

International Journal of High Performance Computing Applications

<http://hpc.sagepub.com/>

Toward Efficient Implementation of Preconditioned Conjugate Gradient Methods On Vector Supercomputers

Rami Melhem and Dennis Gannon

International Journal of High Performance Computing Applications 1987 1: 70

DOI: 10.1177/109434208700100106

The online version of this article can be found at:

<http://hpc.sagepub.com/content/1/1/70>

Published by:



<http://www.sagepublications.com>

Additional services and information for *International Journal of High Performance Computing Applications* can be found at:

Email Alerts: <http://hpc.sagepub.com/cgi/alerts>

Subscriptions: <http://hpc.sagepub.com/subscriptions>

Reprints: <http://www.sagepub.com/journalsReprints.nav>

Permissions: <http://www.sagepub.com/journalsPermissions.nav>

Citations: <http://hpc.sagepub.com/content/1/1/70.refs.html>

TOWARD EFFICIENT IMPLEMENTATION OF PRECONDITIONED CONJUGATE GRADIENT METHODS ON VECTOR SUPERCOMPUTERS

Rami Melhem

THE UNIVERSITY OF PITTSBURGH
PITTSBURGH, PENNSYLVANIA 15260

Dennis Gannon

SUBJECT AREA EDITOR

Summary

We consider large, sparse linear systems that result from the discretization of partial differential equations on regular and irregular domains, and we focus on the application of the preconditioned conjugate gradient (PCCG) method to the solution of such systems. More specifically, our goal is the efficient implementation of the PCCG method on vector supercomputers. The contribution to the above goal is made by the introduction of a data structure that can be effectively manipulated on vector machines, the utilization of preconditioning matrices obtained by incomplete factorization with diagonal update sets, and the introduction of new numbering schemes for both regular and irregular grids.

The International Journal of Supercomputer Applications, Volume 1, No. 1, Spring 1987, pp. 70–98.
©1987 Massachusetts Institute of Technology.

Introduction

Consider the linear system of equations

$$Ax = b$$

where A is a large $n \times n$ symmetric, positive definite matrix. It is widely accepted that the conjugate gradient (CG) method may compete with direct methods for the solution of the equations only if the matrix A is suitably preconditioned. In other words, in order to speed up the convergence of the CG iteration, it is essential to find a preconditioning matrix $M \approx A$, and then apply the CG method to the solution of $M^{-1}Ax = M^{-1}b$. The resulting preconditioned CG method (PCCG) is described as follows:

ALG1: PCCG

Choose an initial guess x_0 ,

$$r_0 = b - Ax_0; h_0 = p_0 = M^{-1}r_0.$$

Repeat for $i = 0, \dots$ until convergence.

1. $\alpha_i = \frac{\langle r_i, h_i \rangle}{\langle Ap_i, p_i \rangle}$
2. $x_{i+1} = x_i + \alpha_i p_i; r_{i+1} = r_i - \alpha_i A p_i$
3. $h_{i+1} = M^{-1}r_{i+1}$
4. $\beta_i = \frac{\langle r_{i+1}, h_{i+1} \rangle}{\langle r_i, h_i \rangle}$
5. $p_{i+1} = h_{i+1} + \beta_i p_i$

where $\langle x, y \rangle$ denotes the inner product of x and y . If the preconditioning matrix is chosen such that $M = U^T D U$, for some upper triangular matrix U and some diagonal matrix D , then it is possible to reduce the linear system $Mh_{i+1} = r_{i+1}$ in step 3 of ALG1 into two triangular systems. The solution of these systems is the price that must be paid for speeding up the convergence of the basic CG method (for which $M = I$).

Unfortunately, this price may be high when the PCCG is implemented on vector supercomputers. More specifically, the solution of a triangular system is a recursive process, which may be very inefficient on vector and pipelined computers. In order to overcome this difficulty, alternative preconditioners have been considered (see, for example, Axelsson, 1984; Concus, Golub, and Meurant, 1985; Meurant, 1984; Saad 1985; and van der Vorst, 1982). Also, multicoloring techniques have been used to reorder the rows and columns of the matrix to minimize the data dependencies, and thus minimize the effect of recursion (see, for example, Kincaid, Oppe, and Young, 1984; Poole and Ortega, 1986; Saad and

Schultz, 1985; and Schreiber and Tang, 1982). However, experimental results show that multicolor orderings decrease the rate of convergence of the PCCG method (Kincaid et al., 1984; Poole and Ortega, 1986; and Schreiber and Tang, 1982). These results raise the following challenge, which was suggested by David Young and coworkers: "Can we find an iterative algorithm which is substantially better on supercomputers than the basic CG method?" (Young et al., 1985).

