

A Review of Numerous Parallel Multigrid Methods*

Craig C. Douglas[†]

May, 1992/August, 1995

Parallel multilevel methods are shown to be the natural precursors to standard multilevel methods based on the personnel computing era of earlier this century. They are also the natural successors to standard multilevel methods in the age of computers. What makes six parallel multilevel methods practical and impractical is discussed in the context of the three algorithms that encapsulate them.

1 Preliminaries

Multigrid methods originated earlier this century, in the *personnel* computing era. Someone who needed to compute an approximation to the solution of a partial differential equation during that era would fill a room with people. After using very simple mechanical calculators to compute parts of the approximation, these people would pass their parts to the other people in the room who needed them. Except for the very different time scales and approximate solution accuracy, this process is similar to computing on today's distributed memory parallel computers.

The most basic model problem for elliptic boundary value problems in multigrid has always been, effectively, the two dimensional Poisson equation

*Appears in slightly edited form in *SIAM News*, V.25, N.3, May, 1992.

[†]IBM Research Division, Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, NY 10598-0218, USA and Department of Computer Science, Yale University, P.O. Box 208285, New Haven, CT 06520-8285, USA and CERFACS, 42 Avenue G. Coriolis, F-31057 Toulouse Cedex, FRANCE. Golly.

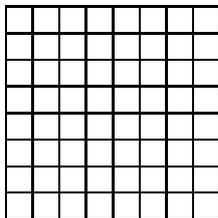


Figure 1: Uniformly meshed domain.

on a square, namely,

$$-\Delta u = f \text{ in } \Omega = (0, 1)^2 \quad ; \quad u = 0 \text{ on } \partial\Omega. \quad (1)$$

A $N \times N$ uniform mesh with mesh spacing $h = 1/(N - 1)$ is placed over Ω as shown in Figure 1.

The grid points are x_{ij} , where $1 \leq i, j \leq N$. A central difference discretization is applied to (1) to get a set of linear equations

$$AU = F \quad (2)$$

where $F_{(i-1)N+j} = h^2 f(x_{ij})$ and $A = [-I, T, -I]$ is the block tridiagonal matrix defined using the $N \times N$ matrices I , the identity matrix, and $T = [-1, 4, -1]$, a tridiagonal matrix.

In the early part of the twentieth century, partial differential equations were approximately solved by relaxation techniques. (In fact, many types of equations are still solved in this way today.) Solving (2) by a Gauss-Seidel method could take the rest of the personnel computers' lives if N was large enough and the required accuracy was strict enough. Hence, new tricks were needed to reduce the computation time.

During the 1920's, and probably several decades earlier, two-level schemes were used by engineers. An auxiliary grid Ω_1 was used to generate an initial guess to the solution to (2) on Ω_2 .

Let $GS(j, n) = n$ Gauss-Seidel sweeps on level j . Algorithm 1, a *one-way two-level* algorithm is defined by:

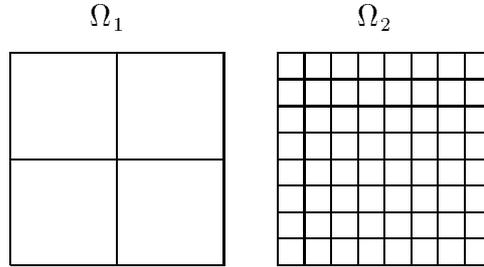


Figure 2: Meshes for the two-level schemes

Algorithm 1 *One-way Two-Level*

- (a) Solve the problem on Ω_1 by any means.
- (b) Interpolate U_1 onto Ω_2 .
- (c) Do $GS(2, \cdot)$ until convergence.

Bilinear interpolation was typically used in step (b). Let N_i be the number of grid points on grid Ω_i and define

$$\sigma_i \approx \frac{N_i}{N_{i-1}}, \quad i > 1. \quad (3)$$

Before digital computers were commonplace, σ_2 was typically between 5 and 10 (see [30]).

During the 1960's, Soviets developed what is now considered true multi-grid (see [3], [4], [19], and [20]). In this case, the auxiliary grid Ω_1 was used to correct approximations to the solution to (2) on Ω_2 . Algorithm 2, a *two-level correction* algorithm is defined by:

Algorithm 2 *Two-Level Correction*

- (a) $U_2 \leftarrow GS(2, n)$.
- (b) $F_1 \leftarrow \sigma^2(F_2 - A_2 U_2)$ where $\Omega_1 \cap \Omega_2$.
- (c) $C_1 \leftarrow GS(1, m)$.
- (d) $U_2 \leftarrow U_2 + \text{Interpolate}(C_1)$.
- (e) Repeat (a)–(d) until convergence.

