In case of a system, the h-principal terms are those contributing to the h-principal terms of $\det L^h$. Other terms are not important in relaxation; namely, they need not satisfy any stability conditions, they can actually be transferred to the right-hand side of the relaxed equations, and they need not even be updated each sweep (only each multigrid cycle; cf. Sec. 10.3).

The h-principal terms all come from discretizing subprincipal terms of L . These are defined as the principal terms (the terms contributing to the highest order derivatives in the determinant of the linearized operator) plus the principal terms of the reduced operator (the operator without singular perturbation terms). Thus, in discretizing any differential operator L , we can confine our attention to its subprincipal part. See the examples in Secs. 19.1 and 20.1.

Regular discretizations of elliptic systems should, and usually do have good (i.e., $O(1)$) h-ellipticity measures. (But see a counter-example in Sec. 17.2.) Singular perturbation or non-elliptic systems can also have such good measures, e.g., by using artificial viscosity or by upwind (upstream) differencing. (Note that a regular elliptic system with lower-order terms may be a singular-perturbation problem on a sufficiently coarse grid.) If, however, characteristic or subcharacteristic directions (i.e., characteristic directions of the reduced equations, in case of singular perturbation problems) coincide with grid directions, upwind differencing schemes are only semi h-elliptic. They have, that

is, bad h-ellipticity measure E^h , but they still have a good semi h-ellipticity measure in the characteristic direction. This measure is defined as follows.

Let $S \subset \{1, \dots, d\}$ be a subset of grid directions. The measure of semi h-ellipticity in directions S , or briefly S-h-ellipticity, of a difference operator L^h is

$$E_S^h(L^h) = \min_{\delta\pi \leq |\theta|_S \leq \pi} |\tilde{L}^h(\theta)| / |L^h|, \quad (2.2)$$

where $|\theta|_S = \max_{j \in S} |\theta_j|$. Full h-ellipticity is the special case $S = \{1, \dots, d\}$. If $S_2 \subset S_1$ then clearly $E_{S_1}^h \leq E_{S_2}^h$, hence S_1 -h-ellipticity entails S_2 -h-ellipticity.

In case (sub)characteristics are aligned with grid directions, full h-ellipticity is not needed for stability. The corresponding S-h-ellipticity is enough; it allows large local oscillations perpendicular to the characteristics, but those oscillations are also allowed by the differential equations.

Fully h-elliptic approximations can be constructed even for non-elliptic or semi-elliptic differential equations by using isotropic artificial viscosity. In various cases, however, semi h-elliptic approximations are preferable, since they entail much less cross-stream smearing. These are mainly cases of strong alignment, that is, cases where (sub)characteristic lines are non-locally (i.e., for a length of many meshsizes) aligned with a gridline, and where this non-local alignment occurs either for many gridlines, or even for one gridline if that line is adjacent to a boundary layer or a similar layer of sharp change in the solution. (For a method to obtain strong alignments and thus avoid smearing -- see Sec. 9.3.)

A convenient way of constructing h-elliptic and semi h-elliptic operators is by term-by-term R-elliptic or semi R-elliptic approximations [B10, Sec. 5.2], [B19, Sec. 3.6]. Another, more physical way is to regard the given boundary-value problem as a limit of an elliptic problem (usually this is physically so anyway), and enlarge the elliptic singular perturbation to serve as artificial-viscosity terms [B17], [B16]. When solving a steady-state problem of originally time-dependent equations, the artificial elliptic terms should conform to the original time-dependent problem, i.e., with those terms that problem should still be well posed. This requirement often determines the sign of the artificial terms. Such physical artificial viscosity terms ensure that computed solutions will exhibit (as $h \rightarrow 0$) only those discontinuities

allowed physically. (Hence, with this approach, explicit entropy conditions are not needed.) A proper amount of anisotropic artificial viscosity gives the correct upstream differencing whenever desired.

The optimal amount of artificial viscosity (either isotropic or anisotropic) is mainly determined not by stability considerations but by the smoothing properties of relaxation. Below a certain level of viscosity, more costly, distributive relaxation will have to be used. Increasing the artificial viscosity slightly beyond the minimum required for convergence of the simplest scheme makes the relaxation ordering-free (see Sec. 3.6 and [B17, Sec. 5.7], [B12, Sec. 4.2]), which is desirable, except perhaps near discontinuities. Considerably larger artificial viscosity makes the algebraic smoothing faster, but impedes the differential smoothing (cf. Sec. 12).

Interior difference equations which are not even semi h-elliptic should be used with care. Their solutions may show large numerical oscillations (giving nice solutions only in the average), and their fast multigrid solvers must have more complicated fine-to-coarse interactions (see for example Sec. 4.2.2). Some quasi-elliptic equations, i.e., cases where $\tilde{L}^h(\theta) / |L^h|$ does vanish for some $|\theta| = \pi$, but not for other $|\theta| \neq 0$, can be solved without much trouble. All that is needed is to average out the bad components (see for example Secs. 18.6, 19.5, 20.6).

2.2 Boundaries, discontinuities

We have so far discussed the stability conditions related to the interior difference equations, away from boundaries. To gain overall stability some additional conditions should be placed at the boundaries. These can be analyzed by mode analysis in case the boundaries are parallel to grid directions (cf. Sec. 7.5). But more general boundaries are difficult to analyze. Usually, however, h-elliptic approximations consistent with a well-posed problem and employing low-order approximations to boundary conditions are stable. The order can then be raised in a stable way by one of the methods of Sec. 10. At any rate, the boundary stability is not related to the stages of developing the main (interior) multigrid processes.

More critical than discretization near boundaries is the treatment of *discontinuities*, whether at boundaries (e.g., boundary layers) or in the interior (e.g., shocks). The basic rule, in multigrid as in unigrid processes, is to try not to straddle the discontinuity by any difference operator (in relaxation as well as in residual transfers. And the rule also applies to the interpolation operators). More precisely, the rule

is not to difference any quantity which is discontinuous in the interval of differencing. This can fully be achieved only in cases where the location of the discontinuity is known or traced (at such discontinuities the above rule overrides upstream differencing if they happen to conflict), or when the discontinuity is more or less parallel to grid directions (so that upstream differencing will automatically satisfy the rule). Captured discontinuities which are not in grid directions must perhaps be smeared; a multigrid way to get high accuracy then is by local refinements (see Sec. 9). Generally, multigrid procedures for discontinuities are now under active investigation (see [B17, Sec. 4] and a remark in Sec. 8.5 below).

3. INTERIOR RELAXATION AND SMOOTHING FACTORS

The crucial step in developing multigrid solvers is the design of interior relaxation schemes with high error-smoothing rates. Namely, the crucial question is how to reduce non-smooth error components for as little computational work as possible, neglecting interactions with boundaries. This is the crucial question, from the point of view of solution efficiency, since reducing *smooth* error components will require less computational work (being done on coarser grids), and since reducing non-smooth error components *near boundaries* will require much less work because it involves local work only near the boundary (which is a lower-dimensional manifold). Also, relaxation is the most problem-dependent part of the algorithm -- other parts are usually quite standard. (The relaxation of boundary conditions is discussed in Sec. 5.3.)

3.1 Local analysis of smoothing

To reduce non-smooth error components is basically a *local* task; it can be done in a certain neighborhood independently of other parts of the domain. This is why it can efficiently be performed through relaxation, which is basically a local process (the information propagates just few meshsizes per sweep). Hence also, the efficiency of this process can accurately be measured by local mode analysis.

That is, one can assume the problem to be in an unbounded domain, with constant (frozen) coefficients, in which case the algebraic error $u^h - \tilde{u}^h$ (where u^h is the exact solution to the discrete equations and \tilde{u}^h is the computed approximation) is a combination of Fourier components $\exp(i\theta \cdot \underline{x}/h)$. For each such Fourier component and any proposed relaxation scheme one can easily calculate the amplification factor

$\mu(\underline{\theta})$, defined as the factor by which the amplitude of that component is multiplied as a result of a relaxation sweep (see simple examples in Sec. 1.2 above and in [S4, Sec. 3.2]). The smoothing factor $\bar{\mu}$ of the relaxation scheme, defined by

$$\bar{\mu} = \max_{\pi/2 < |\underline{\theta}| < \pi} |\mu(\underline{\theta})| \quad (3.1)$$

can then easily be computed (usually by a standard computer routine, such as the SMORATE routine in [M3]). This is indeed the measure we need: $\bar{\mu}$ is the worst (largest) factor by which all high-frequency error components are reduced per sweep, where we define the frequency to be high if the component is not visible (aliases with a lower component) on the next coarser grid (grid 2h).

In case of a system of q grid equations in q unknown grid functions (i.e., q unknowns and q algebraic equations are defined per mesh cell), each Fourier amplitude is a q -vector, hence $\mu(\underline{\theta})$ is a $q \times q$ amplification matrix. $\bar{\mu}$ is still defined as in (3.1), except that $|\mu(\underline{\theta})|$ is replaced by $\rho(\mu(\underline{\theta}))$, where $\rho(\mu)$ is the spectral radius of μ .

For L^h with non-constant coefficients, $\bar{\mu}$ defined by (3.1) depends on the location. In case of nonlinear L^h , the analysis is made for the linearized operator, hence $\bar{\mu}$ depends also on the solution around which linearization is made. The quality of relaxation is then determined by the worst $\bar{\mu}$, i.e., the maximum of $\bar{\mu}$ over all possible coefficients of L^h for any solution which may evolve in the calculations. (One may disregard $\bar{\mu}$ on small regions: See Sec. 3.3.)

A major simplification in calculating $\bar{\mu}$ for complicated systems is to look at subprincipal terms only (see Secs. 2.1 and 3.4 and examples in Secs. 19.2 and 20.2).

Some relaxation schemes do not transform each Fourier component of the error to a multiple of itself. Instead, they couple several (ℓ , say) Fourier components at a time (even in the infinite domain). For example, if relaxation is performed in red-black (checker-board) ordering (cf. Sec. 3.6), the $\underline{\theta}$ component is coupled to the $\underline{\theta} + (\pi, \dots, \pi)$ component. Instead of the $q \times q$ amplification matrix $\mu(\underline{\theta})$ we then have the $(q\ell) \times (q\ell)$ matrix $\mu(\underline{\theta}^1, \dots, \underline{\theta}^\ell)$, describing the transformation of the ℓ q -vector amplitudes corresponding to the coupled components $\underline{\theta}^1, \dots, \underline{\theta}^\ell$. Definition (3.1) is extended to such cases by *defining*

$$\bar{\mu}_v = \max[\rho(C(\underline{\theta}^1, \dots, \underline{\theta}^\ell)\mu(\underline{\theta}^1, \dots, \underline{\theta}^\ell)^v)]^{1/v} \quad (3.2)$$

where the \max is taken over all coupled ℓ -tuples $(\underline{\theta}^1, \dots, \underline{\theta}^\ell)$, C is an $\ell \times \ell$ matrix of $q \times q$ blocks C_{ij} such that $C_{ij} = 0$ for $i \neq j$, $C_{ii} = I_q$ (the $q \times q$ identity matrix) if $|\underline{\theta}^i| > \pi/2$, and $C_{ii} = 0$ otherwise. v is the number of sweeps performed on the finest grid per multigrid cycle; only in the simple case ($\ell = 1$) $\bar{\mu}$ does not depend on v . Examples with $\ell > 1$ see in [B14, Sec. 3.3], [L2, Sec. 2.3.1].

The smoothing factor is the first and simplest quantitative predictor of the obtainable multigrid efficiency: $\bar{\mu}^v$ (or $\bar{\mu}_v^v$) is an approximation to the asymptotic convergence factor obtainable per multigrid cycle. Usually this prediction is more accurate than needed. There are still more accurate predictors (see Sec. 4.1). But the main importance of $\bar{\mu}$ is that it separates the design of the interior relaxation from all other algorithmic questions. Moreover, it sets an ideal figure against which the performance of the full algorithm can later be judged (see Secs. 4, 5).

The analysis of relaxation within multigrid is thus much easier than the analysis of relaxation as an independent iterative solver. The latter is not a local process, and its speed depends on smooth components, which are badly approximated by mode analysis, due to boundaries and variable coefficients. For multigrid purposes, however, wherever the equations (or their linearized version) do not change too much within few meshsizes, the smoothing factor can be used as a standard measure of performance. A general routine for calculating $\bar{\mu}$, called SMORATE, is available [M3].

3.2 Work, robustness and other considerations

In comparing several candidate relaxation schemes we should of course take into account not only their smoothing factors, but also the amount of work per sweep. The aim is generally to have the best high-frequency convergence rate per operation, i.e., the largest $w_0^{-1} \log(1/\bar{\mu})$, where w_0 is the number of operations per gridpoint per sweep. But other considerations should enter as well: The rate should be robust, that is, $\bar{\mu}$ should not depend too sensitively on problem parameters or on a precise choice of various relaxation parameters. Also, between two schemes with similar values of $w_0^{-1} \log(1/\bar{\mu})$ but with very different w_0 , the simpler scheme (where w_0 is smaller) should be preferred, because very small factors $\bar{\mu}$ cannot fully be obtained in practice (owing to the inability of the coarse-grid correction to obtain such small factors for the smooth components, and owing to interactions with

boundaries). Moreover, large values of w_0 leave us with less flexibility as to the amount of relaxation work to be performed per cycle. Very small $\bar{\mu}$ may in fact be below what we need in the Full Multigrid (FMG) algorithm (see Sec. 7).

An important consideration, sometimes overlooked, is that each relaxation sweep should of course be stable. The most familiar schemes are stable, but distributive schemes (Sec. 3.4) for example, can be unstable exactly in cases showing the best $\bar{\mu}$. A trivial example: satisfy each difference equation in its turn by changing its *latest* unknown (in the sweeping ordering) instead of its usual corresponding unknown. $\bar{\mu}$ will then vanish, but the process will be unstable. Stability analysis can in each case be performed as Von-Neumann analysis for time dependent problems, taking the main relaxation marching direction as the time-like direction.

Also, do not forget that relaxation has a certain effect on smooth (low frequency) components, too. Usually this effect is slow: $\mu(\theta)$ is close to 1 for small $|\theta|$. But sometimes schemes which show spectacularly small values of $\bar{\mu}$ also show either bad divergence ($|\mu(\theta)| \gg 1$) or fast convergence ($|\mu(\theta)| \ll 1$) for *low* frequencies. This for example may happen in relaxing hyperbolic (relative to some time-like direction) equations using upstream differencing and marching with the stream (the time-like) direction. Schemes with bad divergence should clearly be rejected (see an example in Sec. 20.3.4, the super-fast smoothing case). Those with fast convergence may also have some disadvantage (in case high-order corrections, as in Sec. 10.2, are desired; see for example [B17, Sec. 2.2]).

It is therefore advisable to add to the program of calculating $\bar{\mu}$ also a routine for checking the stability of the scheme examined, and to calculate, together with (3.1), also the value of $\max_{|\theta| < \pi} |\mu(\theta)|$. It is also useful to calculate weighted mean squares of $\mu(\theta)$ for high-frequency θ 's. Such quantities predict the error decrease in a given number of multigrid relaxation sweeps for a given initial error [B19, Sec. 4.5]. Some of these measures are supplied by SMORATE [M3].

The value of local mode analysis becomes dubious at places of strong discontinuities, e.g., where the coefficients of the differential equation change their order of magnitude discontinuously (or within few meshsizes). This usually happens along manifolds of lower dimensionality, hence more computational work per gridpoint can there be afforded, hence an accurate measure of efficiency is not so needed. But some basic rules, outlined below, must still be followed. One can also employ local R-E

analysis (cf. Sec. 1.1, and more details in [B26]), which yields good quantitative information even in strongly discontinuous cases.

3.3 Block relaxation rule. Semi smoothing

The most basic rule in devising relaxation schemes is that a *locally strongly coupled block of unknowns which is locally decoupled from (or weakly coupled with) the coarser-grid variables should be relaxed simultaneously*. The reason is that a point-by-point relaxation smoothes only along the strongest couplings whereas block relaxation also smoothes along second-strongest couplings (providing the strongest ones are included in the blocks. See for example the case $a \ll C$ in Sec. 1.2, and a more general analysis in [B26, Sec. 3.5 and Sec. 4.6]).

This rule is of course important whether or not the equations are continuous. In case of persistent S_1 -h-ellipticity (i.e., a difference operator with good S_1 -h-ellipticity throughout a substantial subdomain, but without uniformly good S_2 -h-ellipticity measure for any $S_2 \notin S_1$), the rule implies either the use of block relaxation in suitable directions (line relaxation, plane relaxation, etc.), or the use of suitable "semi coarsening" (see below), or both.

Generally, for any set S of grid directions (usually $S \subset \{1, 2, \dots, d\}$, but sometimes including a special bisecting direction could be advantageous), a relaxation scheme is called an S-block relaxation if it relaxes simultaneously all (or many contiguous) equations defined on the same S-subspace. (Two points (x_1, \dots, x_d) and (y_1, \dots, y_d) are in the same S-subspace if $x_j = y_j$ for all $j \notin S$.) For example, the line relaxation in Sec. 1.2 is a y-block (vertical line) relaxation.

Semi coarsening, or more specifically S-coarsening, means that $H_j = 2h_j$ for $j \in S$ and $H_j = h_j$ otherwise, where H_j and h_j are the meshsizes of the coarse grid and the fine grid, respectively, in the x_j direction ($j = 1, \dots, d$). In such a case we need to smooth the error only in directions S . The definition of the smoothing factor should accordingly be modified. We generalize (3.1) to any coarsening situation by defining

$$\bar{\mu} = \max\{\rho(\mu(\theta)) : |\theta| < \pi, \max_{1 \leq j \leq d} |\theta_j| H_j / h_j \geq \pi\}. \quad (3.3)$$

Similarly we generalize (3.2) by defining $C_{ii} = I_q$ if $\max_j |\theta_j^i| H_j / h_j \geq \pi$, and $C_{ii} = 0$ otherwise. The S-smoothing factor is defined as (3.3), or the generalized (3.2), for S-coarsening. (A way for sometimes defi-

ning an improved $\bar{\mu}$, by allowing for example π in (3.3) to be replaced by any smaller positive number, is explained in Sec. 12.)

If point (not block) relaxation is to be used, then S-h-ellipticity is a necessary and sufficient condition for the existence of relaxation schemes with good (i.e., bounded away from 1) S-smoothing factors. This is an easy generalization of a theorem proved in [B12, Sec. 4.2]. The more general situation, with block relaxation, is summarized by the following theorem:

THEOREM 3.1 Let S and S' be two sets of directions: $S, S' \subset \{1, \dots, d\}$. A necessary and sufficient condition for the existence of an S-block relaxation scheme with good S'-smoothing rates is that the discrete operator L^h is uniformly coupled in S' modulo S; that is,

$$E_{S',S}^h(L^h) = \min_{\substack{\pi/2 \leq |\theta|_S \leq \pi \\ \theta'_j = 0_j \text{ for } j \in S}} |\tilde{L}^h(\theta)| / |\tilde{L}^h(\theta')| = O(1). \quad (3.4)$$

$E_{S',S}^h$ is called the measure of uniform coupling in S' modulo S. The theorem states, in other words, that the S'-smoothing factors, produced from L^h by any suitable S-block relaxation, are bounded away from 1 by a quantity which depends only on $E_{S',S}^h(L^h)$.

In this context, the role of block relaxation can sometimes be played by simple point relaxation. This is when the relaxation marching direction conforms with the downstream time-like direction of a hyperbolic-like system, so that a relaxation sweep nearly solves the equations (especially when upstream differencing is used). The role of block relaxation can also be played, more automatically and in more situations, by ILU smoothers (see Sec. 3.8).

Variable coefficients. When the coefficients of L^h (or of its linearization, in case it is nonlinear) are not constant, a perfect smoother is a relaxation scheme whose formal $\bar{\mu}$ (calculated at each point by assuming the coefficients there to extend as constant coefficients throughout) is good at all points. Such a perfect smoother can sometimes require quite costly block relaxation (e.g., plane relaxation) in several directions, because of varying semi-h-ellipticity directions. It is therefore important to realize that such perfect smoothers are not really needed because accidental semi h-ellipticity need not be taken into account. Explanation:

First, the above block-relaxation rule itself suggests that the formal $\bar{\mu}$ may be allowed to be bad (close to 1) at some isolated points.

The multigrid convergence rates will still be good. More importantly, however, consider the common case of a (nearly) non-elliptic differential operator whose (sub)characteristic directions continuously vary over the domain, so that in some particular small region they happen to approach some grid directions. As a result, in that particular region the discretization may be semi h-elliptic, smoothing there will be bad, hence the asymptotic multigrid convergence will slow down. Notice, however, that the components slow to converge are very special ones: They are necessarily high-frequency characteristic components in that particular region; i.e., components smooth in the (sub)characteristic directions but not smooth in all directions. Such components exist on the grid only when (sub)characteristic directions approach grid directions. Elsewhere, such components are not represented at all by the finite-difference solution; they are truncated. Hence, if one is interested in solving the discrete equations only to the level of truncation errors (and hence eventually uses an FMG algorithm -- cf. Sec. 7), such components can be ignored: Their slow algebraic convergence in regions of accidental semi h-ellipticity does not matter, because similar components are not approximated at all in other regions.

Only in cases of strong alignment (see Sec. 2.1) the corresponding block relaxation must be used. But it may be confined to the region of strong alignment. If, for example, the strong alignment is due to grid alignment of boundary layers, it is enough to perform line (or plane) relaxation only at the very lines adjacent to such boundaries (and sometimes not even there -- see [B17, Sec. 3.3]), with just point relaxation elsewhere. If the alignment is strong because it occurs throughout a major subdomain, line (or plane) relaxation of only that special direction is needed there. Alternating-direction block schemes may be needed only if errors far below truncation errors are for some reason desired.

A general way to avoid any need for block relaxation, even when solving far below truncation errors, is to use semi coarsening, as mentioned above. In case of variable coefficients, causing variable coarsening directions, this leads to AMG processes (see Sec. 13.1).

3.4 Distributive, weighted, collective and box Gauss-Seidel. Principal linearization

To obtain efficient smoothing, a selection should be made from an abundance of available relaxation schemes. The choice depends on experience and on some physical insight, with $\bar{\mu}$ calculations serving for final quantitative judgement. We list here some important types of

schemes. Each of those can be operated pointwise or blockwise (see Sec. 3.3) and in different orderings (see Sec. 3.6 below). Some simple schemes are described in more detail in [54]. We first describe the successive displacement species, then we mention their simultaneous-displacement counter-parts (Sec. 3.5). In case of equations with many or complicated linear or nonlinear lower-order terms, it may pay to apply any of these schemes with the scaled principal terms only (see Sec. 10.3). In case of complicated systems, see the general approach in Sec. 3.7.

The most basic scheme is the Gauss-Seidel (GS) scheme, in which all the discrete equations are scanned one by one in some prescribed order. Each of them in its turn is satisfied by changing the value of one corresponding discrete unknown. This is easy to do if the problem is linear and if there is a natural one-to-one correspondence between equations and unknowns (if, e.g., the matrix of coefficients is definite). In case the problem is nonlinear, each discrete equation may be a nonlinear equation in terms of the corresponding unknown. It is then usually best to make just one Newton step toward solving each equation in its turn. This Gauss-Seidel-Newton (GSN) scheme is not related to any global linearization of the system of equations, it just linearizes one discrete equation in terms of one discrete unknown, yielding usually a very simple scheme which does not require any storage other than the storage of the (approximate) solution.

Principal linearization. Moreover, it is actually enough to relax an equation through an approximate linearization of its *h-principal* terms, corresponding to the (sub)principal terms of the differential operator (see Sec. 2.1). Thus, for example, if the equation is $\mu \Delta u + uu_x + \dots$ and the current approximation just before relaxing at some point is \tilde{u} , then the relaxation at that point can simply be the same as if relaxing the equation $\mu \Delta u + \tilde{u}u_x + \dots$. A full linearization would in addition include the term $(u-\tilde{u})\tilde{u}_x$, but on the scale of the grid, hence in relaxation, that term is negligible. This "principal linearization" is as good as full linearization for purposes of relaxation, at least as long as differences of \tilde{u} at adjacent gridpoints are small compared with \tilde{u} itself. Note that for quasi-linear equations, which include almost all practical equations, the principal linearization involves no linearization at all, just trivially retarding the lower-order derivatives in each term, as in the above example.

When relaxation is used as the prime solver, much may be gained by Successive Over Relaxation (SOR), in which the GS correction calculated for each unknown is multiplied by a relaxation parameter ω . The situa-

tion is different when relaxation is used only as a smoother in multi-grid solvers. The best smoothing (lowest $\bar{\mu}$) is usually obtained for the natural value $\omega = 1$, so that GS is not only cheaper (per sweep), but also at least as effective (per sweep) as SOR. Lower $\bar{\mu}$ may be obtained by other parametrizations (e.g., the distributive GS described below), but for regular second-order elliptic equations this gain hardly justifies the extra work involved; Simple GS is probably the best known smoother (especially with red-black ordering -- see Sec. 3.6).

If block relaxation is required (cf. Sec. 3.3), block GS can be used. This means that the blocks are scanned one by one, the equations of each block are simultaneously satisfied by changing the corresponding block of unknowns. In the two-dimensional plane (x,y) , if the blocks are lines parallel to x (lines with constant y), the relaxation is called *x-line GS*, or *xLGS*. Similarly *yLGS* is defined (see example in Sec. 1.2).

When there is no natural one-to-one correspondence between discrete equations and unknowns (the matrix is not approximately definite; e.g., non-elliptic and singular perturbation equations, or elliptic systems which are not strongly elliptic [B19, Sec. 3.6]), simple GS should be replaced either by Distributive Gauss-Seidel (DGS) or by Weighted Gauss-Seidel schemes. In DGS, with each discrete equation we associate a "ghost" unknown, with some prescription being selected for the dependence of regular unknowns on ghost unknowns. Usually, each regular unknown is written as a prescribed linear combination of neighboring ghost unknowns. Then, as in GS, the equations are scanned one by one, each being satisfied by changing the corresponding ghost unknown. This means in practice that a certain pattern of changes is distributed to several neighboring regular unknowns (hence the denomination "distributive" GS); the ghost unknowns do not explicitly appear, nor stored in any way, they just serve for the description of DGS. (In fact their values are never known -- only changes in their values are calculated to induce changes in the regular unknowns.) In case of block (e.g., line) DGS relaxation, a block of ghost unknowns is simultaneously changed to simultaneously satisfy the corresponding block of equations. In two dimensions we thus have *xLDGS* and *yLDGS* schemes. A special case of DGS is the Kaczmarz relaxation (see Sec. 1.1). Other examples are described in detail in Secs. 17.3, 18.3, 19.3 and 20.3. The smoothing analysis of DGS schemes is best executed in terms of the ghost unknowns; see Sec. 3.7.

In Weighted GS (WGS) schemes, with each discrete unknown we associated a ghost equation, which is a preassigned linear combination of

neighboring equations, and we perform GS in terms of the ghost equations. Taking work into account, WGS is usually inferior to DGS, since each equation is calculated several times per sweep, unless the ghost equations explicitly replace the original equations -- which is just a linear transformation of the discrete system of equations. It pays to transform the system (as against performing DGS) only if the resulting system is not more complicated than the original, which is seldom the case: A transformation that yields a simpler system could usually be done already in terms of the differential equations, giving a simpler differential system. (Exceptions are cases where the simplifying transformation gives a worse system for discretization; e.g., a system not in conservation form as in [B16, Sec. 2.1]. To relax in that case in conservation form, a combination of WGS with the DGS scheme of [B16, Sec. 4.1] is indeed needed.)

For systems of equations ($q > 1$), simple GS is appropriate only in case the system is strongly elliptic [B19, Sec. 3.6]. Otherwise collective GS or DGS schemes should be employed. Collective Gauss-Seidel (CGS) is performed when the grid is not staggered, i.e., all the q grid equations and q unknown functions are defined on the same gridpoints: The grid points are scanned one by one, at each point we change simultaneously ("collectively") its q unknowns so as to simultaneously satisfy its q equations. In case of a staggered grid one can divide the domain into (usually overlapping) small boxes. The boxes are scanned, for each one we change simultaneously all unknowns interior to it so as to simultaneously satisfy all equations interior to it. This is called Box GS (BGS. See relevant remarks in Sec. 5.6). DGS schemes are generally more efficient for staggered grids than BGS (except sometimes in very coarse grids; cf. Sec. 6.3), because they do not couple the equations (see Sec. 3.7).

For *nonlinear* equations, all these methods can be used, but instead of fully satisfying an equation (or a collective of q equations, or a box of equations), only one Newton step (or just principal linearization) is made in terms of the corresponding (regular or ghost) unknown (or collective of q unknowns, or box of unknowns). For semi h -elliptic cases, block CGS (e.g. line CGS, meaning simultaneous solution of all equations on a line through changing all that line's unknowns), or block DGS, or block BGS, may be performed (after principal linearization, if needed).

Higher-order equations are sometimes most efficiently relaxed by writing them as systems of lower order equations. For example, the bi-

harmonic can be written as a pair of Poisson equations. Relaxing this system involves less work (per complete sweep) and yields better smoothing (per sweep) than relaxing the biharmonic. But special care should be taken in relaxing the boundary conditions for this system (see Sec. 5.3).

3.5 Simultaneous displacement (Jacobi) schemes

The GS schemes described above are *successive-displacement* schemes: The new value of an unknown (or block of unknowns) replaces the old one as soon as it is calculated, and is immediately used in relaxing the next equations. In *simultaneous displacement* schemes new values replace old ones only at the end of the sweep, after all of them have been calculated; hence each of them is calculated explicitly in terms of old values only. Corresponding to each of the schemes above we have a simultaneous-displacement scheme, called: Jacobi-relaxation, Jacobi-Newton, distributive Jacobi, weighted Jacobi, collective Jacobi, box Jacobi, line Jacobi, line distributive Jacobi, etc. -- corresponding to GS, GSN, DGS, WGS, CGS, BGS, line GS, line DGS, etc., respectively.

Unlike GS, Jacobi schemes often require under-relaxation ($\omega < 1$) in order to provide good smoothing. But with relaxation as a smoother (not an independent solver), good and optimal values of ω are independent of the domain, and can easily be calculated by local mode analysis.

Distributive and weighted Jacobi (under-)relaxation amounts actually to the same thing. An example of an optimized weighted Jacobi scheme is analyzed in [B7, Sec. 3.3].

All experience so far shows that Jacobi schemes are inferior to the corresponding GS schemes. They not only require more work (for operating the relaxation parameter) and more storage (for storing the new values separately), but their smoothing factors are in fact worse. For the 5-point Poisson equation, for example, Jacobi under-relaxation ($\omega_{\text{optimal}} = .8$) yields $\bar{\mu} = .6$, while GS gives $\bar{\mu} = .5$ and $.25$ for lexicographic and red-black orderings, respectively. The situation is similar in all cases so far examined. The advantage of simultaneous-displacement schemes is in their being more amenable to certain rigorous analyses (but there seems to be no practical value to this -- see Sec. 14) and their vectorizability and parallelizability (but red-black GS and similar schemes are also fully parallelizable -- see Sec. 3.6).

3.6 Relaxation ordering. Vector and parallel processing

For successive-displacement schemes, the order in which the equations (or blocks of equations) are relaxed has an important effect on the smoothing factors. The main orderings used are the usual lexicographic (LEX) order (in which the equation at grid point (i_1, \dots, i_d) is relaxed before (j_1, \dots, j_d) if $i_k = j_k$ for $1 \leq k < l$ and $i_l < j_l$), and related orders (LEX order for some permutation of the coordinates, some of them possibly reversed); symmetric relaxation (lexicographic sweep followed by a sweep in the reversed order); red-black (RB) ordering (in which all "red" gridpoints are relaxed before all "black" ones, where the coloring is similar to that of a checkerboard, namely a point (i_1, \dots, i_d) is red if $i_1 + \dots + i_d$ is odd, and black if it is even); and more general pattern relaxation (similar to RB, but with different coloring and possibly more colors). For difference equations involving more than nearest neighbors, RB schemes still depend on the ordering of points within each color. If such points are displaced simultaneously, the scheme is called Jacobi-RB; similarly LEX-RB, etc. The performance (i.e., smoothing per operation) of Jacobi-RB is more like GS schemes than like Jacobi, and like them it does not generally require under-relaxation.

Each of these orderings has its block-relaxation versions. xLGS (or xLDGS) can be done lexicographically forward (increasing y) or backward (decreasing y), or symmetrically (forward alternating with backward). Or, corresponding to RB, we can first relax the even lines, then the odd lines. This is called zebra xLGS (or x-zebra) relaxation. Similarly, yLGS (or yLDGS) can be done upward, downward, symmetrically or zebra. Particularly robust schemes are the Alternating-Direction Zebra (ADZ = x-zebra alternating with y-zebra) and Alternating-Direction Symmetric LGS (ADS = symmetric xLGS alternating with symmetric yLGS). Many more block GS schemes are similarly defined in higher dimensions. The choice of blocks is governed by the rule in Sec. 3.3. Concerning the choice of ordering we have the following remarks.

It has been found that GS with RB ordering is the best for the 5-point Poisson equation [F2]. Similarly, DGS with RB ordering within each of its passes (called briefly Distributive RB, or DRB) is the best for many *systems*, such as Cauchy-Riemann and compressible and incompressible Navier-Stokes equations (see Secs. 17.3, 18.3, 19.3 and 20.3). For 5-point Poisson, RB-GS provides $\bar{\mu}_1 = .25$, $\bar{\mu}_2 = .25$ and $\bar{\mu}_3 = .32$ (cf. Eq. (3.2)), as against $\bar{\mu} = .5$ for LEX-GS. Moreover, RB-GS can be executed with only four operations per grid point, whereas lexicographic GS requires five. Similar comparisons hold for the more complicated

elliptic systems.

In addition, the mentioned RB schemes (more precisely Jacobi-RB) and similar pattern relaxation schemes are fully vectorizable and parallelizable: All the equations of the same color can be relaxed in parallel, thus taking full advantage of computers having vector or parallel processing capabilities. The zebra schemes are similarly parallelizable. (See more about parallelization of all multigrid processes in [B14].)

For non-elliptic equations or for elliptic equations with large non-isotropic lower-order terms (singular perturbation problems, in particular), the first approach ([B6], [S2], [B7], [B10]) was to employ "downstream" ordering, in which the equation at a point A is relaxed before (or simultaneously with) that at point B if the solution at B depends more heavily on the solution at A than vice-versa (e.g., if the fluid flows, or the convection transports, from A to B). This provides very good smoothing factors (better than those for regular elliptic problems). If different "downstream" directions exist at different parts of the domain, this may require a sequence of several relaxation sweeps in several directions. If for example line relaxation is also required, ADS relaxation may be needed, i.e., four passes over the domain. Each pass may be effective in only part of the domain, but the combined sweep will give excellent smoothing everywhere, for any combination of semi h-elliptic approximations in two dimensions (and also in three dimensions, if the grid is coarsened in only two directions (cf. Sec. 4.2.1)). In some particular cases (when the reduced equation is hyperbolic in some time-like direction, and upstream differencing is employed) such schemes yield not only great smoothing but also great convergence, making coarse-grid corrections superfluous.

Our preference today, however, is away from these downstream marching schemes. First, because they are not so good for vector and parallel processing. Also, because in cases where several downstream directions are required, the programming is complicated and the multi-direction procedure is not fully efficient, since it requires several passes over the grid where one or two (efficient) passes is all that would be needed at each multigrid stage. Hence ordering-free schemes were developed, with which good smoothing is obtained for any ordering, including RB and/or zebra (the block-relaxation rules should still be kept). Such ordering-free schemes are obtained either by distributive relaxation [B10, Sec. 6], or by using slightly more artificial viscosity than that required for upstream-differencing [B12, Sec. 4.3], [B17, Secs. 5.7, 6.3, 7.2]. Actually, even these devices (distributive relaxation and/or increased artificial viscosity) are not usually needed, unless one wants a

"perfect smoother" in order to reduce algebraic errors far below truncation errors: If all that matters is the fast approximation of the differential (not the discrete) solution and an FMG algorithm is employed, then the simplest direction-free schemes, such as RB, can do [B17]. Moreover, the tendency of downstream marching schemes to yield fast convergence (not just fast smoothing) may sometimes be disadvantageous (see Sec. 3.2).

3.7 Principle of relaxing general PDE systems

To obtain a good smoother (i.e., a good combination of discretization and relaxation, yielding good differential smoothing, as defined in Sec. 12) for a given system of q partial differential equations $Lu = f$, it is important to understand in advance what smoothing rates are obtainable. The guiding principle here is the following.

The smoothing rate for a given PDE operator L can be as good as the smoothing rates obtainable for the factors of the subprincipal part of $\det L$.

To explain this, we first show how a smoother for L can be constructed in terms of a smoother for the scalar operator $\det L$. One way to do it is through distributive relaxation (cf. Sec. 3.4). Such a relaxation is defined by considering the vector of unknown functions u (or their discrete counterparts) to depend affinely on a "ghost" vector of functions w , i.e., $u = Mw + u^0$, where M is a $q \times q$ differential (or finite-difference) operator, and the vector u^0 is immaterial (since we are interested in changes in w , through which changes in u are defined; w itself is not explicitly present). In terms of w we then devise a suitable relaxation for the operator LM . It is easy to see that the smoothing rate of this relaxation in w will automatically be taken over by the resulting distributive relaxation in u .

One particular choice is to take M to be the transposed matrix of cofactors of L , in which case LM equals $\det L$ times the $(q \times q)$ identity operator. One can thus devise for each component of w any relaxation suitable for $\det L$; the corresponding distributive relaxation for u will have the same smoothing rate.

As mentioned earlier (Sec. 2.1), the only part of L which really participate in devising the smoother is the subprincipal part of the linearized operator; the smoothing rates obtained for L are the same as for that part. Thus, for the discussion here, we may think of L as having subprincipal terms only. In that case we can often factor $\det L$ into simpler factors. Typically, in many physical problems, the factors

are either the Laplacian Δ or the convection-diffusion operator $\Delta + \underline{a} \cdot \underline{\partial}$. Smoothing rates for the latter are discussed for example in [B17], [K2] and [S4].

One general way to devise the relaxation of L in terms of the factorization of $\det L$ is to correspondingly factorize L itself, using the following theorem.

THEOREM 3.2 If $\det L = \ell_1 \ell_2$, where each ℓ_i is a (scalar) differential operator, then one can factorize the $q \times q$ operator L into $L = L_1 L_2$, where L_i are $q \times q$ matrix operators, such that $\det L_i = \ell_i$.

The proof, for which we thank Anthony Joseph, is based on Theorem 5 on page 393 in [L4]. A nice example is the factorization of the elasticity operator

$$\begin{pmatrix} \mu\Delta + \lambda\partial_{xx} & \lambda\partial_{xy} \\ \lambda\partial_{xy} & \mu\Delta + \lambda\partial_{yy} \end{pmatrix} = \begin{pmatrix} \partial_x & \partial_y \\ \partial_y & -\partial_x \end{pmatrix} \begin{pmatrix} \lambda + \mu & 0 \\ 0 & \mu \end{pmatrix} \begin{pmatrix} \partial_x & \partial_y \\ \partial_y & -\partial_x \end{pmatrix}$$

which is indeed useful for its relaxation (through the scheme of Sec. 17.3), especially in case $\mu \ll \lambda$, where simpler schemes fail.

To relax the factorized system $L_1 L_2 u = f$, one can simply introduce the auxiliary vector of unknown functions $v = L_2 u$ and alternating relax the two systems: $L_1 v = f$ and $L_2 u = v$. The combined smoothing rate is no worse than the worst of the rates of the two systems. (If these two rates are very different, the system with the slower rate can be relaxed more times.) Special care should of course be exercised in relaxing near and on boundaries (see Sec. 5.3).

In many cases there is a simpler distributive relaxation which meets the goal of the above guiding principle. It is not necessary that LM be diagonal as in the general approach above; it is enough to get LM to be triangular. Moreover, if the operators on one of the columns of M all have a common divisor, that divisor can be omitted. In this way one can often have each term on the diagonal of LM to be just one separate factor of $\det L$ (some factors possibly appearing in more than one diagonal term), in which case no auxiliary functions (such as v above) are needed. The relaxation schemes in Secs. 17-20 are all of this kind. Also, instead of distributive relaxation one can obtain the same goal by weighted relaxation schemes. The important upshot in any case is that the goal is known in advance, so do not settle for any substantially slower rates. Note that the above smoothing discussion assumes frozen operators, hence may not apply to coarse grids.

3.8 ILU smoothers

The above list of relaxation schemes, although including some of the most efficient smoothers, does not exhaust all possibilities. Of special interest is the use of incomplete LU decomposition (ILU), and related schemes, as smoothers. Such smoothers, first introduced in [W2], have been shown to be very robustly efficient for a wide range of 5-point and 9-point difference equations. For an extensive treatment, see [K2] and the more recent [S9].

The basic ILU process can be described as a Gaussian elimination truncated so as to preserve a certain pattern of sparsity, simply by ignoring (i.e., replacing immediately by zero) any term produced by the elimination process at any matrix position designed to remain zero (e.g., any matrix position which is originally inherent zero). In case of non-linear equations one can apply this process with the principal linearization (see Sec. 3.4).

