

The Fast Adaptive Composite Grid (FAC) Method for Elliptic Equations*

By S. McCormick and J. Thomas

Abstract. The fast adaptive composite grid (FAC) method is a systematic process for solving differential boundary value problems. FAC uses global and local uniform grids both to define the composite grid problem and to interact for its fast solution. It can with little added cost substantially improve accuracy of the coarse grid solution and is very suitable for vector and parallel computation. This paper develops both the theoretical and practical aspects of FAC as it applies to elliptic problems.

1. Introduction. The need for local resolution in physical models occurs frequently in practice. Special local features of the forcing function, operator coefficients, boundary, and boundary conditions can demand resolution in restricted regions of the domain that is much finer than the required global resolution. It is important that the discretization and solution processes account for this locally, that is, that the local phenomena do not precipitate a dramatic increase in the overall computation. Unfortunately, this objective of efficiently adapting to local features is often in conflict with the solution process: equation solvers can degrade or even fail to apply in the presence of varying discretization scales; data structures that account for irregular grids can be cumbersome; the computer architecture may not be able to effectively account for grid irregularity (e.g., “vectorizability” may be inhibited); etc. In fact, even the discretization process itself may find difficulty with this objective: for finite differences, it is problematic to develop accurate difference formulae for irregular grids; for finite elements, this objective is reflected in the substantial overhead costs needed to automate the discretization.

The fast adaptive composite grid method (FAC [11]) is a discretization and solution method designed to achieve efficient local resolution by constructing the discretization based on various regular grids and using these grids as a basis for fast solution. Its basic computational objective is to solve a “good” discretization on an irregular grid by way of regular grids only. Its basic assumption is that both discretization and solution on regular grids are easy by comparison.

Although FAC is in essence very similar to multi-level adaptive techniques (MLAT; cf. [1], [3], [5], [10]) and local defect correction (LDC; cf. [8]), it differs in several simple but important respects. First, FAC is a systematic approach that

Received July 10, 1984; revised August 2, 1985.

1980 *Mathematics Subject Classification*. Primary 65N20; Secondary 65N30.

*This work was supported by AFOSR grant number FQ8671-83-01322 and the National Bureau of Standards.

©1986 American Mathematical Society
0025-5718/86 \$1.00 + \$.25 per page

develops the discrete operator equation explicitly, not as an implicit result of the technique itself. The importance of this is that FAC is generally more robust than MLAT and more efficient than LDC (cf. [11] and [8]), for example. Moreover, its convergence properties are well founded in theory (cf. [11] and Section 3 below). Second, we believe that FAC is conceptually simpler, especially in its most basic form as a correction scheme with exact solvers (see Subsections 3.1 and 4.5). This is partly because FAC does not presume that the coarse grids are used for fast solution of the fine grid equation and thereby clearly distinguishes between this role of the coarse grid and that in transmitting fine grid accuracy throughout the domain. Third, FAC is somewhat more general by description in that it allows for general solvers in the treatment of each grid level, a necessity if FAC is to be patched into existing applications software.

The point here is that the differences are not fundamental—in fact, simple modifications lead to versions of MLAT, LDC, and FAC that are completely equivalent—but that these differences can be critical to numerical performance.

There are other related multigrid-type methods (cf. [2], [9]), including the method of aggregation (cf. [6], [12]). However, these schemes are not applied to local grid discretizations. FAC is also related to other local grid techniques (cf. [4], [7]), although these latter methods are designed for applications that allow local-global communication through the emerging solution only. (As with other multigrid-type techniques, FAC uses the equations for this communication and is usually able to achieve an optimal computation complexity.)

After describing the basic FAC method (Section 2), we expand on the theory of [11] (Section 3), treat several practical issues (Section 4), and demonstrate the typical numerical performance of FAC on two sample problems (Section 5). This development is restricted to the “variational” case (symmetric nonnegative-definite discretizations) with the local grids already in place. We defer nonvariational and adaptive considerations to later work. Note that we borrow freely from multigrid terminology and its constructs.

2. Method. We describe in this section the basic FAC algorithm, restricting our attention to the case of two grids, a *global* and a *local* one. (See Subsection 4.1 for comments on the multi-level case.) For concreteness, we begin with a discussion of our prototype case which we then quickly generalize in Subsection 2.2.

2.1. The Prototype Case (Poisson’s Equation on a Staggered Grid). Consider the two-dimensional cell-centered grid structure depicted in Figure 1. (An alternate but somewhat more complex case can be developed for the vertex-centered structure depicted in Figure 2. See Subsection 4.3.) We place in the region $\Omega = [0, 1] \times [0, 1]$ a uniform *global coarse grid* G with m cells on a side; on the subregion $\Omega_F = [0, \alpha] \times [0, \alpha]$ with $\alpha = k/m$ and $1 \leq k \leq m$, we place a uniform *local fine grid* \mathcal{G}_F with n cells on a side. (In Figure 1, $m = n = 4$ and $k = 2$. Note that small cells are $1/4$ the size of the large ones.) The *composite grid* \mathcal{G} is the union of the cells of \mathcal{G}_F and the cells of G in $\mathcal{C}\Omega_F$, the complement of Ω_F in Ω . We further partition the composite grid so that $\mathcal{G} = \mathcal{G}_C \cup \mathcal{G}_I \cup \mathcal{G}_F$ where \mathcal{G}_I consist of the n coarse cells bounding \mathcal{G}_F and \mathcal{G}_C consists of the remaining coarse cells. Note that \mathcal{G}_C is insulated from \mathcal{G}_F by \mathcal{G}_I . We similarly partition $G = G_C \cup G_I \cup G_F$.




FIGURE 1
(*Staggered Grids*)

The coarse grid G is represented by the cells defined by solid lines, the composite grid \mathcal{G} by cells defined by solid and dotted lines, and the actual boundary by bold lines. The interface $\mathcal{G}_I = G_I$ between the local grid \mathcal{G}_F and the global grid G consists of the *four* large cells with exactly one side common to cells of \mathcal{G}_F .




FIGURE 2
(*Aligned Grids*)

The coarse grid G is represented by \times 's, the composite grid \mathcal{G} by \times 's and \bullet 's, and the actual boundary points by o 's. The interface between the coarse and fine levels consists of the three \times 's bordering \mathcal{G}_F .

Consider the singular Poisson equation

$$(2.1) \quad \begin{aligned} -\Delta\Psi &= \phi && \text{in } \Omega, \\ \Psi_n &= 0 && \text{on } \partial\Omega, \end{aligned}$$

where Ψ and ϕ are functions on Ω and Ψ_n denotes the partial of Ψ in the direction normal to the boundary $\partial\Omega$. As a discretization of (2.1) on \mathcal{G} , let \mathcal{L} be the symmetric nonnegative-definite matrix constructed according to Figure 3. Thus, \mathcal{L} is based on the usual five-point stencil in \mathcal{G}_C and \mathcal{G}_F (except for the latter it reaches to \mathcal{G}_I points with weight scaled by $1/4$). At points of \mathcal{G}_I , \mathcal{L} reaches to its immediate $\mathcal{G}_C \cup \mathcal{G}_I$ neighbors with value $-1/8h^2$, but to both of its \mathcal{G}_F neighbors it reaches with value $-1/4h^2$ (i.e., its weight to both totals 4 times its weight to a point in $\mathcal{G}_C \cup \mathcal{G}_I$). See Subsection 4.2 for a discussion of how \mathcal{L} was derived.