Projection of the style of step(b), known as *injection*, is in disfavor today; once again, bilinear interpolation has become typical. In this case, σ_2 is typically 4. (Note that this changes Ω_1 in the diagram.)

Algorithm 2 is a notoriously slow algorithm for solving (2). The best trick for speeding it up is to use more than two grids. Define Ω_j , $j = 1, \dots, k$ according to Fig. 3 and place uniform meshes Ω_j over Ω , with $h_j = \sigma^{k-j} h_k$, for some $\sigma \in \mathbb{R}$. Then solve

$$A_k U_k = F_k$$

using the U_j , $j < k$, in some prescribed manner.

Other tricks include

- Using different orders of interpolation (e.g., bicubics instead of bilinears) or projection to transfer information between meshes.
- Using better discretization methods (e.g., finite elements or volumes) for the partial differential equation.
- Solving the level 1 problem by a direct method.
- Using a more sophisticated solver, such as conjugate residuals or Bi-CGSTAB.
- Starting the computation on level 1 instead of level k . (This is commonly referred to as a *nested iteration* in multigrid lingo.)

Thus, for a general linear differential equation on some domain Ω that has been discretized into a sequence of problems $j = 1, \dots, k$, we have

$$A_j U_j = F_j, \quad U_j, F_j \in \mathcal{M}_j, \quad A_j \in \mathcal{L}(\mathcal{M}_k). \quad (4)$$

The solution spaces \mathcal{M}_j are typically function or vector spaces based on the discretization method (typically, finite differences, elements, or volumes) and \mathcal{L} stipulates that A_j is a linear operator on \mathcal{M}_j . We assume that there exist mappings between the neighboring spaces:

$$R_j : \mathcal{M}_j \rightarrow \mathcal{M}_{j-1} \quad \text{and} \quad P_{j-1} : \mathcal{M}_{j-1} \rightarrow \mathcal{M}_j.$$

While not necessary (see [15]), we also assume that there are mappings $Q_j : \mathcal{M}_j \rightarrow \mathcal{M}_{j-1}$ such that

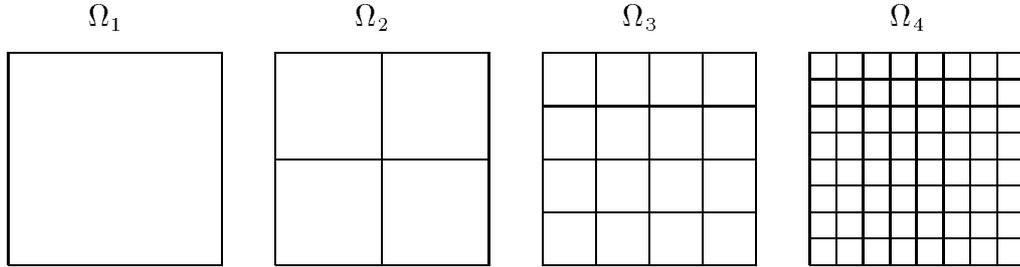


Figure 3: Multiple grid meshes

$$A_{j-1} = Q_j A_j P_{j-1}.$$

This leads to Algorithm 3, a truly abstract correction algorithm of the following form:

Algorithm 3 $MG(j, \{\mu_\ell\}_{\ell=1}^j, x_j, F_j)$

(1) If $j = 1$, then solve $A_1 x_1 = F_1$

(2) If $j > 1$, then repeat $i = 1, \dots, \mu_j$:

(2a) Solving: $x_j \leftarrow \text{PreSolve}_j^{(i)}(x_j, F_j)$

(2b) Residual Correction:

$$x_k \leftarrow x_j + P_{j-1} MG(j-1, \{\mu_\ell\}_{\ell=1}^{j-1}, 0, R_j(F_j - A_j x_j))$$

(2c) Solving: $x_j \leftarrow \text{PostSolve}_j^{(i)}(x_j, F_j)$

(3) Return x_j

This definition assumes that $\mu_1 = 1$. Common values for μ_j , $2 < j \leq k$, are 1 (*V cycle*) and 2 (*W cycle*). It is quite common for either the PreSolver or the PostSolver (but not both) to be the identity operator. Common solvers are relaxation methods and conjugate gradient-like methods. On the coarsest grid, a direct solver may be used.

The purpose of this article is not to analyze Algorithm 3 (see [14] and [15] for extensive analysis) but to discuss various methods for parallelizing it. Further, no mention will be made of nonlinear and/or nested iteration versions of Algorithm 3—the commentary applies to these cases verbatim.

There are two major themes in parallel multigrid today; telescoping and nontelescoping methods. Telescoping methods ($\sigma > 1$) include domain decomposition methods and a method which computes on all levels simultaneously. Nontelescoping methods use multiple coarse subspaces in which the sum of the unknowns on any level equals the sum on any other level.