The robustness of the ILU smoother can be explained by its ability in many cases to automatically produce an approximation to desired block relaxations in varying grid directions. If for example the system contains any sequence of unknowns each of which is strongly coupled only to its predecessor and successor in the sequence, and if the ILU ordering of unknowns conform with the ordering of that sequence (i.e., it gives that sequence upon omitting all other unknowns), then, ignoring weak couplings, the equations of the sequence appear to the ILU process as a separate tridiagonal system, which it automatically solve simultaneously (since Gaussian elimination for a tridiagonal system does not produce new non-zero terms). This is exactly what the block relaxation rule (Sec. 3.3) requires. Provided the weak couplings do not somehow accumulate and spoil this picture. There are special situations where that may happen. Indeed, for some special anisotropic equations there are some very special high-frequency error components which are even considerably *magnified* by the simple ILU process. More advanced "block ILU" (see [K2, Sec. 3.2.3]) should then be used.

One can produce systems where various sequences of strong couplings are so ordered as to contradict any ordering chosen for the ILU process. But this rarely happens in practice; especially in two dimensions, such systems are artificial. Thus, using ILU we need to worry much less about all directions of relaxation required for a perfect smoother. (But such a perfect smoother is usually needed only for solving far below truncation errors; see Sec. 3.3.) A careful comparison, in which the total amount of operations in an FMG algorithm is counted

taking into account the ILU set-up operations, shows ILU schemes to be quite comparable to suitable GS schemes [T5]. In many cases, however, ILU requires much more storage: One needs to store all the non-zero matrix elements, whereas the GS schemes often require storing just the approximate solution itself (e.g., when the problem, whether linear or not, is autonomous, as in all fluid dynamics cases). Also, GS schemes (in red-black ordering) are much faster on vector machines.

Both an advantage and a disadvantage is the fact that ILU is a "package deal", automatic prescription: It tells you exactly what to do in complicated situations, near boundaries for example; but it does not allow you to change the scheme to deal with special needs, such as local relaxation near re-entrant corners and other singularities [B24], other special local treatments, separate smoothing of coupled differential equations, etc. One can however combine ILU smoothers with sophisticated distributive schemes. For example, within the distributive relaxation described in Sec. 19.3 for the Navier-Stokes equations, one can use ILU for relaxing each set of momentum equations ((19.4b) for one j) in terms of the corresponding velocity function u_j^h .

4. INTERIOR TWO-LEVEL CYCLES

Having computed the smoothing factor $\bar{\mu}$, one should expect the asymptotic convergence factor per multigrid cycle to approach $\bar{\mu}^v$, where v is the number of relaxation sweeps (on the fine grid h) per cycle. This ideal figure does not take into account the exact nature of the inter-grid transfers. The next task then is to design those transfers so as to approach the ideal figure. To separate their design from questions related to boundary conditions (which are taken up at the next chapter), we still think in terms of fully-periodic or infinite-space problems; we still, that is, restrict our attention to *interior* processes, because it is there that most of the computational work is invested. Furthermore, we simplify the multigrid situation at this stage by restricting our attention to two grids only, the finest grid $h = (h_1, \dots, h_d)$ and the next coarser grid $H = (H_1, \dots, H_d)$, where usually $H = 2h$. That is, we assume in our analysis that the grid- H equations are solved (exactly) each time the algorithm gets to that grid, without analyzing how that solution is obtained, hence without involving grids coarser than H in the analysis.

These assumptions indeed simplify our studies very much. First, the error can be expanded in a Fourier integral (or series) and the transformations of the amplitudes of different Fourier components by multigrid operations can be calculated. Indeed, for linear systems with constant coefficients only few Fourier components at a time are coupled by these two-level interior processes, hence transformations of Fourier amplitudes are expressed as small matrices (Sec. 4.1). In case of non-constant coefficients, we usually freeze them at some values (treated then as parameters of the analysis). In case of nonlinear equations, their Newton linearization is analyzed (although no such linearization is needed in the actual processes; see Sec. 8.3). The parameters of the analysis then depend on the solutions around which linearization is made.

This freezing of coefficients is reasonable as long as the real coefficients do not change too drastically over a meshsize. Where they do, we can sometimes model them as changing periodically, again making mode analysis with small matrices possible [B19, Sec. 4.7].

When mode analysis becomes too difficult or dubious, or if one simply wishes to skip it, experiments with periodic boundary conditions can be used instead. One can in fact do such experiments even when mode analysis is available, and compare the analysis with the experiments. This is an accurate debugging technique, completely separating away issues related to boundary conditions. Moreover, such periodic-boundary-

condition program could serve as an excellent preliminary stage in developing your real multigrid program. For this preliminary program, the various advices given in Sec. 5.1 concerning the full program could be used, including in particular the use of *multigrid* program to simulate a *two-grid* algorithm, and the trick of near-linearization of non-linear equations.

Some warnings, however, concerning this use of periodic boundary conditions: First, relaxation should better be restricted to simultaneous-displacement (Jacobi, red-black Jacobi, zebra Jacobi, etc.) schemes, to avoid the special non-smoothness created along the starting line (or termination line) of the relaxation sweep. (This non-smoothness would be particularly bad if downstream ordering were used.) Second, periodic boundary conditions often require additional global conditions to be added so as to make the problem well posed. In some cases these extra conditions are easy to implement, involving for example just an adjustment of the solution average by adding a constant. (Multigrid treatment of global conditions is discussed in Sec. 5.6.) Other cases, especially nonlinear ones, are less straightforward. A full section could, and perhaps will, in fact be written about the art of using periodic problems. It is also important to realize that interior studies in general have more limited value for non-elliptic problems, where the interplay with boundaries is more essential (see end of Sec. 4.1).

One should also make sure that at this stage (whether mode analysis or periodic numerical experiments are used) both grid h and grid H are fine enough, and grid- H equations are solved accurately enough (without taking into account the work this accurate solution requires), in order to separate away questions related to coarser grids (see Sec. 6.3). Do not forget, however, in the process of optimizing the grid- H operator L^H , that this is a modeling for a multigrid solution, hence your model must be recursive: The H equations should have the same general form as the original h equations, with the same range of possible parameters.

In addition to the relaxation scheme, studied above, the main issues to be studied at this interior-two-level stage are when to switch (under what criteria, or after how many relaxation sweeps) from grid h to grid H ; what should be the coarse-grid variables; and the type (in the interior) of three multigrid operators: The fine-to-coarse transfer of residuals I_h^H , the coarse grid operator L^H , and the coarse-to-fine interpolation of corrections I_H^h . These issues are one-by-one discus-

ssed in the subsections below. They are later reviewed again, from a more general perspective, in Sec. 11. Relevant to these issues are also Secs. 8.5 and 10.2 (nonlinear problems and higher-order techniques).

4.1 Two-level cycling analysis. Switching criteria

Details of the two-level mode analysis are described in [B19, Secs. 4.6-4.8] and in [S4, Secs. 3.3-3.5, 7, 8, 9]. The former also discusses modifications of the analysis to account for the fact that in practice the grid-H equations are only approximately solved, modification for the case of equations with highly oscillatory coefficients, and ways to make precise comparisons between mode analysis and numerical experiments (for debugging purposes). The main things to know are the following.

On grid $2h$ the Fourier mode $\exp(i\theta \cdot x/h)$ aliases (coincides with) the mode $\exp(i\theta' \cdot x/h)$ if $|\theta_j - \theta'_j| = 0$ or π for $j = 1, \dots, d$. Hence each set of so aliasing components usually includes 2^d components $\{\theta^1, \dots, \theta^{2^d}\}$, called harmonics of each other. They are coupled to each other by the two-level processes. (The special sets with less than 2^d different components do not require special analysis, since they are limits of regular sets.)

We define the two-level cycle as follows: Make ν_1 relaxation sweeps on grid h , then transfer the residual problem to grid H and solve it there exactly, then interpolate that grid-H solution to grid h and add it as a correction to the former grid- h solution, then make ν_2 more relaxation sweeps on grid h . It is easy to see that in the infinite space, if L^h, L^H, I_h^h and I_H^h are all constant operators, and if the error in the solution before such a cycle has the form $\sum_j A_j \exp(i\theta^j \cdot x/h)$, where the sum is over a set of 2^d harmonics, then the error after the cycle will have a similar form, and the new A_j 's will be linear combinations of the old ones. If we deal with a system of q grid equations then each amplitude A_j is a q -vector, hence the overall transformation of the 2^d amplitudes by the two-level cycle is a $(2^d q) \times (2^d q)$ matrix M , which can be denoted $M(\theta)$ where θ is the lowest harmonic ($|\theta| < \pi/2$).

This matrix $M(\theta)$ is called the two-level amplification matrix. The easiest and most modular program for calculating it is to write a different routine for the general matrix-element of each of the five involved processes: relaxation, L^h, I_h^h, L^H and I_H^h . Their respective matrices S^h, L^h, I_h^h, L^H and I_H^h have dimensions $(2^d q) \times (2^d q)$

$(2^d q) \times (2^d q)$, $q \times (2^d q)$, $q \times q$ and $(2^d q) \times q$, respectively, and each of their elements is a function of θ . Then program

$$M(\theta) = (S^h)^{\nu_1} [I - I_H^h (L^H)^{-1} I_h^h L^h] (S^h)^{\nu_2}. \quad (4.1)$$

The main performance measure of the two-level cycle is the two-level asymptotic convergence factor (per cycle)

$$\bar{\lambda} = \max_{|\theta| \leq \pi/2} \rho(M(\theta)). \quad (4.2)$$

Note that $\bar{\lambda}$ depends on the sum $\nu = \nu_1 + \nu_2$, but not on the separate values of ν_1 and ν_2 . In fact, when many cycles are performed the separate values are immaterial. Various other performance measures can similarly be defined. (See [S4, Secs. 3.4-3.5], where the notation M_h^{2h} and $\rho(M_h^{2h})$ is used for our M and $\bar{\lambda}$, respectively. Additional two-level measures will be discussed in Sec. 7.4.)

Using the two-level analysis we try to (roughly) optimize the involved processes; namely, the objective is to maximize $w^{-1} \log(1/\bar{\lambda})$, where $w = A(\nu w_0 + \nu_1 w_1 + \nu_2 w_2)$, w_0 is the work in one relaxation sweep, w_1 is the work of calculating and transferring the residuals, w_2 is the work of the I_H^h interpolation, and A is a factor through which the work on coarser grids is taken into account. For our objective here the value of A is really immaterial, but to have a definite value in later uses we observe that for V cycles we can assume similar operations on each of the grids, hence $A = (1 - \hat{\rho}_1 \dots \hat{\rho}_d)^{-1}$, where $\hat{\rho}_j = h_j/H_j$ (usually $\hat{\rho}_j = .5$) while for W cycles $A = (1 - 2\hat{\rho}_1 \dots \hat{\rho}_d)^{-1}$. To avoid the laborious count of operations and the arbitrary assignment of proper weights to different arithmetic and non-arithmetic operations (which are really machine-dependent), one can use the work of a standard relaxation sweep as the work unit. In complicated problems, where calculating L^h outweighs interpolations, one can then neglect w_2 and take $w_1 = 1$ for full residual weighting and $w_1 = 2^{-d}$ for residual injection. The convergence factor per work unit is then denoted by $\mu = \bar{\lambda}^{1/w}$. As above (Sec. 3.2), in addition to the goal of minimizing μ we should take robustness and simplicity into account.

One can also partly separate the study of I_h^h, L^H and I_H^h from that of relaxation by the Coarse-Grid Correction (CGC) mode analysis, as in [B7, Sec. A.1]. But this is not simpler than the full two-level analysis, especially since relaxation schemes have already been selected in the previous stage. We use a CGC analysis in Sec. 4.3 below.

The ideal factor $\bar{\lambda} = \bar{\mu}^v$ is not always obtainable. If $\bar{\mu}^v$ is too small we will get $\bar{\lambda} > \bar{\mu}^v$, because of significant high-frequency amplitudes generated from low ones by interpolation or by RB-type relaxation (see Sec. 4.3). Even when obtainable, too small values of $\bar{\lambda}$ will require too precise interpolations, hence too much investment in w_1 and w_2 , and will at a later stage be frustrated by other interactions (boundaries and non-constant coefficients). Also, such small $\bar{\lambda}$ will not usually be needed in the final FMG algorithm (see Secs. 7.2-7.3). Hence, the optimal cycle always employs a small v , typically $v \leq 3$.

In regular elliptic problems $v = 1$ is too small to be optimal (unless the sweep includes several passes, as in symmetric and alternating-direction schemes), since the overhead of w_1 and w_2 weights too much against it. Hence usually the optimal number is $v = 2$ for very efficient smoothers ($\bar{\mu} \leq .3$ or so), and $v = 3$ otherwise. A small change in v does not disturb the overall efficiency very much. Considerably larger v are less efficient, because they bring the process into the range of larger feeding from low to high frequencies, while not much more is gained in reducing the overhead (already at $v = 3$, $w_1 + w_2$ is quite small compared with $v\omega_0$).

A possible approach is accommodative: do not fix v in advance, but continue relaxation as long as it exhibits the fast convergence of high frequencies, e.g., as long as the convergence factor (some norm of the residuals divided by the same norm a sweep earlier) is smaller than the smoothing factor $\bar{\mu}$. For non-scalar ($q > 1$) systems, such a criterion can separately be applied to each equation, possibly resulting in more passes for part of the equations. Similarly it may separately be applied at different subdomains (since smoothing is a local process), possibly giving partial relaxation sweeps.

In case of non-elliptic and singular perturbation problems there are some particular smooth error components (smooth characteristic components of the differential operator or the reduced differential operator) for which L^H is a bad approximation to L^h , hence $\bar{\lambda}$ cannot be much smaller than .5, no matter how small $\bar{\mu}^v$ is [B17, Sec. 5.1], [B3]. But for exactly the same components and the same reason, L^h itself is not a good approximation to the differential operator L . Hence, exactly for these components, we do not need much algebraic convergence (convergence to the discrete solution), since the discrete solution itself is far from the differential solution. Hence, for such cases the asymptotic convergence factor $\bar{\lambda}$ is not really the measure we need. The

one we need is obtained by the two-level FMG analysis (see Sec. 7.4). Moreover, for non-elliptic or singular perturbation problems the usual assumption that high-frequency components are local does not hold. It is violated by high-frequency characteristic components in cases of strong alignment (Sec. 2.1). The interior mode analysis should then be supplemented with a half-space analysis (Sec. 7.5).

4.2 Choice of coarse grid

When the fine-grid, with meshsize $h = (h_1, \dots, h_d)$, is given, the choice of a coarse-grid, with meshsize $H = (H_1, \dots, H_d)$, is often straightforward: Take every other line (every other hyperplane, for $d > 2$) of the fine grid in each direction. The coarsening ratio $H_j/h_j = 2$ is usually optimal: it is the smallest recursively convenient number, and it is already big enough to make the coarser-grids work quite small relative to the fine-grid work; larger H_j/h_j will not save significantly more work, but will significantly degrade the smoothing factors (see Eq. (3.3)). The smaller ratio $H/h = 2^h$ may be as efficient (trading larger A for smaller v), and it is recursively convenient in some two-dimensional problems with rotatable operators; see [S4, Sec. 2.5], [R1].

When the fine-grid discretizations are done in terms of "cells" with the discrete variables defined at certain cell positions (e.g., cell centers, or centers of vertical cell boundaries, etc.), and especially when the grid is staggered (different grid functions are defined at different cell positions), it is more convenient to coarsen in terms of the cells: Take every 2^d fine cells as a coarse cell, and then place coarse-grid variables at coarse-cell positions analogous to their positioning in the fine cells. The coarse grid points then are not a subset of the fine grid points. See examples in Secs. 17.4, 18.4 and another approach in [D1].

In some cases the "fine-grid" is not a well-organized grid at all; e.g., a general finite-element triangulation, not based on any grid lines. Then one can still construct the coarse grid as a uniform grid, placed over the domain with no particular relation to the fine grid. Another approach is to base the choice of coarse-grid variables on purely algebraic considerations (Sec. 13.1). Mode analysis is of course not very suitable for analyzing such situations.

4.2.1 Semi coarsening. Semi coarsening, or more specifically S-coarsening, is the technique of using grid H which is coarser than h in only a partial set S of coordinates; i.e., $H_j = 2h_j$ for $j \in S$ and $H_j = h_j$ for $j \notin S$. This means somewhat more work on coarse grids; but

either this or block relaxation are needed in some cases -- see the rule in Sec. 3.3. Semi coarsening is sometimes preferable to block relaxation. For example, in three-dimensional problems where there are two fixed coordinates with stronger couplings, full coarsening would require plane relaxation, which is inconvenient. (Solving these plane equations approximately by one multigrid cycle, if done simultaneously at all planes, will look very much like semi coarsening.) Also, exactly in those cases, semi coarsening involves relatively small work on coarser grids, since two coordinates are still coarsened, hence the total number of points on all coarser grids is at most one third the number of points on the finest grid.

Sometimes, a combination of block relaxation and semi coarsening may be the best. For example, the equation $aU_{xx} + bU_{yy} + cU_{zz}$ with $a \ll b \ll c$, discretized on a cubic grid ($h_x = h_y = h_z$), will best be solved by z-line relaxation and y-z semi coarsening. Generally, rough calculations of S-smoothing factors (Sec. 3.3) immediately show what procedures can be taken.

In some cases block relaxation is of course preferable to semi coarsening. For example, when directions of strong alignment are different at different subdomains. To change accordingly the directions of semi coarsening would be much messier than changing block directions. (But note that strong alignment, as defined in Sec. 3.3, can rarely at all change directions.)

4.2.2 Modified coarse-grid functions. When a difference operator L^h is given which has no good h-ellipticity or semi-h-ellipticity measure, then no relaxation can be efficient in reducing all high-frequency error components. To reduce all components efficiently we will then apply modified coarse-grid correction functions.

Suppose for example that the slow components (i.e., the components for which relaxation is inefficient) are all clustered around some known modes $\varphi_j(\underline{x})$, ($j = 1, \dots, J$). This means that the error $v^h = \tilde{u}^h - u^h$ can be written as $v^h(\underline{x}) = \sum v_j^h(\underline{x})\varphi_j(\underline{x})$, where v_j^h are smooth functions. It is then these smooth functions which we try to approximate by coarse-grid functions v_j^{2h} . See [B15, Sec. 3.2] and [B21]. Sometimes, each of these functions can most efficiently be approximated on a different (e.g., differently rotated) grid.

4.3 Orders of interpolations and residual transfers

The most important aspect of the coarse-to-fine correction interpolation I_h^h and the residual transfer I_h^H is their orders, defined as

follows: The order of I_H^h is m if the interpolation of the low-frequency Fourier component $\exp(i\theta \cdot \underline{x}/h)$, with amplitude 1 on the coarse grid H , creates on the fine grid h , high-frequency components (the harmonics of the low frequency) with amplitudes $O(|\theta|^m)$. It also reproduces the θ component itself on grid h with an amplitude $1 + O(|\theta|^m)$. The order of the fine-to-coarse transfer I_h^H is said to be m , and its secondary order \bar{m} , if a high-frequency harmonic with amplitude 1 on grid h contributes $O(|\theta|^m)$ to the amplitude of the corresponding low frequency θ when transferred to grid H , while a low frequency with amplitude 1 on grid h contributes $1 + O(|\theta|^m)$ to its grid- H amplitude. Thus, linear and bilinear interpolations have order 2, while cubic interpolation is fourth order. Residual transfer by injection ($I_h^H \equiv 1$) has order 0 and infinite secondary order, whereas the usual full-weighting residual transfer (Eq. (4.6) below) is of order 2 and secondary order 2.

What orders should be used in the multigrid cycle? This depends on the orders of derivatives appearing in our equations. Suppose we have a system of q differential equations in q unknown functions, and let m_{ij} be the highest order of differentiation (or differencing) of the j -th unknown in the i -th equation, ($i, j = 1, \dots, q$). We assume, and this is usually the case, that the q unknown functions are interpolated independently of each other and that the residuals of each of the q grid equations are transferred separately from the others. Denote by m^j the order of I_H^h used in interpolating the j -th correction (correction to the j -th unknown function) and by m_i and \bar{m}_i the order and secondary order, respectively, of the I_h^H used in transferring the i -th residual (residuals of the i -th equation).

What m^j , m_i and \bar{m}_i ($i, j = 1, \dots, q$) should be used? Examining orders of magnitude in the CGC mode-analysis operator (the operator in brackets in (4.1)), under the assumption that all $m^j > 0$, we find the following basic rules and observations:

(A) The high-frequency harmonics of the lowest frequencies (those with $|\theta| = O(h)$), are amplified by the CGC operator by a factor with a spectral radius $1 + O\left(\sum_{i,j} h^{m_i + m^j - m_{ij}}\right)$. Hence, to avoid large magnification of high-frequencies, we should have

$$m_i + m^j \geq m_{ij}, \tag{4.3}$$

preferably even $m_i + m^j > m_{ij}$. On the other hand, larger values

$(m_i + m_j > m_{ij} + 1)$ would not significantly further reduce the spectral radius, hence they are asymptotically (when many cycles are made) not needed.

(B) Every high-frequency harmonic (before the CGC cycle) contributes to the corresponding low-frequency (after the cycle) through a $q \times q$ transformation matrix $(L^H)^{-1}B$, where $B_{ij} = O(h^{m_i - m_{ij}})$. This is usually not important asymptotically (for many cycles), but if only one cycle is performed (as in FMG algorithms) that transformation may cause large errors unless

$$m_i > m_{ij} . \tag{4.4}$$

For relaxation schemes with interactions between high and low frequencies (e.g., RB schemes), this transformation may also cause asymptotic degradation unless $m_i > \sum_k (m_{ij} - r_{kj})$, where $O(h^{r_{kj}})$ is the size of the high-frequency errors in the k -th function generated by relaxation from an $O(1)$ low-frequency error in the j -th function. RB and zebra schemes for $q = 1$ give $r_{11} = m_{11}$, hence the rule requires $m_1 > 0$, i.e., full weighting (see Sec. 4.4). This requirement can however be slackened by a more precise look at the nature of these particular schemes (allowing the use of simpler transfers such as the "half injection" $I_h^H \equiv .5$ or "half weighting"; see [F2, Sec. 2], [S4, Sec. 8.1]).

(C) The low-frequency error components themselves are reduced by a factor $O(h^{\tilde{m}})$, where $\tilde{m} = \min(\tilde{p}, \tilde{m}_1, \dots, \tilde{m}_q, m^1, \dots, m^q)$ and \tilde{p} is the lowest of the approximation orders on levels h and H . Hence \tilde{m} must be positive, which is indeed the case for any consistent differencing and interpolation schemes. Larger values of \tilde{m} may of course give better cycle performance. Our experience indicates that $\tilde{m} = 2$ gives considerably better $\bar{\lambda}$ than $\tilde{m} = 1$. Since this is a low-frequency matter, hence non-local, higher \tilde{m} may be effective only if they are carefully matched by corresponding high-order approximations and interpolations at boundaries. But one usually does not have to go into the trouble of $\tilde{m} > 2$. Rather, employ more cycles with $v \leq 3$ (see Sec. 4.1). As a result the factor $O(h^{\tilde{m}})$ will usually be dominated by \bar{u}^v in determining $\bar{\lambda}$.

(D) We also note that every low-frequency error component (before the CGC cycle) contributes to every one of its harmonics (after the j cycle) through a $q \times q$ transformation matrix D , where $D_{jj} = O(h^{m_j})$ and for $i \neq j$ D_{ij} has higher orders in h . This tells us something about the range where relaxation should be efficient (see Sec. 12).

4.4 Variable operators. Full weightings

The above mode-analysis rules are insufficient in case L^h is highly-varying, i.e., its coefficients substantially change between two neighboring gridpoints. For such L^h the residuals after relaxation are also highly varying, hence to represent them correctly on grid H full residual weighting should be used, i.e., I_h^H should satisfy, for any residual function r^h ,

$$(H_1 \dots H_d) \int_{x^h} (I_h^H r^h)(x^H) = h_1 \dots h_d \int_{x^h} r^h(x^h) , \tag{4.5}$$

where x^h are the fine-grid points and x^H are the coarse-grid points. In other words, full weighting "preserves integrals". (Throughout this discussion it is assumed that the difference equations on all grids are written in their divided form, analogous to the differential equations. If, however, they are multiplied through by factors which depend on the meshsize, then one should not forget to have those factors in (4.5), too.) One can regard full weighting as a scheme in which each residual $r^h(x^h)$ on the fine grid is being distributed to several coarse grid points, with weights whose sum is $\tilde{\omega} = h_1 \dots h_d / (H_1 \dots H_d)$. Hence each residual r^h is a weighted average of its transferred values on grid H , times $\tilde{\omega}$. This weighted average represents a certain interpolation, \hat{I}_H^h say. Thus every full weighting I_h^H is the adjoint (or, in matrix terminology, the transpose) of an interpolation \hat{I}_H^h , times $\tilde{\omega}$. We denote this by $I_h^H = \tilde{\omega} \hat{I}_H^{h*}$. The normal (9-point symmetric) full weighting, defined by

$$\left(I_h^{2h} r^h \right) (x^{2h}) = \sum_{\max |v_j| \leq 1} 2^{-d-|v_j|} r^h(x^{2h} + (v_1 h_1, \dots, v_d h_d)) , \tag{4.6}$$

is for example the adjoint of bilinear interpolation, times 2^{-d} .

The requirement (4.5) is equivalent, in terms of the Fourier analysis, to the requirement that I_h^H has a positive order (see Sec. 4.3).

Such full weightings should perhaps be used in almost any case. Only in some particular cases non-full weightings happen to be asymptotically somewhat better. An example is injection in case of the standard 5-point Poisson operator, which yields lower $\bar{\lambda}$ as well as lower w_1 than the full weighting (4.6) (see [B19, Sec. 4.8]). But even in those cases, for the purpose of Full Multigrid (FMG) algorithms (see Sec. 7), full weightings may be preferable. (See rule (4.4) above and [S4, Sec. 3.6].)

4.5 Coarse-grid operator. Variational and Galerkin coarsening

The coarse grid operator L^H should be a proper homogenization of the fine-grid operator L^h . In smooth problems this is easily obtained by good discretizations of both L^h and L^H . In nonlinear problems this is effectively obtained by a suitable FAS averaging of the fine-grid solution (see I_h^H in Sec. 8.5).

Sometimes one needs to derive L^H from L^h , not from the differential operator L , either because L is not available or because one wants an automatic program, for some general class of L^h , or without having to treat separately boundary conditions and whatever other features of the differential problem. That is, one wants to regard the fine grid equations simply as a matrix equation

$$\underline{L}^h \underline{u}^h = \underline{f}^h, \quad (4.7)$$

where the underlines signify matrix notation: \underline{L}^h is an $n^h \times n^h$ matrix, where n^h is the number of unknowns on grid h . The geometry of the grids is only used to construct the inter-grid transfers I_h^h and I_h^H . (A multigrid treatment without any geometrical structures is discussed in Sec. 13.1.)

In case \underline{L}^h is symmetric, a general way to derive L^H is to regard (4.7) as the equivalent problem of finding \underline{u}^h which minimizes the functional

$$\phi^h(\underline{u}^h) = \frac{1}{2} \underline{u}^{h*} \underline{L}^h \underline{u}^h - \underline{f}^{h*} \underline{u}^h, \quad (4.8)$$

where stars stand for transposing. Given now an approximate solution $\tilde{\underline{u}}^h$, with a (smooth) error $\underline{v}^h = \underline{u}^h - \tilde{\underline{u}}^h$ which is to be approximated by the coarse-grid correction $\underline{I}_H^h \underline{v}^H$, the equations for \underline{v}^H are fully specified by requiring it to yield a correction which reduces ϕ^h as far as possible, i.e., \underline{v}^H should minimize $\phi^h(\underline{u}^h + \underline{I}_H^h \underline{v}^H)$. This immediately gives the coarse grid equations

$$\underline{L}_H^H \underline{v}^H = \underline{I}_H^H (\underline{f}^h - \underline{L}^h \underline{u}^h), \quad (4.9)$$

where

$$\underline{I}_H^H = \underline{I}_H^{h*} \quad (4.10)$$

and

$$\underline{L}^H = \underline{I}_H^H \underline{L}^h \underline{I}_H^h. \quad (4.11)$$

Equation (4.9) has the general form (1.14). The specific prescription (4.10)-(4.11) is called *variational coarsening*, since it results from the variational formulation of the problem. It is automatically determined as soon as the interpolation operator I_H^h is selected. I_H^h given by (4.10) is automatically "full" in the sense of Sec. 4.4. (By (4.11), the coarse grid equations (4.9) are not basically changed if I_H^h is multiplied by any scalar, such as $\tilde{\rho}$ in Sec. 4.4.)

In case \underline{L}^h is not symmetric, (4.10) is not always advisable (see Sec. 4.6), but (4.11) can generally be used. This L^H is called the *Galerkin operator*, since it is equivalent to requiring the coarse-grid correction to be a *projection*, i.e., requiring the residuals of the corrected solution $(\tilde{\underline{u}}^h + \underline{I}_H^h \underline{v}^H)$ to vanish when transferred back to the coarse grid:

$$\underline{I}_H^H (\underline{f}^h - \underline{L}^h (\tilde{\underline{u}}^h + \underline{I}_H^h \underline{v}^H)) = 0.$$

The reason Galerkin operators and variational coarsening are not always advisable is the amount of work involved in the construction (4.11), which could be considerably larger than the entire solution work (e.g., when solving by the algorithm in Fig. 1.2). Also, once constructed, the Galerkin operator is often much more complicated than the simpler L^H derived directly from L (e.g., 9-point instead of 5-point formulae), and requires much larger storage for storing all its coefficients, whereas the simpler L^H may require no storage (e.g., whenever L is autonomous, whether linear or not; this includes all fluid dynamic equations, if they are not linearized). See more about this issue in Sec. 11.

4.6 Strongly discontinuous, strongly asymmetric operators

As long as the fine-grid operator does not vary drastically, the above rules for I_h^H , L^H and I_H^h work fine. A more difficult case is that of a strong discontinuity in L^h , i.e., where its coefficients change their *order of magnitude* within a meshsize. Orders of interpolation are not so important then; rather, special forms should be used which take into account the particular nature of the discontinuity. The rule is first to analyze the behavior, near the discontinuity, of the error which is inefficiently reduced by relaxation. This error is approximately a solution to the homogeneous equations. (If it is not, then it has large residuals and therefore there locally exists a relaxation

scheme for which it will be reduced efficiently; cf. Sec. 1.1). Hence its general behavior is like that of solutions to the homogeneous differential equations. The interpolation I_H^h of corrections should take this behavior into account. For example, if we have a diffusion problem $\nabla(a\nabla u) = F$, near a strong discontinuity of the diffusion coefficient $a(x)$ the derivatives of the solution to the homogeneous equation are not continuous; instead, $a\nabla u$ is continuous there, and this can be used to design good interpolation schemes [A1]. In case of singular perturbation or non-elliptic problems, solutions to the homogeneous equations are continuous along (sub)characteristics, hence interpolation should be as much as possible in the characteristic directions. This is possible exactly where it is most important, namely, cases of intended strong alignment (cf. Sec. 2.1). Thus, for example, avoid interpolating across a boundary layer which is intended to be sharply reproduced.

It is less clear how to generally design the residual transfers I_H^h and the coarse grid operators L^H near a strong discontinuity. In the symmetric case the variational rules (4.10)-(4.11) are most robust [A1], even though expensive. For cases which are not essentially symmetric the Galerkin operator (4.11) can still be used, but instead of (4.10) one should take

$$I_{-h}^H = \frac{*h}{-H} \quad (4.12)$$

where $\frac{*h}{-H}$ is an interpolation appropriate in the above sense for L^{h*} , the adjoint of L^h . See [D4].

For non-elliptic and singular perturbation problems, the considerations and experiments in [B17] indicate that improved results are obtained by a full residual weighting in which residuals, on being transferred from a fine gridpoint to a different point (or points) on the coarse grid, are transferred roughly in the downstream direction. As for correction interpolation for such problems, however, it seems that the symmetric schemes are preferable to schemes with upstream bias. Coarse-grid operators identical with the fine-grid ones (hence much cheaper than (4.11)) were used, with excellent FMG results (even where the asymptotic rates were slow).

A general perspective on these questions of coarsening a problem (designing I_h^H , L^H , I_H^h) is given in Sec. 11 below.

5. BOUNDARY CONDITIONS AND TWO-LEVEL CYCLING

The theoretical two-level mode analysis described above (Sec. 4.1), and/or the numerical experiments with periodic boundary conditions (Sec. 4), give us the ideal convergence factor per cycle $(\bar{\lambda})$, or per work-unit $(\bar{\nu})$. These are the *interior convergence factors*, obtained in the absence of boundary interference. The next stage is to construct an actual multigrid program for an actual, bounded domain, and in particular to decide on the special treatment the various processes should take at points near or on boundaries. The goal is to attain or approach the interior convergence factors. For elliptic problems this is generally possible, since smoothing away from the boundary is decoupled from the boundary and since the boundary neighborhood itself is a lower-dimensional set of grid points, hence we can allow there more work (per point) than in the interior, without changing the total work by much. The comparison to the interior factors is a very important tool in debugging the program or finding conceptual mistakes, especially mistakes in treating boundary conditions or interior equations at points adjacent to boundaries. On the other hand, approaching the interior factors is not *all-important*; in fact, optimal performance of the full multigrid (FMG) algorithm may well be obtained without it, especially in non-elliptic or small-ellipticity problems (see end of Sec. 4.1).

In secs. 5.2-5.5 below we mention some rules related to the multigrid processes near or on boundaries. The general remarks of Sec. 11 and the curved-boundary treatment in Sec. 9.3 are also relevant here.

In addition to boundary conditions, some problems have global conditions. These should also be incorporated at this stage. Their multigrid implementation is discussed in Sec. 5.6.

5.1 Simplifications and debugging

It is advisable to start with a program for rectangular domains whose boundaries coincide with grid lines at all levels. This will make the programming much easier (the program in Sec. 1.5 can serve as a model), and will separate away various difficulties related to more general domains. In fact, the first stage in constructing such a program could be the case of periodic boundary conditions (separating away boundary considerations altogether) discussed in Sec. 4.

Having made rectangular models work satisfactorily, one can then proceed to other domains. At this point one has to decide whether to write a general-domain or a specific-domain program. Experience shows

general-domain multigrid programs to be considerably less efficient (typically requiring twice the CPU time). One can model one's general-domain program after MGO1 [S4, Sec. 10], or after MUGPACK, or actually use the MUGPACK or GRIDPACK software [M3]. But the efficiency of this software, too, is still considerably below the efficiency of specific-domain programs (where the efficiency of rectangular domains can be approached). The reason is the many checks that should be made to distinguish between various possible positions of gridpoints with regard to the boundary, especially in interpolation routines, where two grids are simultaneously involved.

It is advisable to start programming cycling algorithms, before proceeding to the additional questions related to the full multigrid (FMG) algorithm (taken up in Sec. 7). Cycling algorithms start with some arbitrary approximation on the finest grid and reduce its error by cycling between that grid and coarser grids. At this stage, one can avoid the question of what cycle to use: For debugging purposes it is best to start with comparing the theoretical two-level asymptotic convergence factor ($\bar{\lambda}$) with the experimental one by an algorithm which simulates a two-level algorithm. This is done by returning from the next coarser grid H back to the finest grid h only when the H equations have been solved to a very good accuracy (e.g., by taking large γ or very small δ in the cycles of Sec. 6.2). In this way we still separate away questions particular to too-coarse grids or related to three or more levels (delaying them to Sec. 6).

Another major simplification is to experiment first with particularly convenient known solutions. Even for complicated nonlinear systems, one can engineer the tested problem so that it has a *known* solution u . This is done simply by planting suitable right-hand sides, both for the interior differential equations and for the boundary conditions. (Even if the original problem is homogeneous, the program should anyway be written for *general* right-hand sides, because right-hand sides on *coarser* grids are obtained from finer grids residuals.)

For many nonlinear problems it is especially useful to experiment with solutions of the form $u = u_0 + \eta u$, where u_0 is a constant (or a constant vector, if u is a vector of functions), and $\eta \ll 1$ is employed in the first experiments. Taking u_0 as the first approximation for a cycling algorithm, the behavior of the solution process should almost identically (identically in the limit $\eta \rightarrow 0$) be the same as for a linear problem with constant coefficients. For such problems precise comparison can be made with mode analysis. The comparison can be pushed to be even more precise by choosing u_1 to be a particular Fourier component.

Afterwards, η can gradually be increased to study the effect of nonlinearity. This effect is in this way easily and precisely separated away from other effects, including effects which are often confused with nonlinearity because they appear in nonlinear terms; e.g., convection, whose magnitude in fluid problems depend on the solution itself. (For the algorithm used in the nonlinear cases -- see Sec. 8.)

Debugging of multigrid programs can generally benefit from relations between the levels. Most bugs and conceptual errors immediately show as irregular behavior in the standard multigrid output (listing the history of the dynamic residual norms for every relaxation sweep on every level, as in Sec. 1.5). A preliminary error-detection table, based on that output, is provided in [B9, Lecture 18]. *Troubles related to treatment of boundaries often show in the following way:* The first couple of cycles exhibit the expected (interior) convergence factor, since the relative weight of errors near the boundaries is small. Later, however, the errors near the boundaries start to dominate and the convergence factor degrades. The coarser is the basic (finest) grid, the sooner this degradation appears.

5.2 Interpolation near boundaries and singularities

The coarse-to-fine interpolation of corrections $I_{H,h}^{h,H}$ should use the boundary conditions on v^H even when they are not explicitly shown on the grid (sometimes they are only implicit in the program). Otherwise extrapolation formulae would be needed for $I_H^{h,H}$, giving slower asymptotic convergence factors [O1]. Exception is the case of discontinuity on the boundary, such as a boundary layer thinner than the mesh-size in which case boundary data should *not* be used (see Sec. 4.6).

Near boundary singularities, such as reentrant corners, the interpolation can be improved by using the asymptotic behavior, whenever known. That is, if the correction v^h to be interpolated from the coarser grid is expected to be of the form $v^h = w^h \psi$, where ψ is a known singular function and w^h is smooth, then polynomial interpolation should be used to interpolate w^h , not v^h . But such improvements are hardly needed (see Sec. 5.7).

5.3 Relaxation of boundary conditions

Except for some simple Dirichlet problems, discrete boundary conditions should generally be relaxed and transferred to the coarser grid in the same way interior difference equations do. It is important to notice that the boundary relaxation may spoil very much the smoothness of inte-

rior residuals near the boundary. Indeed, for a smooth error function, the interior residuals formed near the boundary by relaxing the boundary conditions are $O(h^{l-m})$ times the typical magnitude of other interior residuals, where m is the order of the interior differential equation and l is the order of the boundary condition (usually $l < m$).

One way around this difficulty is immediately realized by looking at the one-dimensional case. It is clear in that case that boundary conditions need not be relaxed at all. Their errors are not functions that can be smoothed out in any way; they are just isolated values, which can always very well be represented on the coarser grid. Analogously in higher dimensional cases, the role of relaxation should not be to impose the boundary conditions, but only to smooth their error along the boundary. Instead of Gauss-Seidel-type relaxation for the boundary condition $Bu = g$, say, one can make a Gauss-Seidel relaxation of the equation $\Delta_S Bu = \Delta_S g$, where Δ_S is an approximation to the Laplace operator along the boundary; e.g., in two-dimensional problems, $\Delta_S = \partial^2 / \partial s^2$, where s is the boundary arclength.

In practice this means that, instead of satisfying the given condition at each boundary point, we only change its error to be equal to an average of the errors at neighboring boundary points. This increases the above l by 2, making the perturbation to the interior smoothness negligible. In case the boundary smoothing factor is not as good as the interior one, a couple of boundary sweeps may be performed per each interior one.

Another way around the above difficulty is to ignore it and rely on more precise residual transfers (Sec. 5.4).

5.4 Residual transfers near boundaries

Relaxation seldom leaves quite smooth residuals near the boundaries, where the normal succession of relaxation steps breaks off (in lexicographic schemes, for example). And this is especially so when relaxation of boundary conditions is not done in the above (Sec. 5.3) manner. Thus, in many cases it is important that each individual fine-grid residual is correctly represented on the coarse grid. This is what we called *full* residual weighting. The full weighting near boundaries, and also near interfaces, is considerably more complicated than the interior full weighting (described in Sec. 4.4). This is because the influence of the residual on the solution depends on its distance from the boundary; e.g., in Dirichlet problems for m -order elliptic equations the influence is proportional to the $(m/2)$ -th power of the distance.

Thus the weight used in transferring a residual from a fine-grid point to a coarse-grid point depends on the distance of both points from the boundary. Near boundary corners the dependence is even more involved.

Hence, near boundaries the interior full-weighting rule (4.5) is modified to the requirement that

$$\sum_{x^H} (I_{h^H}^H r^h)(x^H) W^H(x^H) G(x^H) = \sum_{x^h} r^h(x^h) W^h(x^h) G(x^h) \quad (5.1)$$

is satisfied for any given $r^h(x)$, where $\sum f(x^H) W^H(x^H)$ and $\sum f(x^h) W^h(x^h)$ are discrete approximations, on grids H and h respectively, to the integral $\int f dx$ for any function f , and where $G(\xi)$ has the behavior of the Green function near the boundary. That is, for two neighboring ξ_1 and ξ_2 , the ratio $G(\xi_1) / G(\xi_2)$ roughly gives the ratio between the solutions of $Lu(x) = \delta_{\xi_1}(x)$ and $Lu(x) = \delta_{\xi_2}(x)$, with homogeneous boundary conditions. Usually one can take $G(\xi) = d_{\xi}^{\alpha}$, where d_{ξ} is the distance of the point ξ from the boundary, and $\alpha = m - l - 1$, where l is the order of the highest normal derivative in the neighboring boundary condition.