The composite right-hand side f represents approximate integrated values of ϕ over appropriate squares. We assume here the crudest definition of f where its

		$1/4$		$1/8$		$3/8$		$1/8$		$3/8$		$1/8$		$1/4$	
		$-1/8$				$-1/8$				$-1/8$				$-1/8$	
		$3/8$				$1/2$				$1/2$				$3/8$	
		$-1/4$		$-1/4$		$-1/4$		$-1/4$				$-1/8$		$-1/8$	
		3		4		1		4		1		4		$-1/4$	
		-1		-1		-1		-1				$1/2$		$3/8$	
		3		4		1		4		1		4		$-1/4$	
		-1		-1		-1		-1				$-1/8$		$-1/8$	
		3		4		1		4		1		4		$-1/4$	
		-1		-1		-1		-1				$3/8$		$1/8$	
		2		3		1		3		1		3		$-1/4$	

FIGURE 3
(Stencil Elements for \mathcal{L} -Staggered Grids)

The entries of $h^2\mathcal{L}$ are depicted here, where $h = 1/2m = 1/8$. The diagonal of $h^2\mathcal{L}$ corresponding to a cell has the value depicted in the cell's center. Its off-diagonal entries connecting it to each neighboring cell with a common side have values depicted on that side. For example, the stencil of $h^2\mathcal{L}$ on the southeast corner of \mathcal{G}_F is $(-i \ 3^{-1} \ -1/4)$ and the stencil for its \mathcal{G}_I neighbor is

$$\begin{pmatrix} & & -1/8 \\ -1/4 & & \\ & 3/8 & -1/8 \\ -1/4 & & \end{pmatrix}.$$

$1/4$	$1/8$	$3/8$	$1/8$	$3/8$	$1/8$	$1/4$
$-1/8$	$-1/8$	$-1/8$	$-1/8$	$-1/8$	$-1/8$	$-1/8$
$3/8$	$1/8$	$1/2$	$1/8$	$1/2$	$1/8$	$3/8$
$-1/8$	$-1/8$	$-1/8$	$-1/8$	$-1/8$	$-1/8$	$-1/8$
$3/8$	$1/8$	$1/2$	$1/8$	$1/2$	$1/8$	$3/8$
$-1/8$	$-1/8$	$-1/8$	$-1/8$	$-1/8$	$-1/8$	$-1/8$
$1/4$	$1/8$	$3/8$	$1/8$	$3/8$	$1/8$	$1/4$

FIGURE 4
(Stencil Elements for L-Staggered Grids)

The entries of h^2L are depicted here, again with $h = 1/2m = 1/8$. All off-diagonal connections of L are $-1/8h^2$; only the diagonal varies (according to boundary effects).

entries are the (approximate) values of ϕ in the cell centers, scaled by a factor of h^2 for \mathcal{G}_F cells and by $4h^4$ for $\mathcal{G}_C \cup \mathcal{G}_I$ cells. The target equation is then the composite grid problem given by

$$(2.2) \quad \mathcal{L}u = f, \quad u \in \mathcal{G}.$$

(Throughout this paper we identify grids with their associated function spaces.)

To apply FAC to (2.2), first let the coarse grid operator L associated with G be defined as in Figure 4. Let the *prolongation* or *interpolation* operator $I: G \rightarrow \mathcal{G}$ be based on piecewise constant interpolation so that $u = Iu$ is such that $u_C = u_C$, $u_I = u_I$, and the value of u_F for a given cell of \mathcal{G}_F is $1/4$ the value of u for the corresponding G_F cell. For the *restriction* operator we choose the transpose $I^T: \mathcal{G} \rightarrow G$. Note that our operators satisfy the *variational condition*

$$(2.3) \quad L = I^T \mathcal{L} I.$$

Now partition each vector and operator according to the partitions of \mathcal{G} and G . For example, we write

$$u = \begin{pmatrix} u_C \\ u_I \\ u_F \end{pmatrix} \quad \text{and} \quad L = \begin{pmatrix} L_{CC} & L_{CI} & 0 \\ L_{IC} & L_{II} & L_{IF} \\ 0 & L_{FI} & L_{FF} \end{pmatrix}.$$

Letting id denote the appropriate identity operator, we then have

$$\begin{aligned}
 & \text{(i)} \quad L_{FC}, L_{CF}, \mathcal{L}_{FC}, \text{ and } \mathcal{L}_{CF} \text{ are all zero;} \\
 (2.4) \quad & \text{(ii)} \quad I = \begin{pmatrix} \text{id} & 0 & 0 \\ 0 & \text{id} & 0 \\ 0 & I_{FC} & I_{FF} \end{pmatrix}; \\
 & \text{(iii)} \quad G_C = \mathcal{G}_C \text{ and } G_I = \mathcal{G}_I; \text{ and} \\
 & \text{(iv)} \quad L_{CC} = \mathcal{L}_{CC}, L_{CI} = \mathcal{L}_{CI}, \text{ and } L_{IC} = \mathcal{L}_{IC}.
 \end{aligned}$$

Actually, $I_{FC} = 0$ for our prototype, but we do not expect this in general.

One two-level exact solver cycle of FAC applied to an approximation, u , of \mathcal{U} in (2.2) with right-hand side f is then denoted by $u \leftarrow \text{FAC}(u, f)$ and defined by:

Step 1. $u \leftarrow L^{-1}I^T(f - \mathcal{L}u)$ (coarse grid approximation).

Step 2. $u \leftarrow u + Iu$ (coarse grid correction).

Step 3. $u_F \leftarrow \mathcal{L}_{FF}^{-1}(f_F - \mathcal{L}_{FI}u_I)$ (local fine grid improvement).

Actually, the inverse in Step 1 is taken loosely here for our prototype since both L and \mathcal{L} are singular with null spaces $\eta(L)$ and $\eta(\mathcal{L})$ spanned by $e \equiv 1$ and $e_F \equiv 1$, respectively. Thus by $v = L^{-1}g$ we really mean $v = L^\dagger g$, the minimal norm solution of $Lv = \hat{g} \equiv g - \langle g, e \rangle e / \langle e, e \rangle$. (Here, L^\dagger is the Moore-Penrose generalized inverse of L and $\langle v, w \rangle = v^T w$.) We can in this way analyze FAC on the spaces orthogonal to the ‘‘physical constants’’ e and

$$e = \begin{pmatrix} e_C \\ e_I \\ e_F/4 \end{pmatrix}.$$

See Subsection 3.3 for a general discussion.

Note that FAC attempts to solve the composite grid problem (2.2) by way of uniformly rectangular discrete Poisson equations. Without the component of L in G_F , then FAC would constitute a *block* Gauss-Seidel relaxation scheme with a block each corresponding to $\mathcal{G}_C \cup \mathcal{G}_I$ and to \mathcal{G}_F . The component of L in G_F thus represents an attempt to improve the approximate inverse of \mathcal{L} in order to eliminate the error remaining after Step 3; this is the key to making FAC a fast solver.