2 Telescoping Parallelizations

Domain decomposition techniques [5] offer a relatively easy path to parallelizing Algorithm 3. Each processor computes on the block of data per level that is assigned to it.

Because almost no work is required to get good convergence bounds, this is certainly the most attractive parallelizing technique from a theoretical standpoint. The method is merely a block iterative method used for the approximate solver combined with standard multigrid. Hence, many standard multigrid convergence theorems apply verbatim.

Some communications and processor scheduling issues arise with this method. The amount of data communication involved and the location of the processors in a network will determine whether or not the problem is converted into an input-output problem rather than a computational one.

The factor σ_j in (3) determines the telescoping of unknowns and thus plays a key part in processor scheduling. When the number of unknowns assigned to each processor falls below some threshold (which is a function of the problem, processor speed, and communications bandwidth and latency), some of the processors may be idle some of the time. To avoid this, an agglomeration of unknowns is typically performed and a certain number of processors become completely idle or compute the same thing.

Having idle processors may seem like a waste of computer time on a parallel processor, but it actually is not necessarily the case. If the problem can be solved faster (as timed by a stopwatch) with some processors out of the computation some of time, then it is clearly a good approach, even if it is a bit wasteful. The principal reason for actually solving problems on parallel computers is that the results are perceived to be needed much sooner than just running many problems, one per processor.

As an aside, more and more parallel computing seems to be done on clusters of workstations (i.e., *distributed computing*) rather than on explicitly

parallel machines. The workstation approach means that when a processor is out of the computation, its task scheduler can assign it to work on something else; thus, it may not actually be idle. This also occurs naturally on a parallel machine that has a multi-user or multitasking operating system on each node of the machine.

During the approximate solve steps, information must be passed between processors along neighboring data regions. If a solver like conjugate gradients is used, then information from dot products and matrix-vector multiplies will also clutter the communications' mechanism.

An asynchronous relaxation method would seem to be ideal in this environment. With this approach, each processor does its local relaxation method and uses the information obtained most recently from other processors. This is fast, and if the standard relaxation procedure converges, so does the asynchronous version, just not quite as quickly [8]. Unfortunately, asynchronous relaxation methods have never caught on in the parallel multigrid community, mostly because the small number of smoothing iterations (2–6 iterations) makes it impossible to ignore changes in neighboring data.

Another approach to telescoping multigrid involves massively parallel computers (at least as many processors as unknowns on all of the meshes). Gannon and Van Rosendale [22] proposed what is referred to as a *concurrent multigrid* method. Typically, this means there should be $N \log_d N$ processors for a d dimensional problem.

The concept is that all operations should be performed simultaneously on all unknowns on all levels. An initial approximation of zero is assumed for the solution. Two sets of vectors, q_j and d_j , are used to hold information about right hand sides and data on the spectrum of levels in the sense that

$$F_k \approx \sum_{j=1}^k q_j + d_j.$$

A third set of vectors, x_j , contains the approximations to the solutions to each problem on each level. This information percolates to the finest level, k , to finally provide the approximate solution to the real problem. Algorithm 4, $CMG(k, \mu, F_k)$, is as follows:

Algorithm 4 $CMG(k, \mu, F_k)$ (1) *Initialize in parallel:*

$$x_j = q_j = 0, 1 \leq j < k ; q_k = F_k, x_k = 0.$$

(2) *Repeat* $i = 1, \dots, \mu$:(2a) *Smoothing in parallel*

$$x_j \leftarrow \text{PreSolve}_j^{(i)}(x_j, q_j), 1 \leq j \leq k$$

(2b) *Compute data corresponding to* x_j *in parallel:*

$$d_j \leftarrow A_j x_j, 1 \leq j \leq k.$$

(2c) *Compute residuals in parallel:*

$$q_j \leftarrow q_j - d_j, 1 \leq j \leq k.$$

(2d) *Project* q *onto coarser levels in parallel:*

$$q_1 \leftarrow q_1 + R_2 q_2 ;$$

$$q_k \leftarrow (I - P_{k-1} R_k) q_k ;$$

$$q_j \leftarrow (I - P_{j-1} R_j) q_j + R_j q_{j+1}, 1 < j < k.$$

(2e) *Inject* x *into finer levels in parallel:*

$$x_1 \leftarrow 0 ; x_k \leftarrow x_k + P_{k-1} x_{k-1} ;$$

$$x_j \leftarrow P_{j-1} x_{j-1}, 1 < j < k.$$

(2f) *Inject* d *into finer levels in parallel:*

$$d_1 \leftarrow 0 ; d_k \leftarrow d_k + P_{k-1} d_{k-1} ;$$

$$d_j \leftarrow P_{j-1} d_{j-1}, 1 < j < k.$$

(2g) *Put all the data back into* q *in parallel*

$$q_j \leftarrow q_j + d_j, 1 \leq j \leq k.$$

(3) *Return* x_k

This definition assumes that μ evenly divides twice the number of levels.