Relation (5.1) need not of course be kept very precisely. Residual weightings $I_{h^H}^H$ that deviate from it by 20% may easily still show the same convergence rates. Another way of deriving residual weighting near boundaries is by variational rules, like (4.10) in essentially-symmetric cases. And still other ways exist. It may all seem complicated, but, as explained in Sec. 11, it is in principle no more complicated than discretizing the original differential equations near the boundaries.

5.5 Transfer of boundary residuals

Residuals are defined and are transferred (with some averaging) to the coarser grid H , not only with respect to the interior equations, but also with respect to the boundary conditions. In order to do it in the right scale, the *divided* form of the boundary conditions (the form analogous to the differential conditions, without multiplying through by a power of h) should be used to calculate residuals, average them and transfer. For this purpose a clear conceptual separation should be made between boundary conditions and neighboring interior equations. Incorporating the former into the latter is often convenient to do, but it may easily lead to wrong transfers. (To do it right, one should assume the given boundary condition is incorporated on the finest grid, while the corresponding homogeneous condition is incorporated on all coarser grids.)

Even when correctly done, however, this is equivalent to imposing the boundary condition at relaxation, which, as explained in Sec. 5.3, will sometimes result in large neighboring residuals and hence slower convergence, unless more precise residual weighting is used.)

In symmetric problems one can consistently use the variational relation (4.10) without ever distinguishing between interior equations and boundary conditions, provided good interpolation I_H^h is defined. For some classes of problems this interpolation may be based on the difference equations, interior and boundary alike [D1].

5.6 Treatment and use of global constraints

In addition to boundary conditions many problems also specify some global conditions, such as integral relations, etc. For example, the pressure $p(x)$ in incompressible Navier-Stokes equations is determined only up to an additive constant; for its unique determination one should add an integral condition like

$$\int p(\underline{x}) d\underline{x} = 0 \quad (5.2)$$

(integrating over the entire flow field), or a pointwise condition such as

$$p(\underline{x}_0) = 0. \quad (5.3)$$

Both conditions are in fact "global" in the sense we like to consider here, even though (5.3) does not look so global: One should generally consider as global any single discrete condition which has a large global effect on the solution. Boundary conditions in one-dimensional problems are also of this type (cf. Sec. 5.3). The normalization condition $(u, u) = 1$ in eigenproblems is a nonlinear condition of this type. In that case, unlike (5.2) or (5.3), together with the global condition there also comes a global unknown, the eigenvalue. This is often the case. The one-dimensional boundary conditions are associated with unknown boundary values, which should be considered as global unknowns.

A very useful device is indeed to add global constraints to a problem to make it better posed, freeing as many global parameters. For example, a physical continuation parameter γ for solving a nonlinear problem (cf. Sec. 8.3.2) may be a bad parameter near some "limit points"; i.e., with γ fixed the problem is ill-posed (or nearly ill-posed, hence not approximable on coarse grids). By converting γ into an unknown and adding instead another global characterization of solutions (e.g., the

arclength along the continuation path, as in [K1]), the problem is made well posed (and, in particular, coarse grid well approximate fine-grid errors). Or one can free some accuracy parameters (e.g., allow a small constant error to be added to all equations) and add some known global information (e.g., total mass, or energy, or vorticity, when solving a time-step problem as part of a long-range evolution where such quantities must be strictly conserved). A particularly useful device is to free such accuracy parameters only on coarser grids, using global constraints only for coordinating solutions on those grids with the current fine-grid solution. In this way for example the above bad physical parameter γ can still be fixed on the finest grid. Such a device is important in controlling highly indefinite problem. Still another example of an added constraint with an added parameter (α) appears in Sec. 5.7.

An important advantage of multigrid processing is the easy and natural way with which such global conditions and global unknowns can be handled. To be sure, it is often done in a wrong way: For example, misguided by the practice in relaxation solvers, one would tend to treat (5.3) at the relaxation phase. Imposing such a pointwise global condition just by changing p at \underline{x}_0 is really harmful to the multigrid solution, since it frustrates the error-smoothing processes near \underline{x}_0 .

Global conditions need not be treated at all on the fine grid. There can be no error-smoothing related to such single conditions. All one has to do is to transfer the residual of the condition to serve as the right-hand side for a similar condition on the next coarser grid. In case of a nonlinear condition, FAS should be used (Sec. 8.3). A condition like $(u^h, u^h) = b^h$, for example, will be transferred to the condition $(u^H, u^H) = b^H$, where

$$b^H = b^h + (I_h^H u^h, I_h^H u^h) - (u^h, u^h), \quad (5.4)$$

which is a special case of (8.5). The global nature of a condition like (5.3) becomes increasingly transparent as it is transferred to coarser grids by proper approximations.

The global condition must of course be operated in solving the coarsest-grid problem (cf. Sec. 6.3). Sometimes it should be operated on several of the coarsest grids. For example, approximations to a condition like $(w, u) = b$, where w is a given weight function which changes signs in the domain, must perhaps be operated on a grid fine enough to resolve these sign changes (or at least crudely simulate them). Similarly, the condition $(u^H, u^H) = b^H$ should be operated on a grid fine enough to crudely resolve the sign changes in the solution u .

When a global condition is treated in relaxation (on a coarse but not the coarsest grid, for example) this should be done in a global way. For example, the condition should be satisfied (at the end of each sweep, say) by adding a constant (or a smooth function) to the entire solution, or by multiplying the entire solution by a constant (or a smooth function), so that the error-smoothing process is not frustrated.

There are sometimes conditions which are neither completely global nor quite local, but have some intermediate scale. For example, sometimes, when one global control would not do, several constraints are added, each controlling the solution on one subdomain. Any such condition should not be treated in relaxation wherever the meshsize is small compared with the scale of the condition.

In some relaxation schemes, the global condition seems to be needed in the local relaxation. For example, in the BGS schemes (Sec. 3.4) one solves in small boxes little problems similar to the given boundary-value problem. For the solution in the box to be uniquely determined, a condition like the global condition is needed there. In solving discrete incompressible Navier-Stokes equations in a small local box, for example, a pressure condition similar to (5.2) or (5.3) is needed. The best then is to use in each box a "no change" kind of condition. That is, to require, for example, that some discrete approximation to $\int p(x)dx$ (integration being over the small box) retains its value from before the relaxation step.

5.7 Structural singularities. Reentrant corners. Local relaxation

Structural singularities are singularities in the problem other than those caused simply by singular forcing terms (singular right-hand sides in either the differential equations or the boundary condition). Boundary reentrant corners and singularities in the (left-hand side) differential operators are structural singularities. Such singularities cause the asymptotic multigrid convergence factor to degrade (although it is still bounded away from 1 independently of the meshsize). The reason is that the error components slow to converge in relaxation, which are approximate solution to the homogeneous differential equations (cf. Sec. 4.6), have a singular behavior around the structural singularity, therefore they are not well approximated by a coarse-grid correction, unless this singular behavior is taken into account in formulating I_H^h , L^H and I_h^H (see Secs. 5.2 and 4.6).

It is easy to overcome this difficulty in several other, perhaps simpler, ways. First, the degradation of the asymptotic factors may not

be a two-level feature at all, appearing only because errors of the above type accumulate by cascading through many levels. Indeed, the degradation almost disappeared (in cases of reentrant corners studied in [B24]) when W cycles replaced V cycles. Secondly, even with V cycles, the degraded asymptotic factors little affect the ability of the FMG algorithm to solve the problem to within truncation errors by just one cycle per level. This is explained by the fact that the same singularity causing the degraded convergence also causes large truncation errors, in exactly the same components which are slow to converge. Thus, exactly these components need not converge much in order to be approximated below truncation error. Numerical experiment with reentrant corners ([O1], [B24]) show this to be true, unless the number of levels is very large indeed (as only in cases where local refinements are used).

Finally, a general way to deal with structural singularities, so that even the many-levels- V -cycle convergence factors are not degraded at all, is by *local relaxation sweeps*, i.e., special relaxation passes over a small number of points near the singularity. Experiments with reentrant corners [B24] show that just one such pass over a small number (typically less than 1%) of points before each full sweep is enough to completely restore the interior (i.e., regular) convergence factors.

Another question of course is how to obtain small truncation errors (i.e., good approximation to the differential solution) near structural (and other) singularities. Two general ways are local refinements and subtraction of the singularity. The former is a very general approach, discussed in Sec. 9; the latter can be used when the singular behavior of the solution u is known and simple. For example, near a reentrant corner $u = \alpha\psi + w$, where ψ is a known singular function (e.g., $\psi = r^Y \sin \gamma\theta$ in case of Poisson equation and a corner of π/γ radians), α is an unknown constant, and w is an unknown non-singular function. The procedure then is to rewrite the problem in terms of w . In the new problem α will appear as a global unknown, and a new constraint should be added to express the requirement that w is smooth at the singularity. This constraint is "global", since it controls the size of α , so its treatment, and the determination of α , should follow the rules in Sec. 5.6. When the singularity is subtracted off like this, the degraded convergence factor discussed above should disappear.

6. MANY-LEVEL CYCLES

Having obtained satisfactorily performing two-level cycling algorithms, one needs next to turn on the complete sequence of grids, using now the two-level techniques in recursion. The new algorithmic questions which arise are discussed below. Some of them could theoretically be investigated by three-level mode analysis, but this trouble is neither needed nor normally taken.

6.1 Multigrid cycles. Initial and terminal relaxation

For any grid h , finest or intermediate, a multigrid h-cycle can recursively be defined as follows: Make v_1 relaxation sweeps on grid h , then transfer the residual problem to the next coarser grid H ($= 2h$) and solve it there approximately, using γ H-cycles (unless H is the coarsest grid), then interpolate the grid- H solution and add it as a correction to the grid- h solution, and finally make v_2 more sweeps on grid h . On the coarsest grid the problem is solved either directly or by v_0 relaxation sweeps (cf. Sec. 6.3).

In two-level cycles only the sum $v = v_1 + v_2$ matters of course. When h is an intermediate grid the separate values of v_1 and v_2 do make some difference, although not a big one. In regular elliptic solvers experience shows that $v_2 = \lfloor v/2 \rfloor$ is probably the best prescription (see for example [S4, Tables 3.3]). In double-discretization schemes (Sec. 10.2) it is important to use $v_2 = 0$. In "accommodative" algorithms (see Secs. 4.1, 6.2) the values of v_1 and v_2 vary and they are determined internally.

Note also that the several passes of a complex relaxation sweep (such as ADZ) can be divided between the initial and the terminal stages of the cycle [S4, Sec. 8.2].

6.2 Switching criteria. Types of cycles

The criteria when to switch from a fine grid h to the next coarser grid $H = 2h$ were examined in a previous stage (Sec. 4.1). These same criteria can be used recursively, i.e., not only when h is the finest grid. We need in addition some criteria for switching from any grid H back to the next finer grid h . Two kinds of switches are used: Fixed and accommodative.

Fixed algorithms switch from H back to h after a preassigned number γ of H-cycles. The h-cycle is recursively defined to be of the type $C(v_1, v_2)^\gamma$, if all the H-cycles are of this same type. It is defi-

ned to be of type $F(v_1, v_2)$ if $\gamma = 2$ and the first H-cycle is an $F(v_1, v_2)$ cycle itself while the second H-cycle is a $C(v_1, v_2)^1$ cycle. See flowcharts and operation counts in [B17, Sec. 6.1]. The cycle $C(v_1, v_2)^1$ is also called a V cycle and denoted $V(v_1, v_2)$; see Fig. 1.1. The cycle $C(v_1, v_2)^2$ is also called a W cycle and denoted $W(v_1, v_2)$.

The W cycles are the safest, because the H-grid problem is always solved to a much better accuracy than the accuracy of the h-cycle, hence they perform practically as well as an exact solution of the H problem. Larger values of γ are not normally used, except for simulating two-level algorithms (see Sec. 5.1). F cycles are somewhat less expensive than W cycles in one-dimensional problems with many levels, and also in higher dimensions when semi coarsening is used, but otherwise they perform practically the same. V cycles may save considerable fraction (1/3, in two-dimensional problems) of the work. They are safe to use when the two-level convergence factor $\bar{\lambda}$ is small (e.g., $\bar{\lambda} \leq .15$, as in regular elliptic cases), in which case the convergence per V cycle will be close to $\bar{\lambda}$. If $\bar{\lambda} = .5$, on the other hand, a V cycle may not even attain that factor, because the coarse grid equations themselves will only be solved crudely, the error thus cascading through the levels. This situation arises in cases of severe singularities (cf. Sec. 5.7) or in non-elliptic and singular perturbation problems. In the latter cases, for example, the artificial viscosity on grid kh is k times larger than on grid h , hence visiting grid kh only once per cycle would give an asymptotic convergence factor no better than $1 - 1/k$ [B17, Sec. 5.1]. Since on the coarsest grid $k = O(h^{-1})$, the asymptotic factor will be $1 - O(h)$, which is very poor indeed. In this situation W cycles are absolutely necessary for good asymptotic factors. V cycles may however still work quite satisfactorily in FMG algorithms (see Sec. 7.4 below and the numerical experiments in [B17, Sec. 7.1]).

Accommodative algorithms switch from grid H back to grid h when a certain norm of the residuals on grid H drops below some factor η times the latest value of the corresponding norm on grid h . The parameter η is not a sensitive one. A good general prescription seems to be $\eta = 1.1 \bar{\lambda}$. If $\bar{\lambda}$ is not approximately known, take $\eta = 2^{-d}$, a value related to exchange-rate considerations (cf. Sec. 9.6).

Generally, accommodative algorithms may be troublesome at program development stages, since they may cause more complex interactions between the internal checks and the real questions one likes to examine. Their flexibility may prevent us from seeing some of the troubles, and

they are not suitable for precise comparisons. In the production stages, accommodative algorithms have the disadvantage that they require the extra work of calculating the residual norms. On the other hand, accommodative algorithms are more robust. Also, in complicated problems (which is where this robustness is needed), the residual norm calculation is inexpensive relatively to other calculations, assuming dynamic residuals (calculated anyway in the relaxation process) are used.

6.3 Coarsest grids. Inhomogeneous and indefinite operators

When the multigrid h-cycle performs considerably poorer than expected, it is first important to distinguish between fine-grid and coarse-grid troubles. This distinction is easy to make, by simulating two-level algorithms (taking large γ or small η) and examining whether this improves the convergence factor (per h-cycle), and how much this improvement depends on the size of h . Also examine whether reasonable convergence is obtained on your coarsest grid. If not, or if the trouble is confined to coarse h , the following remarks may be relevant.

Inhomogeneous operators are the main source for the special troubles appearing only on sufficiently coarse grids. On such grids, lower order terms of the operator start to affect, or even dominate, the smoothing and convergence factors. If we have neglected them in designing the fine-grid relaxation, we should now take them into account.

A typical example is the equation $-\Delta u + \sigma u = f$ with purely Neumann boundary conditions. If σ is positive but small, the smoothing factor of a GS relaxation is essentially the same as for Poisson equation, but the convergence factor is roughly $4/(4+h^2\sigma)$, which may be very slow even on the coarsest grid. Hence the coarsest-grid equations should be solved either directly (e.g., by elimination, which is inexpensive since the coarsest grid can contain just few points), or by relaxation, where after each sweep a suitable constant is subtracted from the approximate solution [A1, Sec. 4]. If $\sigma = 0$ everywhere except in some small subdomain, that constant subtraction should be employed on all grids which are not fine enough to resolve that small subdomain.

Indefinite case. If σ is negative, the situation is much worse, whatever the boundary conditions: For the coarse grid to approximate the slowly converging fine-grid component, its meshsize must be fine enough: For large $|\sigma|$, the coarsest meshsize must satisfy $H \leq O(R^{-1/p}(-\sigma)^{-5(p+1)/p})$, where R is the radius of the domain and p is the approximation order. In many cases this H is smaller than the finest affordable meshsize. To use coarser H , a global constraint

should be added (cf. Sec. 5.6). When $|\sigma|H^2 \gg O(1)$, different inter-grid transfers should be employed (cf. Sec. 4.2.2 and [B21]).

In designing the relaxation schemes for complex systems of equations, e.g. in fluid dynamics, we can take only subprincipal terms into account (Secs. 2.1, 3.1). On very coarse grids, however, this is no longer fully justified, and the smoothing factors may deteriorate. We may then have to use either more sweeps (by increasing ν and/or γ , or by using accommodative algorithms), or more sophisticated relaxation. In solving Navier-Stokes equations, for example, improved results were obtained by using the high-speed DGS scheme (see Sec. 19.3) on all finer grids, while employing BGS (see Sec. 3.4) on the two coarsest grids.

Even for homogeneous operators, convergence of h-cycles can sometimes be slower on very coarse grids, because the convergence factor $\bar{\lambda}$ cannot be smaller than $O(h^m)$; see (C) in Sec. 4.3. In such cases one can make more h-cycles, by increasing γ or switching accommodatively, which is inexpensive since h is coarse.

Sometimes troubles seen on coarse grids are only indications of bad procedures at some special, restricted regions, such as boundaries (see Sec. 5.1), or they may signal the need to operate some global conditions, which are not enforced on finer grids (see Sec. 5.6).

Of special concern is the coarsest grid itself. Relaxation there should be *converging*, not just *smoothing* as on other grids. Various conditions not enforced on finer grids must be enforced on the coarsest one, calling for special procedures. If nothing better is known, one can always use either a direct solver or a slow but safe iterative process such as Kaczmarz relaxation (cf. Sec. 1.1); on the coarsest grid they are affordable. Finally note that the coarsest grid cannot efficiently contribute to convergence if all its points happen to lie too close (much closer than a meshsize) to boundaries with Dirichlet boundary conditions.

7. FULL MULTI-GRID (FMG) ALGORITHMS

The cycling algorithms developed in the previous stages are easily converted into full multigrid (FMG) programs. The main difference is that instead of starting with an arbitrary approximation (e.g., $u_0^h \equiv 0$) on the finest grid, the first approximation u_0^h is obtained by an interpolation Π_H^h from a coarse-grid (approximate) solution u^H . Namely, $u_0^h \equiv \Pi_H^h u^H$, where $H = 2h$ and where u^H has been calculated by a similar FMG process with H as its finest level. The full algorithm can be either "fixed" (as for example in Fig. 1.2 above), or "accommodative" (as in [B8, Sec. 1.3], [Bl2, Fig. 1], [Bl1, Sec. 3.6 and Fig. 1], [Bl9, Sec. 2.2]). Both versions are available in the model program FMG1 [M3].

FMG algorithms are in a sense *easier* to program than cycling algorithms. Their main deriving routine is some lines longer, they may include an additional interpolation routine (Π_H^h), and they involve several more algorithmic questions (dealt with in the following subsections) -- but on the other hand they are much more *forgiving*. Their basic performance, which is to solve all problems to the level of truncation errors in just one or two cycles (see Sec. 7.3), is undisturbed by various little mistakes (conceptual mistakes or even programming bugs, especially in treating boundaries) which may degrade very much the *asymptotic* convergence of cycling algorithms. These mistakes may still be important in other situations, hence it is safer to detect them by perfecting the multigrid cycling (as in Secs. 4, 5 and 6) before turning FMG on. But it is important to understand all those cases in which, for good reasons, the FMG results are absolutely satisfactory despite the necessarily bad asymptotic convergence factors. Examples are numerous; some are emphasized in Secs. 3.3, 5.7, 7.4, 7.5 and 18.6. In many of these cases the analyses of FMG described in Secs. 7.4 and 7.5 can be useful.

It is also worth mentioning at this point that the FMG algorithm can incorporate into it continuation processes (see Sec. 8.3.2), grid adaptation processes (see Sec. 9.6), and, generally speaking, any process aimed at solving original "outer" problems (see Sec. 13).

7.1 Order of the FMG interpolation

The FMG interpolation operator Π_H^h is not necessarily the same as the correction interpolation operator I_H^h used in the multigrid correction cycles. Often the order of Π_H^h should be higher than the order of I_H^h , since the first approximation is smoother than the corrections: In the right-hand side of the latter (i.e., in the residuals) the amplitude of high-frequency components is usually comparable to that of low-frequency

quency components.

The optimal order of Π_H^h depends on the purpose of calculations. If one desires ultimately to get the algebraic error (i.e., the errors in solving the *difference* equations) to be very small (far below truncation errors), then Π_H^h should exploit all the smoothness of u^h in order not to produce unnecessary high-frequency algebraic errors. (High-frequency errors are the most expensive to liquidate in the multigrid cycling, since they are processed on the finest grid.) In fact, in such a case the first few correction interpolations I_H^h should also be of suitable high orders. The precise rules for scalar elliptic equations are given in [B7, App. A.2]. Note that these rules assume that the order of smoothness is known in advance.

Usually, however, the smoothness order is not known in advance. More importantly, we are not interested in solving to arbitrarily small algebraic errors; we like them only to be smaller than the truncation errors. The optimal order depends then on the norm by which we measure errors. Suppose we solve a $q \times q$ system of differential equations, and assume our error norm includes difference-quotients up to order ℓ_j in the j -th unknown function ($1 \leq j \leq q$). Then the order \hat{m}^j of the first interpolation of that function should not be less than $p + \ell_j$, where p is the order of approximation. Otherwise, the $O(h^{\hat{m}^j - \ell_j})$ high-frequency errors produced by interpolation would be much larger than the $O(h^p)$ (low-frequency) truncation errors.

In case of equations with strongly discontinuous coefficients, the higher-order interpolation Π_H^h should be of a different form, taking into account the different sense of smoothness in the solutions (cf. Sec. 4.6. A higher-order interpolation of this sort is presented in [A], Eq. (5.12)). The remarks of Sec. 5.2 apply here as well.

In some programs, especially general-domain programs, the higher-order interpolation Π_H^h turned out to cost more CPU time than the rest of the algorithm [O1]. An interpolation of an order smaller than indicated above may then be more practical. In case of rotatable differential operators, simpler higher-order interpolations can be used, based on the equations themselves [H3], [F3, Sec. 3].

7.2 Optimal switching to a new grid

In designing the FMG algorithm one should decide how well the equations on level $H = 2h$ should be solved before the solution is interpolated for the first time to grid h and the h -cycles start. The optimal point to switch is when the work of h -cycles becomes as efficient as the

work of H-cycles in reducing the *differential* error (the difference between the differential solution u and our current computed solutions, \tilde{u}^H or $u_0^H = \Pi_H \tilde{u}^H$). This happens when the algebraic error on grid H , namely $e^H = \|u^H - \tilde{u}^H\|$, is about 2^{-d} times the algebraic error on grid h , $e^h = \|u^h - \tilde{u}^h\|$, where d is the dimension, u^H is the exact solution of the H-equations and u^h is the exact solution of the h-equations. This is because h-cycles are about 2^d times as expensive as H-cycles. The switching point $e^H \approx 2^{-d} e^h$ is roughly equivalent to

$$e^H \approx \beta E^H, \quad \beta = (1-2^{-p}) / (2^d-1), \quad (7.1)$$

where $E^H = \|u^H - u\|$ is the truncation error on grid H and p is the order of approximation (see [B15, Sec. 5.2] and also [F3, App. A]).

In practice the values of e^H and E^H are of course not known, but we can derive from (7.1) the algebraic reduction needed on level H before switching. Namely, denoting by e_0^H the value of e^H when H-cycles are started and by e_*^H its value at the switching point (7.1), and assuming that the switching from the $2H$ -cycles to the H -cycles has been made when a relation similar to (7.1) was reached on level $2H$, we find [B15, Sec. 5.2] that the algebraic reduction on grid H is roughly

$$e_*^H / e_0^H \approx 2^{-p-d}. \quad (7.2)$$

This can be obtained by about

$$(p+d) / \log_2(1/\bar{\lambda}) \quad (7.3)$$

H-cycles, where $\bar{\lambda}$ is the convergence factor per cycles (Sec. 4.1). The switch (7.2) can also be used in an accommodative algorithm, replacing e_*^H/e_0^H by the corresponding ratio of residuals, which of course can be measured. The number of H-cycles usually turns out to be 1 or 2.

7.3 Total computational work. Termination criteria

Suppose that on the finest grid h we wish to obtain an algebraic error smaller than a specified factor α times the truncation error: $e^h < \alpha E^h$. Suppose also that the switch from level $H = 2h$ is made roughly when (7.1) is met; i.e., when $e^H \approx 2^p \beta E^h$. Then the algebraic error reduction required on grid h is roughly $\alpha_1 \approx \alpha(1-2^{-d})/(2^p-1)$. The number of work units to obtain such a reduction is about

$\log \frac{1}{\alpha_1} / \log \frac{1}{\mu}$, where μ is the interior convergence factor per work unit (see Sec. 4.1) and is usually just modestly larger than the interior smoothing factor $\bar{\mu}$. Counting also the work for the reduction (7.2) on coarser grids, we find that the total number of work units theoretically required by the Full MultiGrid algorithm is about

$$\left\{ \log \frac{2^p-1}{\alpha(1-2^{-d})} + \frac{p+d}{2^d-1} \log 2 \right\} / \log \frac{1}{\mu}. \quad (7.4)$$

The actual total number of work units is usually larger than (7.4), because of the need to make *integral* numbers of relaxation sweeps and coarse grid corrections. Typically one V or W cycle for each level, with $v = 2$ or 3 , yields e^h considerably below E^h .

The observation that in FMG algorithms one cycle on each level is needed, and is also basically enough to reduce the algebraic errors to the level of truncation errors (even though sometimes two shorter cycles may be more efficient), can heuristically be understood as follows. The first approximation on grid h , obtained by interpolating the grid- $2h$ solution, necessarily contains two types of errors: (A) High-frequency errors, i.e., errors having oscillations invisible and hence unapproximable on the coarser grid. (B) Aliasing errors, i.e., smooth errors introduced by high-frequency data because on the coarse grid high-frequency data is mistaken for smooth data. Relaxation on grid h can be efficient in removing the high-frequency errors (because of their local nature: At each point they are essentially determined by the neighboring residuals). Having removed the high-frequency errors we have also removed the high-frequency data from the residual problem, hence we can then go back to grid $2h$ with the residual problem to remove the aliasing errors (which are smooth errors, hence not determined by neighboring residuals, hence inefficiently treated by relaxation).

The algorithm may indeed be terminated after a fixed number of cycles on the finest grid h . This number is roughly $\log \frac{1}{\alpha_1} / \log \frac{1}{\bar{\lambda}}$, and in practice it is one or two. Or else, especially if an estimate for $\bar{\lambda}$ is not known, termination can be done when a certain norm of the residuals on grid h becomes smaller than a corresponding norm of $\alpha_1 E^h \approx \alpha(2^p-1)^{-1} \tau_{1h}^{2h}$ (see Sec. 8.4).

One should of course check numerically, using a problem with a known solution or a solution computed on a finer grid, that with these termination procedures $e^h < \alpha E^h$ is indeed obtained. Better still, one can observe the behavior of e^h as function of h . This can approxi-

mately be done even in (the real, production) runs where the solution is not known (see Sec. 1.6. And see more on this in Sec. 13). Quite often these checks reveal mistakes in the discretization schemes, not just in the multigrid algorithm.

7.4 Two-level FMG mode analysis

Instead of developing full multigrid (FMG) programs from the cycling programs, including boundary conditions, one can first develop the FMG algorithm still within the framework of two-level mode analysis (immediately following the stage of Sec. 4). This may again serve to separate away questions related to boundary conditions (questions discussed in Sec. 5), and questions related to many levels and to very coarse grids (Sec. 6) from the particular questions of the FMG algorithm (Secs. 7.1-7.3). The latter can then be examined in the interior, without boundary interference (or also with ideal boundaries -- see Sec. 7.5), and the performance figures so calculated can serve as ideals against which the actual program can be developed and debugged. Such an analysis is particularly useful in cases the usual two-level analysis (that of Sec. 4.1 above) is too pessimistic because of the existence of different components with markedly different convergence properties. For example, in case of nearly non-elliptic or nearly semi-elliptic problems there are smooth characteristic components which converge slower than others, since for such components L^H is not a very good approximation to L^h . But exactly for the same components and for the same reason, L^h itself is not a good approximation to L , hence these components do not need much algebraic convergence, and the fact that they have slow asymptotic rates does not matter [B17, Secs. 5.1, 5.2]. What we need then is an analysis which does not tell us the worst asymptotic rate, but tells us, separately in each mode, how well we solve the problem by an FMG algorithm with a given number of prescribed cycles.

To analyze the FMG solution of $Lu = f$, where L is a $q \times q$ system and has constant coefficients (or frozen local values, in case the original L was variable), we first analyze a single component $u(\underline{\theta}) = \exp(i\underline{\theta} \cdot \underline{x}/h)$. We calculate the corresponding f , and hence also the solution $u^H = u^H(\underline{\theta})$ to the coarse-grid equation $L^H u^H = f^H = I^H f$, where I^H is the local averaging used in our discretization for transferring a continuum function to grid H . The interpolation of u^H to the fine grid gives an approximation $u_0^h = \Pi_H^h u^H$ made up of 2^d Fourier components (the harmonics of $\underline{\theta}$, i.e. all components $\underline{\theta}'$ such that $\underline{\theta}' = \underline{\theta} + (v_1, \dots, v_d)\pi$, v_j integer and $-\pi < \theta'_j \leq \pi$). To the set of 2^d amplitudes we then apply the usual (cycling) two-level mode analysis

(Sec. 4.1); i.e., using (4.1) we calculate the $2^d q \times 2^d q$ matrix $M(\underline{\theta})$ describing the transformation of these 2^d amplitudes by one cycle. The result of applying k such cycles on grid h we denote by $u_k^h(\underline{\theta}) = M(\underline{\theta})^k u_0^h(\underline{\theta})$. Having calculated $u_k^h(\underline{\theta})$ we can then examine its qualities by several measures.

One measure is how well below truncation errors u_k^h is. This is measured for example by

$$\max_{|\underline{\theta}| < \pi} \frac{\|u_k^h(\underline{\theta}) - u(\underline{\theta})\|}{\|u(\underline{\theta}) - u^h(\underline{\theta})\|}, \quad (7.5)$$

where $u^h(\underline{\theta})$ is the exact solution of the grid h equations, and $\|\cdot\|$ is any norm under which we want to guarantee convergence. Note that u_k^h is made of 2^d components, while u^h and u are made of only one of those; the norms can be taken anyway.

Another, perhaps more direct and important measure, is how well we have solved the differential equations. That is, we directly measure $\|u_k^h - u\|$, thus evaluating not only the performance of our fast solver, but also the quality of our discretization scheme itself: We evaluate the total quality of our procedures in solving the differential equations at a given amount of work. In measuring the error $\|u_k^h - u\|$ we should of course give smaller weights to high-frequency $\underline{\theta}$'s; we cannot and need not solve for them as accurately as for low frequencies. Thus, if we aim at an approximation order p , good measures of performance are

$$\max_{|\underline{\theta}| < \pi} \|u_k^h(\underline{\theta}) - u(\underline{\theta})\| / |\underline{\theta}|^p, \quad (7.6)$$

or

$$\left\{ \int_{|\underline{\theta}| < \pi} |\underline{\theta}|^{-2p} \|u_k^h(\underline{\theta}) - u(\underline{\theta})\|^2 d\underline{\theta} \right\}^{\frac{1}{2}}, \quad (7.7)$$

etc. Several such measures can easily be produced, approximately, by the program that calculates $u_k^h(\underline{\theta})$. The program is an easy extension of the usual (cycling) two-level mode-analysis program.

All the issues examined by the two-level cycling analysis (relaxation, the number $v = v_1 + v_2$ of sweeps per cycle, and the interior operators I_h^H , L^H and I_H^h -- see Sec. 4) can further (more accurately) be optimized by the FMG mode analysis. In addition we can examine by this analysis the effect (in the interior) of various Π_H^h interpolation procedures, various values of v_1 , v_2 and k , and, most importantly, various interior discretization procedures. An important advantage is

that this analysis can be used even in cases where no algebraic convergence is desired or obtained (cf. Sec. 10.2).

Moreover, the predictions of the FMG-mode analysis are more robustly held than those of the cycling mode analysis when real boundaries are added. For example, take a problem with singularities in the boundary, such as reentrant corners. The effect of such singularities (similar to the effect of singular perturbations mentioned above) is to make the asymptotic convergence factors per cycle worse than predicted by the interior cycling mode analysis. But this occurs for particular components with large truncation errors (see Sec. 5.7), so that predictions like (7.5) are likely to be roughly held up.

A very simple example of two-level FMG mode analysis for a singular perturbation equation is given in [B17, Sec. 5.2]. A similar analysis holds for semi-elliptic cases.

7.5 Half-space FMG mode analysis. First differential approximations

Another advantage of the two-level FMG mode analysis is the possibility to make it also near a piece of boundary, modelled by a grid hyperplane, so that the entire domain is modelled by a half space. This is particularly important in non-elliptic or singular perturbation problems, where the high-frequencies far in the interior can still strongly be affected by the boundary.

A simple example is given in [B17, Sec. 5.3] for a singular perturbation equation. Its upshot is that both the algebraic error (after a one-cycle two-level FMG algorithm) and the truncation error increase as functions of the distance from the boundary, both eventually becoming as large as the solution itself; but at points where the truncation error is still small (compared with the solution), the algebraic error is smaller yet: the latter is at most a quarter the size of the former. (Again, a similar analysis can be made, with very similar results, for semi-elliptic equations.) Interior analysis could not of course describe this situation, and its relevance in such cases is therefore questionable.

Those examples in [B17] illustrate another technique which can be used whenever one wants to focus one's analysis on smooth components. One can then simplify the analysis very much by using the first-differential approximation (the first terms in a Taylor expansion) of the difference operators, instead of the difference operators themselves. For example, the first differential approximation to the 5-point Laplacian is the differential operator $\partial_{xx} + \partial_{yy} + \frac{1}{12} h^2 (\partial_{xxxx} + \partial_{yyyy})$. The ana-

lysis proceeds in the continuous domain, without mentioning grids except through the quantity h appearing in the operators.

PART II. ADVANCED TECHNIQUES AND INSIGHTS

Having completed the above stages of developing the basic multi-grid solver, one can start introducing various important improvements. Some possibilities are outlined in the following sections, followed by comments of more "philosophical" nature, which can however be readily useful to the practitioner. We then close with general remarks on multi-grid applications to chains of problems and to evolution problems.

8. FULL APPROXIMATION SCHEME (FAS) AND APPLICATIONS

The Full Approximation Scheme (or Full Approximation Storage - FAS) is a widely used version of multigrid inter-grid transfers, explained for example in [B6, Sec. 4.3.1], [B7, Sec. 5], [B8, Sec. 1.2], [B10, Sec. 2.1], [B14, Sec. 2.3]. It has mainly been used in solving nonlinear problems, but it has so many other applications that it should perhaps be used in most advanced programs. The scheme, its programming, and several of its applications are sketched below. Another, perhaps the most important application of FAS is described in Sec. 9, and yet another one in Sec. 15.

8.1 From CS to FAS

Consider first a linear problem $L^h u^h = f^h$ on a certain grid h , with some approximate solution u^h obtained, say, after several relaxation sweeps. Now we like to use coarser grids in order to supply fast approximations to smooth errors. Thus, it is the corrections $v^h = u^h - \tilde{u}^h$ which we try to approximate on the next coarser grid $H = 2h$. In the simpler multigrid programs, such as in Sec. 1.5, the coarse-grid unknown function is indeed v^h , intended to approximate the correction v^h . This multigrid version is therefore called the Correction Scheme (CS). Since $L^h v^h = r^h$, where

$$r^h = f^h - L^h u^h \tag{8.1}$$

is the fine-grid residual, the CS coarse-grid equations are

$$L^H v^H = I_h^H r^h \tag{8.2}$$

where L^H approximates L^h on the coarse grid. Once (8.2) has been approximately solved, its approximate solution \tilde{v}^H is interpolated to

the fine grid and serves as a correction to the fine-grid solution:

$$\tilde{u}_{NEW}^h = \tilde{u}^h + I_H^h \tilde{v}^H \tag{8.3}$$

In the Full Approximation Scheme (FAS) we perform exactly the same steps, but in terms of another coarse-grid variable. Instead of v^h we use

$$\hat{u}^H = I_h^H \tilde{u}^h + v^H \tag{8.4}$$

as the coarse-grid unknown function, where I_h^H is some fine-to-coarse transfer which need not be similar to I_h^H in (8.2). (They are in principle defined on different spaces.) This coarse-grid variable \hat{u}^H approximates $\hat{I}_h^H u^h$, the full intended solution represented on the coarse grid, hence the same "Full Approximation Scheme". The FAS coarse-grid equations, derived from (8.2) and (8.4), are

$$L^H \hat{u}^H = \hat{f}^H \tag{8.5a}$$

where

$$\hat{f}^H = L^H (I_h^H \tilde{u}^h) + I_h^H r^h \tag{8.5b}$$

Having obtained an approximate solution \tilde{u}^H to (8.5), the approximate coarse-grid correction is of course $\tilde{v}^H = \tilde{u}^H - I_h^H \tilde{u}^h$, hence the FAS interpolation back to the fine grid, equivalent to (8.3), is

$$\tilde{u}_{NEW}^h = \tilde{u}^h + I_H^h (\tilde{u}^H - I_h^H \tilde{u}^h) \tag{8.6}$$

To use directly

$$\tilde{u}_{NEW}^h = I_H^h \tilde{u}^H \tag{8.6'}$$

would be worse, of course, since it would introduce the interpolation errors of the full solution u^h instead of the interpolation errors of only the correction v^h (but see end of Sec. 8.5). Notice that $I_h^H \tilde{u}^h$ in (8.6) and in (8.5b) must be identically the same: A common programming mistake is to have a slight difference between the two. This difference may dominate the calculated correction function and hinder convergence. Also it is important to start with identically the same $I_h^H \tilde{u}^h$ as the first approximation in solving (8.5).

Note: The FAS equations (8.5) are a model for each interior dif-

differential equation or boundary condition appearing in the problem: each of them is transferred, separately, by this prescription. Other side conditions, global constraints, etc., are transferred in exactly the same way. See for example Eq. (5.4). In case a double discretization is employed on all levels (see Sec. 10.2), two functions \hat{f}^H should be calculated, one for each coarse-grid operator L^H [B17, Sec. 2.1].

For linear problems, the FAS steps (8.5-6) are fully equivalent to the CS steps (8.2-3). Indeed, the safest way to construct a correct FAS program is to start with a linear subcase, write first a CS program, then convert it to FAS and check that, at each iteration, the FAS results on the finest grid are identically the same as the CS results, except for round-off errors. Common mistakes in implementing FAS, especially in treating boundary conditions, are thus easily detected. The conversion of a CS program to FAS can be done by a trivial addition of three routines [B9, Lecure 12]. It can be done either for a cycling program or for an FMG program. The simplest examples are the cycling program FASCC and the program FMG1 [M3].

8.2 FAS: dual point of view

To see why FAS is preferable to CS in many situations, we rewrite (8.5), using (8.1), in the form

$$L^H \hat{u}^H = f^H + \tau_h^{2h} \quad (8.7)$$

where

$$\tau_h^{2h} = L^H(\hat{I}_h^H \hat{u}^h) - \hat{I}_h^H(L^h \hat{u}^h) \quad (8.8)$$

and $f^H = \hat{I}_h^H f^h$. Observe that (8.7) without the τ_h^{2h} term is the original coarse-grid equation (with the particular discretization $f^H = \hat{I}_h^H f^h$), and that \hat{u}^H approximates $\hat{I}_h^H u_{NEW}^h$, and at convergence $\hat{u}^H = \hat{I}_h^H u^h$. Hence τ_h^H is the fine-to-coarse defect correction, a correction to the coarse-grid equation designed to make its solution coincide with the fine-grid solution.

We can now reverse our point of view of the entire multigrid process: Instead of regarding the coarse grid as a device for accelerating convergence on the fine grid, we can view the fine grid as a device for calculating the correction τ_h^H to the coarse-grid equations. Since this correction depends on the non-smooth components of the solution, we obtain it by interpolating the solution to the fine grid and correcting

its non-smooth components by relaxation. Having obtained the correction τ_h^H by such a "visit" to grid h , we continue the solution process on grid H . Later we may "revisit" grid h , in order to update τ_h^H . In such a case the interpolation (8.6) should better be used if we do not want to lose the non-smooth components of the solution already obtained by relaxation in previous visits. This entire process will then yield a solution on grid h , which we can improve by inserting into it visits to grid $h/2$. Etc.

Since \hat{u}^H is just an improvement to u^H , we can omit the \hat{u}^H , and just understand that the meaning of u^H changes as soon as an approximation \tilde{u}^h exists on the next finer grid h .

This point of view, and the fact that the full fine grid solution is represented on all coarser grids, open up many algorithmic possibilities, as we shall see below.

8.3 Nonlinear problems

The Correction Scheme is not applicable to nonlinear problems, since the correction equation $L^h v^h = r^h$ is valid only for linear L^h . In case L^h is nonlinear, the correction equation can instead be written in the form

$$L^h(\tilde{u}^h + v^h) - L^h(\tilde{u}^h) = r^h \quad (8.9)$$

Transferring this equation to the coarse grid (replacing L^h by L^H , \tilde{u}^h by $\hat{I}_h^H \tilde{u}^h$, v^h by v^H and r^h by $\hat{I}_h^H r^h$) we get the FAS equations (8.5). Thus, one important advantage of FAS over CS is its direct applicability to nonlinear problems. (This is a general property of defect-correction schemes -- see for example [L1], [S3].)

The CS scheme can of course be applied to the Newton linearization of L^h around the current approximation \tilde{u}^h . But FAS is usually preferable, because:

(i) No global linearization is needed. The only linearization is the local one used in relaxation, and even this is seldom needed (see Sec. 3.4). Hence, no extra storage for coefficients of the linearized equation is required.