2.2. *The General Case.* The extension of the notation and definitions of Subsection 2.1 to our basic but more general case is now immediate. Specifically, following the development in [11], consider the composite grid problem (2.2) where now \mathcal{G} is, in general, a space of real-valued functions defined on the composite grid (which again we identify with \mathcal{G}) contained in a region Ω of finite-dimensional Euclidean space. Assume that $\mathcal{L}: \mathcal{G} \rightarrow \mathcal{G}$ is a symmetric positive-definite linear operator and that \mathcal{G} admits a coarse grid G (perhaps $G \subset \mathcal{G}$) in Ω and an intergrid transfer $I: G \rightarrow \mathcal{G}$ that is a full rank linear operator. For concreteness, we assume that the space \mathcal{G} is finite-dimensional and that G is of smaller dimension, so that the coarse grid operator $L: G \rightarrow G$ defined by (2.3) is symmetric, positive-definite and linear. Here, again we choose $I^T: \mathcal{G} \rightarrow G$ as the restriction operator.

As in Subsection 2.1, we assume that the partitions $G = G_C \cup G_I \cup G_F$ and $\mathcal{G} = \mathcal{G}_C \cup \mathcal{G}_I \cup \mathcal{G}_F$ satisfy (2.4). Note that we assume no special properties like connectedness of the local region Ω_F covered by \mathcal{G}_F or any particular dimension of the space in which Ω lies.

The *exact solver* version of FAC is exactly as given in Subsection 2.1.

3. Theory. We provide here foundation for various two-level versions of FAC. (See Subsection 4.1 for the multi-level case.) We begin with the version described in the previous section where both the coarse and fine grid equations are solved exactly.

3.1. *Basic Theory.* Let the *energy* inner product be defined by $\langle u, \omega \rangle_{\mathcal{G}} = \langle u, \mathcal{L}\omega \rangle$ and the induced *energy* norm by $\|u\|_{\mathcal{G}}$ and similarly for L . Let $\delta = \delta(L, \mathcal{L})$ be defined by $\delta = \rho(L)\rho(\mathcal{P}_L\mathcal{L}^{-1})$ where ρ denotes spectral radius and $\mathcal{P}_L\mathcal{L}^{-1} = \mathcal{L}^{-1} - IL^{-1}I^T$. Let $C = C(L, \mathcal{L})$ denote the maximum per cycle energy convergence factor of FAC, i.e., the smallest C so that

$$\|\text{FAC}(u, f) - \mathcal{U}\|_{\mathcal{G}} \leq C\|u - \mathcal{U}\|_{\mathcal{G}}$$

for any u, f in \mathcal{G} .

THEOREM 1. $C \leq (\delta/(1 + \delta))^{1/2}$.

Proof. See [11].

Remark. In collaboration with Vance Faber of the Los Alamos National Laboratory, we were able to improve the bound C in this theorem by showing that we may take $\delta = \rho(L_{II}Q) \leq \rho(L_{II})\rho(Q)$, where $Q = P(\mathcal{P}_L\mathcal{L}^{-1})P$ and

$$P = \begin{pmatrix} 0 & 0 & 0 \\ 0 & \text{id} & 0 \\ 0 & 0 & 0 \end{pmatrix}$$

is the orthogonal projector of \mathcal{G} onto $\mathcal{G}_I = G_I$.

To guarantee an optimal complexity of FAC, it is necessary (though not always sufficient) to establish that $C \ll 1$, by which we mean that a bound for C less than 1 exists that is independent of the mesh sizes of either L or \mathcal{L} . But to use Theorem 1, we must therefore show that $\delta \ll \infty$. As we noted in [11], this is easy to prove for variationally formulated boundary value problems, provided the usual operator regularity and approximation properties hold. Such is the case, for example, when (2.2) is a finite element discretization of a second-order elliptic partial differential equation with smooth coefficients and boundary, with essentially uniform grids G and \mathcal{G}_F , and with piecewise polynomial basis functions of degree at least two. In fact, $\delta \ll \infty$ uniformly in the ratio of the mesh sizes of G and \mathcal{G}_F . But this theorem is not yet sufficiently general to cover many of our practical applications.

First, it does not completely cover our prototype. The essential difficulty here is that the prototype discretization, though variational, is based on piecewise *constant* elements which are inadmissible functions for (2.1). Thus, to use Theorem 1, we must derive estimates for δ by, say, algebraic means, but we defer this to a future paper. A second difficulty here is that it does not cover the case that \mathcal{L} is singular; we will treat this case in Subsection 3.3. A third difficulty is that we cannot guarantee $\delta \ll \infty$ uniformly in the *grade* of G , that is, the ratio of its largest to smallest mesh size. In fact, we can easily construct G_C in general with a large enough

mesh size relative to, say, G_I so that δ is as large as desired. C does not really degrade for such cases, however, as we will deduce from our next theorem. (Other limitations of Theorem 1 will be treated in subsequent sections.)

THEOREM 2. *Suppose G'_C is a nonempty subspace (subgrid) of G_C and let $G' = G'_C \cup G_I \cup G_F$ and $\mathcal{G}' = G'_C \cup G_I \cup \mathcal{G}_F$. Assume $\mathcal{I}': \mathcal{G}' \rightarrow \mathcal{G}$ is a full rank linear operator of the form*

$$\mathcal{I}' = \begin{pmatrix} \mathcal{I}'_{CC} & \mathcal{I}'_{CI} & 0 \\ 0 & \text{id} & 0 \\ 0 & 0 & \text{id} \end{pmatrix}$$

and let $\mathcal{L}' = \mathcal{I}'^T \mathcal{L} \mathcal{I}'$ and $L' = I'^T \mathcal{L} I'$, where I' is I restricted to G' . Then $C(L', \mathcal{L}') \leq C(L, \mathcal{L})$.

Proof. See [11].

The significance of this theorem is that it yields bounds by way of Theorem 1 for the case that the coarse grid is badly graded by virtue of the coarseness of G'_C . More precisely, although $\delta' = \delta(L', \mathcal{L}')$ may be large when the ratio of the mesh sizes in G'_C and $G_I \cup G_F$ is large, Theorem 2 shows that $C' = C(L', \mathcal{L}')$ may not be adversely affected by this ratio. For example, if we imagine G' itself to be a composite grid created by placing a local fine grid patch (that includes $G_I \cup G_F$) on a very coarse global grid, where the coarse and fine grids are both uniform, then G' may be badly graded. Yet the convergence bound C' is no worse than the case of uniform G_C where G'_C is refined to be in balance with $G_I \cup G_F$.

Theorem 2 does not cover the case where a large grade of G arises from the ratio of the mesh sizes of $G_C \cup G_I$ and G_F . This can develop in practice if G is itself a composite grid where its local grid is exactly G_F . The next theorem covers this case.

THEOREM 3. *Suppose now that G'_F is a nonempty subspace of G_F . Let $G' = G_C \cup G_I \cup G'_F$ and $L' = I'^T \mathcal{L} I'$ where I' is I restricted to G' . Then*

$$C(L, \mathcal{L}) \leq C(L', \mathcal{L}').$$

Proof. This follows trivially from the proof of Theorem 1 by appealing to the variational setting and concluding that the accuracy in the coarse grid solution of Step 1 cannot be better for G' than it is for G .