Algorithm 4 has a number of noteworthy aspects. First, communication between adjacent levels is twice the amount that would be expected. This is absolutely required in order for the algorithm to be consistent (without which it diverges).

Second, this algorithm limits what types of iterative methods that qualify as approximate solvers. The cost of one iteration on any level must be the same as the cost on any other level, assuming one processor per unknown. If the matrices A_j are similar enough to each other, Jacobi and conjugate gradients are fine, but Gauss-Seidel and SSOR are not. The latter two methods assume that the data is traversed in a particular order rather than all at once. Hence, the length of time to complete each iteration is dependent on the number of unknowns, violating the requirement of identical time per level.

Third, the number of levels that information must traverse to move from the finest level to the coarsest one and back is twice the usual number. Hence, all work estimates will be $O(2 \log N)$ (with the 2 being part of the constant) asymptotically in the number of levels, the usual multigrid method of estimating the problem complexity. This may seem high, but it is actually equivalent to the complexity of a standard V cycle.

Finally, there is a similar, but quite different approach to parallel multigrid: apply a domain decomposition method to the finest grid and then use a serial computer multigrid method on each of the subdomains. Studies have shown (e.g., [23]) that this is not as fast an approach as using domain decomposition methods on each level of the multigrid algorithm.

3 Nontelegraphing Parallelizations

Several competing methods have the same number of unknowns on each level. Each method has advantages and disadvantages.

The concept of using multiple subspaces to solve a problem whose solution lies in a particular space is hardly new. In fact, no one from the era in which it was invented is alive today. We will never know who really invented it, but we can be certain that it was introduced no later than in 1869 [29]. Further, it has been an active area of research in the field of symmetry groups (see [18]) for many years.

Assume a rooted tree of problems (see (4)) that are arbitrarily numbered. For a given problem k , it either has a set C_k of coarse space correction problems or it has none at all (i.e., $C_k = \emptyset$). When $C_k \neq \emptyset$, there are restriction and prolongation operators for each coarse space problem $\ell \in C_k$ such that

$$R_\ell : \mathcal{M}_k \mapsto \mathcal{M}_\ell \quad \text{and} \quad P_\ell : \mathcal{M}_\ell \mapsto \mathcal{M}_k.$$

We also assume that there are mappings

$$Q_\ell : \mathcal{M}_k \rightarrow \mathcal{M}_\ell \quad \text{such that} \quad A_\ell = Q_\ell A_k P_\ell.$$

These mappings are defined much as in the serial case.

A multiple coarse space correction multigrid scheme is defined by algorithm 5, $MCSMG(j, \mu_j, C_j, x_j, F_j)$, as follows:

Algorithm 5 $MCSMG(j, \mu_j, C_j, x_j, F_j)$

(1) If $C_j = \emptyset$, then solve $A_j x_j = F_j$

(2) If $C_j \neq \emptyset$, then repeat $i = 1, \dots, \mu_j$:

(2a) Smoothing: $x_j \leftarrow PreSolve_j^{(i)}(x_j, F_j)$

(2b) Residual Correction:

$$x_j \leftarrow x_j + \sum_{\ell \in C_j} P_\ell MCSMG(\ell, \mu_\ell, C_\ell, 0, R_\ell(F_j - A_j x_j))$$

(2c) Smoothing: $x_j \leftarrow PostSolve_j^{(i)}(x_j, F_j)$

(3) Return x_j

The performance of any variant of Algorithm 5 is dependent mainly on how small δ_j is, where

$$\|(I - \sum_{\ell \in C_j} Q_\ell^{-1} R_\ell)u\| \leq \delta_j \|u\|, \quad u \in \mathcal{M}_j. \quad (5)$$

This is a measure of how many of the error components in \mathcal{M}_j are not completely represented in the subspaces.

Ta'asan [31] introduced this method to the multigrid community when standard multigrid failed to converge for a class of problems with highly oscillatory solutions. He uses standard interpolation and projection methods and a Kaczmarz relaxation method in his examples.

Using a different set of interpolation and projection methods Hackbusch [26] developed a variant of Ta'asan's method using standard smoothers. Ta'asan's and Hackbusch's methods are both referred to as *robust multigrid*, which adds confusion (and heated discussions) to the field. More recently, Hackbusch's method has been referred to as *frequency decomposition multigrid* [27].