(ii) The programming is very convenient since the FAS equations (8.5) are exactly the same as the original discrete equations, except for a different right-hand side. Hence only one relaxation routine, one residual transfer routine, one boundary relaxation routine, etc., are ever needed in the program.

(iii) In the linearized problems, where sight is lost of the original problem, it is much more difficult to employ the most efficient multigrid approaches: One tends to solve far below truncation errors and therefore Taylor unnecessarily complicated and less vectorizable "perfect smoothers" (cf. Sec. 3.3), getting into unnecessary troubles connected with small ellipticity (cf. Secs. 7.4-5) or boundary singularity (Sec. 5.7). More importantly, one cannot then integrate the FMG algorithm with various processes outside the solver, such as continuation (see Sec. 8.3.2), grid adaptation (Sec. 9.6), and others (Sec. 13).

(iv) The FAS multigrid rate of convergence is not constrained by the convergence rate of Newton iterations. It is still mainly determined by the interior smoothing rate. Solving the nonlinear problem is no more expensive than solving, just once, the corresponding linearized problem. (In many cases, though, Newton convergence rate is fast enough to impose no real constraint to an FMG algorithm.)

(v) FAS is useful in various other ways. Particularly important for nonlinear problems are its applications in solving chains of problems (Sec. 15), near discontinuities (Sec. 8.5) and in automatic local refinement procedures (Sec. 9).

On the other hand, Newton linearization may still be preferred in those cases where it is an essential part of the discretization process, as in some finite-element formulations. This, however, may make such formulations less attractive in multigrid environments.

Although not employed in the FAS multigrid processing, we still use Newton linearizations in the local mode analysis, to estimate smoothing and convergence factors (Secs. 3, 4). See also Sec. 5.1 for a debugging technique related to nonlinear problems.

Examples: Various nonlinear problems have been solved by FAS, including transonic flow problems [S2], [J1], [M4]; steady-state compressible and incompressible Navier-Stokes equations (Secs. 19 and 20); the Bratu equation $\Delta u + \lambda e^u = 0$ (see Sec. 8.3.2); and complementary problems arising from free boundary problems. A simple example of the latter is to calculate the nonnegative function u which minimizes the Dirichlet integral $\int (\nabla u \cdot \nabla u + 2fu) dx$. Without the $u \geq 0$ constraint the problem is of course equivalent to a problem with the Poisson equation $-\Delta u = f$. But the nonnegativity constraint introduces nonlinearity. Using a FAS-FMG algorithm this nonlinear problem is solved with essentially the same amount of work as Poisson problems [B18].

8.3.1 Eigenvalue problems can simply be regarded as nonlinear problems. They are nonlinear since the unknown eigenvalue λ_j multiplies the

corresponding unknown eigenfunction u_j . Also, to fix the eigenfunctions nonlinear orthonormality conditions $(u_i, u_j) = \delta_{ij}$ are added as global constraints. The solution algorithm proceeds as a usual FAS-FMG multigrid, with the global constraints treated basically as in Sec. 5.6. The eigenvalues are updated once per cycle, together with a more precise determination of the individual eigenfunctions within the space spanned by them, by a Rayleigh-Ritz process. Experiments for model problems [B20] show that an FMG algorithm with one $V(2,1)$ cycle on each level gives a discrete eigenfunction with algebraic errors much smaller than truncation errors. Similar work is needed for each additional eigenfunction.

8.3.2 Continuation (embedding) techniques. Nonlinear problems usually have many discrete solutions, and the desired one is obtained only if we start from a close enough first approximation. A common way to obtain a good first approximation, or generally to trace branches of solutions, is by continuation (also called "embedding" or "Davidenko method"): The problem, including its discretization and its approximate solution is written as depending on some parameter

$$\gamma_0 \ll \gamma \ll \gamma_*$$

such that for γ_0 the problem is easily solvable (e.g., it is linear), while for γ_* it is the problem we really need to solve. We advance γ from γ_0 to γ_* in steps $\delta\gamma$ small enough to ensure that the solution to the γ problem can serve as a good first approximation in solving the $\gamma + \delta\gamma$ problem. Sometimes γ is a physical parameter; sometimes the solutions are better defined in terms of a non-physical parameter, such as the arclength of the solutions path [K1]; cf. Sec. 5.6.

As for the relation between multigrid and continuation, several situations arise. Sometimes, the FMG algorithm is a good continuation process by itself. In particular, in non-elliptic and singular perturbation problems where relaxation adds $O(h)$ artificial viscosity, the FMG algorithm starts from highly viscous solutions (since h is large) and gradually eliminates viscosity as it proceeds to finer grids. This is a natural continuation path since problems with large viscosity terms are well-defined and easier to solve. This continuation is carried even much further when the FMG algorithm is continued to include local refinements around thin viscous layers (see Sec. 9).

Sometimes, however, an explicit continuation over the problem path $\gamma_0 \ll \gamma \ll \gamma_*$ should be made, either because the intermediate problems are interesting themselves, or because they are necessary for reaching, or even defining, the desired γ_* solution. When the intermediate problems

are not of interest, they can of course be solved to a lower accuracy, using coarser grids only. The grids cannot all be too coarse, however; the meshsize \hat{h} must participate in the continuation process if components of wavelengths comparable to \hat{h} are needed to keep the solutions in the "attraction region" of the desired solution path.

Even when components with $O(\hat{h})$ wavelengths are needed in the continuation process, in each $\delta\gamma$ step they do not usually change much. We can therefore employ the "frozen τ " techniques described below (Sec. 15), and perform most of the $\delta\gamma$ steps using, on most parts of the domain, very coarse grids, with only few "visits" to grid \hat{h} : Such a continuation process will often require less computational work than the final step of solving (to a better accuracy) the γ_* problem.

A favorite example where these techniques are put to test is the Bratu problem $\Delta u + \lambda e^u = 0$ in the unit square, with $u = 0$ on the boundary. In collaboration with Klaus Stueben we have solved the problem using FAS, freezing λ and adding the value $u(P)$, where P is the center of the square, as a global constraint, treated basically as in Sec. 5.6. In this formulation, the multigrid solver had no problem going around the "limit point" ("turning point") of the solution curve (e.g., the curve of $u(P)$ as function of λ , which is not a unique function). It gave solutions at the same efficiency as corresponding algorithms for Poisson equations. In fact, we could solve by one-cycle FMG algorithm to the level of truncation errors, even problems on the upper branch of the solution curve, and even without continuation at all. The only region where more lengthy calculations were needed was a region where the discrete solution bifurcated, a phenomenon the algorithm was not designed to deal with efficiently. See more details in [S4, Sec. 5.5.1]. A further work along these lines, also to other problems, is reported in [S5].

8.4 Estimating truncation errors. τ extrapolation

As with other defect-correction schemes, the defect can serve also as an error estimator. That is, τ_h^H is an approximation to the local truncation error τ^H (on the coarse grid H), defined by

$$\tau^H = L^H(\hat{I}^H u) - I^H(Lu) \quad (8.10)$$

where u is the true differential solution and I^H and \hat{I}^H are two continuum-to- H transfer operators, defined as follows. I^H is the operator used in our grid- H discretization for $f^H = I^H f$, and \hat{I}^H represents the sense in which we want u^H to approximate u : We want u^H

to actually approximate $\hat{I}^H u$. The injection $(\hat{I}^H u)(x^H) = u(x^H)$ is usually meant, but other local averagings are sensible too.

Note the analogy between (8.10) and (8.8). τ^H is the correction to grid- H right-hand side that would make the grid- H solution coincide with the true differential solution $\hat{I}^H u$, while τ_h^H is the correction that would make it, at convergence, coincide with the fine-grid solution $\hat{I}_h^H u^h$. It is hence clear that at convergence

$$\tau^H \approx \tau^h + \tau_h^H, \quad (8.11)$$

where τ^h is the fine-grid local truncation error, defined with \hat{I}^h such that $\hat{I}^H = \hat{I}_h^H \hat{I}^h$. The sign \approx means equality up to a higher order in h . Relation (8.11) means, more precisely, that if $\tau^h + \tau_h^H$ were used to correct the H -equations, then u^H would be a higher-order approximation to $\hat{I}^H u$; namely, $u^H - \hat{I}^H u$ would equal $w^H - \hat{I}_h^H w^h$, where $L^H w^H = I_h^H \tau^h$ and $L^h w^h = \tau^h$.

Relation (8.11) can be used to inexpensively raise the approximation order. If the local approximation order (order of consistency) at the point x is p , i.e., if $\tau^h(x) \approx c(x)h^p$ where $c(x)$ is independent of h , then $\tau^H(x) \approx 2^p c(x)h^p$, hence $\tau_h^H(x) \approx (2^p - 1) c(x)h^p$, and hence $\tau^H(x) \approx 2^p (2^p - 1)^{-1} \tau_h^H(x)$. To raise the approximation order all we have then to do is to change the grid- H equations (8.5) by writing them as in (8.7) and multiplying τ_h^H by the fixed factor $2^p (2^p - 1)^{-1}$. This operation is called τ extrapolation. It resembles Richardson extrapolation, but it can profitably be done even in cases the latter cannot (e.g., in cases $p = p(x)$ is not constant), because it extrapolates the equation, not the solution. The τ extrapolation can be shown to be a special case of the higher-order techniques of Sec. 10.2 below, but it is especially simple and inexpensive. It costs only one multiplication on the coarser grid. It is probably best to use it in an FMG algorithm with $W(v,0)$ cycles, since a terminal relaxation with the lower order discretization would impair the approximation (even though its order would remain higher). An option for τ extrapolation exists in the model program FMG1 [M3].

Because of the analogy to the local truncation errors, τ_h^H is also called the relative local truncation error — the local truncation error of grid H relative to grid h . It is a by-product of the FAS processing which can be used to estimate the true local truncation errors: $\tau^h \approx (2^p - 1)^{-1} \tau_h^H$. Hence it can be used in FMG stopping criteria (see Sec. 7.3) and in grid adaptation criteria (Secs. 9.5, 15).

8.5 FAS interpolations and transfers

The consideration in determining the residual transfer I_h^H , the correction interpolation I_H^h and the FMG interpolation Π_H^h in FAS are basically the same as in CS, but there are some additional possibilities and we should also specify now the FAS solution transfer \hat{I}_h^H .

For linear problems the choice of \hat{I}_h^H does not matter of course; all choices will give identically the same results. The solution efficiency of many nonlinear problems is also insensitive to the exact choice of \hat{I}_h^H . The choice does matter only where the problem coefficients drastically vary over a meshsize. By the "problem coefficients" we mean those of the linearized problem. In practice, using FAS, we do not linearize the problem, but the transferred solution $\hat{I}_h^H u^h$ implicitly determines the problem coefficients on the coarse grid H — determining the coefficients may in fact be regarded as the purpose of this transfer (although, unlike the CS situation, the coefficients can change on grid H , with the changing approximation). When the problem coefficients are highly variable, it is important to have each coarse-grid coefficient a suitable average of neighboring values of the corresponding fine-grid coefficients. The coarse-grid problem, in other words, should be a proper "homogenization" of the fine-grid problem. Such homogenization is usually obtained by using full weighting for \hat{I}_h^H (as for I_h^H in Eqs. (4.5) - (4.6)).

In some, very special situations the dominant solution-dependent term in the coefficients may have the form $g(u)$, where g is a sensitive function; large changes in g are caused by more-or-less normal changes in u over a meshsize. In such a case the weighting \hat{I}_h^H should have the special form

$$\hat{I}_h^H \bar{u}^h = g^{-1} \bar{I}_h^H(g(\bar{u}^h)), \quad (8.12)$$

where \bar{I}_h^H is a normal full weighting, such as (4.6), $g(\bar{u}^h)$ is a grid-function such that $(g(\bar{u}^h))(x^h) = g(\bar{u}^h(x^h))$ for every fine-grid point x^h , and g^{-1} is the inverse function of g ; that is, $g^{-1}(g(\alpha)) = \alpha$ for any value α . If several sensitive functions such as g appear in the coefficients, several \hat{I}_h^H may correspondingly have to be used. (So far we have not seen a practical problem where this was required.)

An important possibility offered by FAS is the interpolation near an interior discontinuity, such as a shock or an interface. The grid- H solution, introducing smooth changes to the grid- h solution, may change

the location of such a discontinuity, moving it a meshsize or two. Near the discontinuity the correction $\bar{v}^H = \bar{u}^H - \hat{I}_h^H u^h$ will then be highly non-smooth; it will look like a pulse function. Interpolating it as a correction to the fine grid will introduce there unintended high oscillations. To avoid this, the FAS interpolation (8.6) should be replaced by (8.6') near the discontinuity. This is easy to implement, by adopting the following, more general rule.

Use (8.6) everywhere except near points where $\bar{u}^H - \hat{I}_h^H u^h$ is comparable to \bar{u}^h , where (8.6') should be used.

8.6 Application to integral equations

When the integral equation

$$\int K(x,y) u(y) dy = f(x,u(x)) \quad (8.13)$$

is discretized in a usual way on a grid with $n = O(h^{-d})$ points, the unknowns are all connected to each other; the matrix of the (linearized) discrete system is full. A solution by elimination would require $O(n^2)$ operations. An FMG solution would require $O(n^2)$ operations, since each relaxation sweep costs $O(n^2)$ operations. In case (8.13) is nonlinear in u , FAS-FMG would be used. Using the FAS structure, even for linear problems, we can often reduce this operation count very much, by exploiting smoothness properties of K .

In most physical applications $K(x,y)$ has a singularity at $y = x$. So usually $K(x,y)$ becomes either smaller or smoother as the distance $|y - x| = (\sum_1^d (y_j - x_j)^2)^{1/2}$ increases. "Smaller" and "smoother" are in fact related: The former is obtained from the latter by differentiations of (8.13) with respect to $x = (x_1, \dots, x_d)$ (possibly replacing the integral equation by an integro-differential equation). In either case, one can obtain practically the same accuracy in the numerical integration using meshsizes that increase with $|y - x|$, cutting enormously the work involved in relaxation. Usually $u(y)$ is much less smooth than $K(x,y)$ for large $|y - x|$. The integration with increasing meshsizes cannot then use point values of \bar{u}^h , but should use local averages of \bar{u}^h , taken over boxes whose size increases with $|y - x|$. Exactly such averages are supplied by the sequence of coarser grids in the FAS structure. The FAS solution transfers \hat{I}_h^H should of course represent full weighting. One can increase the accuracy of integration by using, in addition to the full-weighting averages, higher local moments, represented on additional coarser grids.

8.7 Small storage algorithms

Various effective methods for vastly reducing the storage requirement of the multigrid algorithm, without using external storage, can be based on the Full Approximation Scheme. One simple method [B10, §2.2] is to use the fact that a problem whose finest grid is h can satisfactorily be solved by an FMG algorithm with only one h -cycle (see Sec. 7.3). This means that only one visit is needed to grid h , including the FMG interpolation Π_h^H , a couple of relaxation sweeps, and the residuals and solution transfers I_h^H and \hat{I}_h^H , back to grid H . All these operations can be made "wave-like" by just one pass over grid h , requiring no more than few columns at a time kept in memory. (When the operations of one relaxation sweep have been completed up to a certain column, the operations of the next sweep can immediately be completed on the next column, etc.) This visit is enough to supply the corrected right-hand side \hat{f}^H on grid H (cf. Sec. 8.2), hence enough to calculate \bar{u}^H , without any storage allocated to grid h , except for the few mentioned continuously shifted columns.

\bar{u}^H is as precise as \bar{u}_{NEW}^h . The usual terminal sweeps of the h -cycle are only done if we need the solution on grid h , their role is to smooth the interpolation errors, not to reduce the error. Moreover, suppose that what we really need from our calculations is some functional of the solution, $\phi(u)$ say, so we would like to calculate $\phi^h(\bar{u}_{NEW}^h)$. All we have to do is to calculate, incidentally to transferring the solution \bar{u}^h back to grid H , the values of both $\phi^h(\bar{u}^h)$ and $\phi^H(I_h^H \bar{u}^h)$. Then, having later obtained \bar{u}^H , we can calculate

$$\phi^H(\bar{u}^H) = \phi^H(\bar{u}^H) + \phi^h(\bar{u}^h) - \phi^H(I_h^H \bar{u}^h), \quad (8.14)$$

which is practically as accurate as $\phi^h(\bar{u}_{NEW}^h)$, since $\bar{u}^H - I_h^H \bar{u}^h$ is small and smooth.

This simple procedure reduces the required storage typically by the factor 2^d , without increasing the computational work. Other procedures can reduce the storage much farther by avoiding the storage of coarser grids too, except for a certain $(n_k h_k) \times (n_k h_k)$ box on each grid $h_k = 2^k h$. The h_k box is shifted within the h_{k-1} box to supply the h_{k-1} corrections. The amount of work increases since on "re-visiting" the h_k box we need to reproduce its own h_k corrections. This can be done only in the interior of the box, distance $O(h_k |\log \epsilon|)$ from the boundary of the box, where $\epsilon = O(h^p)$ is the desired accuracy

on the finest grid, because closer to that boundary the h_{k+1} high-frequency errors are larger than the desired accuracy. Hence we must have $n_k \geq O(|\log \epsilon|)$. The overall required storage can therefore be reduced to $O(|\log \epsilon|^d |\log h|)$ (not just $O(|\log h|)$ as mistakenly calculated in [B7, Sec. 7.5]). Such procedures are called segmental refinement techniques.

Another small-storage multigrid algorithm, not based on FAS, is described in [H1]. It is a region dissection procedure, particularly suited for elongated domains.

9. LOCAL REFINEMENTS AND GRID ADAPTATION

Non-uniform resolution is needed in many, perhaps most, practical problems. Increasingly finer grids are needed near singularities, near non-smooth boundaries, at boundary layers, around captured shocks, etc., etc. Increasingly coarser grids are needed for exterior problems on unbounded domains. The multi-level FAS approach gives a convenient way to create non-uniform adaptable structures which are very flexible, allowing fast local refinements and local coordinate transformation, and whose equations are still solved with the usual multigrid efficiency. Moreover, the grid adaptation can naturally be governed by quantities supplied by the FAS multigrid processing, and it can naturally be integrated with the FMG algorithm to give increasingly better approximations to the true differential solution, at a fast, nearly optimal rate. These techniques, outlined below, are described in more detail in [B7, Secs. 7, 8, 9], [B8, Secs. 2, 3, 4] or [B10, Secs. 3, 4]. An application to three-dimensional transonic flows is described in [B23].

Another highly flexible discretization using a multigrid solver is described in [B2]. It is based on a finite element approach, which makes the program simpler, especially for complicated structures, but the execution is less efficient. The techniques outlined below are also applicable to finite element formulations as in [B11], [B9, Lecture 4].

9.1 Non-uniformity organized by uniform grids

Our non-uniform discretization grows from the simple observation that the various grids (levels) used in usual multigrid algorithms need not all extend over the same domain. The domain covered by any grid may be only a proper part of the domains covered by coarser grids. Each grid h can be extended only over those subdomains where the desired meshsize is roughly less than $2h$. In such a structure, the effective

meshsize at each neighborhood will be that of the finest grid covering it: see Fig. 9.1.




FIGURE 9.1: A piece of non-uniform grid (A) and the uniform levels it is made of (B, C, D, E). The heavy line shows a piece of the boundary, with a reentrant corner calling for the local refinements produced by the local patches (C, D, E).

This structure is very flexible, since local grid refinements (or coarsening) is done in terms of extending (or contracting) uniform grids, which is relatively easy and inexpensive to implement. A scheme named GRIDPACK [M3] has been developed for constructing, extending and contracting general uniform grids, together with many service routines for such grids, including efficient sweeping aids, interpolations, displays, treatment of boundaries and boundary data, etc. It is fully described in [B30]. One of its advantages is the efficient storage: The amount of logical information (pointers) describing a uniform grid is proportional to the number of strings of points (contiguous sets of gridpoints on the same gridline), and is therefore usually small compared with the number of points on the grid. Similarly, the amount of logical operations for sweeping over a grid is only proportional to the number of strings. Changing a grid is inexpensive too. One can easily add finer levels, or extend existing ones, thus effecting any desired refinements.

Moreover, this structure will at the same time provide a very efficient solution process to its difference equations, by using its levels also as in a multigrid solver. For this purpose the Full Approximation Scheme must be used, because in parts of the domain not covered by the finer grid h , the coarser grid $H = 2h$ must certainly show the full solution, not just a correction. Indeed, the FAS approach naturally fit here: We use on grid H the equations $L^H u^H = \hat{f}^H$, where \hat{f}^H is given by (8.5b) wherever $L^h u^h$ is well defined (i.e., in the interior of grid h), and $\hat{f}^H = f^H$ otherwise. In other words (cf. Eqs. 8.7-8), the fine-to-coarse correction τ_h^H is simply cancelled wherever it is not defined. Applying an FMG algorithm with these structures and equations, we will get a solution that in each subdomain will satisfy its finest-grid equations, while at interior boundaries (of fine levels not covering the entire domain) the solution will automatically be as interpolated from the coarser grid. Note that the coarse-grid solution is influenced by the finer grid even at regions not covered by the latter, since the coarse-grid equations are modified in the refined region.

In other words, a patch of the next finer level h can be thrown on any part of a given grid $H = 2h$, correcting there the latter's equations to the finer-grid accuracy. Moreover, several such patches may be thrown on the same grid. Some or all of the patches may later be discarded, but we can still retain their τ_h^H corrections in the grid- H equations.

An important advantage is that difference equations are in this way defined on uniform grids only. Such difference equations on equi-

distant points are simple, inexpensive and standard, even for high-order approximations, whereas on general grids their weights would have to be calculated by lengthy calculations separately for each point. Relaxation sweeps are also made on uniform grids only. This simplifies the sweeping and is particularly useful for line relaxation schemes.

9.2 Anisotropic refinements

It is sometimes desired to have a grid which resolves a certain thin layer, such as a boundary layer. Very fine meshsizes are then needed in one direction, namely, across the layer, to resolve its thin width. Even when the required meshsize is extremely small, not many gridpoints are needed, since the layer is comparably thin, provided, of course, that fine meshsizes are used only in that one direction. We need therefore a structure for meshsizes which get finer in one direction only.

In case the thin layer is along coordinate hyperplane ($x_j = \text{const.}$), this is easily done by semi refinements: Some levels H are refined by the next level h only in the j coordinate, $h_j = H_j/2$ whereas $h_i = H_i$ for $i \neq j$. See Fig. 9.2. In fact, different patches may have different refinement directions. Thus, the set of all grids is arranged logically in a tree, each grid having a unique next-coarser grid, but possibly several next-finer grids, instead of the former linear ordering. All these grids are still uniform, and can still easily be handled by GRIDPACK.

Note that the next-finer grids of a given grid H may geometrically have some overlap. All that is needed in such cases is to set priority relations, to tell which correction t_h^H applies at each point of grid H . Such priority relations are simply set by the order in which the corrections are transferred to grid H .

In case the thin layer is not along coordinate lines, the methods of the following sections could be used.

9.3 Local coordinate transformations

Another dimension of flexibility and versatility can be added to the above system by allowing each of the local patches to have its own set of local coordinates.

Near a boundary or an interface, for example, the most effective discretization is made in terms of coordinates in which the boundary (or interface) is a coordinate line. In such coordinates it is much easier to formulate high-order approximations near and on the boundary, and to introduce meshsizes which are much smaller across than along the boun-

dary layer (Sec. 9.2); etc. In the interior, local patches of coordinates aligned with characteristic directions (along streamlines, for instance) can greatly reduce the cross-stream numerical viscosity (cf. Sec. 2.1), thus yield superior approximations to non-elliptic equations.

Each set of coordinates will generally be used with more than one grid, so that (i) local refinements, isotropic or anisotropic, in the manner described above, can be made within each set of coordinates; and (ii) the multigrid processing retains its full efficiency by keeping the meshsize ratio between any grid and its next-coarser one properly bounded.




FIGURE 9.2: A piece of non-uniform, boundary-layer type grid (A) and the uniform rectangular subgrids it is made of (B, C, D, E). The meshsize in the local patches (C, D, E) is halved horizontally only.

Since local refinement can be made within each set of coordinates, the *only* purpose of the coordinate transformation is to provide the grid with the desired orientation, i.e., to have a given manifold (such as a piece of the boundary) coincide with a grid hyperplane. Since, furthermore, this needs to be done only locally, it can be obtained by a simple and standard transformation. For example, in two-dimensional problems, let a curve be given in the general parametric form

$$x = x_0(s), \quad y = y_0(s), \quad (s_1 < s < s_2) \quad (9.1)$$

where s is the arclength, i.e.,

$$x_0'(s)^2 + y_0'(s)^2 = 1. \quad (9.2)$$

To get a coordinate system (r, s) in which this curve coincides with the grid line $(r=0)$, we use the standard transformation

$$x(r, s) = x_0(s) - ry_0'(s), \quad y(r, s) = y_0(s) + rx_0'(s). \quad (9.3)$$

Locally, near $r=0$, this transformation is isometric (simple rotation).

The main advantage of this transformation is that it is fully characterized by the single-variable functions $x_0(s)$, $y_0(s)$. These functions, together with $x_0'(s)$, $y_0'(s)$ and $q(s) = x_0''/y_0' = -y_0''/x_0'$ can be stored as one-dimensional arrays, in terms of which efficient interpolation routines from (x, y) grids to (r, s) grids, and vice versa, can be programmed once for all. The difference equations in (r, s) coordinates are also simple to write in terms of these stored arrays, since, by (9.2-3),

$$\frac{\partial}{\partial x} = -y_0' \frac{\partial}{\partial r} + \frac{x_0'}{1+rq} \frac{\partial}{\partial s}, \quad \frac{\partial}{\partial y} = x_0' \frac{\partial}{\partial r} + \frac{y_0'}{1+rq} \frac{\partial}{\partial s}. \quad (9.4)$$

A different kind of multi-level procedure using a combination of cartesian grids and grids curved along boundaries is described in [S4, §11]. The main difference is that all levels, from coarsest to finest, are used there both for the cartesian and for the curved grids, and at each level the relaxation includes interpolations between the two types of grids, while the present approach is to regard the curved grids as a finer level which correct the finest cartesian grid near the boundary. The present approach is perhaps more economic and flexible, but it re-

quires a (crude) approximation to the boundary conditions to be given on the cartesian grids, too.

9.4 Sets of rotated cartesian grids

Another variant of this procedure is required in case the location of the thin layer (interface, shock, etc.) is not fully defined. For this purpose, each level will be a set of rotated cartesian grids, possibly overlapping. The finer the level, the finer (richer) is also the set of rotations. See Fig. 9.3. The self-adaptive criteria (see Sec. 9.5) can be employed to decide where to refine the set of rotations (together with refining the meshsize in one direction). Hence the scheme can capture discontinuities (thin layers), without defining them as such. The stronger the discontinuity, the better its resolution.

Since only rotated cartesian grids are needed in this scheme, the finite difference equations are as simple as ever. Hence this method is sometimes preferable even in cases where the location of the thin layer is known.

9.5 Self-adaptive techniques

The flexible organization and solution process, described above, facilitate the implementation of variable mesh-size $h(x)$ and the employment of high and variable approximation order $p(x)$. How, then, are mesh-sizes and approximation-orders to be chosen? Should boundary layers, for example, be resolved by the grid? What is their proper resolution? Should high-order approximations be used at such layers? How does one detect such layers automatically? In this section we survey a general multigrid framework for automatic selection of $h(x)$, $p(x)$ and other discretization parameters in a (nearly) optimal way. This system automatically resolves or avoids from resolving thin layers, depending on the goal of the computations, which can be stated through a simple function. (For more details see [B7, §8], [B8, §3]).

As our directive for sensible discretization we consider the problem of minimizing a certain error estimator E subject to a given amount of solution work W (or minimizing W for a given E . Actually, the control quantity will be neither E nor W , but their rate of exchange). This optimization problem should of course be taken quite loosely, since full optimization would require too much control work and would thus defeat its own purpose.

The error estimator E has generally the form

(A)




Fig. 9.3A. (Caption on the following page).

(B)




FIGURE 9.3: Grid orientation around an interior thin layer. The two coarsest levels (A) have the usual orientation 0 . The next level (B) has 3 orientations: 0 , $\frac{\pi}{4}$ and $-\frac{\pi}{4}$ (the later is not applied here). The next level (not shown) would have 7 orientations: 0 , $\pm\frac{\pi}{8}$, $\pm\frac{\pi}{4}$, $\pm\frac{3\pi}{8}$; etc. The successors (refinements) of a grid will always have either the same orientation or one of the two closest ones (e.g., each successor of the $\frac{\pi}{4}$ -oriented grid in B will have orientation $\frac{\pi}{4}$, $\frac{\pi}{8}$ or $\frac{3\pi}{8}$).

$$E = \int_{\Omega} G(x) \tau^h(x) dx, \quad (9.5)$$

where $\tau^h(x)$ is the local truncation error (cf. (8.10)) at x . $G(x) > 0$ is the error-weighting function. It should in principle be imposed by the user, thus defining his goal in solving the problem. In practice $G(x)$ serves as a convenient control. It is only the relative orders of magnitude of $G(x)$ at different points x that really matter, and therefore it can be chosen by some simple rules. For example, if it is desired to compute l -order derivatives of the solution up to the boundary then $G(x) \approx d_x^{m-1-l}$, where d_x is the distance of x from the boundary, and m is the order of the differential equation.

The work functional W is roughly given by

$$W = \int_{\Omega} \frac{w(p(x))}{h(x)^d} dx, \quad (9.6)$$

where d is the dimension and h^{-d} is therefore the number of grid-points per unit volume. (Replace h^d by $h_1 \dots h_d$ in case of anisotropic grids.) $w = w(p)$ is the solution work per grid-point. In multigrid processing, this work depends mainly on the approximation order (consistency order) $p(x)$. If the high-order techniques of Sec. 10 are used then usually $w(p) \approx w_0 p$, although sometimes $w(p) = O(p^3)$ for unusually high p (see Sec. 10.1).

Treating $h(x)$ as a continuous variable, the Euler equations of minimizing (9.5) for fixed (9.6) can be written as

$$G \frac{\partial \tau}{\partial h} - \lambda d w(p) h^{-d-1} = 0, \quad (9.7)$$

where λ is a constant (the Lagrange multiplier), representing the marginal rate of exchanging optimal accuracy for work: $\lambda = -dE/dW$.

In principle, once λ is specified, equation (9.7) determines, for each $x \in \Omega$, the local optimal values of $h(x)$, provided the truncation function $\tau^h(x)$ is fully known. In some problems the main behavior of $\tau^h(x)$ near singularities or in singular layers is known in advance by some asymptotic analysis so that approximate formulae for $h(x)$ can a priori be derived from (9.7). (Near source-type singularity (9.7) should be modified for that purpose, since τ^h has an essential and singular sign reversal at the source [B24].) More generally, however, equation (9.7) is coupled with, and should therefore be solved together with, the given differential equations. Except that (9.7) is solved to a cruder approximation. This is done in the following way:

In the FAS solution process we readily obtain the quantity τ_h^H . By (8.11) and (9.5), the quantity $-\Delta E(x) = G(x) \tau_h^H(x)$ can serve as an estimate for the decrease in E per unit volume owing to the refinement from H to h in the vicinity of x . By (9.6), this refinement requires the additional work (per unit volume) $\Delta W = w(p)h^{-d}(1-2^d)$. The local rate of exchanging accuracy for work is $Q = -\Delta E/\Delta W$. If Q is much larger than the control parameter λ , we say that the transition from H to h was highly profitable, and it pays to make a further such step, from h to $h/2$. So we will next establish grid $h/2$ in the neighborhood of x , as in any other neighborhood where there are points with $Q \gg \lambda$.

The computer work invested in the test is negligible compared with the solution work itself, since Q is calculated by just a couple of operations per point on the coarser grid H once per cycle.

A similar test can be used to decide on changing the local approximation order $p(x)$, with τ_h^H being replaced by the p_1 -to- p_0 defect-correction (10.1) and correspondingly $\Delta W = (w(p_1) - w(p_0))h^{-d}$. Or, if we treat p as a constant over the domain, but we like to optimize that constant, we can measure ΔE globally; i.e., measure directly the change in some quantity of interest (e.g., some functional of the solution we are most interested in), due to the transition from p_1 to p_0 . Correspondingly, global ΔW will be used. Whether locally or globally, the order will be increased beyond p_1 if $-\Delta E/\Delta W \gg \lambda$. Other discretization parameters, such as the computational boundaries (when the physical domain is unbounded), or refinements in grid orientations (see Sec. 9.4), can be decided by similar tests all based on comparing some $-\Delta E/\Delta W$ to the exchange-rate parameter λ . How to control λ and coordinate it with the solution algorithm is discussed in the next section.

9.6 Exchange rate algorithms. λ -FMG

Near a severe singularity many levels of increasingly finer grids on increasingly narrower subdomains may be needed, and formed by the above criteria. If the usual FMG algorithm were applied to these levels, too much work would be spent, since too many passes on coarser grids would be made. Only when all grids cover the same domain is it true that the coarse-grid work is small compared to the next-finer grid work, since the latter deals with about 2^d as many points. This is no longer so when local refinements are used: Finer grids may include less points than some coarser grids. The amount of work in a usual FMG algorithm would therefore be much greater than proportional to the total

number of gridpoints; (9.6) would not hold.

A better procedure is to decrease the accuracy-to-work exchange rate λ in a gradual sequence $\lambda_1 > \lambda_2 > \dots$, and to use the solution obtained for λ_{j-1} as the first approximation to the solution on the grids formed for λ_j . In the absence of singularities, and for uniformly p -order approximations, this process with the ratio $\lambda_j/\lambda_{j-1} = 2^{-p-d}$ would yield the regular FMG algorithm, so generally it is called λ -FMG. It was tested [B24] for the Poisson equation $\Delta u = f$, with severe singularity in f (e.g., $f = r^{-3.5}$, where r is the distance to a certain boundary point), or with 2π reentrant corner. The ratio $\lambda_j/\lambda_{j-1} = 1/16$ was used, and for each λ_j the collection of local refinements was determined by (9.7), using the roughly known behavior of r^h as function of the distance from the singularity. The 5-point Laplacian was employed, with red-black Gauss-Seidel relaxation. In the reentrant corner case, it was essential to use the local relaxation technique (Sec. 5.7). Results were invariably excellent: Algebraic errors $\|\tilde{u}^\lambda - u^\lambda\|$ smaller than the truncation errors $\|u^\lambda - u\|$ were obtained by one $V(1,1)$ cycle per λ . More importantly, the differential error $E = \|\tilde{u}^\lambda - u\|$ as function of the total work W (measured as the total number of points traversed in all the relaxation sweeps) behaves the same as in regular cases: In regular problems with second-order approximations on two-dimensional uniform grids $E = O(h^2) = O(W^{-1})$, and the above experiments indeed clearly yield $E = O(W^{-1})$, for several orders of increase in W . This confirms the validity not only of the λ -FMG algorithm, but also of the adaptation criterion (9.6). Had we used only uniform grids, the singularity would severely cripple $E(W)$; e.g., in the reentrant corner case it would yield $E = O(h) = O(W^{-1/2})$.

Switching criteria based on the exchange rate λ could also be used in the multigrid cycles themselves. Typically, the algorithm would continue relaxing on a given grid wherever $\hat{Q}(x) \equiv -G(x)\Delta r(x)/\Delta W > a_1\lambda$ where $-\Delta r(x)$ is the local decrease in the residual $|r(x)|$ per grid-point per sweep and ΔW is the corresponding work. Wherever $\hat{Q}(x) < a_1\lambda$ but $-G(x)|R(x)|/\Delta W > a_2\lambda$, a switch would be made to the next coarser grid. The local relaxation mentioned above (and in Sec. 5.7), could be effected by applying such criteria locally.

Ultimately, such exchange-rate criteria unify all the multigrid switching and adaptation criteria, integrating them into one algorithm, in which λ is gradually decreased. The process can be continued indefinitely, with increasingly finer levels created, globally and locally, in an almost optimal way. It can be terminated when either E or

W or λ reach preassigned limits. The above techniques of anisotropic grids, local transformations and rotations, and adaptation of approximation orders can all be integrated into such an exchange-rate algorithm.

10. HIGHER-ORDER TECHNIQUES

A sound way of constructing high-order approximations to a given differential problem $LU = F$, is first to construct a multigrid program with a low approximation order, and then convert it into a high-order program. The lower order is easier to develop and is also useful as a component in the higher-order program. Such programs are usually more efficient than programs which use high-order difference operators throughout. We mostly recommend the method of Sec. 10.2 below, especially for non-elliptic and singular perturbation problems.

10.1 Fine-grid defect corrections. Pseudospectral methods

Given a program for solving the (linear or nonlinear) low-order (order p_0) discrete system $L_0^h u^h = f^h$, an obvious multigrid approach for raising the approximation order is by high-order "deferred" (or "defect") corrections introduced once per cycle on the currently-finest grid [B10, §3.4]. That is, we add to f^h the correction

$$\sigma_{1,0}^h(x^h) = L_0^{h-h}(x^h) - L_1^h u^h(x^h), \quad (10.1)$$

where L_1^h is the higher-order operator, its approximation order (consistency order) being $p_1 > p_0$, and u^h is the current approximate solution. A similar correction is of course introduced to the discrete boundary conditions, too. To save h -cycles one should employ an FMG algorithm (Sec. 7), and use corrections like (10.1) at all the FMG stages (i.e., for every currently-finest grid). The total amount of work is then still basically given by (7.4). Note that that work is proportional to the approximation order p_1 . However, this count does not take into account the calculation of (10.1) once per cycle. For lower p_1 this extra work may be less than the other work within the cycle (a couple of sweeps on each level), but for high p_1 it becomes dominant and makes the amount of work per cycle proportional to p_1 (assuming spectral-type methods cannot be used and the complexity of calculating L_1^h is thus proportional to p_1), hence the total work is in principle $O(p_1^2)$. Furthermore, for higher p_1 we have in principle to use higher computer

precision, making the work of each arithmetic operation (in calculating (10.1)) again proportional to p_1 , bringing the total work to $O(p_1^3)$. This can be reduced to just $O(p_1^2)$ by a method of *compound deferred corrections*, i.e., taking $p_0 = p_1/2$ and solving for L_0^h by deferred corrections to a system of order $p_0/2$, and so on recursively. In the normally used range of p_1 , however, the work of even the uncompounded deferred correction is often still dominated by relaxation and hence still proportional to p_1 .

This technique can in particular be applied to pseudospectral approximations L_1^h , i.e., approximations attaining very high order (proportional to $1/h$) through a discrete spectral (Fourier, or Chebyshev, etc.) representation of the solution, using fast numerical transformers (e.g., FFT -- the fast Fourier transform) to obtain that representation and to calculate from it the approximate derivatives at gridpoints (cf. e.g. [G3]). Using reasonably high order in L_0^h (e.g., $p_0 = 4$, itself calculated by deferred corrections to a second-order operator) one can attain the spectral approximation order with just few applications of the spectral operator. (Other spectral multigrid methods are described in [B35], [S8], [Z1], [Z2].)

The deferred correction technique (suggested by L. Fox) is a special case of the concept of defect corrections (see [L1], [S3], [A2]). An important advantage of such a technique is that the higher-order operator L_1^h (and the corresponding higher order boundary conditions) need not be stable. This gives much freedom in the relatively difficult task of calculating L_1^h . This freedom is especially welcome in non-elliptic and singular perturbation cases, where convenient central approximations are unstable.

The reason L_1^h need not be stable is that the convergence of the defect correction iterations, to the solution corresponding to L_1^h , is fast only in the smooth components (for which L_0^h is a good approximation to L_1^h) and is very slow in the high-frequency components. Since instability is a property of high-frequencies, it can creep in only very slowly. The growth of unstable modes within the few cycles made is not too damaging.

The whole purpose of defect corrections is in fact to correct low-frequency components; only for such components higher-order approximations, such as L_1^h , are much better than lower-order approximations like L_0^h . Recognizing this and the fact that in multigrid processes low frequencies are converged via the coarse-grid corrections, we see that the main effect of the defect corrections can be obtained by applying them

only at the stage of transferring residuals to the coarser grids. This would save about two work-units per cycle, and would give better approximations in case L_1^h is unstable. This idea, from a different point of view, is described in the next section.

10.2 Double discretization: High-order residual transfers

On any given grid participating in multigrid interactions, discrete approximations to the continuous operator L are used in two different processes: in relaxation sweeps, and in calculating residuals transferred to coarser grids. The two discretization schemes need not be the same [B19, Sec. 3.11]. The discretization L_0^h employed in the relaxation sweeps must be stable (see Sec. 12), but its accuracy may be lower than the one we wish to generate. The discretization L_1^h used in calculating the transferred residuals determines the accuracy of our numerical solution, but it need not be stable. This "double discretization" scheme is especially useful in dealing with non-elliptic and singular perturbation problems: One can use the most convenient (but sometimes unstable) central differencing for L_1^h , and add artificial viscosities (see Sec. 2.1) only to L_0^h . This will ensure stable solutions which still have the accuracy of the central differencing.

Note that such a multigrid process will not converge to zero residuals, since it uses two conflicting difference schemes. The very point is, indeed, that the solution produced may be a better approximation to the *differential* solution than can be produced by either scheme.

The lack of algebraic convergence makes the usual two-level mode analysis irrelevant for double discretization schemes. Instead they can be analyzed by the two-level FMG mode analysis (Sec. 7.4).

Double discretization schemes can of course similarly be applied to boundary conditions; e.g., to Neumann conditions: Simple first-order schemes can be used in relaxation, while second-order Neumann conditions (which are sometimes complicated and may sometimes be unstable) can be used to transfer boundary-condition residuals to coarser grids.