3.2. Approximate Solvers. For the theory here to be of general practical value, it must allow for the use of inexact solvers in Steps 1 and 3 of FAC. It may be that iterative methods are used in Step 3 of many practical algorithms. But more critically, FAC will often be used in a multi-level setting by placing several increasingly finer local grids about isolated phenomena. This has the effect that all but the very coarsest level is itself a composite grid and the solver corresponding to Step 1 for this level is just FAC. In short, a multi-level version of FAC necessarily involves approximate solvers.

To treat such cases, suppose $0 \leq \alpha, \beta < 1$ where $\beta < 1 - C$ and $C = C(L, \mathcal{L})$. Then one *approximate solver* cycle of FAC is denoted by $u \leftarrow \text{FAC}_{\alpha, \beta}(u, f)$ and defined by:

Step 1. Let $r = I^T(f - \mathcal{L}u)$ and set $u \leftarrow L^{-1}r + q$ where $\|q\|_L \leq \beta \|L^{-1}r\|_L$.

Step 2. $u \leftarrow u + Iu$.

Step 3. Let $\ell_F = f_F - \mathcal{L}_{FI}u_I$ and set $v_F \leftarrow \mathcal{L}_{FF}^{-1}\ell_F + q_F$, where $\|q_F\|_{\mathcal{L}_{FF}} \leq \alpha\|u_F - \mathcal{L}_{FF}^{-1}\ell_F\|_{\mathcal{L}_{FF}}$. Set $u_F \leftarrow v_F$.

Step 1 requires that $LU = r$ be solved to within a relative energy error of β . This means that the error in the initial vector u must be reduced by a fixed factor β . Step 3 requires a similar condition, namely, that u_F be improved as an approximate solution of $\mathcal{L}_{FF}u_F = \ell_F$ at least by a factor of α . Note that $\text{FAC}(u, f) = \text{FAC}_{0,0}(u, f)$.

THEOREM 4. *Suppose the energy convergence bound $C = C(L, \mathcal{L})$ for FAC satisfies $C < 1$ and let $\beta < 1 - C$ and $\alpha < 1$. Then, a bound on the energy convergence factor $C_{\alpha,\beta}$ for $\text{FAC}_{\alpha,\beta}$ is given by $C_{\alpha,\beta} \leq ((1 - \alpha^2)(C + \beta)^2 + \alpha^2)^{1/2} < 1$.*

Proof. Assume without loss of generality that $f = 0$ and that u is arbitrary. Note that u is now the actual error in (2.2). We first consider the case $\alpha = 0$ so that

$$\text{FAC}_{\alpha,\beta}(u, f) = \text{FAC}(u, f) + Qq$$

where Qq is just $\text{FAC}(q, 0)$ with Step 1 suppressed. Since Q cannot increase energy, then

$$\begin{aligned} \|Qq\|_{\mathcal{L}} &\leq \|Iq\|_{\mathcal{L}} = \|q\|_L \leq \beta\|L^{-1}I^T\mathcal{L}u\|_L = \beta\langle \mathcal{L}L^{-1}\mathcal{L}^T\mathcal{L}u, \mathcal{L}u \rangle^{1/2} \\ &\leq \beta\langle \mathcal{L}^{-1}\mathcal{L}u, \mathcal{L}u \rangle^{1/2} = \beta\|u\|_{\mathcal{L}}. \end{aligned}$$

The last inequality follows from noting as in [11] that $\mathcal{P}_L = \mathcal{L}^{-1} - IL^{-1}I^T$ is nonnegative. The conclusion for $\alpha = 0$ now follows from noting that

$$\|\text{FAC}_{\alpha,\beta}(u, f)\|_{\mathcal{L}} \leq \|\text{FAC}(u, f)\|_{\mathcal{L}} + \|Qq\|_{\mathcal{L}} \leq (C + \beta)\|u\|_{\mathcal{L}}.$$

Now lift the restriction that $\alpha = 0$ and let $v^\alpha = \text{FAC}_{\alpha,\beta}(u, f)$ and $v^0 = \text{FAC}_{0,\beta}(u, f)$. We must show that

$$(3.1) \quad \|v^\alpha\|_{\mathcal{L}} \leq C_{\alpha,\beta}\|u\|_{\mathcal{L}}.$$

To this end, note first that $v_F^0 = \mathcal{L}_{FF}^{-1}\ell_F = -\mathcal{L}_{FF}^{-1}\mathcal{L}_{FI}u_I^0$ where u^0 is $\text{FAC}_{\alpha,\beta}(u, 0)$ with Step 3 suppressed. Thus,

$$(3.2) \quad \|q_F\|_{\mathcal{L}_{FF}} = \|v_F^\alpha + \mathcal{L}_{FF}^{-1}\mathcal{L}_{FI}u_I^0\|_{\mathcal{L}_{FF}} \leq \alpha\|u_F^0 + \mathcal{L}_{FF}^{-1}\mathcal{L}_{FI}u_I^0\|_{\mathcal{L}_{FF}}.$$

A little calculation shows that

$$\|v^\alpha\|_{\mathcal{L}}^2 \leq \|v_F^\alpha + \mathcal{L}_{FF}^{-1}\mathcal{L}_{FI}u_I^0\|_{\mathcal{L}_{FF}}^2 + \|v^0\|_{\mathcal{L}}^2$$

and that

$$\|u_F^0 + \mathcal{L}_{FF}^{-1}\mathcal{L}_{FI}u_I^0\|_{\mathcal{L}_{FF}}^2 = \|u^0\|_{\mathcal{L}}^2 - \|v^0\|_{\mathcal{L}}^2.$$

Thus, by (3.2) we have

$$(3.3) \quad \|v^\alpha\|_{\mathcal{L}}^2 \leq \|v^0\|_{\mathcal{L}}^2 + \alpha^2(\|u^0\|_{\mathcal{L}}^2 - \|v^0\|_{\mathcal{L}}^2).$$

But

$$(3.4) \quad \|v^0\|_{\mathcal{L}}^2 \leq (C + \beta)^2\|u\|_{\mathcal{L}}^2$$

and

$$(3.5) \quad \|u^0\|_{\mathcal{L}}^2 \leq \|u\|_{\mathcal{L}}^2.$$

(3.3)–(3.5) now combine to prove (3.1) and the theorem is proved.

3.3. *Singular and Nearly Singular Equations.* When \mathcal{L} is singular, as for our prototype case, FAC may be modified to apply without difficulty provided the singularity is global and exactly represented in L . Specifically, we require that

$$(3.6a) \quad \mathcal{L}_{FF} \text{ is nonsingular,}$$

$$(3.6b) \quad \eta(\mathcal{L}) \subset \mathcal{R}(I),$$

where η denotes null space and \mathcal{R} denotes range. (Note that (3.6b) is equivalent to the assumption that $\eta(\mathcal{L}) = I\eta(L)$.) The modification of FAC is achieved by replacing L^{-1} in Step 1 by L^\dagger . Replacing L^{-1} by L^\dagger also in $\delta = \delta(L, \mathcal{L})$, we then have the following result.

THEOREM 5. *Suppose $f \in \mathcal{R}(\mathcal{L})$ so that (2.2) has a solution. Then, the estimates in Theorems 1–4 apply here to the singular case provided (3.6) holds.*

Proof. A careful examination of the proof of Theorem 1 in [11] shows that it applies here with the only modification being the replacement of the inverse by the generalized inverse. (Note, for example, that our assumption on f allows us again to assume without loss of generality that $f = 0$.) Extending the proofs of Theorems 2–4 is completely straightforward.