Frederickson and McBryan [21], using an approach different from that of Gannon and Van Rosendale, also investigated ways to keep all of the processors busy on a massively parallel single instruction, multiple data (SIMD) machine. They used standard interpolation and projection methods and an elaborate smoother on each level. Unless great care is taken, this method computes the correction in one of the correction spaces while the corrections in the remaining spaces add up (pointwise) to zero. Their method is referred to as *parallel superconvergent multigrid*. A comparison of this method with that of Gannon and Van Rosendale might make an interesting student exercise.

The methods of Ta'asan, Hackbusch, and Frederickson-McBryan all use interleaved grids. For a problem on a square, this translates into the following, where the numbers refer to which subproblem the unknowns belong:

3	4	3	4	3	4	3	4
1	2	1	2	1	2	1	2
3	4	3	4	3	4	3	4
1	2	1	2	1	2	1	2
3	4	3	4	3	4	3	4
1	2	1	2	1	2	1	2

(The motivation is similar to that for multicolored orderings [1] for standard iterative methods.) Each of these methods requires that all of the matrices associated with the spaces be generated, except in trivial cases, thus doubling the memory requirements expected for solving boundary value problems. In addition, the coarse space operators are more difficult to compute using Hackbusch's variant than with either of the other two methods. In general, these are all space-wasteful methods.

A fourth, very different approach that my colleagues and I developed, is referred to as either *constructive interference* [17] or, more recently, *domain reduction* (many references can be found in [9] and [16]). Our original motivation for using a multiple coarse space parallel multigrid algorithm was to eliminate the approximate solve step from standard multigrid algorithms. The solver takes most of the computational time but contributes almost nothing to the convergence rate, whereas coarse grid corrections take little time and reduce the error substantially. A general theory and simple examples were developed for multiple coarse space methods using no solver on the finest grid and mutually orthogonal subspaces which covered all of the error components of the original space (thus, $\delta_j = 0$ in (5)). This leads to very efficient direct methods rather than the expected iterative ones.

A side benefit of this theory is that the fine grid problem and, if a trick is used, most of the coarse space matrices do not need to be generated [9]. This method can use substantially less memory than a standard iterative or multigrid algorithm.

An additional note about domain reduction is that it leads naturally to more than 2^d subspaces for a d dimensional problem. An eight way decomposition of a problem on a square can be constructed, leading to problems

defined on squares, rectangles, and triangles [6]. Both 60- and 64-way decompositions of a problem on a cube can be constructed with moderate difficulty [9]. In theory, a 192-way decomposition of a problem on a cube is possible. The entire problem would be solved 2660 times faster, if each of the 192 subproblems were solved by sparse Gaussian elimination, than if the original problem is solved by the same method (the latter is not advised, however).

To see how each of the variants operates, consider various projection operators in one dimension (where $x_0 = c_{N+1} = 0$):

- Linear projection: for $1 \leq i \leq N$,

$$y_i = x_{i-1} + 2x_i + x_{i+1}.$$

Let one subspace consist of the odd numbered y_i 's and another subspace for the even numbered y_i 's. This is used by both Ta'asan and Frederickson-McBryan.

- Linear-linear orthogonal complement projection: for $1 \leq i \leq N$,

$$y_i = \begin{cases} x_{i-1} + 2x_i + x_{i+1} & i \text{ even,} \\ -x_{i-1} + 2x_i - x_{i+1} & i \text{ odd.} \end{cases}$$

The subspaces are defined as in the linear projection case. The values at odd numbered y_i 's correspond to the orthogonal complement of the traditional space defined by linear projection. This is used by Hackbusch and is actually a set of pre-wavelets.

- Symmetric-antisymmetric projection: for $1 \leq i \leq N/2$,

$$y_i = \frac{x_i + x_{N-i+1}}{2} \quad \text{and} \quad \hat{y}_i = \frac{x_i - x_{N-i+1}}{2}.$$

One subspace consists of y_i 's and the other of \hat{y}_i 's. This defines a domain folding (or reduction) where even and odd functions are annihilated in exactly one subspace and exactly reproduced in the other. This is the method that my colleagues and I used.

The two dimensional definitions of the above are defined in the obvious manner using tensor products.

Unfortunately, these methods have not been directly compared on a non-trivial example problem. The closest is a simple problem from the literature:

$$-10^5 u_{xx} - 10^{-5} u_{yy} = f \text{ in } (0, 1)^2 \quad ; \quad u = 0 \text{ on } \partial(0, 1)^2.$$

To make things comparable with known results, a uniform mesh is used, a standard (simple) discretization, the energy norm, one Jacobi iteration in the analysis for multigrid (MG) and robust multigrid, and a direct solve on the coarsest level(s). The contraction factors are the following:

Method	Contraction factor	
MG	0.97	(17×17 grid)
Robust MG	0.33	(h independent)
Domain reduction (iterative)	ϵ	(h independent)
Domain reduction (direct)	0.00	(h independent)
Parallel superconvergent	“small”	(h independent)

The solver in the domain reduction method is either iterative (solving each problem to an accuracy of ϵ) or direct. If the smoother called for in parallel superconvergent can be constructed, then the contraction factor missing from the table will be very small, on the order of 0.05. Note that a line relaxation method, instead of point Jacobi, would make multigrid work well.