The double discretization scheme need not be confined to the currently finest level; it can also be used on coarser levels. This will give better coarse-grid corrections, and hence faster algebraic convergence. (In non-elliptic and singular perturbation cases the algebraic convergence is usually determined by the quality of the coarse-grid correction [B17, §5.1].) It is also more convenient to program, since the same residual transfer routine, based on L_1^h , is used on all levels.

Moreover, if only L_0^h is used on coarser levels, the gain in approximation order per cycle cannot be more than p_0 ; hence the final approximation order cannot exceed $2p_0 + r_0$, where r_0 is the convergence order of relaxation [B17, §2.2]. Such a restriction does not exist if L_1^h is used for residual transfers on all levels. The approximation order p_1 can then be attained, perhaps even in one cycle, no matter how high p_1 is. In particular, pseudo-spectral approximations can be used in L_1^h , yielding very high approximation orders in few cycles.

In order to obtain the high approximation orders several rules should be observed: Suitable interpolation orders and residual-transfer orders should be employed. The right orders can be derived by crude mode analysis, as in Sec. 4.3, but with particular attention to boundary (see in particular rule (C) in Sec. 4.3). FMG algorithms with $W(v,0)$ cycles should be used (see Sec. 6.2), to ensure accurate enough solution of the course-grid equations and to avoid degradation of the approximation by terminal relaxation. Also, when double discretization is used on all levels together with the Full Approximation Scheme (see Sec. 8), notice that two different right-hand sides should be used on coarser grids, one for relaxation and a different one for residual transfers [B17, §2.1].

In case L_1^h is a better approximation than L_0^h not only for smooth components but also in the high-frequency range, the method of fine-grid defect corrections (Sec. 10.1) will eventually give smaller errors than the coarse-grid defect correction described here. But the gain will hardly justify the extra work involved in calculating (10.1) separately from the calculation of residuals. In problems where L_1^h is unstable, the present method is both faster and more accurate.

Double discretization schemes have already been used successfully in various cases, including fourth and sixth order approximations to Poisson equation [S1]; second-order approximations to simple singular perturbation problems [B17, §7], [B3, §7]; and second-order approximations to incompressible Navier-Stokes equations with high Reynolds numbers. Also, the λ extrapolation (Sec. 8.4) can be viewed as a special case, where $L_1^h = (2^{p_0} L^h - L^{2h}) / (2^{p_0} - 1)$.

1.3 Relaxation with only subprincipal terms

A particularly useful application of the above techniques is to employ a simple relaxation operator L_0^h where non-principal terms are neglected; more precisely, to employ the simplest stable L_0^h which approximate the subprincipal terms of the differential operator (see Sec. 2.1). Other terms need to be approximated only in L_1^h . For some fluid-dynamics systems this procedure can save a substantial amount of work. The techniques of either Sec. 10.1 or 10.2 can be used with this relaxation; more work is saved by the latter, but the former is safer. On very coarse grids this type of relaxation may give worse performance. In such cases use more sweeps or reintroduce the neglected non-principal terms.

1. COARSENING GUIDED BY DISCRETIZATION

The term "coarsening" is used here for the entire process of transferring a residual problem $L^h v^h = r^h$ from a fine grid h to the next coarser grid $H (= 2h)$. This includes the formulation of the coarse-grid problem $L^H v^H = r^H$, where the coarse-grid operator L^H and the fine-to-coarse transfer I_h^H should be determined both in the interior and near boundaries, and similar equations should be transferred for the boundary conditions themselves, and for any other side conditions the problem may have.

Any sufficiently general method of coarsening implies a discretization method, in the following sense: If the differential problem $u = f$ is discretized by any method, giving the problem $L^h u^h = f^h$ on grid h , and if this problem is then successively coarsened to $2h$, $4h$, $8h$, etc., by successive applications of the same coarsening method, then in the limit (for a sufficiently coarse grid) we obtain a discretization of $Lu = F$ which does not depend on the original discretization $L^h u^h = f^h$, but only on the method of coarsening. The limit discretization, in this sense, is the fixed point of the coarsening method. (In practice the limit is almost fully established after just a couple of coarsening levels).

This quite trivial observation has important implications. It implies that coarsening is at least as rich and difficult as discretization. It implies that controversies and competing techniques will emerge concerning coarsening similar to the ones in the field of discretization. Indeed such competitions have already surfaced. For example, the competition between finite-difference and finite-element methods, a dispute

which in fact consists of several separate issues: The variational derivation of discrete equations (or coarse-grid equations) vs. direct differencing; the interpolation issue (finite-elementers insist on using the same interpolation — the same "element" — as used in deriving the discrete equations, while finite-differencers allowing more freedom in interpolation, sometimes gaining higher accuracy in some error norms); the issue of general triangulation vs. uniform grids; and the issue of compactness of high-order approximations. These issues should not be confused with each other: variational derivation is possible and natural even without the use of elements [F4, §20.5]. Uniform grids can be used with finite-element solutions, too, changing the elements only near boundaries, a structure much more efficient computationally, especially in conjunction with multigrid methods [B11]. High-order compact operators arise quite naturally in the finite-element method, but such operators can also be derived by finite-difference approaches, such as the operator compact implicit method [C1] and also the Hodge method [L3].

All these and other issues arise as well with regard to coarsening, and the competing approaches are generally successful in coarsening wherever they are successful as discretization procedures — which is usually in problems where they are more natural. Variational approaches ([N2], [B7, App. A.5], [H2]) are natural for self-adjoint problems, and provided the most robust and automatic coarsening procedures for such problems [A1], [D1], although they can be replaced by much less expensive procedures (analogous to direct differencing) if the self-adjoint problem is not particularly complicated (cf. Secs. 4.5-4.6). In singular perturbation problems, such as those arising in fluid dynamics, discretization as well as coarsening are most successfully guided by physical understanding (artificial viscosity, upstream differencing, etc.). And so on.

The attempt to devise general fine-to-coarse transfers, good for all problems, is as hopeless (or as hopeful) as the attempt to have general, completely problem-independent discretization procedures.

But if this argument tells us how complicated coarsening can be, it also shows a general way to handle this difficulty. Namely: the coarsening method can always be guided by the discretization scheme.

Indeed, conversely to the statement above (that every coarsening implies discretization), one can say that every discretization scheme can be used to derive a coarsening procedure. This is done by imitation or analogy: Think about *discretizing* the problem $Lv = r$ on the coarse grid H ; then replace the operations done on the continuous domain by analogous operations done on the fine-grid h ; e.g., replace

integrations by analogous summations (or by integrations by elements, in case v^h is given in terms of finite elements). Galerkin discretization schemes, for example, are easily translated in this way into analogous coarsening formulae of the type (4.11) [N2, Sec. 3].

A coarsening procedure analogous to the finest-grid discretization scheme is called compatible coarsening. It is not necessary to use compatible coarsening, but it almost always makes a good sense to do so. In case the discretization scheme is a bad one, this would give a bad coarsening and hence slow asymptotic convergence rates of the multigrid cycling. But experience with several such cases (e.g., boundary singularities improperly treated) show that, if compatible coarsening is used, this slowness does not matter, because the source of slowness (bad discretization) is also, and for the same components, a source for larger truncation errors, hence an FMG algorithm (with the same discretization scheme on all currently-finest levels) still solves *below truncation errors* in the usual number of cycles (one or two, depending on the interior processes). Moreover, the slower *asymptotic* rates can in this way serve as a detector for the bad discretization, which otherwise may be passed unnoticed.

Compatible coarsening makes sense also from the point of view of computer resources and programming effort. For example, if a great generality and simplicity of programming is obtained by a discretization scheme (e.g., finite elements) which on the other hand spends a lot of computer time and storage to assemble the discrete equations and store them, the coarsening procedure can do the same since the time and storage it spends will be smaller than those already spent on the fine grid.

There are some special cases in which compatible discretizations are not quite available. These are cases where the discretization scheme is not general enough, because it specifically uses features of the finest grid not present on coarser ones. It uses for example a finest grid exactly laid so that its lines coincide with special lines of the problem, such as boundaries or lines of strong discontinuities (as in [A1]). In such situations compatible discretization is not well defined. To define it we must think in terms of a more general discretization scheme. (Again, the coarsening process serves to detect a certain flaw in discretization: In this case the flaw is the lack of generality.)

When double discretization is used (Sec. 10.2), compatible coarsening means the use of such a double discretization on coarser levels, too (as indeed recommended in Sec. 10.2).

12. REAL ROLE OF RELAXATION

The role of relaxation in multigrid processes has often been stated: It is to *smooth the error*; i.e., to reduce that part of the error (the "high-frequency" part) which cannot be well approximated on the next coarser grid. Some elaboration and clarification of this statement is important.

What is the "error" we want to smooth? It is usually thought of as the *algebraic error*, i.e., the difference $u^h - \tilde{u}^h$ between our calculated solution \tilde{u}^h and the discrete solution u^h (the exact solution to the discrete equations). However, in view of the double discretization scheme (Sec. 10.2), where u^h is not well-defined, it becomes clear that what relaxation should really do is to smooth the *differential error*, i.e., the difference $u - \tilde{u}^h$, where u is the solution to the given differential equations. In fact, this is the true role of relaxation even when double discretization is *not* used, if what we want to approximate is u , not u^h : It is the smoothness of $u - \tilde{u}^h$ which permits its efficient reduction via the coarser grid.

Thus, the important measure of relaxation efficiency is not the algebraic smoothing factor $\bar{\mu}$, but the differential smoothing factor, the factor by which the high-frequency part of $u - \tilde{u}^h$ is reduced per sweep. This is not usually recognized because the latter factor is not constant: It approximately equals $\bar{\mu}$ when the high-frequencies in $u - \tilde{u}^h$ are large compared with those in $u - u^h$ (where u^h is the local solution to the discrete equations employed in relaxation), but below this level $\bar{\mu}$ may mislead, and when \tilde{u}^h is closer to u than to u^h in their non-smooth components, the differential factor may even be larger than 1. For example, in solving a singular perturbation problem with strong alignment (see Sec. 2.1), we can reduce the algebraic smoothing factors of point Gauss-Seidel relaxation by taking a larger artificial viscosity and, more importantly, by taking it *isotropically*, instead of anisotropically. This would not however improve the overall performance of our double-discretization FMG algorithm (see the experiments in [B17, §7]), since it would not reduce the *differential* smoothing factors.

The differential smoothing is the purpose of relaxation not only on the finest grid h_* . On any grid h , it's relaxation reduces its range of high frequencies in the error $u - \tilde{u}^{h*}$, where we interpret changes in \tilde{u}^h as changes in \tilde{u}^{h*} via the interpolation relations.

We can here also elaborate on what are those "high-frequency components" (of the differential error) that should be converged by relaxation on grid h . Generally speaking we say that these are the components

12. REAL ROLE OF RELAXATION

The role of relaxation in multigrid processes has often been stated: It is to *smooth the error*; i.e., to reduce that part of the error (the "high-frequency" part) which cannot be well approximated on the next coarser grid. Some elaboration and clarification of this statement is important.

What is the "error" we want to smooth? It is usually thought of as the *algebraic error*, i.e., the difference $u^h - \tilde{u}^h$ between our calculated solution \tilde{u}^h and the discrete solution u^h (the exact solution to the discrete equations). However, in view of the double discretization scheme (Sec. 10.2), where u^h is not well-defined, it becomes clear that what relaxation should really do is to smooth the *differential error*, i.e., the difference $u - \tilde{u}^h$, where u is the solution to the given differential equations. In fact, this is the true role of relaxation even when double discretization is *not* used, if what we want to approximate is u , not u^h : It is the smoothness of $u - \tilde{u}^h$ which permits its efficient reduction via the coarser grid.

Thus, the important measure of relaxation efficiency is not the algebraic smoothing factor $\bar{\mu}$, but the differential smoothing factor, the factor by which the high-frequency part of $u - \tilde{u}^h$ is reduced per sweep. This is not usually recognized because the latter factor is not constant: It approximately equals $\bar{\mu}$ when the high-frequencies in $u - \tilde{u}^h$ are large compared with those in $u - u^h$ (where u^h is the local solution to the discrete equations employed in relaxation), but below this level $\bar{\mu}$ may mislead, and when \tilde{u}^h is closer to u than to u^h in their non-smooth components, the differential factor may even be larger than 1. For example, in solving a singular perturbation problem with strong alignment (see Sec. 2.1), we can reduce the algebraic smoothing factors of point Gauss-Seidel relaxation by taking a larger artificial viscosity and, more importantly, by taking it *isotropically*, instead of anisotropically. This would not however improve the overall performance of our double-discretization FMG algorithm (see the experiments in [B17, §7]), since it would not reduce the *differential* smoothing factors.

The differential smoothing is the purpose of relaxation not only on the finest grid h_* . On any grid h , it's relaxation reduces its range of high frequencies in the error $u - \tilde{u}^{h*}$, where we interpret changes in \tilde{u}^h as changes in \tilde{u}^{h*} via the interpolation relations.

We can here also elaborate on what are those "high-frequency components" (of the differential error) that should be converged by relaxation on grid h . Generally speaking we say that these are the components

"invisible" on the next coarser grid H , i.e., Fourier components $\exp(i\theta \cdot \underline{x}/h)$ which on grid H coincide with lower components, that is to say components with $\pi < \max_j |\theta_j| H_j/h_j$, $\max_j |\theta_j| \leq \pi$ (cf. Eq. (3.3)).

More precisely we should include in the "high-frequency" range all those components that are not efficiently reduced by relaxation on the other grids, which can for example be any range of the form

$$\{(\theta_1, \dots, \theta_d)\}: \alpha_j h_j / H_j \leq \theta_j \leq \alpha_j \text{ for at least one } j \quad (12.1)$$

where each α_j is fixed (assuming H_j/h_j is the same for all levels), $0 < \alpha_j \leq \pi$. That is, we can allow some of the highest frequency components on any intermediate level not to converge efficiently by relaxation at that level ($\alpha_j < \pi$), as long as those components efficiently converge by the next-finer-level relaxation. This may leave the highest frequencies on the finest grid uncontrolled, but they are unimportant and can be eliminated by averaging the final results. Examples where this further understanding of mode analysis is relevant are mentioned in Sec. 18.6 and in [B17, Sec. 5.7].

The range of frequencies to be reduced by relaxation can also be modified by modified coarse-grid functions of the type mentioned in Sec. 4.2.2. In such cases relaxation may not reduce some high-frequency error components; but the unreduced components are very special ones, hence they are described by few parameters. This is a general property of relaxation (see Sec. 1.1). Very generally we can thus say that the role of relaxation is to reduce the *information content* of the error, so that it becomes approximable by a lower dimensional approximation space.

Another important point to clarify is that relaxation should be efficient only as long as the high-frequency error components have relatively large amplitudes: When the high-frequency errors are too small compared with the low-frequency ones, relaxation cannot usually be efficient because of certain *feeding from low to high components*. Such feeding is caused by interaction with boundaries, and by non-constant coefficients, and by the high-frequency harmonics generated when the low-frequency error is corrected via the coarse-grid cycle (see observation (D) in Sec. 4.3). Sometimes such feeding is even caused by the interior relaxation itself; e.g., red-black relaxation of an order- m differential equation produces $O(h^m)$ high-frequency errors from $O(1)$ low-frequency errors. When the size of high-frequency amplitudes approaches the size fed from low frequencies, relaxation should be stopped; this is the point where the coarse-grid correction should be made. If relaxation is stopped in time, then the range of strong interactions with low-frequencies is

not entered. It is also only then that the multigrid convergence rates can accurately be deduced from the smoothing-rate analyses.

Finally, even though smoothing is the main role of relaxation, we should not forget its influence on other components. Some relaxation schemes with extremely good smoothing factors are either unstable or they cause large amplification of some low-frequency errors (see Sec. 3.2).

We can thus say in summary that the role of relaxation is to reduce large amplitudes of certain (usually high-frequency) components of the differential error, while avoiding from significantly amplifying its other components.

Stability of the difference equations used in relaxation is only a tool in performing this role, not an end by itself.

13. DEALGEBRAIZATION OF MULTIGRID

An interesting line in the development of multigrid can be viewed as a gradual "dealgebraization", a gradual liberation from algebraic concepts, and the development of methods that increasingly exploit the underlying *differential* nature of the problems. We'd like to briefly trace this line here, so as to bring out some concepts useful in practical implementations.

As the first step of dealgebraization we can regard the replacement of "acceleration" by "smoothing". The early two-grid and multi-grid approach viewed coarse-grid corrections mainly as a tool for accelerating the basic iterative process - the fine-grid relaxation. Only later it became clear that the only role of relaxation is to smooth the error. (cf. Sec. 12, where a further "dealgebraization" of the smoothing concept is described.) This slight shift in understanding revolutionized the multigrid practice: It made it clear that the fine-grid process is basically local, hence analyzable by local mode analysis. This understanding, together with that analysis, produced the truly efficient multigrid cycles, in which very few sweeps are made on each grid before switching to coarser ones, and in which the fine-to-coarse meshsize ratio assumes the (practically) optimal value of 1:2.

The next dealgebraization steps are related to the trivial understanding that we are not primarily interested in solving the algebraic equations (obtaining u^h), but we are interested in approximating the differential solution u . First, this implies that we have to solve the algebraic equations only "to the level of truncation errors", i.e., only

to the point that our calculated solution \tilde{u}^h satisfies $\|\tilde{u}^h - u^h\| \approx \|u - u^h\|$; further reduction of $\tilde{u}^h - u^h$ is meaningless.

This implies that the asymptotic convergence rate of the multigrid cycle is not important by itself. What counts is the amount of work we need in an FMG algorithm in order to reduce the error from its original value on grid $2h$, which is approximately $\|\pi_{2h}^h u^{2h} - u\| \approx \|u^{2h} - u\|$, to the desired level $\|u^h - u\|$. This is a reduction by a factor of 2^p only, which can usually be achieved in one cycle. (The fundamental reason for this is again non-algebraic: See Sec. 7.3.) Evidently it is then more relevant to think in terms of the FMG analysis (Sec. 7.4) than in terms of asymptotic rates.

Even the later viewpoint, that we want to reduce the errors to the level of truncation errors, is too algebraic-oriented. It is tied too much to one given discretization on one given grid. The optimal moment of switching from a certain currently-finest grid H to a new, finer grid $h = H/2$ is not necessarily when $\|\tilde{u}^H - u^H\| \approx \|u - u^H\|$. Rather, it is determined by comparing H -cycles to h -cycles in their efficiency at driving \tilde{u}^h closer to u (see Sec. 7.2). What really counts is the behavior of the differential error $E = \|\tilde{u}^h - u\|$ as a function of the total accumulated computational work W .

We want $E(W)$ to be as fast-decreasing as possible.

From this as our objective we can derive correct switching criteria, i.e. decide when to establish a new finer grid. The next step is to realize that criteria based on $E(W)$ can be applied locally, to decide not only when to have a finer grid, but also where to have it. This naturally brings us to grid adaptation (Sec. 9.5). Indeed one can integrate the switching and self-adaption criteria (discussed in Secs. 6.2, 7.2, 9.5) into a total multi-level adaptive algorithm, where switching between levels and creating new, or extending existing, levels are all governed by the same exchange-rate criteria (see Sec. 9.6).

Another step away from fixed algebraic concepts is to allow variable discretization schemes, i.e., schemes which can be changed throughout the algorithm to promote faster decreasing $E(W)$. This includes the use of higher-order, variable-order and adaptable order schemes, governed again by $E(W)$ criteria (see Sec. 9.5 and [Blö, §3.6, 4.3]). Using different discretization schemes in relaxation and in residual transfers (Sec. 10.2) is a further step in that direction.

By now we have gone quite far beyond the notion of multigrid as just a fast algebraic solver, toward viewing it as a total treatment of

the original problem. This is proved to be a very beneficial general principle: Always think of multigrid in terms of as original a problem as possible: For example, instead of using Newton iterations, employing multigrid as a fast solver of the linearized problems, apply multigrid directly to the non-linear problem (Sec. 8.3). Instead of solving an eigenproblem by the inverse power method, with multigrid as the fast inverter, you can multigrid directly the original eigenvalue problem (Sec. 8.3.1). Instead of using multigrid for solving each step in some outer iterative process — be it a continuation process, a time-dependent evolution, a process of optimizing some parameters or solving an inverse problem, etc. — apply it directly to the originally given problem (cf. Secs. 8.3.2, 15, 16). Instead of a grid adaptation process where the discrete problem on each grid configuration is completely solved (by multigrid, say) and then used to decide on an improved grid configuration, the whole adaptation process can be integrated into a multigrid solver (Sec. 9.6). And so on.

An illustration to this approach is the solution of optimization problems, where the parameter to be optimized is some continuum function on which the solution u depends. This "parametric function" may for example be the shape of the boundary (e.g., the shape of an airplane section which we want to optimize in some sense), or a certain coefficient of the differential equations (e.g., in inverse problems, where one tries to determine this coefficient throughout the domain so that the solution will best fit some observational data), etc. Multigridding the original problem means that we solve it by some FMG algorithm, where already at the coarser FMG stages we treat the given optimization problem, by optimizing a coarser representation of the parametric function. On the finer grids, incidentally to relaxing the equations, we optimize that function locally (when this makes sense), and then we introduce smooth corrections to the function during the coarse-grid correction cycles. Instead of using the multigrid solver many times, we may end up doing work only modestly larger than just one application of that solver.

13.1 Reverse trend: Algebraic multigrid

Contrary to the above line of dealgebraization, there is a recent trend to develop purely Algebraic Multi-Grid (AMG) algorithms. By this we mean a multi-level algorithm without any geometry, without grids. An algebraic (linear or nonlinear) system of equations is given. To solve it fast, a sequence of increasingly "coarser" levels is created. A coarser level in this context is a related, but much smaller, algebraic sys-

tem. The choice of the coarse-level variables, and of the coarse-to-fine interpolation I_H^h , is based not on geometric positions but on the algebraic equations themselves: The coarse variables are chosen so that each fine-level variable is strongly coupled to one or more of them, and each I_H^h coefficient can for example be chosen to be proportional to the corresponding coupling strength. The fine-to-coarse transfer and the coarse-level matrix are then constructed by prescriptions like (4.12) and (4.11), respectively. The theoretical background directing the various choices is developed in [B26].

Generally, the efficiency that can be achieved by such algebraic algorithms is below that of algorithms built to exploit the geometric information, let alone the further efficiency obtainable by further de-algebraization. On the other hand these algebraic solvers may be used as black boxes for larger classes of problems. They may especially be useful in cases where the geometrical information is too complicated, such as finite-element equations based on arbitrary partitions, or various problems which are not differential in their origin but still lend themselves for fast multi-level solutions. Also, there are cases of finite-difference equations on a uniform grid, in which the usual geometric choice of coarse-grid variables is not good, since too many fine-grid variables happen to depend too weakly on the coarse-grid variables (cf. e.g. [A1, Sec. 8]). Algebraic multigrid can then perform better. Because of its sensitive coarsening, there is in AMG no need for special relaxation schemes, in varying block and marching directions (cf. Sec. 3.3); simple Gauss-Seidel is for example used for all definite problems. Experiments on a wide range of problems, including discretization of regular and degenerate second-order elliptic equations as well as problems with no continuous origin, show that the typical multigrid convergence rates are robustly obtained [B29], [S6]. The AMG set-up time is expensive, but still comparable to the set-up time required by any Galerking coarsening (4.11). Work is underway to generalize AMG to other classes of matrices, such as those arising in discretizing systems of differential equations.

14. PRACTICAL ROLE OF RIGOROUS ANALYSIS

A good deal of the literature on multigrid consists of articles with rigorous analyses of the algebraic convergence. For a growing class of problems the basic multigrid assertion is rigorously proven, namely, that an FMG algorithm will solve the algebraic system of n equations (n unknowns on the finest grid) to the level of truncation errors in less than Cn computer operations; or at least, that $Cn \log \frac{1}{\epsilon}$ operations are enough to reduce the L^2 norm of the error by any desired factor ϵ . The emphasis is on C being independent of n ; it may depend on various parameters of the given differential problem. This is clearly the best one can do in terms of the order of dependence on n , hence the result is very satisfying.

The question discussed below is what role such rigorous analyses can have in the practical development of multigrid techniques and programs. It is an important question for the practitioner, who may wonder how much of those proofs he should try to understand.

The main shortcoming of the rigorous results is that they are usually unrealistic in terms of the size of C . In most of the proofs C is not even determined. This does not change the important fact that the best the proof could do in terms of C is very unrealistic: In most cases the provable constant is many orders of magnitude larger than the one obtainable in practice. In some typical cases the rigorous bound is $C \approx 10^8$, while the practical one is $C \approx 10^2$. Only in the very simplest situation (equations with constant coefficients in a rectangle) one can obtain realistic values of C , by Fourier methods [F5], [B7, App. C], [S4, Sec. 8]. Recently, some analyses have been made which obtain reasonable (although still several times larger than the practical) values of C for more general problems [B4], [V1], [B26].

What can then be the practical value of the Cn results, especially those where C is unreasonably large? Usually in complexity analyses results with undetermined constants are sought in cases where the size of the constants is indeed less important. A typical result would for example be that a solution to some problem, depending on some parameter n , is obtained in Cn operations. Here C may be unimportant, since changing C by orders of magnitude will only slightly increase the range of n for which the problem is solvable. But this spirit of undetermined constants is clearly pushed way too far when the estimate is Cn , the typical constant is $C = 10^8$ and the typical value for n is 10^3 to 10^5 . Here C becomes more important than n . In

the practical range of n , the provable Cn result is then vastly inferior to results obtained by simpler algorithms (such as banded elimination with typically $4n^2$ operations; not to mention drastic improvements obtainable by modern sparse-matrix packages [D3]). Thus, the values of n for which the unrealistic rigorous result can compete with much simpler solution methods is very far out in the range of overkilling the problem. In a sense, one proves efficiency of an algebraic solution process by taking an extremely unreasonable algebraic problem. The efficiency is not in terms of solving the given differential problem.

The usual rigorous theory, being too concerned on making C independent of n , is often careless about its dependence on various problem parameters. This dependence can be hair-raising indeed, something like $\exp(\exp(\dots))$, with as many compounded exponentials as there are stages in the proof. Hence, a very distorted picture is in fact supplied about the real complexity in solving the given differential problem.

The implied intention of "Cn" theorems with unspecified or unrealistic C is sometimes understood as follows: The rigorous analysis only tells us that a constant C exists, its actual value can then be determined empirically. That is, if we have calculated with $n = 10^3$ and solved the problem in 10^5 operations, say, then the rigorous proof guarantees that for $n = 10^4$ we would solve the problem in 10^6 operations. This understanding is wrong: The nature of the rigorous proofs is such that the information for $n = 10^3$ does not help the estimates for $n = 10^4$. The only rigorous estimate is still $C10^4$ operations, with the same unrealistic C . The guess that the number of operations for $n = 10^4$ will be 10^6 is purely non-rigorous. Even heuristically it does not follow in any way from the "Cn" theorem. Nothing in that theorem excludes, even heuristically, an operation count such as $Cn/(1+10^4C/n^2)$, for example, with an astronomically large C .

Thus, if one literally believed these rigorous bounds, one would not use the multigrid method in the first place. This indeed historically happened: The estimates in [F1] are so bad (although only the simplest problem is considered; cf. [37, §10]), and those of [B1] so much worse (even though his constants are undetermined), that nobody was encouraged to use such methods. They were considered to be merely of asymptotic curiosity.

Several other cases from the multigrid history are known where wrong practical conclusions were derived from the asymptotic rigorous

analysis. For example, non-smooth boundaries, reentrant corners in particular, gave troubles in the rigorous proofs. This led to the wrong conclusion that there are real troubles there. The practical fact is that such problems are solved to within truncation errors as easily as regular problems; even the asymptotic algebraic convergence rates in such cases can be made to attain the interior rates (see Sec. 5.7). The difficulties are purely difficulties of the proof, not of the computational process. The proof made us too pessimistic. In other cases similar errors made people too optimistic, because their asymptotic relations do not show the real difficulties encountered in the real range. Some people did not realize, for example, the very real difficulties in solving degenerate and singular perturbation equations (in particular indeterminate problems such as $\epsilon\Delta + k^2u = f$, where ϵ is positive but small), because these difficulties disappear for sufficiently small meshsizes. For such meshsizes are far too small to be used in practice. The terrible growth of C as function of ϵ is not seen if all we are interested in is that C will not depend on n .) Fedorenko had a completely wrong idea about the practical meshsize ratios and the number of grids to be used. He writes: "The proposed method thus consists of solution with the aid of an auxiliary net; if this latter is extremely large, the problem can also be solved on it by using a net of a particular type for the problem, and so on". And several similar historical examples could be given.

It is indeed not reasonable to expect unrealistic performance estimates to be of practical value. In practice we are interested in understanding the difference between one algorithm which solves the problem in a few minutes CPU time and another algorithm which solves it in a few hours, or in days. A rigorous result that tells us that the solution will surely be obtained within a few weeks (even years) of CPU time, does not explain that difference. The factors important in the proof may vary remotely and non-quantitatively be related to those operating in practice. Even in cases of much more reasonable C (such as [B4]), the relative values of C in two competing approaches (e.g., V cycles vs. W cycles) does not point to their relative efficiency in practice. The rigorous proof tells us more about the efficiency of the proof than about that of the actual algorithm. Hence, for all its pure-mathematical interest and intellectual challenge, much of the existing rigorous approach is not a practical tool. The only role is to generally enhance our confidence in the method, a psychological role that should not be slighted.

The purpose of the analysis should be borne in mind. We are not trying here to prove any central mathematical idea. We are engaged in a very practical problem, namely, how to solve the equations *fast*. This is in its nature as practical a problem as, say, building an airplane or understanding nuclear fission. (In fact the only purpose of the fast solvers is to aid solving such engineering and scientific problems.) One would not postpone building airplanes until rigorous proofs of their flight capabilities are furnished. Clinging to rigorous mathematics, like clinging to any secure images, may have wrong contexts. Moreover, in this business of fast solvers what one tries to apriori estimate is nothing but the computer time (and other computer resources), which is after all exactly *known* in each particular case, even though *aposteriori*. The main practical aims of theoretical understanding should therefore be:

(i) To give us realistic and quantitative insights to the important factors affecting the overall efficiency. The insight should be simple enough and still precise enough so that one can use it to improve our algorithms, and perhaps even to debug our programs. The insight should of course be independent of the numerical experiments.

(ii) Even more important than quantitative performance prediction, one urgently needs to know whether the performance (predicted or found empirically) is as good as one could *hope* to get (see the situation described in Sec. 0.2). Hence the main theoretical task is to provide us with ideal performance figures, which the practical algorithm should then attempt to approach.

The local mode analysis is an example of a theory constructed with these aims in mind. This is amply emphasized throughout Part I of the present paper. So is also the (rigorous) theory in [B26], which was developed to guide multigrid algorithms in situations where mode analysis is not available (cf. Sec. 13.1). At any rate, it is strongly recommended not to restrict oneself to numerical experiments only, without *any* supporting theory. The experiments can be, and have been, quite misleading: They may happen to show, for some particular cases, much better results than should generally be expected. More often, they show results much inferior to those that could be obtained, because of some conceptual mistakes and/or programming bugs. Experience has taught us that careful incorporation of (usually non-rigorous) theoretical studies is necessary for producing reliable programs which fully realize the potential of the multigrid method.

15. CHAINS OF PROBLEMS. FROZEN T

We often need to solve not just one isolated problem but a sequence of similar problems depending on some parameter. For example, we may be studying the effect of changing some physical parameters on the "performance" of a system, where the performance is measured in terms of the solution u to a differential problem. We may want to find for what physical parameters the performance is optimal. Or, in "inverse problems", we may desire to find the physical parameters for which the solution best fit some physically observed behavior. Or we may need to solve a sequence of problems in a continuation process (see Sec. 8.3.2). Or, the most familiar case, the parameter may be the time t , and each problem in the sequence may represent the implicit equations of one time step.

The key to a highly efficient multi-level treatment of such a sequence of problems is to understand the behavior of high-frequency components. Most often, the change in one step (i.e., the change from one problem in the sequence to the next) is a global change, governed by global parameters. In some problems, the relative changes in high-frequency components are therefore comparable to the relative changes in low ones. Hence, for such problems, the absolute high-frequency changes in each step are negligible — they are small compared to the high frequencies themselves, and therefore small compared with the discretization errors. In such cases one need not use the finer grids at each step; the finer the level the more rarely it should be activated. Often this is the situation in most parts of the domain, but in some particular parts, such as near boundaries, significant high-frequency changes do take place in every step, hence more refinement levels should more often be activated in those parts only.

The Full Approximation Scheme (FAS) gives us a convenient structure in which to see smooth changes in the solution without (locally) activating finer grids. The way to neglect changes in wavelengths smaller than $O(h)$, without neglecting those components themselves, is to freeze τ_h^{2h} (see Sec. 8.2), i.e., to use on grid $2h$ the values of τ_h^{2h} calculated in a previous step and thus in the present step avoid any visit to grid h . Once in several steps of visiting grid $2h$, a visit can be made to grid h , to update τ_h^{2h} . In visiting grid $2h$ changes in τ_{2h}^{4h} are made, their cumulative values since the last visit to grid h can serve to decide when a new visit to grid h is needed, using exchange-rate criteria (see Secs. 9.5-6 and [B11, §3.9]). Since these criteria can be applied locally, one can decide when and *where* to activate increasingly finer levels.

An obvious but important remark: Whether the above procedures are used or not, and whether FAS or CS is employed, in each step (i.e., for each problem in the chain) it is normally more economic to work on the correction problem, taking the previous-step solution as a first approximation. When FAS-FMG is used, this is easily done, even for non-linear problems, as follows. First, the old values of τ_h^{2h} should be used in the 2h-stage of the FMG algorithm (i.e., before grid h is ever visited in the present step). Secondly, the FMG interpolation (first interpolation to grid h in this step) should be a FAS-like interpolation, using the old values of u^h ; i.e., like Eq. (8.6), but with possibly higher order Π_H^h replacing Π_H (cf. Sec. 7.1).

Solving the chain of problems we usually need to monitor certain solution functionals. In order to calculate such a functional ϕ with finest-grid accuracy even at steps not visiting the finest grid, transfers as in Eq. (8.14) can be used.

16. TIME DEPENDENT PROBLEMS

The experience with multigrid applications to evolution problems is quite sparse, but several possibilities can already be outlined, such as fast solvers to implicit equations, coarse-grid time steps, highly adaptable structures, high-order techniques and global conservation facilities.

One obvious application is to use fast multigrid solvers for solving the set of algebraic equations arising at each time step when implicit time differencing is employed. Such differencing is normally needed whenever the physical signal speed is considerably higher than the speed of substantial changes in the solution. The latter speed determines the size of time steps δt we need to approximate the solution accurately, but with such δt and explicit differencing the numerical signal speed will be slower than the physical one, causing numerical instability. Using implicit equations and solving them by multigrid can be viewed as a way to inexpensively obtain high signal speeds by propagating information on coarse grids. Indeed, with multigrid solvers the cost of an implicit time step is comparable to that of an explicit one.

In many cases one can even do much better using techniques as in Sec. 15 above. For second-order parabolic problems, for example, significant changes in high-frequency components, whose wavelength is $O(h)$, occur only in very particular places such as

- (i) initially, for a short time interval $0 < t \leq O(h^2)$;
- (ii) at distance h and time interval $O(h^2)$ from points where significant changes occur in boundary conditions or in forcing terms (source terms) of the equation.

At all other places significant high-frequency changes are induced by comparably significant low-frequency changes. Hence the frozen- τ technique, with a special control for time-dependent problems [Bl1, §3.9], can give us a solution with the fine-grid accuracy but where most of the time in most of the domain we use coarse grids only. The cost of an average step may then be far smaller than the cost of an explicit time-step. For the heat equation in the infinite space and steady forcing terms, for example, one can show by Fourier analysis that marching from initial state to 90% steady-state, following the solution through-out with close to finest-grid accuracy, can in this way cost computational work equivalent to just 10 explicit time steps. (The finest grid needs to be activated only in the first few time steps, and very rarely later.)

Notice that when we march (calculating smooth changes in the solution) on coarse grids we can also use large time steps with *explicit* differencing. When fully adapted grids are used there is no need for implicit differencing, because each range of components is in effect handled by a meshsize comparable to the wavelength and by a time-step corresponding to the propagation speed, so that no conflict arises between different characteristic speeds.

For problems with *small* parabolicity (e.g., parabolic singular perturbation to a "reduced" hyperbolic system), the above technique can be superposed on an integrator of the reduced system (which may itself be based on a method of characteristics).

The multi-level techniques can also be applied to a parabolic part of the system, such as the implicit pressure equation in integrating Navier-Stokes equations [B23]. Here too, the techniques of Sec. 15 can further save a lot of fine-grid processing.

Whether the fast solver is used or not, the multi-level procedures can also give highly flexible discretization structures. Patches of finer grids with correspondingly finer time steps can be used in any part of the space-time domain, in a manner similar to Sec. 9.1. Anisotropic refinements, local coordinate transformations and rotated cartesian grids can be used as in Secs. 9.2, 9.3 and 9.4, all controlled by exchange-rate criteria (cf. Secs. 9.5, 9.6); but instead of criteria based on the τ_h^H of the approximate solution, criteria here will be based on recently accumulated changes in τ_h^H [B11, Sec. 3.9].

In some problems, especially when integrating over long time periods, certain quantities, such as total mass, must strictly be conserved, otherwise the physics of the system would fundamentally change. Imposing such global constraints, with the corresponding freeing of some accuracy parameters in the difference equations, can easily be incorporated when a multigrid solver is used at each time step (see Sec. 5.7).

Finally, independently of the above techniques, one could use a multigrid procedure similar to Sec. 10.2 above to efficiently increase the approximation order of a stable discretization

$$L_0^h u_0^h(x,t) = 0, \quad (t > 0) \quad (16.1)$$

of a time-dependent system, not necessarily linear, by using coarse-grid defect corrections. Typically L_0^h is a simple low-order implicit operator, allowing simple integration. One wants to use a simple (e.g., central in time) higher-order operator L_1^h , which may be unstable, to

raise the approximation order. This can be done by integrating the defect equation

$$L_0^H v^H = -L_1^h L_1^h u_0^h, \quad (16.2)$$

and then correcting

$$u_1^h = u_0^h + L_1^h v^H, \quad (16.3)$$

where H may either be coarser than h (coarse-grid defect correction) or $H = h$ (defect correction on the same grid). If the order of consistency of L_j^h is p_j , ($j=1,2$), then for low-frequency components

$$|u_1^h - u| = O(h^{p_1} t + h^{p_0} t^2), \quad (16.4)$$

where u is the differential solution and t is the time interval over which (16.1-3) is integrated from initial conditions $u_1^h(x,0) = u_0^h(x,0) = u(x,0)$.

The scheme (16.1-3) is always stable, since only the stable operators L_0^h and L_0^H are integrated. Notice that this is true only if u_0^h is integrated independently of u_1^h . The seemingly similar scheme, in which after each time step u_0^h is reinitialized by being replaced with the more accurate u_1^h , may well be unstable. On the other hand it does pay to reinitialize every $O(1)$ time interval, short enough to make the second term in (16.4) smaller than the first one.

The independent integration of u_0^h and v_0^h requires extra storage. By taking $H = 2h$ this extra storage becomes only a fraction of the basic storage (one time level of u_0^h), and the computational work is also just a fraction more than the work of integrating (16.1). These two-level schemes can be extended to more levels and more approximation orders. The multigrid exploits here the fact that the higher-order approximation L_1^h is desired only for sufficiently low frequencies; for the highest frequencies (where numerical instability occurs) L_0^h is in fact a better approximation than L_1^h .

PART III. APPLICATIONS TO FLUID DYNAMICS

Over the past twelve years (starting with [B5] and [S2]), but mostly during the last six, an ever-increasing number of works have applied multigrid techniques to solve *steady-state* flow problems. This includes works on transonic potential flows, such as [A4], [B31], [B32], [D5], [M5], [J1], [S2], [S8], [B23], [C3], [M4], [S7] and [W4], the latter five treating 3-dimensional flows; works on the Stokes and incompressible Navier-Stokes equations, like [B5], [B13], [B19], [D2], [F6], [T6], [V2], [V3], [W2], [W3]; works on the Euler equations [B16], [J2], [J3], [N4] and on the compressible Navier-Stokes equations [B16]. A survey of all this is not attempted here. Our purpose here is to trace a line of development which gradually leads from very simple equations to the most complicated ones, adding the difficulties step by step, but always maintaining the full multigrid efficiency; i.e., insisting on *solving every problem to $O(h^2)$ accuracy in just few work units*, where the work unit is the minimal amount of computer operations needed to express a discretization of the problem on a grid with meshsize h , and where the operations used can be fully parallelized (or vectorized) over the entire grid. *Minimal computer storage is also maintained, i.e., a storage just a fraction more than needed to store the solution itself on grid h .* Moreover, to show how these goals are achieved for the more complicated systems of equations, our emphasis here is on the treatment of *systems* of differential equations, although the line of development starts of course with simple scalar equations. In particular, the work on the scalar convection-diffusion problem [B17] is a crucial step in that line, as will become clear in Sec. 19.3, not to mention the extensive work on the Poisson equation and on more general diffusion problems.