Two warnings are in order concerning the singular case. First, the restriction $f \in \mathcal{R}(\mathcal{L})$ in Theorem 5 is needed to ensure that the coarse grid approximation in Step 1 is sensible. Specifically, if $f = f_{\mathcal{R}} + f_{\eta}$ where $f_{\mathcal{R}} \in \mathcal{R}(\mathcal{L})$ and $f_{\eta} \in \eta(\mathcal{L}) = \mathcal{R}^\perp(\mathcal{L})$, then the coarse grid approximation would not necessarily approximate $\mathcal{L}^\dagger f_{\mathcal{R}}$. That is, even though $I^T(f_{\mathcal{R}} - \mathcal{L}u)$ is in $\eta^\perp(L) = \mathcal{R}(L)$ because $f_{\mathcal{R}} - \mathcal{L}u$ is in $\eta^\perp(\mathcal{L}) = (Iy(L))^\perp$, generally $I^T f_{\eta}$ is not in $\eta(L)$. More to the point, $IL^\dagger I^T f_{\eta}$ is generally not in $\eta(\mathcal{L})$. Thus, $\eta(\mathcal{L})$ components of f can feed $\eta^\perp(\mathcal{L})$ components of u , thereby diverting u from its approximation of $L^\dagger f$. To avoid this when a basis for $\eta(\mathcal{L})$ is known, f_{η} should be *a priori* eliminated by orthogonalization. For example, before applying FAC to our prototype, we simply set $f \leftarrow P_e f$ where $P_e = \text{id} - ee^T/e^T e$.

The second warning is that the component of u in $\eta(\mathcal{L})$ will generally change during an FAC cycle. This may be of no concern, but when it is, a solution with a particular $\eta(\mathcal{L})$ component can be approximated *a posteriori* again by orthogonalization with respect to a basis for $\eta(\mathcal{L})$ when it is known. For example, the minimal norm solution $u = \mathcal{L}^\dagger f$ for our prototype can be approximated by setting $u \leftarrow P_e u$ after FAC is applied.

We now observe directly that there is also no difficulty with the nearly singular case, provided we have a condition analogous to (3.6b). Specifically, if \mathcal{L} has eigenvalues near zero in some relative sense, then we require that their associated eigenvectors are well approximated on grid G . For example, let \mathcal{L} be our prototype operator and consider the perturbation $\mathcal{L}_\epsilon = \mathcal{L} + \epsilon \text{id}$, $\epsilon > 0$. Note that $L_\epsilon = L + \epsilon I^T I$ and that the eigenvector, e , of \mathcal{L}_ϵ belonging to ϵ satisfies $e = Ie$, that is, this eigenvector is *exactly* represented on G . But since $I^T I e = I^T e$, then $L_\epsilon(e/\epsilon) = I^T e$ from which follows $IL_\epsilon^{-1} I^T e = e/\epsilon$ and, therefore, $\mathcal{P}_{L_\epsilon}(\mathcal{L}_\epsilon^{-1} e) = 0$. From this, we can conclude that $\delta(L_\epsilon, \mathcal{L}_\epsilon) = (1 + O(\epsilon))\delta(L, \mathcal{L})$ or, loosely speaking, that the near singularity of \mathcal{L}_ϵ does not adversely affect our estimate for C . Of course, for $\text{FAC}_{\alpha, \beta}$, we must have an approximate solver in Steps 1 and 3 that is able to deal effectively with nearly singular matrices.

4. Practical Matters. In the following subsections, we discuss in brief several issues that are important for practical implementation of FAC.

4.1. Multi-level FAC. For simplicity, we have thus far restricted our development to the two-level case. Extension to the general multi-level case is, however, fairly straightforward. Specifically, we suppose now that grid \mathcal{G} is a composite of a coarse grid G^0 and a sequence of $p \geq 1$ increasingly finer local grids, where the region Ω_F^i covered by a given local grid G_F^i contains the region covered by the next finer one G_F^{i+1} , $0 \leq i \leq p - 1$. This local refinement sequence thus gives rise to a sequence of global grids $G^0, G^1, \dots, G^p = \mathcal{G}$ that are equipped with interpolation operators $I^{i+1}: G^i \rightarrow G^{i+1}$ and their induced grid operators $L^i: G^i \rightarrow G^i$, $i = 0, 1, \dots, p - 1$. Write $L^p = \mathcal{L}$ and assume that each grid pair G^i and G^{i+1} together with their operators and partitions satisfy the assumptions of Subsection 2.2. Now suppose $f^i \in G^i$ is given and that $u^i \in G^i$ is an approximate solution of the level i problem

$$(4.1) \quad L^i U^i = f^i.$$

With the *cycling parameter* $\mu \geq 1$ given, then one multi-level μ -cycle version of $\text{FAC}_{\alpha,\beta}$ on level i is denoted by $u^i \leftarrow \text{FAC}_{\mu,\alpha,\beta}(u^i, f^i)$ and defined *recursively* by:

Step 1. If $i = 0$, set $u^i \leftarrow L^{i-1} f^i + q$ where $\|q\|_{L^i} \leq \beta \|L^{i-1} f^i\|_{L^i}$ and stop. Otherwise go to Step 2.

Step 2. Set $f^{i-1} \leftarrow I^{i-1}(f^i - L^i u^i)$ and $u^{i-1} \leftarrow 0$. Do $u^{i-1} \leftarrow \text{FAC}_{\mu,\alpha,\beta}(u^{i-1}, f^{i-1})$ μ times.

Step 3. Let $b_F^i = f_F^i - L_{FI}^i u_F^i$ and set $v_F^i \leftarrow L_{FF}^{i-1} b_F^i + q_F^i$ where $\|q_F^i\|_{L_{FF}^i} \leq \alpha \|u_F^i - L_{FF}^{i-1} b_F^i\|_{L_{FF}^i}$. Let $u_F^i \leftarrow v_F^i$.

Note that only the coarsest level G^0 is approximately solved according to the parameter β . The approximate solver for all other levels $i < p$ is in fact $\text{FAC}_{\mu,\alpha,\beta}$ itself. Note also that the computation on each level $i > 0$ is only on the local grids G_F^i .

The simplest cycling scheme is given by choosing $\mu = 1$ which, borrowing from multigrid terminology, yields the *sawtooth* or *slash* cycle version that on level p is written as $u^p \leftarrow \text{FAC}_{\alpha,\beta}(u^p, f^p)$ and given simply by:

Step 1. For $i = p, p - 1, \dots, 1$, set $f^{i-1} \leftarrow I^{i-1}(f^i - L^i u^i)$ and $u^{i-1} \leftarrow 0$.

Step 2. Set $u^0 \leftarrow L^{-1} f^0 + q$ where $\|q\|_{L^0} \leq \beta \|L^{-1} f^0\|_{L^0}$.

Step 3. For $i = 1, 2, \dots, p$, set $u^i \leftarrow u^i + I^i u^{i-1}$. Let $b_F^i = f_F^i - L_{FI}^i u_F^i$ and set $v_F^i \leftarrow L_{FF}^{i-1} b_F^i + q_F^i$ where $\|q_F^i\|_{L_{FF}^i} \leq \alpha \|u_F^i - L_{FF}^{i-1} b_F^i\|_{L_{FF}^i}$. Set $u_F^i \leftarrow v_F^i$.