4 Implementing Parallel Multigrid

4.1 Shared Virtual Memory

I have been distributing a public domain multilevel, aggregation-disaggregation code, Madpack (see [10] and [13]) for some years. Madpack is really a linear algebra package, rather than a package designed specifically for partial differential equations. The user specifies the domain or differential operator to Madpack indirectly, not directly. Interpolations and projections are computed as matrix-vector multiplies. Several sparse matrix formats are allowed including a stencil format which is very efficient on regular meshes.

Madpack2, first distributed in early 1986, has a sparse matrix-vector (and matrix transpose-vector) multiplier, a sparse direct solver, and three

smoothers, namely, Gauss-Seidel, conjugate gradients, Orthomin(1). The latter two are preconditioned by SSOR. It is quite compact.

As an experiment, I parallelized parts of it in 1989 in the spirit of finding out how painlessly it could be done. Since I wanted to run my code on distributed memory (Intel iPSC2) and shared memory (Sequent Symmetry) parallel processors, as well as a network of workstations with a minimal amount of code changes, I started from the C version of Madpack and adapted it to the Linda system [7]. (I rejected proprietary message passing systems at the time on the grounds that I refuse to program in any computer's assembly language, so why would I voluntarily program my communications in exactly that sort of nonportable environment?)

First, I realized that the obvious approach of storing matrices by column strips and vectors by row strips meant that the sparse matrix-vector multiplies were trivial to implement in parallel. I then replaced the three smoothers by two-diagonally preconditioned conjugate gradient and Orthomin(1)—which added a parallel dot product routine. Then I load balanced the operating processors by maintaining the same number of unknowns per processor independent of the level (effectively agglomerating at each level). Finally, I made absolutely sure that if I was doing a direct solve on the coarsest level, that I had only one processor involved. (The point of this exercise was to learn something about implementing something substantial in parallel and get something running quickly, not to produce a product quality code.)

The good news is that it was reasonably efficient in all three machine environments. The first implementation (on the Sequent) took a long time; the second (on the Intel iPSC2), which unfortunately required porting to the distributed memory environments, took an afternoon. With twenty minutes more work, I also had a simple three dimensional domain reduction example running (with 99% parallel efficiency) that produced publishable results. Assuming there was really enough data associated with each active processor to keep them all active computationally, the parallel efficiency could be kept in the 75%–99% range for most problems tried, even accounting for idle processors.

The bad news is that the differential equation front end to Madpack had to determine how to break natural data objects like vectors and sparse matrices into strips. This was not nearly as painful as had been expected, but introduced many subtle bugs into the process that had to be found and fixed.

4.2 Explicit Message Passing

Over the years, Madpack has evolved into an object oriented code based on a combination of C and Fortran [13]. More solvers have been added, temporary memory is dynamically allocated and freed, and the extremely complicated data structures have been hidden from the user.

The latest version, Madpack5, has been parallelized [28] using several message passing systems (IBM's MPL, Intel's NX/2, and MPI). Information concerning the global problem must be entered before the package can determine which processor has which part of the global data. A fairly simple approach was taken once again. The advantage is that the user interface remains almost identical to the serial case: each processor calls the same routines in the same way as in the serial case, but one more routine is called to register local and global information.

By using a small collection of routines for the parallel communications that is not based on any specific message passing system, it is now fairly painless to add interfaces to message passing systems not already addressed. Due to the large number of ports of systems like PVM [24] and MPI [25], this is likely to be the direction that almost everyone in the scientific computing community takes.

4.3 Cache Awareness

Many parallel computers of the middle 1990's use RISC based processors. As is commonly known, RISC processors require very careful use of their memory caches in order to operate at a high percentage of their peak speeds. (A cache is a very fast memory system that is tightly coupled to the processor which duplicates small parts of the main memory.) Most multigrid codes, for parallel or serial computers, do not attempt to optimize for caches. This is a serious flaw since it can be done quite easily, though not completely portably.