Most works mentioned above lag far behind the ideal performance, for various reasons (see discussion in Sec. 1.7). To achieve the goals stated above, many of the principles delineated in the previous two parts of this Guide are, and perhaps must be, used. Other principles described above have not yet been used, but they are available, ready to be added and enhance the power of the flow solvers presented in this part. This includes: methods of flexible local refinements and local coordinate curving (see Sec. 9); higher-order techniques (Sec. 10: the double discretization scheme of Sec. 10.2 is already used to obtain the mentioned $O(h^2)$ approximation in cases of inviscid or small-viscosity problems, but still higher orders are obtainable, if desired, for small extra work);

procedures for further reducing computer storage (Sec. 8.7); the general approach of multigridding directly the real "outer" problem (e.g., the optimization or design problem for which the flow equations are solved; cf. Sec. 13); and the methods for efficiently treating sequences of many boundary-value problems and solving time-dependent problems (Secs. 15 and 16).

The work described in this part has been done in collaboration with Nathan Dinar and Ruth Golubev. Much of it has appeared before in [B19] and [B16].

17. CAUCHY-RIEMANN EQUATIONS

17.1 The differential problem

As a first simple exercise in multigridding a system of partial differential equations we have studied the elliptic system

$$u_x + v_y = F_1 \quad (17.1a)$$

$$u_y - v_x = F_2 \quad (17.1b)$$

in a domain Ω , where $u = u(x,y)$ and $v = v(x,y)$ are the unknown functions, the subscripts denote partial derivatives, and $F_i = F_i(x,y)$ are given functions. All functions are real. The homogeneous system $F_1 = F_2 = 0$ are the usual Cauchy-Riemann equations, which express analyticity of the complex function $u + iv$.

The matrix-operator form of (17.1) is

$$L \begin{pmatrix} u \\ v \end{pmatrix} \equiv \begin{pmatrix} \partial_x & \partial_y \\ \partial_y & -\partial_x \end{pmatrix} \begin{pmatrix} u \\ v \end{pmatrix} = \begin{pmatrix} F_1 \\ F_2 \end{pmatrix} \quad (17.2)$$

where ∂_x and ∂_y are partial derivatives with respect to x and y , respectively. The determinant of L is the Laplace operator $-\Delta = -\partial_x^2 - \partial_y^2$. Hence (17.2) or (17.1) is a second-order elliptic system and its solution is determined by one condition along the boundary $\partial\Omega$. As such a boundary condition we can, for example, require

$$(u(x), v(x))_n = G(x), \quad (x \in \partial\Omega), \quad (17.3)$$

where $(u, v)_n$ denotes the component of the vector (u, v) normal to the boundary in the outward direction. From (17.1a), (17.3) and the divergence theorem (or Stokes, or Gauss, formula) we get the "compatibility condition"

$$\int_{\Omega} F_1 dx dy = \int_{\partial\Omega} G ds. \quad (17.4)$$

If (17.4) holds then equations (17.1) or (17.2), with the boundary condition (17.3), is a well-posed problem: A unique solution exists and depends continuously on the data F_1 , F_2 and G .

17.2 Discrete Cauchy-Riemann equations

Suppose we first try to approximate (17.1) by the central difference equations

$$\frac{u^h(x+h, y) - u^h(x-h, y)}{2h} + \frac{v^h(x, y+h) - v^h(x, y-h)}{2h} = F_1^h(x, y) \quad (17.5a)$$

$$\frac{u^h(x, y+h) - u^h(x, y-h)}{2h} - \frac{v^h(x+h, y) - v^h(x-h, y)}{2h} = F_2^h(x, y). \quad (17.5b)$$

The corresponding difference operator is

$$L^h = \begin{pmatrix} \mu_x^h \partial_x^h & \mu_y^h \partial_y^h \\ \mu_y^h \partial_y^h & -\mu_x^h \partial_x^h \end{pmatrix} \quad (17.6)$$

where the averaging and differencing operators are defined by

$$\begin{aligned} \mu_x^h \phi(x, y) &= \frac{1}{2} [\phi(x + \frac{h}{2}, y) + \phi(x - \frac{h}{2}, y)], \\ \mu_y^h \phi(x, y) &= \frac{1}{2} [\phi(x, y + \frac{h}{2}) + \phi(x, y - \frac{h}{2})], \\ \delta_x^h \phi(x, y) &= \frac{1}{h} [\phi(x + \frac{h}{2}, y) - \phi(x - \frac{h}{2}, y)], \\ \delta_y^h \phi(x, y) &= \frac{1}{h} [\phi(x, y + \frac{h}{2}) - \phi(x, y - \frac{h}{2})], \end{aligned} \quad (17.7)$$

hence

$$\mu_x^h \partial_x^h = \partial_x^{2h}, \quad \mu_y^h \partial_y^h = \partial_y^{2h}$$

and

$$\det L^h = -(\mu_x^h \delta_x^h)^2 - (\mu_y^h \delta_y^h)^2 = -\Delta^{2h}$$

with the symbol (see Sec. 2.1)

$$\tilde{L}^h(\theta_1, \theta_2) = (\sin^2 \theta_1 + \sin^2 \theta_2) / h^2. \quad (17.8)$$

This operator is not h -elliptic, since $\tilde{L}^h(\pi, 0) = \tilde{L}^h(0, \pi) = \tilde{L}^h(\pi, \pi) = 0$.

Indeed, the homogeneous (17.5) equations ($F_1 = F_2 = 0$) have the oscillatory solutions

$$u^h(\alpha h, \beta h) = C_0 + C_1 (-1)^\alpha + C_2 (-1)^\beta + C_3 (-1)^{\alpha+\beta} \quad (17.9a)$$

$$v^h(ah, \beta h) = C_4 + C_5(-1)^\alpha + C_6(-1)^\beta + C_7(-1)^{\alpha+\beta} \quad (17.9b)$$

which corresponds to nothing similar in the solution of the differential equation. Note, however, that solutions like (17.9) vanish in the average, i.e., $M^h u^h = M^h v^h = 0$ for a suitable local averaging operator M^h , such as $M^h = \mu_{xy}^h$ or $M^h = (\mu_{xy}^h)^2$. Hence, the solutions of (17.5) will be good solutions in the average. Such difference operators are called *quasi elliptic* [BD, Sec. 3.4]. See further remarks in Sec. 17.6.

Let us now construct an h -elliptic approximation L^h to (17.1). If the equations are to have the form

$$D_x^1 u^h + D_y^2 v^h = F_1^h \quad (17.10a)$$

$$D_y^3 u^h - D_x^4 v^h = F_2^h \quad (17.10b)$$

where D_x^j and D_y^j are some difference approximations to ∂_x and ∂_y , then $\det L^h = -D_x^1 D_x^4 - D_y^2 D_y^3$ should be an elliptic approximation to the Laplace operator $-\Delta$. The simplest such operator is the five-point operator which is obtained by taking either

$$D_x^1 = D_x^4 = \partial_x^h, \quad D_y^2 = D_y^3 = \partial_y^h, \quad (17.11)$$

or one-sided differences such as

$$D_x^1 = \partial_x^F, \quad D_y^2 = \partial_y^F, \quad D_y^3 = \partial_y^B, \quad D_x^4 = \partial_x^B \quad (17.12)$$

where $\partial^F = \mu^h \partial^h - \frac{h}{2} \partial^h \partial^h$ and $\partial^B = \mu^h \partial^h + \frac{h}{2} \partial^h \partial^h$. Approximations like (17.12) do not give a central approximation to (17.1), and their truncation error is therefore $O(h)$. We thus prefer to use (17.11). This we can do only by using staggered grids for u^h and v^h .

The grid we use and the positioning of the discrete variables are shown in Fig. 17.1. With this positioning we can indeed approximate (17.1) by

$$\partial_x^h u^h + \partial_y^h v^h = F_1^h \quad \text{at cell centers } \textcircled{1} \quad (17.13a)$$

$$\partial_y^h u^h - \partial_x^h v^h = F_2^h \quad \text{at interior vertices } \textcircled{2}, \quad (17.13b)$$

and the symbol is that of the 5-point Laplacian, namely,




FIGURE 17.1 Discretization of Cauchy-Riemann Equations.

A typical part of the grid is shown. The discrete unknown functions u^h and v^h and their computed approximations \tilde{u}^h and \tilde{v}^h (u and v in the figure) are defined at the centers of vertical and horizontal links, respectively. The first equation (17.13a) is centered at cell centers, where its right-hand side, F_1^h is defined and where $\textcircled{1}$ is shown in the figure. The second equation (17.13b) is centered, and F_2^h is defined, at the grid vertices, as shown by $\textcircled{2}$ in the figure.

$$\tilde{L}_h(\theta_1, \theta_2) = \frac{4}{h^2} \left(\sin^2 \frac{\theta_1}{2} + \sin^2 \frac{\theta_2}{2} \right). \quad (17.14)$$

This symbol vanishes only for $\theta_1 \equiv \theta_2 \equiv 0 \pmod{2\pi}$. Thus (17.13) is an elliptic (even R-elliptic -- see [BD, Sec. 3.6]) difference system.

For simplicity we consider here domains with boundaries along grid lines. It is then simple to discretize the boundary condition (17.3). On each boundary link (the heavy lines in Figure 17.1) the variable $(u, v)_n$ is already defined at the center of the link, so (17.3) is discretized to

$$(u^h, v^h)_n = G^h \quad \text{at midpoints of boundary links.} \quad (17.15)$$

Summing (17.13a) over all the cells of our domain we get the compatibility condition

$$\sum_{\text{cell centers}} F_1^h(x, y) = \sum_{\text{boundary midpoints}} G^h(x, y) \quad (17.16)$$

which is the discrete analog of (5.4).

THEOREM 17.1 If (17.16) holds, then the discrete Cauchy-Riemann equations (17.13) with the boundary conditions (17.15) have a unique solution.

Indeed, the total number of equations (17.13), (17.15) equals the total number of cells and vertices in the grid. The number of discrete unknowns is the number of links. Hence, by a well-known formula of Euler, there is one more equation than unknowns. But the equations are dependent, as we saw in constructing the compatibility condition (17.16). Hence, if (17.16) holds, we can remove an equation and have the same number of equations as unknowns. It is therefore enough to prove the theorem for the homogeneous case $F_1^h \equiv 0$, $F_2^h \equiv 0$, $G^h \equiv 0$. In this case (17.13a) implies the existence of a discrete "stream function" ψ^h , defined at the vertices of the grid, such that $u^h = \partial_y \psi^h$, $v^h = -\partial_x \psi^h$. The homogeneous (17.13b) yields $\Delta^h \psi^h \equiv 0$, and the homogeneous (17.15) implies that ψ^h along the boundary vertices is constant. Hence, by the maximum principle, ψ^h is constant everywhere. Thus, in the homogeneous case $u^h \equiv 0$ and $v^h \equiv 0$, which is what we had to show.

17.3 DGS relaxation and its smoothing rate

Most relaxation schemes are based on one-to-one correspondence between equations and unknowns: The basic relaxation step is to satisfy (or over-satisfy, or under-satisfy) one of the discrete equations by changing the corresponding unknown (or satisfy a block of equations by changing the corresponding block of unknowns). Such one-to-one correspondence is not always natural. In our case, it is clear already in the differential equations (17.1) that it would be unnatural to regard (17.1a), say, as the equation corresponding to the unknown u , and (17.1b) as the one corresponding to v . The entire system corresponds to (u,v) . In the difference equations it would be impossible to have even a one-to-one correspondence between pairs of equations and pairs of unknowns, since the number of unknowns is one less than the number of equations.

We will therefore use "distributive relaxation", i.e., a relaxation scheme that satisfies each discrete equation in its turn by distributing changes to several unknowns, in a natural manner.

To derive a natural distributive scheme we note that neither (17.13a) nor (17.13b) are elliptic equations by themselves. It is their combination together which is elliptic. Hence, in relaxing (17.13a), for example, we should take (17.13b) into account. The simplest way to do it, which is also a special case of the general prescription described in Sec. 3.7, is to relax (17.13a) in such a way that equations (17.13b) are not "damaged", i.e., in a way which preserves the residuals of (17.13b). We do this by simultaneously changing four unknowns, in the following way:

Let $(\tilde{u}^h, \tilde{v}^h)$ be the current approximation to (u^h, v^h) . Let (x,y) be the cell center where we next wish to relax (17.13a), and let

$$r_1^h = F_1^h - \partial_x \tilde{u}^h - \partial_y \tilde{v}^h \quad (17.17)$$

be the "dynamic residual" at (x,y) . That is, r is the residual at (x,y) just before relaxing there. The relaxation step of (17.13a) at (x,y) is made up of the following four changes:

$$\begin{aligned} & \tilde{u}^h(x + \frac{h}{2}, y) + \tilde{u}^h(x - \frac{h}{2}, y) + \delta \\ & \tilde{u}^h(x - \frac{h}{2}, y) + \tilde{u}^h(x + \frac{h}{2}, y) - \delta \\ & \tilde{v}^h(x, y + \frac{h}{2}) + \tilde{v}^h(x, y - \frac{h}{2}) + \delta \\ & \tilde{v}^h(x, y - \frac{h}{2}) + \tilde{v}^h(x, y + \frac{h}{2}) - \delta \end{aligned} \quad (17.18)$$

where

$$\delta = \frac{1}{4} h r_1^h \quad (17.19)$$

It is easy to check that the distribution of changes (17.18) is such that the residuals

$$r_2^h = F_2^h - \partial_y \tilde{u}^h + \partial_x \tilde{v}^h \quad (17.20)$$

at all neighboring vertices are not changed, whatever the value of δ . The choice (17.19) for the ghost unknown δ is made so that after the changes the residual $r_1^h(x,y)$ will vanish. This is in the manner of the Gauss-Seidel relaxation, where old values are replaced by new values so as to satisfy one difference equation. Such schemes are therefore called *Distributive Gauss-Seidel* (DGS) schemes. In case k of the four values changed in (17.18) are boundary values ($k=1$ near boundaries, except near corners), then no such change should be introduced in those values, and (17.19) is replaced by

$$\delta = \frac{1}{4-k} h r_1^h \quad (17.21)$$

The relaxation of (17.13b) is made in a similar manner. If (x,y) is the vertex to be relaxed, the relaxation step will include the changes

$$\begin{aligned}
 \tilde{u}^h(x, y + \frac{h}{2}) + \tilde{u}^h(x, y - \frac{h}{2}) + \delta \\
 \tilde{u}^h(x, y + \frac{h}{2}) - \tilde{u}^h(x, y - \frac{h}{2}) - \delta \\
 \tilde{v}^h(x + \frac{h}{2}, y) + \tilde{v}^h(x - \frac{h}{2}, y) - \delta \\
 \tilde{v}^h(x + \frac{h}{2}, y) - \tilde{v}^h(x - \frac{h}{2}, y) + \delta,
 \end{aligned}
 \tag{17.22}$$

where

$$\delta = \frac{1}{4} h r_2^h.
 \tag{17.23}$$

The distribution (17.22) is such that the residuals r_1^h will be preserved, and δ in (17.23) is such that equation (17.13b) at (x, y) will be satisfied by the changed variables.

The above relaxation steps can be taken in various orders. In our programs, each complete relaxation sweep comprised of two passes: The first pass relaxes equation (17.13a) by (17.18-19), letting (x, y) traverse all cell centers, and the second pass similarly scans all the grid vertices, relaxing (17.13b) by (17.22-23). Within each pass the best ordering of points is red-black (RB), although lexicographic ordering was used in the experiments [D2]. Some operations can be saved by passing first on one color of cells and a matching color of vertices, then on the other cell color, and then on the remaining vertices.

In terms of the general formulation of Sec. 3.7 the above scheme is equivalent to introducing the ghost functions w_1^h and w_2^h , where

$$\begin{pmatrix} u^h \\ v^h \end{pmatrix} = \begin{pmatrix} \partial_x^h & \partial_y^h \\ \partial_x^h & -\partial_y^h \end{pmatrix} \begin{pmatrix} w_1^h \\ w_2^h \end{pmatrix} = M^h w^h
 \tag{17.24}$$

(accidentally $M^h = L^h$), and relaxing the resulting system $\Delta^h w_i^h = F_i^h$, ($i = 1, 2$), Δ^h being the 5-point Laplace operator. Hence the smoothing factor of this relaxation is the same as that of Δ^h , that is, for lexicographic ordering $\bar{\mu} = .5$, while in RB ordering $\bar{\mu}_1 = \bar{\mu}_2 = .25$ and $\bar{\mu}_3 = .32$ (cf. Eq. (3.2)).

17.4 Multigrid procedures

Assume now we have a sequence of grids (levels) with mesh-sizes h_1, \dots, h_M , where $h_{k+1} = \frac{1}{2} h_k$. The relative position of the different grids is shown in Fig. 17.2. Instead of $F_1^h, F_2^h, G^h, u^h, v^h, \tilde{u}^h, \tilde{v}^h$, r_1^h and r_2^h used above, the discrete functions on the k -th level will

be denoted by $F_1^k, F_2^k, G^k, u^k, v^k, \tilde{u}^k, \tilde{v}^k, r_1^k$ and r_2^k , respectively. Similarly, u_x^k and u_y^k will stand for $u_x^{h_k}$ and $u_y^{h_k}$.

The coarse-to-fine interpolation can be of first order, since this is the highest order of derivatives in the Cauchy-Riemann operator (see Sec. 4.3). An obvious way of doing such an interpolation (see Fig. 17.2) is

$$I_k^{k+1} \tilde{u}^k(x, y) = \begin{cases} \tilde{u}^k(x, y \pm \frac{1}{2} h_{k+1}) & \text{if } x \text{ is on a coarse-grid line} \\ u_x^{k+1} \tilde{u}^k(x, y) & \text{otherwise,} \end{cases}
 \tag{17.25}$$

and similarly for $I_k^{k+1} \tilde{v}^k$. One can of course use linear interpolations instead.

The Cauchy-Riemann problem is linear. We can therefore make coarse-grid corrections either by the Correction Scheme (CS) or the Full-Approximation Scheme (FAS, cf. Sec. 8). In the latter case we have to define the fine-to-coarse transfer of solution $(I_{k+1}^k \tilde{u}^{k+1}, I_{k+1}^k \tilde{v}^{k+1})$. We use the following averaging (cf. Fig. 17.2):

$$I_{k+1}^k \tilde{u}^{k+1}(x, y) = u_y^{k+1} \tilde{u}^{k+1}(x, y)
 \tag{17.26a}$$

$$I_{k+1}^k \tilde{v}^{k+1}(x, y) = u_x^{k+1} \tilde{v}^{k+1}(x, y).
 \tag{17.26b}$$




FIGURE 17.2 A coarse-grid cell divided into fine-grid cells.

Same notations as in Fig. 17.1, with larger and heavier type being used for the coarse grid and lighter type for the fine grid.

The number of operations in such an algorithm, using the Correction Scheme, is about $40n$, where n is the number of unknowns on the finest grid. Almost all these operations are either additions or shifts (i.e., multiplications by an integer power of 2), and the algorithm is fully parallelizable.

There is a faster way of solving the discrete Cauchy-Riemann equations (17.13): Subtracting from u^h a function u_0^h which satisfies $\partial_x^h u_0^h \equiv F_1^h$, a new system is obtained in which $F_1^h \equiv 0$. The problem can then be rewritten as a Poisson problem for the discrete stream function ψ^h (see Sec. 17.2). Solving that Poisson problem by a similar FMG algorithm, together with the operations of subtracting u_0^h and constructing u^h and v^h would require about $17n$ operations (additions and shifts only). The main purpose of this chapter, however, was to study methods for solving elliptic systems. The techniques developed for the present simple system guided us in developing the more complicated cases described below.

17.6 Remark on non-staggered grids

Any staggered-grid formulation can yield a non-staggered one (on a finer grid) simply by overlaying several staggered grids, properly shifted, on top of each other. For example, shifting equations (17.13) by the four shifts $(0,0)$, $(0,h/2)$, $(h/2,0)$ and $(h/2,h/2)$, the four systems together are equivalent to equations (17.5) for grid $h/2$. In fact, various non-staggered formulations appearing in the literature can be shown to be such interlacing of staggered formulations. They are wasteful in that the same accuracy is already obtained by just one of the interlacing subgrids, for much less work.

Also, the decomposition of the grid into interlacing subgrids locally decoupled from each other introduces a subtle kind of instability (typical to quasi-elliptic operators in general): Certain high-frequency modes (those which look like low frequency modes on each subgrid) are magnified in the discrete solution much beyond their size in the differential solution. This may show as large truncation errors in higher Sobolev norms. It can be corrected by averaging.

Such an averaging, and the multigrid solution of such a quasi-elliptic system, is discussed in Sec. 18.6.

18. STEADY-STATE STOKES EQUATIONS

18.1 The differential problem

As a prelude to the treatment of the full Navier-Stokes equations, we consider now the steady-state Stokes equations in a d -dimensional domain

$$\nabla \cdot \underline{u} = F_0 \tag{18.1a}$$

$$-\Delta \underline{u} + \nabla p = \underline{F} \tag{18.1b}$$

where $\underline{u} = (u_1, \dots, u_d)$ represents the velocity of a fluid and p represents the pressure, $\nabla = (\partial_1, \dots, \partial_d)$ is the gradient operator, $\Delta = \partial_1^2 + \dots + \partial_d^2$ is the Laplace operator, and F_0 and $\underline{F} = (F_1, \dots, F_d)$ are given forcing functions. (18.1) are the equations of "creeping" flows (vanishing Reynolds number). (18.1a) is the "continuity equation" (usually with vanishing source term: $F_0 \equiv 0$), and (18.1b) is the vector of d momentum equations.

The matrix-operator form of (18.1) is

$$L \begin{pmatrix} p \\ u_1 \\ \vdots \\ u_d \end{pmatrix} \equiv \begin{pmatrix} 0 & \partial_1 & \dots & \partial_d \\ \partial_1 & -\Delta & & \\ \vdots & & \circ & \\ \partial_d & & & -\Delta \end{pmatrix} \begin{pmatrix} p \\ u_1 \\ \vdots \\ u_d \end{pmatrix} = \begin{pmatrix} F_0 \\ F_1 \\ \vdots \\ F_d \end{pmatrix} \tag{18.2}$$

and the operator determinant is

$$\det L = (-\Delta)^d \tag{18.3}$$

Hence (18.1) is a $2d$ -order elliptic system and will require d boundary conditions. These are usually given by specifying the velocity on the boundary

$$\underline{u}(x) = \underline{G}(x), \quad (x \in \partial\Omega) \tag{18.4}$$

where $\underline{G} = (G_1, \dots, G_d)$.

Equations (18.1) with the boundary conditions (18.4) constitute a well-posed problem, provided the compatibility condition, obtained from (18.1a) and the divergence theorem,

$$\int_{\Omega} F_0 \, dx = \int_{\partial\Omega} G \cdot ds \quad (18.5)$$

is satisfied, where ds is the boundary element multiplying an outward normal unit vector.

18.2 Finite-difference equations

By arguments similar to those in Secs. 17.2 and 17.6 we find it best to discretize (18.1) on a staggered grid (but see Sec. 18.6). Such a grid, in the two-dimensional case, is shown in Fig. 18.1. In the general d -dimensional case, the grid hyperplanes (planes if $d = 3$, lines if $d = 2$) define cells, each cell with $2d$ faces. The discrete




FIGURE 18.1 Discretization of two-dimensional Stokes equations

A typical part of the grid is shown. The discrete pressure \tilde{p}^h is defined at cell centers $\{p\}$. The discrete velocity \tilde{u}_1^h is defined at centers of vertical links $\{1\}$ (interior centers; \square = boundary and exterior centers), and \tilde{u}_2^h is defined at centers of horizontal links $\{2\}$ and \square . The discrete continuity equations are centered at cell centers $\{p\}$. The j -th momentum equation is centered at interior values of \tilde{u}_j^h $\{1\}$. The exterior values of \tilde{u}_1^h and \tilde{u}_2^h (at \square and \square , respectively, but not on the boundary) are fictitious.

velocity u_j^h and its computed approximations \tilde{u}_j^h are defined at centers of j -faces, i.e., faces perpendicular to the j -th coordinate. The discrete pressure p^h and its computed approximation \tilde{p}^h are located at cell centers. The discrete approximation to (18.1) can then be written as

$$\sum_{j=1}^d \partial_j^h u_j^h = F_0^h \quad \text{at cell centers} \quad (18.6a)$$

$$-\Delta^h u_j^h + \partial_j^h p = F_j^h \quad \text{at centers of } j\text{-faces, } (j = 1, \dots, d) \quad (18.6b)$$

where $\partial_j^h \phi(x) = \frac{1}{h} [\phi(x + \frac{1}{2} h_j) - \phi(x - \frac{1}{2} h_j)]$, h_j is h times the unit vector in direction x_j , and the discrete approximation Δ^h to the Laplace operator is the usual $(2d+1)$ -point approximation $\sum_{l=1}^d (\partial_l^h)^2$. For a point x near a boundary, however, $\Delta^h u_j^h(x)$ may involve an exterior value $u_j^h(x^e)$. This value is defined by quadratic extrapolation from $u_j^h(2x-x^e)$, $u_j^h(x)$ and $u_j^h(x^b) = G_j^h(x^b)$, where x^b is a boundary point on the segment (x, x^e) . This definition is used to eliminate the exterior value from $\Delta^h u_j^h(x)$, so that the discrete Laplacian is modified and includes a boundary value of u_j^h .

The matrix operator of (18.6) is

$$L^h = \begin{pmatrix} 0 & \partial_1^h & \dots & \partial_d^h \\ \partial_1^h & -\Delta^h & & \\ \vdots & & \ddots & \\ \partial_d^h & & & -\Delta^h \end{pmatrix}, \quad (18.7)$$

hence $\det L^h = (-\Delta^h)^d$ and its symbol is

$$L^h(\theta) = h^{-2d} \left\{ \sum_{j=1}^d \left(2 \sin \frac{\theta_j}{2} \right)^2 \right\}^d, \quad (18.8)$$

which is positive for $0 < |\theta| \leq \pi$. The difference system (18.6) is therefore h -elliptic (see Sec. 2.1) and even R -elliptic [B19, Sec. 3.6].

The boundary condition (18.4) is approximated by the above way for treating boundary and exterior values of u_j^h . For simplicity, consider the case of domains whose boundary is contained in grid lines (or grid planes). In this case the velocity normal to the boundary is conveniently defined at the center of boundary faces, and the discrete analog to (18.5) is naturally written as

$$\sum_{\underline{x}} F_0^h(\underline{x}) = \sum_{\underline{y}} G_n^h(\underline{y}), \quad (18.9)$$

where \underline{x} runs over all cell centers, \underline{y} runs over all centers of boundary faces, and $G_n^h(\underline{y})$ is the (given) normal velocity at \underline{y} .

THEOREM 18.1 *The discrete Stokes equations (18.6), with exterior and boundary values determined by the boundary conditions as above, have a unique solution, upto an additive constant in p^h , if and only if (18.9) is satisfied.*

The proof is simple. The number of equations is the same as the number of unknowns, since for each interior $u_j^h(\underline{x})$ there corresponds an equation (18.6b) at \underline{x} , and for each unknown $p^h(\underline{y})$ there corresponds an equation (18.6a) at \underline{y} . The pressure values p^h are determined only upto an additive constant, but, on the other hand, the equations are dependent; summing (18.6a) over all cell centers we get (18.9). That is to say, if (18.9) is not satisfied we get a contradiction. If (18.9) is satisfied, it is enough to show that in the homogeneous case ($F_0^h \equiv 0$, $\underline{F}^h \equiv 0$, $\underline{G}^h \equiv 0$), the only solution is the trivial one ($\underline{u}^h \equiv 0$, $p^h \equiv \text{constant}$). Indeed, if $\underline{F}^h \equiv 0$ it is easy to see from (18.6b) that

$$\begin{aligned} 0 &= \sum_{j=1}^d \Sigma_1^j \{ -\Delta^h u_j^h(\underline{x}) + \partial_j^h p^h(\underline{x}) \} u_j^h(\underline{x}) \\ &= \sum_{j=1}^d h^{-2} \Sigma_2^j \{ u_j^h(\underline{x}) - u_j^h(\underline{y}) \}^2 \\ &\quad + \sum_{j=1}^d h^{-2} \Sigma_3^j \{ u_j^h(\underline{x}) - u_j^h(\underline{z}) \} u_j^h(\underline{x}) + \Sigma_4 p^h(\underline{x}) \sum_{j=1}^d \Delta^h u_j^h(\underline{x}), \end{aligned}$$

where the point \underline{x} in Σ_1^j runs over all interior positions of u_j^h (points \odot in Fig. 18.1); the pair $\{\underline{x}, \underline{y}\}$ in Σ_2^j runs over all pairs of neighboring interior positions of u_j^h ; the pair $\{\underline{x}, \underline{z}\}$ in Σ_3^j runs over all pairs of neighboring positions u_j^h , with \underline{x} being an interior position (\odot in Fig. 18.1), and \underline{z} being a boundary or exterior position (\square in Fig. 18.1); and \underline{x} in Σ_4 runs over all cell centers (\square in Fig. 18.1). The term with Σ_4 vanishes by (18.6a), since $F_0^h \equiv 0$. In the Σ_3^j term, by the way exterior values are defined, we get (for $G^h \equiv 0$) $u_j^h(\underline{z}) = 2u_j^h(\underline{x}) - \frac{1}{3}u_j^h(\underline{y})$, where \underline{y} is the interior neighbor of \underline{x} opposite \underline{z} . Hence,

$$\sum_{j=1}^d \left\{ \Sigma_5^j \{ u_j^h(\underline{x}) - u_j^h(\underline{y}) \}^2 + \Sigma_3^j \left[4u_j^h(\underline{x})^2 + u_j^h(\underline{y})^2 - \frac{7}{3} u_j^h(\underline{x}) u_j^h(\underline{y}) \right] \right\} = 0$$

where Σ_5^j runs as Σ_2^j except for terms added to Σ_3^j . This form is

positive definite, hence $u_j^h \equiv 0$. By (18.6b) $p^h \equiv \text{const}$.

18.3 Distributive relaxation

Designing the relaxation scheme for (18.7) by the general approach of Sec. 3.7, note that L^h is almost triangular. To triangularize it, it is enough to transform its first column by multiplying L^h on the right with

$$M^h = \begin{pmatrix} -\Delta^h & 0 & \dots & 0 \\ -\partial_1^h & 1 & & \\ \vdots & & \ddots & \\ -\partial_d^h & & & 1 \end{pmatrix}, \quad (18.10)$$

(The first column of M^h could be obtained as the cofactors of the first row of L^h , divided by their maximal common divisor $(-\Delta^h)^{d-1}$.) This distribution operator M^h implies the following relaxation scheme.

First, for each $1 \leq j \leq d$, the j -th momentum equation (18.6b) is relaxed by scanning, in some prescribed order (RB ordering is best), all the interior j -face centers, changing at each such point \underline{x} the current value $\tilde{u}_j^h(\underline{x})$ so as to satisfy the momentum equation at \underline{x} . Then, the continuity equation (18.6a) is relaxed by scanning the cell centers (preferably in RB ordering), introducing at each such center \underline{x} the following changes:

$$\tilde{u}_j^h(\underline{x} + \frac{1}{2} \underline{h}_j) + \tilde{u}_j^h(\underline{x} - \frac{1}{2} \underline{h}_j) + \delta, \quad (j = 1, \dots, d), \quad (18.11a)$$

$$\tilde{u}_j^h(\underline{x} - \frac{1}{2} \underline{h}_j) + \tilde{u}_j^h(\underline{x} + \frac{1}{2} \underline{h}_j) - \delta, \quad (j = 1, \dots, d), \quad (18.11b)$$

$$\tilde{p}^h(\underline{x}) + \tilde{p}^h(\underline{x}) + \frac{2d}{h} \delta, \quad (18.11c)$$

$$\tilde{p}^h(\underline{x} + \underline{h}_j) + \tilde{p}^h(\underline{x} - \underline{h}_j) - \frac{1}{h} \delta, \quad (j = 1, \dots, d), \quad (18.11d)$$

$$\tilde{p}^h(\underline{x} - \underline{h}_j) + \tilde{p}^h(\underline{x} + \underline{h}_j) - \frac{1}{h} \delta, \quad (j = 1, \dots, d), \quad (18.11e)$$

where

$$\delta = \frac{h}{2d} r_0^h(\underline{x}) = \frac{h}{2d} \left(F_0^h - \sum_{j=1}^d \partial_j^h \tilde{u}_j^h \right) \text{ before the changes} \quad (18.12)$$

so that after these changes at \underline{x} the new residual $r_0^h(\underline{x})$ vanishes.

The relaxation of the continuity equation, and its modification near a boundary, are shown in Fig. 18.2. Observe that near the boundary, unlike the situation away from it, the relaxation of the continuity equation at a point does introduce slight changes in neighboring momentum residuals, as if locally contradicting the triangularity of L_M^h . But





FIGURE 18.2 Continuity-equation relaxation step in two-dimensional Stokes equations

- (a) The cell at the center of the figure is relaxed by 9 simultaneous changes. The amount of change is displayed at the position of the changed variable (cf. Fig. 18.1). $\delta = h r_0^h(\underline{x})/4$, where $r_0^h(\underline{x})$ is the dynamic residual at the relaxed cell.
- (b) Configuration of changes in a boundary cell. $\delta = h r_0^h(\underline{x})/3$.
- (c) Configuration of changes in a corner cell. $\delta = h r_0^h(\underline{x})/2$.

these slight changes do not cause later (when the momentum equations are relaxed) significant "feed-back", i.e., too large changes back in r_0^h , because near the boundary feed-back changes are partly "absorbed" by the boundary.

The smoothing factor is the same as for GS relaxation for L_M^h , i.e., the same as for Δ^h . Hence, in two dimensions, for lexicographic ordering $\bar{u} = .5$ and for RB ordering, if used in all 3 passes, $\bar{u}_1 = \bar{u}_2 = .25$ and $\bar{u}_3 = .32$ (cf. Eq. (3.2)). For $d = 3$ and lexicographic ordering $\bar{u} = .563$.

18.4 Multi-grid procedures

For multi-grid processing of Stokes equations we use a sequence of grids (levels) with meshsizes h_1, \dots, h_M , where $h_{k+1} = \frac{1}{2} h_k$, and where the grid lines (or grid planes) of level k are every other grid line (plane) of level $k + 1$. Hence, each cell of level k is the union of 2^d cells of level $k + 1$. In two dimensions ($d = 2$) the configuration is shown in Fig. 18.3. Instead of $\underline{F}^h, \underline{u}^h, p^h, \underline{r}^h, \underline{u}_j^h$ and ∂_j^h used in Secs. 18.2-3 and in Eq. (17.7), the discrete functions and operators on the k -th level are now denoted by $\underline{F}^k, \underline{u}^k, p^k, \underline{r}^k, \underline{u}_j^k$ and ∂_j^k , respectively.




FIGURE 18.3 A coarse-grid cell divided into fine-grid cells
Same notation as in Fig. 18.1 is used, with heavy type for the coarse grid and light type for the fine grid.

We have solved Stokes equations both with the Correction Scheme (CS) and with the Full-Approximation Scheme (FAS), getting of course identical results. We describe here the procedures in terms of FAS, since CS is not extendible to the nonlinear Navier-Stokes equations.

Coarse-to-fine interpolations. In the FMG algorithm, to obtain residuals smaller than their truncation errors, the first coarse-to-fine interpolation has to be of order at least four for the velocities and at least three for the pressure (see Sec. 7.1). The design of such interpolations is straight-forward, although it turns out somewhat cumbersome near boundaries. The coarse-to-fine interpolation of corrections has to be of orders at least two for the velocities and one for the pressure (see Sec. 4.3). We used bilinear (i.e., order two) interpolations for both.

The fine-to-coarse transfers are made by averaging. For the FAS transfer of u_j^{k+1} we can use the same averaging as for r_j^{k+1} , ($j = 1, \dots, d$), which can be either the minimal-operation transfer

$$I_{k+1}^k r_j^{k+1} = \mu_1^{k+1} \dots \widehat{\mu_j^{k+1}} \dots \mu_d^{k+1} r_j^{k+1}, \quad (j = 1, \dots, d) \quad (18.13)$$

or the full weighting

$$I_{k+1}^k r_j^{k+1} = \mu_j^{k+1} \mu_1^{k+1} \dots \mu_d^{k+1} r_j^{k+1}, \quad (j = 1, \dots, d), \quad (18.14)$$

where the hat in (18.13) indicates the term to be skipped in the sequence. The residual-weighting (18.13) is less expensive than (18.14), especially since it requires calculating only one half of the fine-grid residuals. But (18.14) is more reliable in the nonlinear case and near boundaries, since it is "full" (see Sec. 4.4).

The FAS transfer of \tilde{p}^{k+1} can be made with the same weighting as the transfer of the continuity-equation residuals

$$I_{k+1}^k r_0^{k+1} = \mu_1^{k+1} \dots \mu_d^{k+1} r_0^{k+1}, \quad (18.15)$$

which is both simplest and full. In fact, if the minimal-operations transfer (18.13) is used for the velocities \tilde{u}_j^{k+1} , then (18.15) need not really be calculated: If the FAS continuity equation on level k is written in the form

$$\sum_{j=1}^d \partial_j^k u_j^k = f_0^k \quad (18.16a)$$

(where $f_0^k \equiv F_0^k$ on the currently finest level l), it is easy to see that (18.15) is equivalent to

$$\tilde{f}_0^k = \mu_1^{k+1} \dots \mu_d^{k+1} f_0^{k+1}, \quad (k < l), \quad (18.16b)$$

which does not depend on \tilde{u}^{k+1} .

The compatibility condition (18.9) is automatically obtained (up to round-off errors) on all levels provided it holds on the finest one. This results directly from (18.15).

18.5 Numerical results

Our early (1978) experiments with the above procedures, in various cycling and FMG algorithms, on two-dimensional rectangular domains (chosen only for programming simplicity), are described in [B19] and in more details in [D2]. The program itself is available [M3]. The experiments were not optimal because we have used lexicographic instead of RB ordering in relaxation (RB was used in our recent experiments with non-staggered grids; see Sec. 18.6). The asymptotic convergence rates were 20% slower than those predicted by the smoothing rate $\bar{\mu}$, but not more than 6% slower than predicted by the two-level analysis. The fact that the smoothing rate is not fully exploited indicates that better convergence rates may be obtainable with better inter-grid transfers, but this does not seem to worth the extra work: The obtained rates are good enough to give a solution below truncation errors by an FMG algorithm with one $v(2,1)$ cycle per level (cf. Fig. 1.2). With RB relaxation, $v(1,1)$ should already suffice (see Sec. 18.6).

18.6 Non-staggered grids. Recent results

Having received complaints about the inconvenience of staggered grids despite their advantages (see Sec. 17.6), we have recently experimented with Stokes (and Navier-Stokes) solvers on conventional, non-staggered grids, where all unknowns are defined at the same gridpoints. The approach is to use a quasi-elliptic approximation, but to properly average the results. Thus, ∂_j^h in (18.6) or (18.7) has throughout been changed into the (long) central differencing $\partial_j^{2h} = \mu_j^h \partial_j^h$, replacing (18.8) by

$$\tilde{L}^h(\theta) = h^{-2d} \left\{ \sum_1^d \left(2 \sin \frac{\theta_j}{2} \right)^2 \right\}^{d-1} \left\{ \sum_1^d \sin^2 \theta_j \right\}. \quad (18.17)$$

This symbol vanishes for some $|\theta| = \pi$, showing L^h to be only quasi-

elliptic. More precisely, there are the high-frequency components

$$\begin{pmatrix} p \\ \underline{u} \end{pmatrix} = \begin{pmatrix} 1 \\ 0 \end{pmatrix} e^{i\theta \cdot \underline{x}/h} \quad (\theta_j = 0 \text{ or } \pi, 1 \leq j \leq d; |\underline{\theta}| = \pi) \quad (18.18)$$

which satisfy the homogeneous difference equations. These components by themselves do not really matter more than adding a constant to the pressure: They do not affect $\partial_j^{2h} p^h$, nor any other term in the difference equations. But other components, in the neighborhood of (18.18), do matter for the solution. They thus constitute high-frequency components which are not locally controlled. They will give rise to a subtle kind of instability (see Sec. 17.6), and they will not efficiently be reduced by relaxation, hence will be slow to converge in conventional multigrid algorithms. (Fast convergence can still be obtained by modified coarse-grid functions or by AMG algorithms; see Secs. 4.2.2 and 13.1.)

Both troubles (instability and slow smoothing) are closely related and easily cured by the same simple device: The bad components can be eliminated by averaging the pressure, i.e., replacing

$$\tilde{p}^h + (\nu_1^h \dots \nu_d^h) \frac{2\tilde{p}^h}{p^h}. \quad (18.19)$$

This operation need to be done only on the final result. Similar bad convergence of some high-frequency components of an intermediate level does not matter, since those components are efficiently reduced by the next-finer-level relaxation (cf. Sec. 12).

To test this approach, without any interference of questions related to boundary conditions, we have studied the two-dimensional ($d = 2$) case, on the square $\{|x|, |y| \leq \pi\}$, with periodic boundary conditions. With such boundary conditions, u_j and p are determined only up to an additive constant each. We have used distributive Gauss-Seidel relaxation based on the distribution operator

$$M^h = \begin{pmatrix} -\Delta^h & 0 & 0 \\ -\partial_1^{2h} & 1 & 0 \\ -\partial_2^{2h} & 0 & 1 \end{pmatrix} \quad (18.20)$$

(cf. (18.10)), with RB ordering in each of its 3 passes (corresponding to the 3 differential equations). Inter-grid transfers were standard: Full weighting (4.6), bilinear interpolation of corrections and cubic FMG interpolation of the first approximation. The coarse-grid operator is the same non-staggered, central Stokes operator as employed on the

finest grid.