Another type of cycling scheme is *accommodative* (again borrowing from multigrid terminology) which is denoted by $u^i \leftarrow \text{FAC}A_{\alpha,\beta}(u^i, f^i)$ and defined by:

Step 1. Same as $\text{FAC}_{\mu,\alpha,\beta}$.

Step 2. Let $f^{i-1} \leftarrow I^{i-1}(f^i - L^i u^i)$ and $u^{i-1} \leftarrow 0$. Do $u^{i-1} \leftarrow \text{FAC}A_{\alpha,\beta}(u^{i-1}, f^{i-1})$ until $\|u_F^{i-1} - L_{FF}^{i-1} f_F^{i-1}\|_{L_{FF}^{i-1}} \leq \beta \|L_{FF}^{i-1} f_F^{i-1}\|_{L_{FF}^{i-1}}$.

Step 3. Same as $\text{FAC}_{\mu,\alpha,\beta}$.

Of course, $\text{FAC}A_{\alpha,\beta}$ requires an estimate of the relative error in the energy norm. When such estimates are not directly available, they may be made by way of the *energy functional* $F(u) = \langle Lu - 2f, u \rangle$ or by appealing to the *residual norm* $r(u) = \|Lu - f\|$. Such estimates are usually acceptable in connection with FAC because of its strong convergence properties. (The behavior of $\|u - U\|_L$ may not be reflected in $F(u)$ for methods that converge only sublinearly or erratically; neither is it

reflected in $r(u)$ for methods like simple relaxation that tend to dispose the error to the smooth end of the spectrum of L .)

An important point here is that the convergence bound in Theorem 4 applies directly to $FAC A_{\alpha,\beta}$. On the other hand, unless we are willing to accept a bound for $FAC \mu_{\alpha,\beta}$ that depends on p , then we must have a condition that relates μ and the two-grid bounds $C^i = C_{\alpha,0}(L^{i-1}, L^i)$. For example, if we let $\bar{C} = \max_{1 \leq i \leq p} C^i$, then a convergence bound on $FAC \mu_{\alpha,\beta}$ is any $\hat{C} \geq \bar{C}, \beta$ that satisfies $\hat{C}^{2\mu} - \hat{C}^2 + \bar{C}^2 = 0$. Note that a μ can always be found that guarantees the existence of such a $\hat{C} < 1$ provided $\beta, \bar{C} < 1$. But only if $\bar{C} \leq 1/2$ will $\mu = 2$ suffice and only if $\bar{C} = \beta = 0$ will $\mu = 1$ suffice! However, our numerical experiments suggest that $FAC/\alpha,\beta$ converges with a bound $\hat{C} < 1$ independent of p ; this is important in terms of complexity, since the coarse grids may swamp the computation for $\mu > 1$. It is just that our theory does not yet cover this case.

Note for the exact solver $FAC/_{0,0}$ that there is no real purpose in having subsequent refinements over precisely the same geometric region; that is, if $\Omega_F^i = \Omega_F^{i-1}$, then the computation on G^i is completely wasted. Of course, for $FAC/\alpha,\beta$ in general, such "intermediate" grids may be *critical* to complexity by providing initial guesses for the local fine-grid problems that are much better than the coarser grids can provide.

4.2. \mathcal{L} Construction. Perhaps one of the more critical aspects of FAC is the construction of a composite grid operator \mathcal{L} that yields acceptable discretization properties, especially with respect to truncation error. We describe one possibility here that assumes that the only serious difficulty is in dealing with the nonuniformity of \mathcal{G} .




FIGURE 5
(*Extended Fine Grid—Staggered Grids*)

The composite grid \mathcal{G} is represented by the cells defined by solid lines, the extended grid G by solid and dotted lines.

To be specific, suppose \mathcal{G} is a subgrid of the *extended fine grid* G , where we have in mind Figure 5 for our prototype in which the local fine grid is extended over the entire domain Ω . Let L be an operator defined on G and suppose that the interpolation operator $I: \mathcal{G} \rightarrow G$ is linear and full rank and that $\mathbf{f} \in G$. (We suppose that acceptable L , I , and \mathbf{f} are easily constructed.) The composite grid problem (2.2) is then constructed by setting $\mathcal{L} = I^T L I$ and $\mathcal{f} = I^T \mathbf{f}$. Note that only the entries of \mathcal{L} involving \mathcal{G}_I are irregular, that is, it may not be necessary to appeal explicitly to L for entries of \mathcal{L} corresponding to \mathcal{G}_C . (See, for example, the prototype \mathcal{L} depicted in Figure 3.) Note also that I^T is used as the transfer from G to \mathcal{G} ; any (global) scalar multiple of I^T will do, but I^T must not be scaled locally.

The assumption here is of course that we do not need the solution U in G of

$$(4.2) \quad LU = \mathbf{f},$$

but rather are willing to accept the best energy approximation to it in $\mathcal{R}(I)$. In fact, this is in tune with the basic assumption of FAC that there is significantly less resolution needed in $\mathcal{E}\Omega_F$ than in Ω_F . The specific assumption here, which is somewhat subtle, is that we want to solve (2.2) to within an energy error tolerance that is some fixed factor $\gamma < 1$ of the error in the coarse grid approximation to its solution. That is, the convergence criterion is that the approximation u must satisfy

$$(4.3) \quad \|u - \mathcal{U}\|_{\mathcal{E}} \leq \gamma \|IL^{-1}I^T\mathcal{f} - \mathcal{U}\|_{\mathcal{E}}.$$

(A more stringent criterion may, however, be needed in such special cases as the removal of singularities described in Subsection 4.7.) We can meet this criterion simply by performing ℓ cycles of some version of FAC with convergence bound C provided $C^\ell \leq \gamma$. The importance of this observation is that, if $C \ll 1$, then the computational complexity of this FAC can then be predicted from the algorithm details and the complexities of the approximate solvers and grid transfers, assuming of course that the cycling scheme is fixed as opposed to accommodative.

4.3. \mathcal{G} and \mathcal{I} Construction. Construction of \mathcal{G} and \mathcal{I} are not entirely straightforward. In our experiments we have appealed to a finite element setting based either on constant or bilinear elements corresponding, respectively, to the staggered and aligned grids in Figures 1 and 2. Though the staggered grid case is very simple conceptually and computationally, convergence is expectedly slow and our theory does not directly imply $C \ll 1$. On the other hand, fast convergence for aligned grids with second degree elliptic boundary value problems is well established in both theory (Section 3) and numerical experiments (Section 5); but complexity is somewhat higher and some care is necessary in constructing \mathcal{G} and \mathcal{I} (e.g., it is inappropriate here to place fine grid points between adjacent points of G_I in Figure 2 by assuming their values are “free”).

Note that FAC is applicable to virtually any dimension of the underlying geometry. In fact, FAC can be very effective for three-dimensional models, since local phenomena can often be “captured” with relatively small grid patches. Note also that FAC need not be restricted to rectangular grid structures; it can be used when the local patches are in boundary fitted coordinates, for example.