For example, consider a red-black ordering of Gauss-Seidel on a grid like that in Fig. 1. The red-black ordering specifies that the red points are relaxed on first, then the black ones. The usual implementation does this. Hence, all of the data (matrix, unknowns, and right hand side) passes through the cache twice. Instead, the following can be done:

Algorithm 6 *Cache Aware Red-Black Gauss-Seidel*

- (a) *Update all of the red points in row 1.*
- (b) *Do $j = 2, N$*
 - (b1) *Update all of the red points in row j .*
 - (b2) *Update all of the black points in row $j - 1$.*
 - (b3) *End Do*
- (c) *Update all of the black points in row N .*

When 4 grid rows of data along with the information from the corresponding rows of the matrix can be stored in cache simultaneously, this is a cache based algorithm (a domain decomposition approach can be used to do this when 4 rows do not fit). The advantage is that all of the data and the matrix pass through cache only once (instead of the usual twice) per iteration of Gauss-Seidel.

On many machines, this improvement alone speeds up a multigrid solver by 25 – 40% alone. Other steps, like residual computation, projection of residuals, and interpolation can be drawn into a complex, cache aware multigrid algorithm [12]. The speedup can be quite startling (e.g., 600% on some machines with particularly badly designed memory systems).

5 Conclusions

Parallel multilevel methods, which originated earlier this century in the personnel computing era, were the natural precursors to standard multilevel methods on single processors. Parallel multilevel methods are also the natural successors, in the age of advancing computers, to standard multigrid methods. What makes six parallel multilevel methods different, practical, and impractical was discussed in the context of the three algorithms that encapsulate them. Because no one has carefully compared all of these methods, on a collection of common problems, on a variety of machine architectures, it is difficult to determine the conditions in which a particular method is really the right or wrong choice.

Using high level programming tools makes implementing these methods much easier, although nontrivial. The use of low level tools, such as explicit message passing methods, should have been dismissed by the scientific computing community as a waste of human time. However, the community has

flocked to message passing systems and this choice has been considered here. In fact, the choice is not as bad as it once was now that a standard exists for message passing systems (i.e., MPI).

Whether or not high level tools are used, parallel multilevel methods scale well assuming there is enough data to make using a parallel computer quite worthwhile. By paying close attention to the cache, dramatic improvements in speed can also be achieved.

There are a number of parallel multigrid codes in existence. In particular, the SUPRENUM project produced a collection of interesting codes [2]. There are several parallel multigrid codes available by anonymous ftp from MGNet.

MGNet

MGNet (see [10] and [11]) is an Internet repository which contains numerous preprints, conference proceedings, codes, a large bibliography, and general information about multigrid methods. It can be accessed through either the World Wide Web or anonymous ftp.

Where	Anonymous ftp	WWW URL
USA	na.cs.yale.edu	http://na.cs.yale.edu/mgnet/www/mgnet.html
Europe	ftp.cerfacs.fr	http://www.cerfacs.fr

A monthly electronic newsletter can be subscribed to by sending a request by e-mail to *mgnet-requests@cs.yale.edu*.

References

- [1] L. M. ADAMS AND H. F. JORDAN, *Is SOR color-blind?*, SIAM J. Sci. Stat Comp., 7 (1986), pp. 490–506.
- [2] M. ALEF, *Implementation of a multigrid algorithm on SUPRENUM and other systems*, Parallel Comput., 20 (1994), pp. 1547–1557.
- [3] G. P. ASTRAKHANTSEV, *An iterative method of solving elliptic net problems*, Z. Vycisl. Mat. i. Mat. Fiz., 11 (1971), pp. 439–448.

- [4] N. S. BAKHVALOV, *On the convergence of a relaxation method under natural constraints on an elliptic operator*, Z. Vycisl. Mat. i. Mat. Fiz., 6 (1966), pp. 861–883.
- [5] A. BRANDT, *Multigrid solvers on parallel computers*, in Elliptic Problem Solvers, M. H. Schultz, ed., Academic Press, New York, 1981, pp. 39–83.
- [6] F. BREZZI, C. C. DOUGLAS, AND L. D. MARINI, *A parallel domain reduction method*, Numer. Meth. for PDE, 5 (1989), pp. 195–202.
- [7] N. CARRIERO AND D. GELERNTER, *Linda in context*, Comm. ACM, 32 (1989), pp. 444–458.
- [8] D. CHAZAN AND W. L. MIRANKER, *Chaotic relaxation*, J. Lin. Alg. Applic., 2 (1969), pp. 199–222.
- [9] C. C. DOUGLAS, *A tupleware approach to domain decomposition methods*, Appl. Numer. Math., 8 (1991), pp. 353–373.
- [10] ———, *MGNet Digests and Code Repository*. Monthly digests subscribed to by sending a message to mgnet-requests@cs.yale.edu and an anonymous ftp site (casper.cs.yale.edu) for codes and papers on multigrid and related topics.
- [11] ———, *MGNet: a multigrid and domain decomposition network*, ACM SIGNUM Newsletter, 27 (1992), pp. 2–8.
- [12] ———, *Caching in with multigrid algorithms: problems in two dimensions*, Tech. Rep. 20091, IBM Research Division, Yorktown Heights, New York, 1995.
- [13] ———, *Madpack: A family of abstract multigrid or multilevel solvers*, Comput. Appl. Math., 14 (1995), pp. 3–20.
- [14] C. C. DOUGLAS AND J. DOUGLAS, *A unified convergence theory for abstract multigrid or multilevel algorithms, serial and parallel*, SIAM J. Numer. Anal., 30 (1993), pp. 136–158.
- [15] C. C. DOUGLAS, J. DOUGLAS, AND D. E. FYFE, *A multigrid unified theory for non-nested grids and/or quadrature*, E. W. J. Numer. Math., 2 (1994), pp. 285–294.