The compatibility condition, analogous to (18.5), seems at first to cause some troubles. In the discrete approximation it breaks up into 4 different conditions, each obtained by summing the discrete continuity equation on one of the four staggered subgrids into which that equation is decoupled. Even if we take the trouble to satisfy these four conditions on the finest grid, they will not be satisfied on coarser grids, unless the coarse-grid equations are adjusted, e.g. by adding 4 suitable constants to F_0^k , one constant on each subgrid. Actually, all this is not needed. The said adjustment is below the level of the coarse-grid truncation, hence it does not improve the quality of the coarse-grid contribution to the fine-grid convergence. The fine-grid compatibility condition is itself similarly unimportant, unless one wants to solve the algebraic system below truncation errors. We simply ignore this condition at all levels.

Two additional compatibility conditions emerge because of our periodic boundary conditions, each obtained by summing one of the momentum equations over the grid. In the differential equations chosen by us these conditions are automatically satisfied, since we calculate a right-hand side \underline{F} from a known solution (so that we know the exact differential solution and can compare our results with it). For reasons similar to the above, we ignored satisfying these compatibility conditions, too, in our discrete approximations, whether fine or coarse.

Components in the neighborhood of (18.18), eventually almost eliminated by (18.19), are still present, even though their influence on the residuals must initially be very small. Hence, asymptotically (after many cycles) we must get slow convergence. Indeed, on large $N \times N$ grids we have obtained asymptotic convergence factors of about $1-160N^{-2}$ per either $v(1,1)$ or $w(1,1)$ cycle. Initially, however, those components are not important and convergence factors are excellent; e.g., in the experiments of Table 18.1 the average convergence factor for the first five cycles was always better than .2 per $V(1,1)$ cycle.

The main point of course is that the convergence rate of high-frequency components is immaterial by itself, since we do not need to converge them much to get below their large truncation errors. What does then matter is how many cycles per level an FMG algorithm needs in order to solve the problem to below truncation. Typical results are summarized in Table 18.1. They show that convergence well below truncation errors is already obtained for an algorithm with only 3.6 RWUs (relaxation work units). Being based on the RB relaxation, the algo-

rithm is fully vectorizable.

FMG Algorithm			Results			
grid size	type of cycle	# cycles per level	$\ u - \tilde{u}^h\ _\infty$	$\ v - \tilde{v}^h\ _\infty$	$\ p - \tilde{p}^h\ _\infty$ without averaging	with averaging
16x16	v	1	.0259	.0379	.4589	.1984
		2	.0217	.0224	.3306	.1671
		3	.0209	.0213	.3247	.1662
	w	1	.0236	.0356	.4400	.1873
		2	.0219	.0219	.3287	.1666
		3	.0211	.0216	.3275	.1662
32x32	v	1	.0056	.0107	.1216	.0821
		2	.0044	.0069	.0774	.0556
		3	.0043	.0068	.0627	.0541
	w	1	.0047	.0090	.1248	.0813
		2	.0045	.0069	.0789	.0564
		3	.0043	.0068	.0630	.0543
64x64	v	1	.0019	.0036	.0345	.0172
		2	.0020	.0029	.0180	.0141
		3	.0020	.0028	.0154	.0141
	w	1	.0020	.0032	.0367	.0187
		2	.0020	.0029	.0184	.0143
		3	.0020	.0028	.0158	.0141

TABLE 18.1 Stokes solutions on non-staggered grid

Results are shown for an FMG solution of a problem whose exact differential solution is $u = v = p = \sin(\cos(x+2y))$, the right-hand sides F_0 and F being calculated accordingly. F_0^h and F^h on all levels are obtained by averaging F_0 and F on $h \times h$ squares. The approximate solution $(\tilde{u}^h, \tilde{v}^h, \tilde{p}^h)$ is obtained by the algorithm explained in the text, with relaxation counts $v_1 = v_2 = 1$ (see Secs. 1.4, 1.6; the v-cycle algorithm is exactly that of Fig. 1.2). The undetermined additive constants in \tilde{u}^h , \tilde{v}^h and \tilde{p}^h were of course properly subtracted off.

19. STEADY-STATE INCOMPRESSIBLE NAVIER-STOKES EQUATIONS

19.1 The differential problem

Using the notation of Sec. 18.1, the steady-state incompressible Navier-Stokes equations in d dimensions can be written in the form

$$\nabla \cdot \underline{u} = F_0 \tag{19.1a}$$

$$\underline{Q}\underline{u} + \nabla p = \underline{F} \tag{19.1b}$$

where

$$Q = -\frac{1}{R} \Delta + \underline{u} \cdot \nabla = -\frac{1}{R} \Delta + \sum_{j=1}^d u_j \partial_j \tag{19.2}$$

R being the Reynolds number; i.e., R^{-1} is a scaled viscosity-over-density coefficient. (It is assumed that time and distance are scaled so that \underline{u} is $O(1)$ and domain dimensions are $O(1)$.) The principal part of this system is the Stokes system (18.1) (rescaling $p \rightarrow Rp$), hence for $R \ll O(1)$ the solution processes of Sec. 18 are directly applicable here (using FAS of course to deal with the nonlinearity). But we will be interested in solving also for large R , hence we will look at the corresponding subprincipal operator (cf. Sec. 2.1)

$$L = \begin{pmatrix} 0 & \partial_1 & \dots & \partial_d \\ \partial_1 & Q & & \\ \vdots & & \circ & \\ \partial_d & & \circ & Q \end{pmatrix} \tag{19.3}$$

Observe indeed that (19.3) is not the full Newton linearization for (19.1); only subprincipal terms are kept. Since $\det L = -\Delta^{d-1}$, (19.1) is again an elliptic system of order $2d$ and therefore requires d boundary conditions. Usually the velocities \underline{u} are given on the boundary as in (18.4), leading again to the computability condition (18.5).

For $R \rightarrow \infty$, $\det L \rightarrow -\Delta(\underline{u} \cdot \nabla)^{d-1}$, hence the streamlines are the subcharacteristics (characteristic lines of the reduced equation), and in the limit only one condition is needed all along the boundary, plus $d - 1$ conditions at points where the flow enter the domain (because $\underline{u} \cdot \nabla$ is a streamwise derivative which, with the singular perturbation $-\frac{1}{R} \Delta$, permits discontinuity only at the exit end of each streamline). Typically, for $R \rightarrow \infty$, the velocities \underline{u} are given at entrance points, and either the normal velocity or the pressure at other boundary points.

For theoretical investigations of this system, see for example [T3].

19.2 Staggered finite-difference approximations

The discretization is carried out on the same staggered grid as before (Fig. 18.1), using the difference equations

$$\sum_{i=1}^d \partial_i^h u_i^h = F_0^h \quad \text{at cell centers} \quad (19.4a)$$

$$Q_j^h u_j^h + \partial_j^h u_j^h = F_j^h \quad \text{at } j\text{-face centers} \quad (j = 1, \dots, d) \quad (19.4b)$$

where Q^h is some difference approximation to Q . (Non-staggered grids are discussed in Sec. 19.6.) Since $\det L^h = -\Delta^h (Q^h)^{d-1}$, it is clear that L^h has good h -ellipticity measure if and only if Q^h does. Hence, all we have to construct is a good approximation to Q . For small to moderate $hR|\underline{u}|$ (i.e., $hR|\underline{u}|$ not much larger than 1) this can be done by central differencing. But for larger $hR|\underline{u}|$ upstream differencing or explicit artificial viscosity terms should be used.

The artificial viscosity terms may be anisotropic, so that the total (physical and artificial) viscosity has the form $-\tau_i \bar{\beta}_i h (\partial_i^h)^2$. For stability of the simplest DGS schemes (Sec. 19.3), $\bar{\beta}_i(x) > .5 \max |u_i(y)|$ is needed, where the maximum is taken over all y neighboring x . Upstream differencing is the same as $\bar{\beta}_i = .5 |u_i|$. For large R , sharp cross-stream changes (large solution changes per cross-stream meshsize) can travel with the stream, and one may like to avoid smearing them by using no cross-stream artificial viscosity. This is only possible by *strong alignment* (cf. Sec. 2.1), i.e., by using a grid (sometimes through the method of Sec. 9.3) such that one of its principal directions is (nearly) aligned with the stream throughout a large subdomain. One can then use $\bar{\beta}_i = O(h|\underline{u}|)$ (necessary for stability) in the stream direction, together with zero (or small) cross-stream artificial viscosity.

In three-dimensional problems it may sometimes be difficult and unnecessary to have the flow (nearly) aligned with one grid direction, but a grid can be used so that each streamline (nearly) stay in one grid plane. We call this *plane alignment*. In this case only sharp changes perpendicular to these planes are resolvable. For that purpose, zero artificial viscosity perpendicular to the planes should be used.

Another, more special case of strong alignment may arise near boundaries. Namely, to obtain sharp numerical boundary layers on grids

which do not resolve the *physical* boundary layer, the Q^h operators near the boundary should be constructed so that tangential components of $Q^h \underline{u}$ do not straddle the boundary layer, i.e., do not include boundary values. This "alignment" of Q^h can be done without aligning the grid.

The artificial viscosity may introduce $O(h)$ error, but such errors can be eliminated by omitting the artificial terms from the residuals calculated for the transfer to the next coarser grid (cf. Sec. 10.2); except that the rule of not straddling the boundary layer should still be kept by those residuals, too.

For a full discussion of the discretization and multigrid procedures for the convection-diffusion operator Q , see [B17].

In case the boundary-layer interaction is important (e.g., in the driven cavity problem, where this interaction determines the entire flow), the boundary layer should be resolved. One can use then anisotropic local refinements (see Sec. 9.2), with local coordinate transformations in case the boundary is not along grid directions (Sec. 9.3).

19.3 Distributive relaxation

Generalizing the scheme in Sec. 18.3 to any operator Q^h , relaxation is guided by the distribution operator

$$M^h = \begin{pmatrix} -Q^h & 0 & \dots & 0 \\ -\partial_1^h & 1 & & \\ \vdots & & \circ & \\ -\partial_d^h & & \circ & 1 \end{pmatrix} \quad (19.5)$$

and thus proceeds as follows:

The j -th momentum equation (19.4b) is relaxed by changing \tilde{u}_j^h only, in any order and manner suitable for the operator Q^h . Thus, for $hR|\tilde{u}| \ll O(1)$ the best perhaps is the red-black (RB) Gauss-Seidel scheme: \tilde{u}_j^h is changed at each point so as to satisfy (19.4b) at that point, the red points being scanned first, the black next. For $hR|\tilde{u}| \gg 1$ this relaxation is still one of the best, except for cases of intended strong alignment described above (Sec. 19.2). In those cases, block relaxation must be used, but only in the specific strong-alignment direction, in its specific subdomain. This means streamwise line relaxation, except for plane relaxation in cases of plane alignment. And to obtain sharp numerical boundary layers one should relax in blocks (i.e., simultaneously, in one or several blocks) exactly those equations where the special Q^h

was constructed for that end (cf. Sec. 19.2). In many cases, the block relaxation can be replaced by a suitable downstream or ILU scheme (cf. Sec. 3.3). And in case of plane alignment the plane relaxation could be replaced by point relaxation if the grid is not coarsened in the direction perpendicular to the planes (cf. Secs. 4.2.1 and 3.3). Complicated alternating-direction line relaxation schemes are needed only if a fast solution is desired with errors far below truncation errors. A simple RB scheme, fully parallelizable and vectorizable, can thus most often be used.

Having relaxed in this way one pass per each momentum equation ($j = 1, \dots, d$), we then make a pass of relaxation for the continuity equation (19.4a), by scanning the cells one by one, preferably in red-black ordering. At each cell the distributive relaxation step resulting from (19.5) is a generalization of Fig. 18.2 above: Denoting by χ^h the characteristic function of the relaxed cell (i.e., $\chi^h = 1$ at that cell center, and $\chi^h = 0$ at all other cell centers), the relaxation step changes all functions \tilde{u}_j^h and \tilde{p}^h by the prescription

$$\tilde{u}_j^h + \tilde{u}_j^h - \delta h \partial_j^h \chi^h, \quad (j = 1, \dots, d) \quad (19.6a)$$

$$\tilde{p}^h + \tilde{p}^h + \delta h Q^h \chi^h, \quad (19.6b)$$

where δ is still given by (18.12).

The smoothing factor $\bar{\mu}$ is the slowest among the factors obtained for the triangular operator

$$L_M^h = \begin{pmatrix} -\Delta^h & \partial_1^h & \dots & \partial_d^h \\ 0 & Q^h & & \\ \vdots & & \text{O} & \\ 0 & & \text{O} & Q^h \end{pmatrix}. \quad (19.7)$$

Hence $\bar{\mu} = \max(\bar{\mu}^\Delta, \bar{\mu}^Q)$, where $\bar{\mu}^\Delta$ is the smoothing factor of RB relaxation for Δ^h , and $\bar{\mu}^Q$ is the smoothing factor of the relaxation of Q^h . Hence $\bar{\mu}^\Delta = .25$ if one or two sweeps are performed per cycle, while $\bar{\mu}^\Delta = .32$ in case of three sweeps per cycle. If downstream relaxation is used for the momentum equations (hence for Q^h), one can obtain $\bar{\mu}^Q \leq \bar{\mu}^\Delta$ and hence $\bar{\mu} = \bar{\mu}^\Delta$. As discussed above, however, especially in cases of widely varying stream directions, it is not important to obtain perfect smoothing by relaxing in all these directions. The relaxation rules specified above, deviating from simple (RB) point rela-

xation only in some cases of strong alignment, ensure fast reduction of all high-frequency components, except for nearly characteristic components which anyway have large truncation errors almost everywhere in the flow field (cf. Sec. 3.3). For this to hold the total viscosity coefficients $\bar{\beta}_i$ (see Sec. 19.2) should be slightly larger than the minimum, e.g. $\bar{\beta}_i = .7|u_i|$. In case the minimum $\bar{\beta}_i = .5|u_i|$ is used, as in upstream differencing, some highest frequencies are not reduced (unless downstream relaxation ordering is everywhere ensured); but this is not important from the point of view of differential smoothing (cf. Sec. 12), which will in fact be damaged if $\bar{\beta}_i$ is increased too much ($\bar{\beta}_i = .7|u_i|$ is still good). In any case, the distribution matrix (19.5) reduces the problem of relaxing the Navier-Stokes system into consideration concerning the relaxation of the scalar convection-diffusion operator Q^h (which is in detail studied in [B17]).

19.4 Multigrid procedures and numerical results

The grids, their relative positions and the interpolation procedures between them are generally the same as for the Stokes equations (Sec. 18.4). Because of the nonlinearity, FAS is of course used (see Sec. 8), and the full weighting (18.14) is preferred over (18.13) in the fine-to-coarse transfers of both the velocities and the momentum residual functions.

For $R > O(h^{-1})$, the momentum residuals themselves can be calculated in two ways, using either the same $O(h)$ approximation Q^h used in relaxation (see Sec. 19.2), or the $O(h^2)$ central approximation to Q_*^h (i.e., Q^h with zero artificial viscosity). The latter is the "double-discretization" scheme (see Sec. 10.2) which exploits the fact that Q^h is a better approximation for the high-frequency components, hence in relaxation, while Q_*^h is a better (higher) approximation for the low-frequency components which converge through the interaction with the coarse grid. Q_*^h should, however, respect intended discontinuities (boundary layers) in the same way that Q^h does, even if it means $O(h)$ local truncation errors; the global discretization errors (e.g., the L_1 -norms of $u - \tilde{u}^h$, $v - \tilde{v}^h$ and $p - \tilde{p}^h$) will still be $O(h^2)$.

For large hR it is also advisable to use W cycles or accommodative algorithms (see Sec. 6.2). On very coarse grids, where velocity changes per meshsize are comparable to the velocity itself, BGS relaxation (see Sec. 3.4 and the end of Sec. 5.6) is safer than the DGS rela-

xation described above.

Our old (1978) numerical results are reported in [D2], with two examples in [B19]. They clearly show convergence to below truncation errors in an FMG algorithm with one accomodative cycle per level, making only two relaxation sweeps on the finest grid. At the time we were worried about asymptotic convergence rates, but we should not; see end of Sec. 4.1. Recent, improved experiments feature the double discretization, BGS relaxation on very coarse grids, and W cycles. The results will be tabulated in a separate report.

19.5 Results for non-staggered grids

The non-staggered approach to the Stokes equations (Sec. 18.6) can also be used for the Navier-Stokes equations. On a non-staggered grid the short-central difference operator ∂_j^h in (19.4) should throughout be replaced by the long-central ∂_j^{2h} . Hence, ∂_j^h should be replaced by ∂_j^{2h} in the distribution operator (19.5), too (yielding M^h as in (18.20) with Δ^h replaced by Q^h). As in Stokes equations, error components around (18.18) are not efficiently reduced by relaxation; but in the Navier-Stokes case, non-linear interactions can actually amplify these components. Therefore, the pressure averaging (18.19) must now be applied not only to the final results but also to any pressure corrections, just before it is interpolated to a finer grid.

This approach was tested with RB ordering, full residual weighting (4.6), bilinear interpolation of corrections, W(2,0) cycles and an FMG algorithm with bicubic interpolation of solutions. For clarity, the table below shows results for the case of periodic boundary conditions with the smooth non-aligned solution $u = v = p = 1 + .2 \sin(\cos(x+2y))$ and $R=\infty$. Compatibility conditions were ignored (see Sec. 18.6), but the correct averages of u, v and p were enforced, so as to make the periodic problem well posed. Double discretization was used, with $\bar{B}_1 = B|u_1|$ in relaxation and $\bar{B}_2 = B_*|u_1|$ in residual transfers. The table shows the error $\|u - \tilde{u}^h\|_1 + \|v - \tilde{v}^h\|_1 + \|p - \tilde{p}^h\|_1$.

grid	(R, R _*)	(1., 1.)		(.7, .7)		(1., 0)
		1	2	1	2	1
16 x 16			.1602		.1101	.1233
32 x 32			.0809		.0583	.0309
64 x 64		.0420	.0416	.0298	.0291	.0067

20. COMPRESSIBLE NAVIER-STOKES AND EULER EQUATIONS

The steady-state Euler equations of inviscid flows will be treated here as a special, limit case of the full steady-state Compressible Navier-Stokes (CNS) equations. Some terms can be dropped in the inviscid limit, but there is no essential difference between the numerical solution of inviscid flows and that of slightly viscous flows, because $O(h)$ artificial viscosity should anyway be used in relaxation, and it should closely resemble the physical viscosity to ensure that only physical discontinuities are admitted at $h \rightarrow 0$. The double-discretization scheme (Sec. 10.2) will be used; in the present context this simply means that the artificial viscosity, when needed, is employed only in the relaxation operators, not in the difference operators by which residuals are calculated in the fine-to-coarse transfers. The latter can generally use simple central differencing, but both types of operators should respect, as much as possible, the flow discontinuities. The multigrid processes will first be described in terms of the simpler quasi-linear form of the equations, discretized on a staggered grid, and then their modification to the conservation form and to non-staggered discretization will be discussed.

20.1 The differential equations

20.1.1 Conservation laws and simplification. The time-dependent Compressible Navier-Stokes (CNS) equations in d dimensions may be written in conservation form as

$$\frac{\partial W}{\partial t} + \sum_{j=1}^d \partial_j F_j = \underline{f} \tag{20.1}$$

where $\partial_j = \frac{\partial}{\partial x_j}$ and

$$W = \begin{pmatrix} \rho u_1 \\ \vdots \\ \rho u_d \\ \rho \\ e \end{pmatrix}, \quad F_j = \begin{pmatrix} \rho u_j u_1 + \tau_{j1} \\ \vdots \\ \rho u_j u_d + \tau_{jd} \\ \rho u_j \\ e u_j + \sum_i u_i \tau_{ij} - \kappa \partial_j \epsilon \end{pmatrix}, \quad \underline{f} = \begin{pmatrix} f_1 \\ \vdots \\ f_d \\ f_\rho \\ f_\epsilon \end{pmatrix}$$

$$\tau_{ij} = \tau_{ji} = -\nu(\partial_i u_j + \partial_j u_i) - \lambda \delta_{ij} \sum_{k=1}^d \partial_k u_k + p \delta_{ij}, \tag{20.2}$$

ρ being the fluid density, $\underline{u} = (u_1, \dots, u_d)$ the velocity vector, e the total energy per unit volume, ϵ the specific internal energy, p the pressure, λ and μ the viscosity coefficients, and κ the coefficient of heat conductivity. (20.1) is a system of $d + 2$ equations: the first d equations are the d momentum equations, next is the continuity (mass conservation) equation, and the last is the energy equation. The $d + 2$ basic unknowns may be considered to be \underline{u} , ρ and ϵ , in terms of which e is given by

$$e = \rho \left(\epsilon + \frac{q^2}{2} \right), \quad q^2 = \sum_{i=1}^d u_i^2, \quad (20.3)$$

and p by the equation of state

$$p = \bar{p}(\epsilon, \rho). \quad (20.4)$$

For a perfect gas, for example, the equation of state is $p = (\gamma - 1)\epsilon\rho$, where γ is the ratio of specific heats. Generally, $p_\epsilon = \partial\bar{p}/\partial\epsilon$ and $p_\rho = \partial\bar{p}/\partial\rho$ are positive. The coefficients λ , μ and κ are given functions, usually, of ϵ (or functions of the fluid temperature, which in turn is a function of ϵ). We will treat these coefficients as constants, since they change slowly. The whole discussion below will remain precisely valid as long as any change in any of these coefficients over a meshsize is small compared with the coefficient itself.

In most aerodynamical applications the right-hand side \underline{f} vanishes, but there are other applications where the external body force (f_1, \dots, f_d) , or the mass source f_ρ , or the energy source f_ϵ do not vanish. A general \underline{f} is assumed here, mainly because many of the numerical experiments are set to have known specified solutions by prearranging \underline{f} accordingly.

The steady-state CNS equations are given by

$$\sum_{j=1}^d \partial_j F_j = \underline{f} \quad (20.5)$$

together with (20.3) and (20.4). It will be convenient below to substitute (20.3) into equations (20.5), but to treat (20.4) as an additional equation and p as an additional unknown. Allowing (20.4) not to be satisfied until convergence will substantially simplify the solution process, and more than justify the additional storage needed for p . Thus we will have a system of $n = d + 3$ differential equations in n unknown functions. The n -vector $U = (u_1, \dots, u_d, \rho, \epsilon, p)$ will serve as our vector of unknowns.

Simplified equations, but not in conservation form, are obtained as follows. The i -th momentum equation is simplified by subtracting from it u_i times the continuity equation. Then the energy equation is simplified by subtracting from it u_i times the i -th (simplified) momentum equation for $i = 1, \dots, d$, and e/ρ times the continuity equation. The resulting system is

$$\rho \sum_j u_j \partial_j u_i + \sum_j \partial_j \tau_{ij} = f_i \quad (i = 1, \dots, d) \quad (20.6a)$$

$$\sum_j \partial_j (\rho u_j) = f_\rho \quad (20.6b)$$

$$\rho \sum_j u_j \partial_j \epsilon - \sum_j \partial_j (\kappa \partial_j \epsilon) + \sum_{i,j} \tau_{ij} \partial_i u_j = f_\epsilon \quad (20.6c)$$

$$p = \bar{p}(\epsilon, \rho) \quad (20.6d)$$

which, in view of (20.2), is a system of n equations for the n unknowns U . In terms of this simpler system we will now study the principal and the inviscid subprincipal parts. This will tell us what boundary conditions are appropriate and which terms are locally dominant, which is important for designing the relaxation scheme and the form of the artificial viscosity terms.

20.1.2 The principal part of (20.6), i.e., the part of the linearized operator which contributes to the highest-order terms of its determinant, is the operator

$$L_p = \begin{pmatrix} -\mu\Delta - \bar{\lambda}\partial_{11} & \cdots & -\bar{\lambda}\partial_{1d} & 0 & 0 & 0 \\ \vdots & \ddots & \vdots & \vdots & \vdots & \vdots \\ -\bar{\lambda}\partial_{d1} & \cdots & -\mu\Delta - \bar{\lambda}\partial_{dd} & 0 & 0 & 0 \\ 0 & \cdots & 0 & \underline{u} \cdot \underline{\nabla} & 0 & 0 \\ 0 & \cdots & 0 & 0 & -\kappa\Delta & 0 \\ 0 & \cdots & 0 & 0 & 0 & 1 \end{pmatrix} \quad (20.7)$$

where $\partial_{ij} = \partial_i \partial_j$, $\Delta = \partial_{11} + \dots + \partial_{dd}$, $\underline{u} \cdot \underline{\nabla} = u_1 \partial_1 + \dots + u_d \partial_d$ and $\bar{\lambda} = \lambda + \mu$. It can be shown by dimensional analysis that on a small enough scale the behavior of solutions to the original system depends only on L_p . Thus, on a sufficiently small scale, viscosity is the main mechanism that determines velocities, convection determines density, diffusion determines the internal energy, and they are all loca-

lly independent. It also follows from (20.7) that the full CNS system requires d boundary conditions for the velocities (usually \underline{u} is given) and one for ϵ (usually given in the form of boundary temperature or temperature gradient) on all boundaries, and an additional condition for ρ (or for p) at one end of each streamline.

20.1.3 Elliptic singular perturbation. Since $\underline{u} \cdot \underline{\nabla}$ is a factor of $\det L_p$, the steady-state CNS system is not elliptic: The streamlines are its characteristic lines (the *only* characteristics, as long as the flow is viscous). This means that an addition of artificial h-ellipticity will be needed in local numerical processes, unless the grid exactly aligns with the stream. Therefore, and for uniform treatment of all artificial terms in the inviscid limit, we will regard already the CNS differential system as a limit, $\nu \rightarrow 0$ say, of an elliptic system, obtained by adding a singular-perturbation term to the continuity equation (20.6b), rewriting it (times ρ) as

$$\sum_{j=1}^d [\rho \partial_j (\rho u_j) - \partial_j (\nu \partial_j \rho)] = f_\rho. \quad (20.8)$$

ν should be positive for the additional term to be compatible (i.e., give a well-posed problem together) with the time-dependent system (20.1), so as to make the limits $\nu \rightarrow 0$ and $t \rightarrow \infty$ interchangeable. With $\nu > 0$, the additional term indeed represents a physical effect, namely, static molecular diffusion, which could normally be neglected.

The system (20.6) with (20.8) replacing (20.6b) is called the augmented CNS (ACNS) system. Its principal-part determinant is $\kappa \nu \mu \Delta^{d+2}$, so it is indeed elliptic. It requires the same $d+1$ boundary conditions (on \underline{u} and ϵ) as before, plus a boundary condition on ρ (or p), on all boundaries. But the sign of ν ensures that the latter condition will affect the solution in the limit $\nu \rightarrow 0$ only at points where the flow enters the domain. At non-entry boundaries, an artificial boundary layer (discontinuity in the limit) would be formed; but it can be avoided by using the original continuity equation (20.6b) as the extra boundary condition at such points. If all boundaries are such, however, we will have only gradient conditions on ρ along the boundaries, hence we will need an extra integral condition to uniquely determine the solution to the ACNS system, and also, in the limit, to the CNS system. This condition is usually the total mass, or some equivalent datum. Indeed, if the flow nowhere enters the domain, rigid walls are all around, then the total mass is determined

only by the initial conditions, and therefore should be added as an extra condition to the steady-state equations.

20.1.4 Inviscid (Euler) and subprincipal operators. The inviscid case (Euler equations) is the system (20.6) with vanishing viscosities and heat conduction: $\lambda = \mu = \kappa = 0$. More precisely, the flow is inviscid (free of viscous and heat-conduction effects) where λ , μ and κ are small compared with $\rho l q$, where l is a typical length of change of \underline{u} and ϵ . Usually there will be some particular narrow subdomains, such as boundary layers, where l will be just small enough to make the flow viscous. Thus, viscosity effects can seldom be completely neglected. Anyway, wherever the flow is inviscid, the scale where viscosity dominates is much smaller than the scale of changes in the flow, which will later also be the scale of our grid h . So we like to isolate the terms which dominate the flow in that intermediate scale (small scale in terms of the flow geometry, but large enough to neglect viscosity and heat conduction). These are the sub-principal terms, defined as all the terms that are either principal or become principal when λ , μ and κ , or some of them, vanish. They form the following sub-principal operator L_s

$$L_s = \begin{pmatrix} Q_\mu - \bar{\lambda} \partial_{11} & \cdots & -\bar{\lambda} \partial_{1d} & 0 & 0 & \partial_1 \\ \vdots & \ddots & \vdots & \vdots & \vdots & \vdots \\ -\bar{\lambda} \partial_{d1} & \cdots & Q_\mu - \bar{\lambda} \partial_{dd} & 0 & 0 & \partial_d \\ \rho^2 \partial_1 & \cdots & \rho^2 \partial_d & Q_\nu & 0 & 0 \\ p \partial_1 & \cdots & p \partial_d & 0 & Q_\kappa & 0 \\ 0 & \cdots & 0 & -p_\rho & -p_\epsilon & 1 \end{pmatrix}, \quad (20.9)$$

where generally

$$Q_a = -\underline{\nabla} \cdot (a \underline{\nabla}) + \rho \underline{u} \cdot \underline{\nabla}. \quad (20.10)$$

This is the operator that should be kept in mind in the local processing, such as relaxation, and in the choice of discretization to be used with relaxation. The coefficients \underline{u} , ρ and p appearing in L_s are actually the values of some solution around which the flow is examined through principal linearization (see Sec. 3.4); they will always be derived from the current approximate solution \tilde{U} (see Sec. 20.3.5). It can always be assumed that the current approximation is

close enough to the solution, by employing continuation if necessary (see Sec. 8.3.2). The determinant of L_S , developed by its last row, is

$$\det L_S = Q_\mu^{d-1} (Q_\kappa Q_\nu (Q_\mu - \bar{\lambda} \Delta) - (\rho^2 p_\rho Q_\kappa + p p_\epsilon Q_\nu) \Delta). \quad (20.11)$$

The reduced (principal) operator L_R is defined as L_S for $\lambda = \mu = \kappa = \nu = 0$, i.e., the principal part of the inviscid limit.

$$L_R = \begin{pmatrix} Q_0 & & & 0 & 0 & \partial_1 \\ & \ddots & & \vdots & \vdots & \vdots \\ & & \ddots & & & \\ & & & Q_0 & 0 & 0 & \partial_d \\ \rho^2 \partial_1 & \dots & \rho^2 \partial_d & Q_0 & 0 & 0 \\ p \partial_1 & \dots & p \partial_d & 0 & Q_0 & 0 \\ 0 & \dots & 0 & -p_\rho & -p_\epsilon & 1 \end{pmatrix} \quad (20.12)$$

$$\det L_R = Q_0^d (Q_0^2 - \rho^2 a^2 \Delta) = \rho^{d+2} (\underline{u} \cdot \underline{v})^d [(\underline{u} \cdot \underline{v})^2 - a^2 \Delta], \quad (20.13)$$

where $a = (p_\rho + \rho^{-2} p p_\epsilon)^{1/2}$ is the speed of sound. (In the time-dependent inviscid case the operator in brackets in (20.13) is replaced by $[(\frac{\partial}{\partial t} + \underline{u} \cdot \underline{\nabla})^2 - a^2 \Delta]$, showing that a is the speed relative to the flow at which small disturbances would be propagated.) The ratio $M = q/a = (\underline{u} \cdot \underline{u})^{1/2}/a$ is called the Mach number. Where $M < 1$ the flow is called subsonic, where $M > 1$ it is called supersonic, and the line where $M = 1$ is the sonic line. We can see from (20.13) that the steady-state inviscid supersonic equations are hyperbolic, regarding the stream direction \underline{u} as the time-like direction, with three families of characteristic lines and three characteristic speeds: q , $|a-q|$ and $|a+q|$. The steady-state inviscid subsonic equations are neither hyperbolic nor elliptic, and have only one family of characteristic lines, namely, the streamlines. In either case the equations are of order $d + 2$, hence require $d + 2$ boundary conditions per streamline. The only restriction on these imposed by (20.13) is that, in the subsonic case, at least one condition should be given all along the boundary (on both sides of each streamline). Actually the situation is more complicated since the flow can be partly subsonic and partly supersonic and, more importantly, acceptable solutions of the inviscid equations are only those obtainable as limits of solutions of

the viscous equations. This latter requirement determines which of the boundary conditions of the full CNS equations will affect the inviscid flow away from the boundary, and which will be ignored in the inviscid limit, creating a discontinuity (boundary layer). It also determines what type of discontinuities (shocks) are permissible in the interior. The derived rules for permissible discontinuities are sometimes expressed as "entropy conditions".

Our approach here to the inviscid case will generally be to imitate the physics. Instead of deriving entropy conditions and then imposing them numerically, and instead of getting fully into the question of correct boundary conditions, we will locally (i.e., in relaxation) use a numerical scheme which contain artificial viscosity exactly analogous to the physical viscosity, thus ensuring correct selection of discontinuities. Our local processes need artificial viscosity anyway, to eliminate high-frequency errors.

The required magnitude of the artificial viscosity coefficients can be seen from $\det L_S$. They should be as effective on the scale of the meshsize h as other terms, hence, regarding each differentiation symbol as $O(h^{-1})$, the coefficients $(\lambda, \mu, \nu, \kappa)$ should be chosen so as to make the order of $\det L_S$ homogeneous in h . It is easy to see that this is obtained if and only if the artificial λ, μ, ν and κ are all $O(h|\underline{u}|)$, in which case $\det L_S$ is homogeneously of order h^{-d-2} .

20.1.5 Incompressible and small Mach limits. The incompressible limit is the case of vanishingly small $\partial \rho / \partial p$, or indefinitely large p_ρ . The main operator in this case is the cofactor of p_ρ in (20.9) (reducing by one the number of equations, corresponding to the fact that ρ is no longer unknown). The resulting system is reducible: the momentum and continuity equations form a separate system of equations for \underline{u} and p , easily simplified to (19.3) above. An exactly analogous situation arises if p_ϵ , instead of p_ρ , is large, and, more generally, whenever the Mach number is small. Thus, if one develops a multi-grid solver for cases which include regions with small Mach, its discretization and relaxation should be efficient in the incompressible limit (19.3).

20.2 Stable staggered discretization

20.2.1 Discretization of the subprincipal part. The stable discretization constructed here is for use in relaxation processes, and will thus (see Sec. 2.1) be based on considerations concerning the subprincipal operator (20.9) with the coefficients \underline{u} , ρ , p , P_ρ , P_ϵ regarded as fixed. For all admissible values of these coefficients, the discretization and relaxation processes should be uniformly effective. In particular, they should accommodate the incompressible limit (cf. Sec. 20.1.5). We will therefore use a staggered grid as in Secs. 18.2 and 19.2 (non-staggered grids are discussed in Sec. 20.6), with ρ^h , ϵ^h and p^h (the discrete approximations to ρ , ϵ and p) defined at

$$x_k^h = x_0^h + (k_1 h_1, \dots, k_d h_d) \quad (\text{cell centers}), \quad (20.14)$$

and u_i^h (approximating u_i) at

$$x_k^{h,i} = x_k^h + .5h_{-i} \quad (i\text{-face center}), \quad (20.15)$$

where $k = (k_1, \dots, k_d)$ are vectors of integers, $h_1 \times \dots \times h_d$ is the size of the grid-h cells, $h_{-i} = h_i e_{-i}$, and e_{-i} is the i -th coordinate unit vector. In two dimensions ($d = 2$) the staggering is depicted in Fig. 20.1.




FIGURE 20.1 Grid staggering for compressible Navier-Stokes discretization

With this staggering (20.9) can indeed be approximated by replacing ∂_i by its *short* central analog ∂_i^h ,

$$\partial_i^h \phi(x) = [\phi(x + .5h_{-i}) - \phi(x - .5h_{-i})] / h_i,$$

replacing ∂_{ij} by $\partial_{ij}^h = \partial_i^h \partial_j^h$, and each Q_α by a proper approximation Q_α^h . Calling the resulting operator L_s^h , we find, similar to (20.11),

$$\det L_s^h = (Q_u^h)^{d-1} \{ Q_k^h Q_v^h (Q_u^h - \bar{\lambda} \Delta^h) - (\rho^2 P_\rho Q_k^h + P P_\epsilon Q_v^h) \Delta^h \}, \quad (20.15)$$

where $\Delta^h = \sum_{jj} \partial_{jj}^h$ is the usual $(2d+1)$ -point discrete Laplacian. Thus, the approximation is h -elliptic provided Q_u^h , Q_v^h and Q_k^h are h -elliptic.

The approximation of these diffusion-convection operators is generally similar to that of Q^h in Sec. 19.2, with ρ/μ replacing R , with $O(h\rho|u|)$ replacing $O(h|u|)$, and with some modifications in case of shocks. Such modifications, introduced of course in the conservative formulation (see Sec. 20.6), have been studied for the nonlinear convection-diffusion equation $-\epsilon \Delta u + \partial_1(u^2) + \partial_2 u = f$, but are not yet implemented in the CNS programs. The main emerging rule, as with boundary layers, is to avoid differencing across discontinuities: The stronger the shock the more precisely rotated and upwinded is the calculation of fluxes, so the more weakly it straddles the shock. Where boundary layer and shocks are not present, we will thus approximate any Q_α at any gridpoint x by

$$Q_\alpha^h(x) = \frac{d}{\sum_{j=1}^d} \left(-\partial_j^h (\alpha_j^h(x) \partial_j^h) + \bar{\rho}^h(x) \bar{u}_j^h(x) \partial_j^{2h} \right), \quad (20.16)$$

where $\bar{\rho}^h(x)$ and $\bar{u}_j^h(x)$ are central averages of ρ^h and u_j^h , respectively, over points nearest to x , and

$$\alpha_j^h(x) = \max[\alpha, \beta h_j \rho^h |u_j^h|]. \quad (20.17)$$

The \max in (20.17) is taken in principle over all values of h_j , ρ^h and u_j^h in some neighborhood of x . Usually any neighboring values, not necessarily exactly maximal ones, can be taken, except near a stagnation point, where \underline{u} nearly vanishes. β is an $O(1)$ parameter: $\beta = \frac{1}{2}$ normally gives upstream differencing, but slightly larger values ($\beta = .6$ or $.7$) may give better results (cf. [B17]).

20.2.2 The full quasi-linear discretization. Guided by the above discretization scheme for the sub-principal part, the discrete approximation to the full CNS system (20.6) on the staggered grid (20.14)-(20.15) will be written as

$$Q_{\mu}^h u_i^h - \bar{\lambda} \sum_j \partial_{ij}^h u_j^h + \partial_i^h p^h = f_i^h \quad \text{at } x_k^{h,i}, \quad (i=1, \dots, d) \quad (20.18a)$$

$$Q_{\nu}^h \rho^h + (\rho^h)^2 \sum_j \partial_j^h u_j^h = f_{\rho}^h \quad \text{at } x_k^h \quad (20.18b)$$

$$Q_{\kappa}^h \epsilon^h + p^h \sum_j \partial_j^h u_j^h - B^h(\underline{u}^h) = f_{\epsilon}^h \quad \text{at } x_k^h \quad (20.18c)$$

$$p^h - \bar{p}(\epsilon^h, \rho^h) = 0 \quad \text{at } x_k^h, \quad (20.18d)$$

where all Q_{α}^h are defined by (20.16)-(20.17) and $B^h(\underline{u}^h)$ is the simplest central approximation to

$$B(\underline{u}) = \mu \sum_{i,j} (\partial_i u_j + \partial_j u_i) \partial_j u_i + \lambda (\sum_i \partial_i u_i)^2. \quad (20.19)$$

The exact form of B^h is not important since it is neither a principal nor a subprincipal term. $f^h = (f_1^h, \dots, f_d^h, f_{\rho}^h, f_{\epsilon}^h)$ are some local averages of f ; injection $f^h = f$ is usually used, except in some cases where this fails to give good enough approximations on the finest grid (relative to the grid-2h solution, whose right-hand side is $\underline{f}^{2h} = \frac{1}{h} \underline{f}^h$).

The scheme (20.18) will be used in the relaxation processes. The same scheme, but without the artificial viscosity terms ($\alpha_j^h = \alpha$ in (20.16)), will be used to calculate residuals transferred to coarse grids, thus making the overall approximation $O(h^2)$.

In the inviscid case ($\lambda, \mu, \kappa \ll \rho h |\underline{u}|$) the term with $\bar{\lambda}$ in (20.18a) and the term $B^h(\underline{u}^h)$ in (20.18c) may be omitted. They do not contribute to the h-ellipticity of the system. The resulting scheme is nothing but Euler equations with a simple form of artificial viscosity, derived from the viscous (Navier-Stokes) equations.

20.2.3 Boundary conditions. At this stage of development, to separate away various algorithmic questions (cf. Sec. 4), the numerical experiments were conducted with known smooth solutions U , employing first periodic boundary conditions and then Dirichlet conditions in two dimensions ($d = 2$). In the periodic case no boundaries are actually present; gridpoint (x_1, x_2) is simply identified with gridpoint

$(x_1+2\pi, x_2+2\pi)$. This enables us later to check that no slow-down is caused by boundary conditions. The Dirichlet case at this stage is the square domain $\{|x|, |y| < \pi\}$, with \underline{u} , ρ and ϵ given on its boundary. The staggered grid is square and is placed so that the boundary of the domain coincides with cell boundaries; i.e., $h_1 = h_2 = h = 2\pi/N$ and $x_0^h = (h/2, h/2)$ (see (20.14)).

Moreover, to simplify the code development, the Dirichlet boundary conditions are at this stage placed not exactly on the boundary but in their natural staggered-grid positions. For example, ρ and ϵ are specified at the cell centers immediately outside the boundary, i.e., on the lines $\{|x|, |y| = 2\pi + h/2\}$. This is easy to do at this stage since the numerical tests are made with known solutions U , whose values are in particular known on that staggered boundary. Ultimately, these staggered-boundary conditions will be obtained by (quadratic) extrapolation from the real boundary conditions as well as interior values; the present type of conditions is only employed in order to separate away questions as to exactly when and how this extrapolation should be made.