4.4. Five-Point Formulae. Our prototype FAC scheme preserves the five-point structure of L (except, of course, at G_I), but the more accurate aligned grid case requires nine-point formulae. As a compromise between accuracy and complexity,

instead of solving the coarse and local grid equations exactly or using iterative schemes based on the nine-point stencil, we can develop approximate solvers based on judiciously chosen five-point formulae. For example, let L and \mathcal{L}_{FF} be based on the nine-point stencil discretization of Poisson's equation for the aligned grid case illustrated in Figure 2. In particular, if we start with L based on the nine-point stencil

$$L \sim -\frac{1}{3h^2} \begin{pmatrix} -1 & -1 & -1 \\ -1 & 8 & -1 \\ -1 & -1 & -1 \end{pmatrix}$$

and use linear interpolation, then we are led to the stencils

$$\mathcal{L}_{FF} \sim -\frac{1}{3h^2} \begin{pmatrix} -1 & -1 & -1 \\ -1 & 8 & -1 \\ -1 & -1 & -1 \end{pmatrix} \quad \text{and} \quad L \sim -\frac{1}{3h^2} \begin{pmatrix} -1 & -1 & -1 \\ -1 & 8 & -1 \\ -1 & -1 & -1 \end{pmatrix}.$$

(Note the scale for L . Other stencil entries for \mathcal{L} are depicted in Figure 6.) Consider the approximating five-point stencils

$$\hat{\mathcal{L}}_{FF} \sim -\frac{1}{h^2} \begin{pmatrix} -1 & -1 & -1 \\ -1 & 4 & -1 \\ -1 & -1 & -1 \end{pmatrix} \quad \text{and} \quad \hat{L} \sim -\frac{1}{h^2} \begin{pmatrix} -1 & -1 & -1 \\ -1 & 4 & -1 \\ -1 & -1 & -1 \end{pmatrix}.$$

Now the coarse grid step in FAC involves (approximately) solving an equation of the form

$$(4.4) \quad LU = f.$$

The point here is that we may replace this equation with the *defect correction* step

$$(4.5) \quad u \leftarrow u - \hat{L}^{-1}(Lu - f).$$

Similar comments hold for \mathcal{L}_{FF} . Note that the residual $Lu - f$ is computed using the nine-point formula, but the subsequent computation of the (approximate) inverse uses the five-point formula. Note also that Theorem 4 applies here directly, even when the equations for \hat{L} are solved only approximately.

4.5. Correction (CS) vs. Full Approximation (FAS) Scheme. MLAT [5] is developed in an FAS mode in the sense that the coarse grid maintains an approximation to the *solution*, not the *correction*. This no doubt offers certain advantages in that it facilitates error estimation, extrapolation, and grid adaption, so it can be useful to implement FAC in such a mode. However, the CS version described here has its own attributes, including conceptual simplicity (although this is quite subjective), some computational simplicity (especially for $\text{FAC}_{0,0}$ since the composite grid residuals are zero in \mathcal{G}_C and \mathcal{G}_F), and roundoff error advantages (the composite grid residual may be accumulated in double precision; carrying the full approximation on all levels would eliminate the advantage of doing this).

4.6. Conservation of Mass. FAC was developed as a response to some difficulties encountered in applying MLAT to the "five-spot" oil reservoir simulation problem described in Subsection 5.2. To be more specific, suppose L is defined so that the solution of the coarse grid equation exhibits some *conservation* property. For example, for cell-centered grids, L can be constructed by using the differences across cell edges to enforce (say) energy conservation within each cell. Such constructions provide an "exact" sense of (discrete) conservation so that exact local

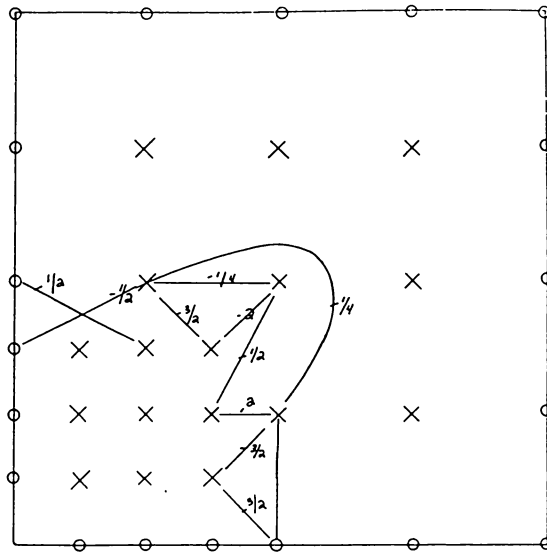


FIGURE 6
(Irregular Stencil Elements for \mathcal{L} -Aligned Grids)

The off-diagonal entries of $\mathcal{L}/3h^2$, with $h = 1/8$, corresponding to each type of irregular neighbor pair are depicted as connections between these pairs. All other types of connections not shown correspond to the regular 9-point stencil and have value -1 . The diagonal entries are such that the row sums of \mathcal{L} are zero (depending of course on the boundary conditions and how they are implemented). Note for example that the stencil at the corner point of \mathcal{G}_1 is

$$\frac{1}{3h^2} \times \begin{pmatrix} -1 & -1 & -1 \\ -1/4 & 17/2 & -1 \\ -1/2 & -2 & -1/2 \\ & -1/2 & -1/4 & -1 \end{pmatrix}.$$

conservation can in principle be achieved, even though the accuracy of the approximation itself is inhibited by discretization error. It may seem questionable to demand that conservation be achieved beyond the accuracy of truncation, but the practical point here is that application engineers often make such demands.

With MLAT, maintenance of conservation is not at all straightforward. If the measure of conservation is to assess the local strengths of the right-hand side (“sources” and “sinks”) on G , then the FAS transfer will generally be in conflict because it has the effect of significantly altering the original strengths (i.e., values of f). It is therefore generally true that coarse grid conservation is damaged, even though the accuracy of its approximation may be well below coarse grid discretization error. Attempts to correct this loss of conservation can be quite *ad hoc* and may inhibit convergence.

Maintaining conservation for FAC is not so problematic. The essential point here is that the measure should be taken via the composite grid equation (2.2). That is, if \mathcal{L} is constructed in conservation form, then the extent to which conservation is

attained on \mathcal{G} can be assessed via the residual $r = f - \mathcal{L}u$. In fact, Step 1 of FAC can be viewed as a mechanism for constructing artificial sources and sinks at points of G_I in order to force the coarse grid equations to approximate the conservation objective of \mathcal{G} . In any event, this objective is not in conflict with, but is precisely the same as, the goal of discretization, so conservation on \mathcal{G} can be achieved to any desired accuracy.

4.7. Removing Singularities. An interesting possibility is to use FAC for removing singularities. More precisely, we imagine that there are advantages to representing (a component of) the solution by a continuous (*explicit* or *implicit*) equation in some local region of the domain. The composite grid would then consist of a global discrete coarse grid G and a local continuous region \mathcal{G}_F ; the composite operator would then consist of the usual discretization on $\mathcal{G}_C \cup \mathcal{G}_I$ and a continuous equation in \mathcal{G}_F . The latter may consist of the differential equation itself, or perhaps a formal, explicit representation of a component of the solution based on local information together with the usual discretization to approximate the remaining component. We will report on this concept elsewhere.