- [16] C. C. DOUGLAS AND J. MANDEL, *A group theoretic approach to the domain reduction method*, Computing, 48 (1992), pp. 73–96.
- [17] C. C. DOUGLAS AND W. L. MIRANKER, *Constructive interference in parallel algorithms*, SIAM J. Numer. Anal., 25 (1988), pp. 376–398.
- [18] A. FÄSSLER AND E. STIEFEL, *Group Theoretical Methods and Their Applications*, Birkhäuser, Berlin, 1992.
- [19] R. P. FEDORENKO, *A relaxation method for solving elliptic difference equations*, Z. Vycisl. Mat. i. Mat. Fiz., 1 (1961), pp. 922–927. Also in U.S.S.R. Comput. Math. and Math. Phys., 1 (1962), pp. 1092–1096.
- [20] ———, *The speed of convergence of one iteration process*, Z. Vycisl. Mat. i. Mat. Fiz., 4 (1964), pp. 559–563. Also in U.S.S.R. Comput. Math. and Math. Phys., 4 (1964), pp. 227–2356.
- [21] P. O. FREDERICKSON AND O. A. MCBRYAN, *Parallel superconvergent multigrid*, in Multigrid Methods: Theory, Applications, and Supercomputing, S. F. McCormick, ed., vol. 110 of Lecture Notes in Pure and Applied Mathematics, Marcel Dekker, New York, 1988, pp. 195–210.
- [22] D. B. GANNON AND J. R. ROSENDALE, *On the structure of parallelism in a highly concurrent pde solver*, J. Parallel Distrib. Comput., 3 (1986), pp. 106–135.
- [23] U. GÄRTEL AND K. RESSEL, *Parallel multigrid: grid partitioning versus domain decomposition*, in Proceedings of the 10th International Conference on Computing Methods in Applied Sciences and Engineering, R. Glowinski, ed., New York, 1992, Nova Science Publishers, pp. 559–568.
- [24] A. GEIST, A. BEGUELIN, J. J. DONGARRA, W. JIANG, R. MANCHEK, AND V. SUNDERAM, *PVM Parallel Virtual Machine*, MIT Books, Cambridge, MA, 1994.
- [25] W. D. GROPP, E. LUSK, AND A. SKJELLUM, *Using MPI: Portable Parallel Programming with the Message-Passing Interface*, Scientific and Engineering Computation, MIT Press, Cambridge, MA, 1994.

- [26] W. HACKBUSCH, *A new approach to robust multi-grid solvers*, in ICIAM'87: Proceedings of the First International Conference on Industrial and Applied Mathematics, A. H. P. Burgh and R. M. M. Mattheij, eds., SIAM, Philadelphia, 1988, pp. 114–126.
- [27] ———, *The frequency decomposition multigrid method, part I: Application to anisotropic equations*, Numer. Math., 56 (1989), pp. 229–245.
- [28] B. MOON, G. PATNAIK, R. BENNETT, D. FYFE, A. SUSSMAN, C. C. DOUGLAS, J. SALTZ, AND K. KAILASANTH, *Runtime support and dynamic load balancing strategies for structured adaptive applications*, in Proceedings of the Seventh SIAM Conference on Parallel Processing for Scientific Computing, D. H. Bailey, P. E. Bjorstad, J. R. Gilbert, M. D. Mascagni, R. Schrieber, H. D. Simon, V. J. Torczon, and L. T. Watson, eds., SIAM Books, Philadelphia, 1995, pp. 575–580.
- [29] H. A. SCHWARZ, *Über einige Abbildungsaufgaben*, Ges. Math. Abh., 11 (1869), pp. 65–83.
- [30] R. V. SOUTHWELL, *Relaxation Methods in Engineering Science*, Oxford University Press, Oxford, 1940.
- [31] S. TA'ASAN, *Multigrid Methods for Highly Oscillatory Problems*, PhD thesis, Weizmann Institute of Science, Rehovot, Israel, 1984.