In the inviscid case some of the above conditions are redundant, but the code can handle this automatically (cf. Sec. 20.1.4).

20.3 Distributed relaxation for the simplified system

20.3.1 General approach to relaxation design. Since the problem at hand is not elliptic, one should not attempt obtaining "perfect smoothers" (see Secs. 3.3, 7 and end of Sec. 4.1). So the question is how to use the usual measure of relaxation performance, namely, the smoothing factor $\bar{\mu}$, in selecting the relaxation scheme. We do it by dividing the relaxation design into the following three stages.

(A) First, a *pointwise* (not-block) and direction-free relaxation, with low $\bar{\mu}$ (per operation, of course, and with the other considerations of Sec. 3.2) is constructed for the fully h-elliptic operator L_e^h obtained from L^h when sufficiently large and *isotropic* artificial viscosity terms are used. This means replacing (20.17) by

$$\alpha_j^h = \max\{\alpha, \beta \rho^h \max_{\ell} h_{\ell} |u_{\ell}^h|\}, \quad (20.20)$$

and choosing β just large enough to make excellent $\bar{\mu}$ obtainable independently of relaxation marching directions. It means that β

should be appreciably larger than the minimal value $\beta = .5$ needed for stability of simple relaxation schemes. For example one can take for this purpose $\beta = 1$; any larger value of β will not substantially change the smoothing factors, and will especially not change the comparison between different relaxation schemes.

(B) Having designed the relaxation scheme, it is then actually used with the anisotropic artificial viscosities (20.17), rather than (20.20). Note that if the flow is not (nearly) aligned with the grid, there is no fundamental difference between the two. If the flow is aligned with the grid, there is one kind of high-frequency error components V which are not deflated in the anisotropic as in the isotropic case because $L^h V$ is much more closely singular than $L_e^h V$: These are the "characteristic components", i.e. high-frequency components which are smooth in the flow direction. When the flow is not aligned with the grid, neither scheme approximate these components well. If there is alignment, only the anisotropic scheme approximates them well, but exactly for this scheme and these components the pointwise relaxation is not effective. (The effectiveness of relaxation in the isotropic case is obtained for a characteristic component at the price of not really approximating its amplitude in the differential solution, which makes the fast convergence for this component meaningless; cf. Sec. 12). It is meaningful to get good smoothing for the characteristic components only if the alignment is strong (i.e., intended and consistent; see Sec. 3.3), and exactly then it is possible to do that, via block relaxation.

Thus, the pointwise relaxation scheme, designed in terms of the isotropic artificial viscosities (20.20), is actually used with the anisotropic viscosities (20.17), and it is modified to the corresponding block scheme in case of strong alignment. "Corresponding" means that the same distribution matrix M^h , and a similar relaxation ordering, are maintained while the "ghost functions" w_1 (see Sec. 3.7) are relaxed in blocks instead of pointwise, where the blocks are in the specific strong-alignment direction.

(C) To this basic relaxation scheme, several improvements can be added. First, the value of β can be lowered, either experimentally, or theoretically through the modified smoothing range (12.1). (As explained there (following Eq. (12.1)), if β is lowered to near its minimal stable value ($\beta = .5$), the final result may be improved by averaging. Also note that near $\beta = .5$, downstream relaxation orde-

ring may become badly divergent for some special smooth component which may not show up in one experiment but badly affect another; see Sec. 20.3.4). Then the equations may more precisely be rotated and upwinded near strong shocks, double discretization schemes may be introduced (see Sec. 10.2), etc.

In the two latter stages, (B) and (C), efficiency should of course mainly be measured not in terms of asymptotic factors, but in terms of FMG performance, whether experimentally (see Sec. 1.6) or theoretically (see Secs. 7.4, 7.5).

20.3.2 Possible relaxation scheme for inviscid flow. In the inviscid case the principal difference operator is L_r^h (cf. (20.12)), and the usual distribution operator (cf. Secs. 3.7 and 19.3) would be

$$M_r^h = \begin{pmatrix} 1 & & & & & -\partial_1^h Q_0^h \\ & \cdot & & & & \vdots \\ & & \cdot & & & -\partial_d^h Q_0^h \\ & & & \circ & & -\rho^2 \Delta^h \\ & & & & 1 & -p \Delta^h \\ & & & & & (Q_0^h)^2 \\ 0 & \dots & 0 & 0 & 0 & \end{pmatrix}, \quad (20.21a)$$

where the last column is made of the cofactors of the last row in L_r^h , divided by their common factor $(Q_0^h)^2$. Since

$$L_r^h M_r^h = \begin{pmatrix} Q_0^h & & & & 0 & 0 & 0 \\ & \cdot & & & \vdots & \vdots & \vdots \\ & & \cdot & & Q_0^h & 0 & 0 \\ & & & \circ & & 0 & 0 \\ \rho^2 \partial_1^h & \dots & \rho^2 \partial_d^h & & Q_0^h & 0 & 0 \\ p \partial_1^h & \dots & p \partial_d^h & & 0 & Q_0^h & 0 \\ 0 & \dots & 0 & -p_\rho & -p_e & (Q_0^h)^2 - \rho^2 a^2 \Delta^h & \end{pmatrix} \quad (20.21b)$$

the relaxation process is essentially decoupled into relaxing the convection operator Q_0^h and, separately, the potential-flow operator $L_{pot}^h = (Q_0^h)^2 - \rho^2 a^2 \Delta^h$. The former has been discussed in Sec. 19.3; the latter is currently under renewed study via numerical experiments with potential flows (see remarks in Sec. 21).

creased viscosities is the special case of large Mach numbers and downstream relaxation, a case which is not direction-free and which is also ruled out below for other reasons.) Hence, examined here is mainly the inviscid case $\mu = \bar{\lambda} = \kappa = \nu = 0$.

By (20.23), the relaxation can be performed as $d + 1$ separate passes, d of them with the operator Q_0^h , and one with the system (20.24). Hence $\bar{\mu} = \max(\bar{\mu}^Q, \bar{\mu}')$, where $\bar{\mu}^Q$ is the smoothing factor of relaxing Q_0^h and $\bar{\mu}'$ is the smoothing factor in the collective relaxation of (20.24). Both depend on the order in which the corresponding passes are made, on the value of β in (20.20), and on the direction of \underline{u} , and $\bar{\mu}'$ also depends on the Mach number M . This dependence is shown in Table 20.1. Notice that, except for RB(2), $\bar{\mu} = \bar{\mu}' \geq \bar{\mu}^Q$, and in most cases $\bar{\mu} \approx \bar{\mu}' \approx \bar{\mu}^Q$ (assuming the same ordering is used in all passes). This results from the fact that Q_0^h is a divisor of the determinant of (20.24).

The table shows it very efficient to use red-black (RB) ordering in all passes. First, because of its usual advantage of being vectorizable (cf. Sec. 3.6). Secondly, for sufficiently large β (and hence also for viscous flows), RB smoothing rates are superior to all others. Only when β approaches its minimal value (e.g., for $\beta \leq 1$), downstream ordering (e.g., lexicographic ordering in case all $u_i \geq 0$) shows better smoothing rates; but this is not direction-free: it is complicated to maintain downstream ordering in case the flow directions (i.e., the signs of u_1, \dots, u_d) change with location. Also, the components for which RB is not so good at smaller β are exactly the characteristic components one needs not care about (see Sec. 20.3.1).

Moreover, the main advantage of the downstream relaxation is shown as super-fast smoothing in case of small β and large M ; indeed, $\bar{\mu} \rightarrow 0$ as $\beta \rightarrow .5$ and $M \rightarrow \infty$. But this small $\bar{\mu}$ shows only the behavior of high-frequency components; there are low-frequency ones for which exactly this relaxation badly diverges: the amplification is about $2^{-4}M$ for a component $(\theta_1, \theta_2) \approx 2^{-4}M^{-1}(u_1, u_2)/q$. This does not happen when the slower schemes are used, such as RB, or even the downstream ordering with larger β ($\beta \geq .8$, say): For these latter schemes some low-frequency components may still diverge, but the divergence is slow and easily corrected by the multigrid coarse-grid corrections.

Table 20.1 also shows symmetric ordering (SGS) to yield a very efficient smoother. The bad behavior of low frequencies for large M

Relaxation ordering	M	$(u_1, u_2) = (1, 1)$			$(u_1, u_2) = (0, 1)$		
		2.0	1.0	0.5	2.0	1.0	0.5
RB(1)	*	.27	.33	1.00	.27	.31	.50
	0.0-0.1	.27	.43	1.00	.27	.31	.50
	1.0-5.0	.40-.41	.46-.48	1.00	.27-.29	.33-.36	.50-.55
RB(2)	*	.39	.50	1.00	.36	.39	.42
	0.0-0.1	.27	.65	1.00	.27	.31	.50
	1.0-5.0	.33-.35	.41-.42	1.00	.27-.29	.33-.39	.50-.55
Lex+	*	.38	.25	0	.46	.42	.47
	0.0-1.0	.44-.50	.42-.50	.42-.50	.48-.50	.47-.50	.49-.50
	2.0	.41	.35	.24	.47	.45	.48
	5.0	.39	.29	.056	.46	.44	.50
Lex-	*	.63	.75	1.00	.58	.66	.72
	0.0-5.0	.63	.75-.76	1.00	.58	.66-.67	.84-.86
Lex+-	*	.54	.63	1.00	.58	.66	.72
	0.0-5.0	.54-.55	.63-.64	1.00	.58	.66-.67	.84-.86
SGS	*	.48	.42	0	.50	.49	.49
	0.0-0.1	.50	.50	.50	.50	.50	.50
	1.0	.51	.54	.60	.51	.54	.60
	5.0	.49	.46	.24	.50	.51	.53

TABLE 20.1 Smoothing factors for two-dimensional Euler equations $\bar{\mu} = \max(\bar{\mu}^Q, \bar{\mu}')$. In the rows where * stands for M , $\bar{\mu}^Q$ is displayed; in the others $\bar{\mu}'$ is displayed. RB(i) is red-black ordering with i sweeps per cycle ($\bar{\mu}'$ derived by Eq. (3.2)). Lex+ is lexicographic ordering, Lex- is the reversed lexicographic, Lex+- is lexicographic with only the y coordinate reversed, and SGS is Symmetric GS (a Lex+ pass alternating with a Lex- one, $\bar{\mu}$ calculated per pass).

and $\beta \approx .5$ is still theoretically present here, although more weakly. One can eliminate this trouble simply by switching to larger β (e.g., $\beta = 1$) wherever M is large and the relaxation marching happens to be downstream. In fact, there is no reason to use the same β at the same location in all passes, especially as still another value ($\beta = 0$) is used for fine-to-coarse residual transfers.

20.3.5 Summary: relaxation of the full system. The scheme outlined above for the subprincipal part L_s^h easily translates to the following relaxation procedure for the full quasi-linear system (20.18). We denote by $\tilde{U}^h = (\tilde{u}^h, \tilde{\rho}^h, \tilde{\epsilon}^h, \tilde{p}^h)$ the dynamic approximation to the solution $U = (u, \rho, \epsilon, p)$, i.e., the approximation stored there just prior to any described relaxation step, while $U^h = (u^h, \rho^h, \epsilon^h, p^h)$ will denote the exact solution of the stable difference equations (20.18). $r^h = (r_1^h, \dots, r_d^h, r_\rho^h, r_\epsilon^h, r_p^h)$ will denote the dynamic residuals, i.e. the left-hand sides of (20.18) applied to \tilde{U}^h instead of U^h and subtracted from the right-hand sides. Thus, at each step of each relaxation pass \tilde{U}^h and r^h change.

The relaxation steps are first described for the case where there is no strong alignment between the grid and the flow direction.

A relaxation sweep consists of $d + 1$ passes. The recommended ordering within each pass is either the red-black ordering (which we used throughout our numerical experiments so far) or the symmetric lexicographic (SGS).

First, one pass is made for each momentum equation. The i -th equation at the point $x_k^{h,i}$ is relaxed by the replacements

$$\tilde{u}_i^h + \tilde{u}_i^h + \psi_{i,k}^h \quad (20.25a)$$

$$\tilde{p}^h + \tilde{p}^h + \bar{\lambda} \beta_i^h \psi_{i,k}^h \quad (20.25b)$$

where $\psi_{i,k}^h$ is a function defined at all i -face centers (see (20.15)), vanishing on all of them except at the relaxed point $x_k^{h,i}$. Its value at $x_k^{h,i}$ is determined so that

$$\tilde{Q}_\mu^h \psi_{i,k}^h(x_k^{h,i}) = r_i^h(x_k^{h,i}) \quad (20.26)$$

where the coefficients of \tilde{Q}_μ^h are evaluated at $x_k^{h,i}$ (see (20.16)), based on \tilde{U}^h . Neglecting non-principal effects this means that after

the changes (20.25), the i -th momentum equation (20.18a) at $x_k^{h,i}$ will be satisfied.

Then, a pass is made on cell centers. The three equations defined at the cell center x_k^h are relaxed simultaneously by replacements of the form

$$\tilde{u}_i^h + \tilde{u}_i^h - \partial_i^h \phi_k^h, \quad (i = 1, \dots, d) \quad (20.27a)$$

$$\tilde{p}^h + \tilde{p}^h + Q_{\mu+\bar{\lambda}}^h \phi_k^h \quad (20.27b)$$

$$\tilde{\rho}^h + \tilde{\rho}^h + \hat{\rho}_k^h \quad (20.27c)$$

$$\tilde{\epsilon}^h + \tilde{\epsilon}^h + \hat{\epsilon}_k^h \quad (20.27d)$$

where ϕ_k^h , $\hat{\epsilon}_k^h$ and $\hat{\rho}_k^h$ are functions defined at cell centers (see (20.14)), vanishing everywhere except at the relaxed point x_k^h . At x_k^h their values are chosen so as to satisfy the three equations

$$(\tilde{Q}_\nu^h \hat{\rho}_k^h - \tilde{\rho}^2 \Delta^h \phi_k^h)(x_k^h) = r_\rho^h(x_k^h) \quad (20.28a)$$

$$(\tilde{Q}_\epsilon^h \hat{\epsilon}_k^h - \tilde{p} \Delta^h \phi_k^h)(x_k^h) = r_\epsilon^h(x_k^h) \quad (20.28b)$$

$$(-\tilde{p}_\rho \hat{\rho}_k^h - \tilde{p}_\epsilon \hat{\epsilon}_k^h + \tilde{Q}_{\mu+\bar{\lambda}}^h \phi_k^h)(x_k^h) = r_p^h(x_k^h) \quad (20.28c)$$

where $\tilde{\rho}$, \tilde{p} , \tilde{p}_ρ , \tilde{p}_ϵ and the coefficients of \tilde{Q}_ν^h , \tilde{Q}_ϵ^h and $\tilde{Q}_{\mu+\bar{\lambda}}^h$ are all evaluated at x_k^h , based on \tilde{U}^h .

The functions $\psi_{i,k}^h$, ϕ_k^h , $\hat{\rho}_k^h$ and $\hat{\epsilon}_k^h$ mentioned above do not actually appear in the program, of course. They just serve to concisely describe the relaxation steps.

Instead of the separate $d + 1$ passes just described, they could of course be merged in any desired fashion.

In case of strong alignment, i.e., if one grid direction nearly coincides with the flow direction throughout a major subdomain, relaxation should be done in the corresponding blocks (lines in the alignment direction). This means that each unknown function (ψ^h in case of (20.25) and ϕ^h , $\hat{\rho}^h$ and $\hat{\epsilon}^h$ in case of (20.27)) are free (i.e., not fixed to be zero) not just at one gridpoint at a time, but at all gridpoints of the relaxed block, thus giving exactly the number of parameters needed to simultaneously satisfy the equations ((20.26) or

(20.28)) at all gridpoints of the block. In case of *plane alignment* (cf. Sec. 19.2) it is advisable to coarsen the grid only in that plane directions (cf. Sec. 4.2.1), in which case no block relaxation will be needed.

The total work of the relaxation sweep is only a fraction (20% or so) larger than the work of expressing the difference equations (20.18), or calculating their residuals, at all gridpoints.

20.4 Multigrid procedures

The grids, their relative positions and the interpolation procedures between them are generally as for Stokes equations (Sec. 18.4), with ρ and ϵ transferred similarly to p , and the residuals of the energy and state equations transferred similarly to the residuals of the continuity equations. Because of the nonlinearity, FAS is of course used (see Sec. 8), and the full weighting (18.14) is preferred over (18.13) in the fine-to-coarse transfers of both the velocity u_i and the i -th momentum residual.

An option for *double discretization* (cf. Sec. 10.2) is included. Namely, the artificial viscosity terms may be omitted in calculating the residuals to be transferred from any grid to the next coarser one. More generally, the artificial-viscosity coefficient β (see (20.17)), may have different values at different stages of the algorithm.

Whatever the value of β , in any stage of the solver, the discretization should also attempt not to straddle strong discontinuities, by calculating one-sided fluxes. As mentioned above, such features have been tested in simpler problems, but are not yet incorporated in our CNS system.

Note that in FAS, Dirichlet boundary conditions appear the same on all grids, whether the grid is the currently-finest or a correction grid. This can also be the case in the staggered boundary conditions employed at this stage (see Sec. 20.2.3): The exact differential solution is enforced at the staggered boundary points of the coarser, correction grid H , too. On those points, the value of $\hat{u}_h^{H,h}$ is also defined to be the exact solution, hence, for the purpose of interpolating back to the finer grid (like Eq. (8.6)), the difference $\tilde{u}_h^H - \hat{u}_h^{H,h}$ is defined to be zero at the staggered boundary points.

As in other non-elliptic cases, w cycles are generally preferred to V cycles (see Sec. 6.2).

21. REMARKS ON SOLVERS FOR TRANSONIC POTENTIAL EQUATIONS

21.1 Multigrid improvements

The multigrid solution of transonic potential flows has been studied by the author long ago in collaboration with Jerry South [S2], [B7, Sec. 6.5]. At that time multigrid research was less advanced, and many of the improved approaches described in this Guide were not implemented. Collected below is a list of important improvements that should be incorporated in future works.

(A) The Neumann boundary condition and the constant-potential-jump condition in the wake should not be enforced in relaxation, only smoothed (see Sec. 5.3). Thus, in relaxation, the potential jump at each wake point should be just set to be the average of the jumps at the neighboring wake points. The conditions should only be enforced on the coarsest level. Likewise, Kutta condition should only be applied at the coarsest level (cf. Sec. 5.6); the far-field conditions on each level should accordingly be adjusted at the stage of interpolating-and-adding the coarse-grid correction.

(B) Improved rates can be obtained if, before each full relaxation sweep, special local sweeps are made around singular points, such as trailing edges (see Sec. 5.7).

(C) Near a strong shock it may be better to use interpolation procedures which take jumps in ρ into account (see Sec. 4.6).

(D) Instead of stretching coordinates (to cover large exterior domains) and other transformations, a better multigrid procedure is to use increasingly coarser grids on increasingly larger domains, possibly with local refinements, anisotropic refinements and local coordinates (see Sec. 9). On such grids simple relaxation schemes can be used; block relaxation is only needed in directions of strong alignment (see Sec. 3.3).

(E) Most importantly, because of the non-elliptic nature of the problem, perfect smoothers and good asymptotic convergence rates should not be attempted; much simpler and vectorizable schemes can be used if all one attempts is to get below truncation errors (see end of Sec. 4.1 and Secs. 3.3, 7 and 20.3.1). Correspondingly, the performance of the algorithm should be ascertained through direct measurements of $\|\tilde{\phi}^h - \tilde{\phi}^{2h}\|$ (see Sec. 1.6) and of $|F^h(\tilde{\phi}^h) - F^{2h}(\tilde{\phi}^{2h})|$, where $F(\phi)$ is any solution functional one wants to get as the end result of the computations. If the norm measures discontinuous quan-

tities, such as velocities, it should be an L_1 norm [B17, Sec. 4.5].

21.2 Artificial viscosity in quasi-linear formulations

The transonic potential equation can be written in the quasi-linear form

$$[(\underline{u} \cdot \underline{\nabla})^2 - a^2 \Delta] \phi = 0 \quad (21.1)$$

where $\underline{u} = \underline{\nabla} \phi$. This operator has appeared above in the discussion of Euler and Navier-Stokes equations (see, e.g., Eq. (20.13), or (20.21)). This physical origin of the operator suggests a physical form for the artificial viscosity which should be added to it, different from the Murman-Cole-type forms currently in use. Namely, the artificial viscosity should be added to $\underline{u} \cdot \underline{\nabla}$, before it is squared, using generally the form of Q_0^h (see (20.16)-(20.17) for $\alpha = 0$). In particular, if upwinding is desired, it is the operator $\underline{u} \cdot \underline{\nabla}$ that should be upwinded, before it is squared.

This form of artificial viscosity (or upwinding) is not only smoother and more physical, it is also more straightforward than the usual scheme where the operator should be rotated before it is upwinded. Also, the new scheme requires no distinction between subsonic and supersonic points. The main difference between the two schemes is near sonic points ($M \approx 1$), or near shock transition from supersonic to subsonic. The Murman-Cole artificial viscosity vanishes there, and may therefore give rise to non-physical solutions.

In deeply subsonic regions, where (21.1) is uniformly elliptic, the form of the artificial viscosity does not matter of course, and one can switch to fully central approximations. In multigrid processing it is not important to do that, because $O(h^2)$ approximations can be obtained, even in the supersonic regions, by omitting the artificial viscosity (or the upwinding) from the operator used in the fine-to-coarse residuals transfer.

This latter operator should however respect shocks, as far as possible. Namely, it should be written in conservation form, and the stronger the shock, the weaker should it be straddled by flux calculations.

REFERENCES

- [A1] R.E. Alcouffe, A. Brandt, J.E. Dendy, Jr. and J.W. Painter: The multi-grid methods for the diffusion equation with strongly discontinuous coefficients. *SIAM J. Sci. Stat. Comput.* 2 (1981), 430-454.
- [A2] W. Auzinger and H.J. Stetter: Defect corrections and multigrid iterations. In [H4], 327-351.
- [A3] A. Arakawa: Computational design for long-term numerical integration of the equations of fluid motion: Two-dimensional incompressible flow. Part I. *J. Comput. Phys.* 1 (1966), 119-143.
- [A4] B. Arlinger: Multigrid techniques applied to lifting transonic flow using full potential equation. SAAB-SCANIA Rep. L-0-1 B439, 1978.
- [B1] N.S. Bakhvalov (Bahvalov): Convergence of a relaxation method with natural constraints on an elliptic operator. *Ž. Vyčisl. Mat. i Mat. Fiz.* 6 (1966), 861-885. (Russian)
- [B2] R.E. Bank: Multi-level Iterative Method for Nonlinear Elliptic Equations. *Elliptic Problem Solvers* (M. Schultz, ed.), Academic Press, New York 1981, 1-16.
- [B3] C. Börgers: Mehrgitterverfahren für eine Mehrstellendiskretisierung der Poisson-Gleichung und für eine zweidimensionale singular gestörte Aufgabe. Diplomarbeit, Universität Bonn, Bonn, 1981.
- [B4] D. Braess: The convergence rate of a multigrid method with Gauss-Seidel relaxation for the Poisson equation. In [H4], 368-386.
- [B5] A. Brandt: Multi-level adaptive technique (MLAT) for fast numerical solution to boundary value problems. *Proc. 3rd Internat. Conf. on Numerical Methods in Fluid Mechanics* (Paris, 1972), Lecture Notes in Physics 18, Springer-Verlag 1973, 82-89.
- [B6] A. Brandt: Multi-level adaptive techniques. IBM Research Report RC6026, 1976.
- [B7] A. Brandt: Multi-level adaptive solutions to boundary-value problems. *Math. Comp.* 31 (1977), 333-390. ICASE Report 76-27.
- [B8] A. Brandt: Multi-level adaptive techniques (MLAT) for partial differential equations: Ideas and software. *Mathematical Software III* (John R. Rice, ed.), Academic Press, New York 1977, 273-314. ICASE Report 77-20.
- [B9] A. Brandt et al.: Lecture notes of the ICASE Workshop on Multigrid Methods. ICASE, NASA Langley Research Center, Hampton, VA., 1978.
- [B10] A. Brandt: Multi-level adaptive solutions to singular-perturbation problems. *Numerical Analysis of Singular Perturbation Problems* (P.W. Hemker and J.J.H. Miller, eds.), Academic Press 1979, 53-142. ICASE Report 78-18.
- [B11] A. Brandt: Multi-level adaptive finite-element methods: I. Variational problems. *Special Topics of Applied Mathematics* (J. Frehse, D. Pallaschke, U. Trottenberg, eds.), North Holland 1980, 91-128. ICASE Report 79-8.

- [B12] A. Brandt: Numerical stability and fast solutions to boundary-value problems. *Boundary and Interior Layers - Computational and Asymptotic Methods* (J.J.H. Miller, ed.), Boole Press, Dublin 1980, 29-49.
- [B13] A. Brandt: Multi-level adaptive computations in fluid dynamics. *AIAA J.* 18 (1980), 1165-1172.
- [B14] A. Brandt: Multigrid Solvers on Parallel Computers. *Elliptic Problem Solvers* (M. Schultz, ed.), Academic Press, New York 1981, 39-84.
- [B15] A. Brandt: Stages in developing multigrid solutions. *Numerical Methods for Engineering* (E. Absi, R. Glowinski, P. Lascaux, H. Veysseyre, eds.), Dunod, Paris 1980, 23-44.
- [B16] A. Brandt: Multigrid solutions to steady-state compressible Navier-Stokes equations. *Proc. Fifth International Symposium on Computing Methods in Applied Sciences and Engineering*, Versailles, France, December 14-18, 1981.
- [B17] A. Brandt: Multigrid solvers for non-elliptic and singular-perturbation steady-state problems. Weizmann Institute of Science, Rehovot, Israel, 1981. *Computers and Fluids*, to appear.
- [B18] A. Brandt and C.W. Cryer: Multigrid algorithms for the solution of linear complementary problems arising from free boundary problems. MRC Technical Summary Report #2131, University of Wisconsin, Madison, Oct. 1980. *SIAM J. Sci. Stat. Comp.* 4 (1983), No. 4.
- [B19] A. Brandt and N. Dinar: Multigrid solutions to elliptic flow problems. *Numerical Methods for Partial Differential Equations* (S. Parter, ed.), Academic Press 1979, 53-147. ICASE Report 79-15.
- [B20] A. Brandt, S. McCormick and J. Ruge: Multigrid algorithms for differential eigenproblems. *SIAM J. Sci. Stat. Comp.* 4 (1983), 244-260.
- [B21] A. Brandt and S. Ta'asan: Multigrid methods for highly oscillatory problems. Research Report, August 1982.
- [B22] A. Brandt, J.E. Dendy, Jr. and H. Ruppel: The multi-grid method for the pressure iteration in Eulerian and Lagrangian hydrodynamics. *J. Comp. Phys.* 34 (1980), 348-370.
- [B23] J.J. Brown: A multigrid mesh-embedding technique for three-dimensional transonic potential flow analysis. Boeing Commercial Airplane Company, Seattle, Washington, April 1981. NASA Conf. Publ. 2202, 1982.
- [B24] D. Bai and A. Brandt: Local mesh refinement multigrid techniques. Report, Appl. Math., Weizmann Inst., Rehovot, Israel, 1984.
- [B25] D. Barkai and A. Brandt: Vectorized multigrid Poisson solver for the CDC Cyber 205. *Applied Math. and Comput.* 13 (1983), 215-228.
- [B26] A. Brandt: Algebraic multigrid theory: The symmetric case. *Preliminary Proc. Int. Multigrid Conf.*, Copper Mountain, Colorado, April, 1983.
- [B27] A. Brandt: Guide to multigrid development. In [H4], 220-312.
- [B28] A. Brandt, S.R. Fulton and G.D. Taylor: Improved spectral multigrid methods for periodic elliptic problems. *J. Comput. Phys.* (Submitted).
- [B29] A. Brandt, S. McCormick and J. Ruge: Algebraic multigrid (AMG) for sparse matrix equations. *Sparsity and Its Applications* (D.J. Evans, ed.), Cambridge University Press, 1984.
- [B30] A. Brandt and D. Ophir: GRIDPACK: Toward unification of general grid programming. *PDE Software: Modules Interfaces and Systems* (B. Engquist, ed.), North Holland, 1984.
- [B31] J.W. Boerstael: A multigrid algorithm for steady transonic potential flows around aerofoils using Newton iteration. *J. Comp. Phys.* 48 (1982), 314-344.
- [B32] J.W. Boerstael and A. Kassies: Integrating multigrid relaxation into a robust fast-solver for transonic potential flows around lifting surfaces. AIAA Paper 83-1885, 1983.
- [B33] A. Brandt: Videotape lectures on multigrid methods. Production Numbers L8101-L8110, M8111-M8114, M8166-M8175. Office of Instructional Services, Colorado State University, Colorado 80523, Phone 303-491-5416.
- [B34] A. Brandt, S. McCormick and J. Ruge: Algebraic multigrid (AMG) for automatic multigrid solution with application to geodetic computations. Report, Colorado State University, Fort Collins, Colorado, 1983.
- [B35] A. Brandt, S.R. Fulton and G.D. Taylor: Improved spectral multigrid methods for periodic elliptic problems. Colorado State University, Fort Collins, Colorado, September 1983.
- [C1] M. Ciment, S.H. Leventhal and B.C. Weinberg: The operator compact implicit method for parabolic equations. *J. Comp. Phys.* 28 (1978), 135-166.
- [C2] F.W.C. Campbell and J. Robson: Application of Fourier Analysis to the visibility of gratings. *J. Physiol. (Lond.)* 197 (1968), 551-566.
- [C3] D.A. Caughey: Multigrid calculation of three-dimensional transonic potential flows. *Applied Math. and Computations* 13 (1983), 241-260.
- [D1] J.E. Dendy, Jr.: Black box multigrid. *J. Comp. Phys.* 48 (1982), 366-386.
- [D2] N. Dinar: Fast methods for the numerical solution of boundary-value problems. Ph.D. Thesis, The Weizmann Institute of Science, Rehovot, Israel (1979).
- [D3] I.S. Duff: Sparse matrix software for elliptic PDE's. In [H4], 410-426.
- [D4] J.E. Dendy: Black box multigrid for nonsymmetric problems. *Appl. Math. Comput.* 13 (1983), 261-284.

- [D5] H. Deconinck and C. Hirsch: A multigrid method for the transonic full potential equation discretized with finite elements on an arbitrary body fitted mesh. *J. Comp. Phys.* 48 (1982), 344-366.
- [F1] R.P. Fedorenko: On the speed of convergence of an iteration process. *Ž. Vyčisl. Mat. i Mat. Fiz.* 4 (1964), 559-564. (Russian)
- [F2] Ha. Foerster, K. Stueben and U. Trottenberg: Non-standard multigrid techniques using checkerboard relaxation and intermediate grids. *Elliptic Problem Solvers* (M. Schultz, ed.), Academic Press, New York, 1981, 285-300.
- [F3] Ha. Foerster and K. Witsch: On efficient multigrid software for elliptic problems on rectangular domains. Preprint No. 458, SFB 72, Universität Bonn, Bonn, 1981.
- [F4] G.E. Forsythe and W.R. Wasow: *Finite-Difference Methods for Partial Differential Equations*. Wiley, New York 1960.
- [F5] P.O. Frederickson: Fast approximate inversion of large sparse linear systems. Math. Report 7-75, Lakehead University, Ontario, Canada, 1975.
- [F6] L. Fuchs: Multi-grid solution of the Navier-Stokes equations on non-uniform grids. NASA Conf. Publ. 2202, 1982.
- [G1] S. Geman and D. Geman: Stochastic relaxation, Gibbs distribution, and the Bayesian restoration of images. Division of Appl. Math., Brown University, Providence, Rhode Island, September 1983.
- [G2] U. Ghia, K.N. Ghia and C.T. Shin: High-Re solutions for incompressible flow using the Navier-Stokes equations and a multigrid method. *J. Comp. Phys.* 48 (1982), 387-411.
- [G3] D. Gottlieb and S.A. Orszag: Numerical Analysis of Spectral Methods: Theory and Applications. CBMS-NSF Regional Conference Series in Applied Mathematics, Society for Industrial and Applied Mathematics, 1977.
- [H1] W. Hackbusch: The fast numerical solution of very large elliptic differential schemes. *J. Inst. Maths. Applics.* 26 (1980), 119-132.
- [H2] W. Hackbusch: Multi-grid convergence theory. In [H4], 177-219.
- [H3] J.M. Hyman: Mesh refinement and local inversion of elliptic partial differential equations. *J. Comp. Phys.* 23 (1977), 124-134.
- [H4] W. Hackbusch and U. Trottenberg (eds.): *Multigrid Methods*, Lecture Notes in Math. 960, Springer-Verlag, 1982.
- [J1] A. Jameson: Acceleration of transonic potential flow calculations on arbitrary meshes by the multiple grid method. AIAA Paper 79-1458, 1979.
- [J2] A. Jameson: Solution of the Euler Equations for two dimensional transonic flow by a multigrid method. *Applied Math. and Computations* 13 (1983), 327-356.
- [K1] H.B. Keller: Numerical solution of bifurcation and nonlinear eigenvalue problems. *Applications of Bifurcation Theory* (P. Rabinowitz, ed.), Academic Press, New York, 1977, 359-384.
- [K2] R. Kettler: Analysis and comparison of relaxation schemes in robust multigrid and preconditioned conjugate gradient methods. In [H4], 502-534.
- [K3] L. Kauffman: Matrix methods for queuing problems. Bell Laboratories, Murray Hill, New Jersey, August 1982.
- [L1] B. Lindberg: Error estimation and iterative improvement for the numerical solution of operator equations. Report UIUCDCS-R-76-820, University of Illinois, Urbana, Illinois, 1976.
- [L2] J. Linden: Mehrgitterverfahren für die Poisson-Gleichung in Kreis und Ringgebiet unter Verwendung lokaler Koordinaten. Diplomarbeit, Bonn 1981.
- [L3] R.E. Lynch and J.R. Rice: The Hodge method and its performance. *Recent Advances in Numerical Analysis* (C. de Boor, ed.), Academic Press, New York, 1978, 143-179.
- [L4] S. Lang: *Algebra*, Addison-Wesley, 1965.
- [M1] J. Meinquet: Multivariate interpolation at arbitrary points made simple. *J. Applied Math. Phys.* (ZAMP) 30 (1979), 292-304.
- [M2] P. Meissle: A priori prediction of roundoff error accumulation in the solution of a super-large geodetic normal equation system. NOAA Professional Paper 12, National Oceanic and Atmospheric Administration, Rockville, Maryland, 1980.
- [M3] MUGTAPE 84, A tape of multigrid software and programs, including GRIDPACK; MUGPACK; simple model programs (CYCLEV, CYCLE C, FASCC, FMGL and an eigenproblem solver); Stokes equations solver; SMORATE; BOXMG [D1]; MGOO and MGO1 [S4]. Available at the Department of Applied Mathematics, The Weizmann Institute of Science, Rehovot, Israel.
- [M4] D.R. McCarthy and T.A. Reyhner: Multigrid code for transonic potential flow about inlets. *AIAA J.* 20 (1982), 45-50.
- [M5] D.R. McCarthy and R.C. Swanson: Embedded mesh multigrid treatment of two-dimensional transonic flows. *Applied Math. and Computations* 13 (1983), 399-418.
- [M6] W.J.A. Mol: Numerical solution of the Navier-Stokes equations by means of a multigrid method and Newton-iteration. *Proc. Seventh Int. Conf. on Num. Methods in Fluid Dynamics* (W.C. Reynolds, R.W. MacCormack, eds.), Lecture Notes in Physics 141, Springer-Verlag, 1981.
- [N1] K.A. Narayanan, D.P. O'Leary and A. Rosenfeld: Multi-resolution relaxation. TR-1070 MCS-79-23422, University of Maryland, College Park, Maryland, July 1981.
- [N2] R.A. Nicolaides: On multiple grid and related techniques for solving discrete elliptic systems. *J. Comp. Phys.* 19 (1975), 418-431.
- [N3] R.A. Nicolaides: On the l^2 convergence of an algorithm for solving finite element equations. *Math. Comp.* 31 (1977), 892-906.
- [N4] R.H. Ni: A multiple grid scheme for solving the Euler equations. AIAA Paper 81-1025, June 1981.

- [O1] D. Ophir: Language for processes of numerical solutions to differential equations. Ph.D. Thesis (1979), The Weizmann Institute of Science, Rehovot, Israel.
- [R1] M. Ries, U. Trottenberg and G. Winter: A note on MGR methods. Preprint 461, Universität Bonn, SFB 72, 1981. *Lin. Alg. Appl.* 49 (1983), 1-26.
- [S1] S. Schaeffer: High-order multigrid methods. Ph.D. Thesis, Colorado State Univ., Fort Collins, Colorado, 1982.
- [S2] J.C. South and A. Brandt: Application of a multi-level grid method to transonic flow calculations. *Transonic Flow Problems in Turbomachinery* (T.C. Adamson, M.F. Platzer, eds.), Hemisphere, Washington, 1977, 180-207. ICASE Report 76-8.
- [S3] H.J. Stetter: The defect correction principle and discretization methods. *Num. Math.* 29 (1978), 425-443.
- [S4] K. Stüben and U. Trottenberg: Multigrid methods: fundamental algorithms, model problem analysis and applications. In [H4], 1-176.
- [S5] H. Schwichtenberg: Die Erweiterung des Mehrgitterprogramms MG01 auf nichtlineare Aufgaben. Diplomarbeit, Institut fuer Angewandte Mathematik, Universitaet Bonn, 1983.
- [S6] K. Stueben: Algebraic multigrid (AMG): Experiences and comparisons. *Applied Math. and Computations* 13 (1983), 419-452.
- [S7] A. Shmilovich and D.A. Caughey: Application of the multigrid method to calculations of transonic potential flow about wing-fuselage combinations. *J. Comp. Phys.* 48 (1982), 462-484.
- [S8] C.L. Streett, T.A. Zang and M.Y. Hussaini: Spectral multigrid methods with applications to transonic potential flow. ICASE Report 83-11, 1983.
- [T1] S.L. Tanimoto: Template matching in pyramids. *Computer Graphics and Image Processing* 16 (1981), 356-369.
- [T2] K. Tanabe: Projection methods for solving a singular system of linear equations and its applications. *Num. Math.* 17 (1971), 203-214.
- [T3] R. Temam: *Navier-Stokes Equations: Theory and Numerical Analysis*. North-Holland, 1977.
- [T4] D. Terzopoulos: Multilevel computational processes for visual surface reconstruction. *Computer Vision, Graphics and Image Processing* 24 (1983), 52-96.
- [T5] C.A. Thole: Beiträge zur Fourieranalyse von Mehrgittermethoden: V-Cycle, ILU Glättung, anisotrope Operatoren. Diplomarbeit, Universität Bonn.
- [T6] T. Thunell and L. Fuchs: Numerical solution of the Navier-Stokes equations by multi-grid techniques. *Numerical methods in laminar and turbulent flow* (C. Taylor, J.W. Schrefler, eds.), Pineridge Press, 1981.
- [V0] I.Y. Vakhutinsky, L.M. Dudkin and A.A. Ryvkin: Iterative aggregation -- a new approach to the solution of large-scale problems. *Econometrica* 47 (1979), 821-841.
- [V1] R. Verfürth: The contraction number of a multigrid method with ratio 2 for solving Poisson's equation. Ruhr-Universität Bochum, West Germany, 1981.
- [V2] R. Verfürth: Two algorithms for mixed problems. *Proc. Int. Multigrid Conf.*, Copper Mountain, Colorado, April 1983.
- [V3] R. Verfürth: A multilevel algorithm for mixed problems. Preprint, Bochum 1983.
- [W1] P. Wesseling: A robust and efficient multigrid method. In [H4], 614-630.
- [W2] P. Wesseling and P. Sonneveld: Numerical experiments with a multiple grid and a preconditioned Lanczos type method. *Approximation Methods for Navier-Stokes Problems* (R. Rautmann, ed.), Lecture Notes in Math. 771, Springer-Verlag, 1980, 543-562.
- [W3] P. Wesseling: Numerical solution of stationary Navier-Stokes equation by means of a multiple grid method and Newton iteration. Report NA-18, Dept. of Math., Delft University of Technology, 1977.
- [W4] A.J. van der Wees, J. van der Vooren and J.H. Meelker: Robust calculation of 3D transonic potential flow based on the nonlinear FAS multigrid method and incomplete LU-decomposition. AIAA Paper 83-1950, 1983.
- [W5] H.R. Wilson and J.R. Bergen: A four mechanism model for spatial vision. *Vision Res.* 19 (1979), 19-32.
- [Z1] T.A. Zang, Y.S. Wong and M.Y. Hussaini: Spectral multigrid methods for elliptic equations. *J. Comp. Phys.* 48 (1982), 485-501.
- [Z2] T.A. Zang, Y.S. Wong and M.Y. Hussaini: Spectral multigrid methods for elliptic equations II. NASA CR 172131, 1983.

Additional References

- [S9] P. Sonneveld, P. Wesseling, P.M. de Zeeuw: Multigrid and conjugate gradient methods as convergence acceleration techniques. *Proc. Multigrid Conf. Bristol 1983* (H. Holstein and D. Paddon, eds.).
- [J3] D.C. Jespersen: Design and Implementation of a multigrid code for the Euler equations. *Applied Math. and Computations* 13 (1983), 357-374.