4.8. Large-Scale Computation. The basic FAC premise is that it can increase accuracy without a commensurate increase in cost. Assuming that special local phenomena precipitate the need for higher local resolution, then FAC allows this resolution to remain local. Furthermore, it allows the computation to be restricted to uniform discretizations or, more precisely, to problems that have characteristics similar to the basic grid G . Both of these features make FAC attractive in a large-scale computational environment, the first by expanding capabilities beyond machine limitations and the second by taking maximum advantage of the host machine architecture. Concerning the latter, FAC lends itself to both vector and multiprocessor computation. In fact, FAC may be used in a multiprocessor machine as a vehicle for communicating between processors that have been assigned in correspondence to a partition of a fine *global* grid.

5. Numerical Results. We report here on numerical experiments with FAC applied to two types of problems, one a modification of our prototype with a re-entrant boundary and the other a singular, aligned grid discretization of the so-called five-spot problem in oil reservoir simulation. All experiments were performed on a Cyber 205.

5.1. Modified Prototype. Inspired by the model of a fire in a structure of rectangular rooms with adjoining windows and doors, we consider the simplified geometry depicted in Figure 7. On this staggered grid we discretize (2.1) as before, except that the boundary conditions on the external boundary are now Dirichlet. (The re-entrant boundary condition is Neumann.)

Table 1 depicts convergence rates for two sizes of problems and two types of fine grid solvers. The size is determined by the values of m and n (here, G and \mathcal{G}_F have $2m \times m$ and $2n \times n$ cells, respectively). In both cases, the door is $m/4$ cells wide. The solvers are all Gauss-Seidel with the number of sweeps being $NS = 1, 2, 3$ and 100 (the latter to simulate an exact solver.) The coarse grid is solved by 100 Gauss-Seidel sweeps in each case. The asymptotic rate C_R is determined as the apparent limit of the ratios of the composite grid *residuals* after subsequent cycles of




FIGURE 7
(Modified Prototype)

This is similar to Figure 1 except that the region now consists of two squares adjoined by a “door”, that is, the region is a rectangle with a re-entrant boundary almost bisecting it. The local fine grid \mathcal{G}_F has the same structure as G (except that the door is relatively larger) and is placed to increase resolution about the opening.

FAC. (This was checked against the energy errors which were known by construction.) Note that the “exact” version achieves a rate of about 1/4 and that the rate for one sweep Gauss-Seidel is about double that. Interestingly enough, there is almost nothing to gain by a few more sweeps, suggesting that the exact solver achieves its increased error reduction by eliminating smooth components.

5.2. *Five-Spot Problem.* Consider the so-called *five-spot* problem given by (2.1) with $\phi = \sigma^{SW} - \sigma^{NE}$, where σ^{SW} and σ^{NE} are Dirac delta functions at the respective southwest and northeast corners of Ω . This represents the discrete pressure equation in a subregion of an oil field (of constant permeability) with an injection and extraction well at opposing corners of Ω . This problem is similar to our prototype except that here we use aligned grids and construct $\mathcal{G}_F = \mathcal{G}_F^{SW} \cup \mathcal{G}_F^{NE}$ by placing a fine grid patch in the southwest *and* northeast corners of Ω . Thus, \mathcal{G}_F^{SW} is as depicted in Figure 2 and \mathcal{G}_F^{NE} is exactly the same but in the opposing corner. Let m (n^{SW} and n^{NE}) be the number of grid points of G (\mathcal{G}_F^{SW} and \mathcal{G}_F^{NE} , respectively) in each coordinate direction. In Figure 2, $m = 5$ and $n^{SW} = n^{NE} = 4$.

TABLE 1
(Staggered Grids)

FAC per cycle convergence factors for various values of the number (NS) of fine grid Gauss-Seidel relaxations.

M	N	CR			
		NS = 1	NS = 2	NS = 3	NS = 100
16	16	0.44	0.43	0.42	0.23
32	32	0.47	0.46	0.45	0.25

TABLE 2
(Aligned Grids)

Convergence factors scaled by 10^3 for successive cycles of FAC. All subsequent factors were 0.044.

Cycle Number																						
2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	
32	36	27	40	25	48	23	58	26	60	31	54	36	50	40	47	42	45	43	45	43	44	

In our experiments, we used the “nine-point” discretization of Figure 6 with ghost points and one-sided differencing at the boundary. Using direct elimination on both the coarse and fine levels, we restricted our tests to the case $n^{SW} = n^{NE} = 4$ but varied $m = 5, 9, 17, 33,$ and 65 . Table 2 depicts the ratios of composite grid residual norms on successive FAC cycles for $n = 65$. (Results for the other values of n were almost identical.) Note that the asymptotic convergence factor is 0.044 and that this is approximately achieved at the outset.

Acknowledgments. We are indebted to Achi Brandt, Richard Ewing, Vance Faber, John Gary, and Roland Sweet for many inspiring discussions on FAC and related topics.

Department of Mathematics
Colorado State University
Fort Collins, Colorado 80523

1. D. BAI & A. BRANDT, *Local Mesh Refinement Multilevel Techniques*, Research report, Dept. of Appl. Math., Weizmann Institute of Science, Rehovot, Israel, 1983.
2. R. E. BANK & C. DOUGLAS, “Sharp estimates for multigrid rates of convergence with general smoothing and acceleration,” *SIAM J. Numer. Anal.*, v. 4, 1985, pp. 617–633.
3. M. BERGER & T. JAMESON, *An Adaptive Multigrid Method for the Euler Equations*, Lecture Notes in Physics, Proc. 9th Internat. Conf. Numer. Methods in Fluid Dynamics, June, 1984, Saclay, France.
4. M. BERGER & J. OLIGER, “An adaptive mesh refinement for hyperbolic partial differential equations,” *J. Comput. Phys.*, v. 53, 1984, pp. 484–512.
5. A. BRANDT, “Multilevel adaptive solutions to boundary value problems,” *Math. Comp.*, v. 31, 1977, pp. 333–390.
6. F. CHATELIN & W. MIRANKER, “Acceleration by aggregation of successive approximation methods,” *Linear Algebra Appl.*, v. 43, 1982, pp. 17–47.
7. W. G. GROPP, “Local uniform mesh refinement for elliptic partial differential equations,” *SIAM J. Sci. Statist. Comput.*, to appear; also available as Res. report no. 278, Dept. of Comput. Sci., Yale University, New Haven, CT, 1983.
8. W. HACKBUSCH, “Local defect correction method and domain decomposition techniques,” in *Defect Correction Methods: Theory and Applications* (K. Böhmer and H. J. Stetter, eds.), *Computing Supplementum 5*, Springer-Verlag, Wien, 1984, pp. 89–113.
9. J. MANDEL, S. McCORMICK & J. RUGE, “An algebraic theory for multigrid methods for variational problems,” *SIAM J. Numer. Anal.* (To appear.)
10. S. McCORMICK, *A variational theory for multilevel adaptive techniques* (MLAT), Proc. Multigrid Conference, Bristol, Sept., 1983, *IMA J.*
11. S. McCORMICK, “Fast adaptive composite grid (FAC) methods,” in *Defect Correction Methods: Theory and Applications* (K. Böhmer and H. J. Stetter, eds.), *Computing Supplementum 5*, Springer-Verlag, Wien, 1984, pp. 115–121.
12. W. MIRANKER & V. PAN, “Methods of aggregation,” *Linear Algebra Appl.*, v. 29, 1980, pp. 231–257.