Numerical solution techniques of partial differential equations are major sources of large linear systems. In any of these systems, the coefficient matrix A is generated by the discretization of a partial differential equation on a finite grid G such that each row of A corresponds to a node in G . The specific discretization used defines a neighboring relation between the nodes of G . For example, if finite element analysis is used, then two nodes are neighbors if they belong to the same element. This neighboring relation determines the sparsity structure of A . More specifically, an element a_{ij} of A is nonzero only if nodes i and j are neighbors in G .

For rectangular domains covered by regular grids, the nonzero elements in A are confined to few diagonals. In this case, good preconditioning matrices may be obtained by the incomplete Cholesky factorization of A (Gustafsson, 1978; Manteuffel, 1980; Meijerink and van der Vorst, 1977)—namely, a Cholesky factorization in which only selected elements of A are modified during the factorization process. The positions of these selected elements are specified by an update set P , and the corresponding factorization is denoted by $IC(P)$. For example, if $P = P_a = \{(i,j); a_{ij} \neq 0\}$, the set of positions which contain nonzeros in A , then the corresponding Cholesky factorization, $IC(P_a)$, does not allow any fill-ins in the matrix A . The conjugate gradient method, which uses a preconditioned matrix obtained by an incomplete Cholesky factorization with update set P , is denoted by $IC(P)/PCCG$.

If, on the other hand, the matrix A is generated from irregular grids, then its nonzero elements are not confined to any regular pattern. In this case, a suitable preconditioner for A may be obtained by choosing M to be the SSOR matrix (Hageman and Young, 1981).

In order to obtain efficient implementations of the $IC(P)/PCCG$ methods on supercomputers, advances have to be made in three fronts: (1) suitable data structures which are efficiently manipulated on vector computers

"Numerical solution techniques of partial differential equations are major sources of large linear systems."

"We present a data structure that may be used efficiently for the manipulation of general sparse matrices on vector computers."

have to be chosen; (2) better preconditioning matrices have to be used; and (3) renumbering schemes have to be found that satisfy a balance between the advantages of increasing the recursion span in step 3 of ALG1 and the disadvantages of slowing down the convergence of the PCCG method.

We present a data structure that may be used efficiently for the manipulation of general sparse matrices on vector computers. Then we examine the class of preconditioning matrices obtained by incomplete Cholesky factorizations. By using the notion of update sets rather than fill-in sets, we show that the SSOR matrices are incomplete factorization matrices with an empty update set. Also, the zero extension factorizations *ICCG(0)* (Meijerink and van der Vorst, 1981) and *MICCG(0)* (Gustafsson, 1978) are special instances of an incomplete factorization $IC(P_d)$ in which the update set is given by $P_d = \{(i,i); i = 1, \dots, n\}$. This $IC(P_d)$ factorization is, in fact, general, and may be applied to matrices generated from irregular grids.

Column-wise multicolor schemes are introduced for rectangular grids. Numbering schemes for pierced rectangular grids are described, and their effect on the vectorization potential of the PCCG method is analyzed. The multicolor numbering schemes are generalized and applied to irregular grids, and finally, some experimental results obtained on the CRAY X-MP vector computer are presented.

Results and Discussion

1. THE STRIPE DATA STRUCTURE

In Melhem (1986a), stripe data structures are introduced as a means to include all the nonzero elements of a matrix in a structure suitable for parallel processing on linear computational arrays. We will demonstrate that this same structure can be used effectively to process sparse matrices on pipelined and vector computers.

Very briefly, a stripe in an $n \times n$ matrix A is a set of positions $S = \{(i, \sigma(i)); i \in I \subseteq I_n\}$, where $I_n = \{1, \dots, n\}$ and σ is a strictly increasing function—that is, if $(i, \sigma(i))$ and $(j, \sigma(j))$ are in S , then

$$i < j \rightarrow \sigma(i) < \sigma(j). \quad (1)$$

Note that a stripe need not include a position for every row of the matrix. If, however, it does contain a position for every row, then it is called a complete stripe. Two

stripes, $S_1 = \{(i, \sigma_1(i))\}$ and $S_2 = \{(i, \sigma_2(i))\}$, are ordered by $S_1 < S_2$ if, for any $(i, \sigma_1(i)) \in S_1$ and $(j, \sigma_2(j)) \in S_2$, we have

$$i \leq j \rightarrow \sigma_1(i) < \sigma_2(j). \quad (2)$$

With this, a stripe structure of A is defined as a set of stripes, $\Sigma_A = \{S_{-\pi_1}, \dots, S_{\pi_2}\}$, such that (1) $S_{-\pi_1} < \dots < S_{\pi_2}$ and (2) the position (i, j) of any nonzero element a_{ij} of A is in some stripe S_k , $-\pi_1 \leq k \leq \pi_2$. The number of stripes $\pi = \pi_1 + \pi_2 + 1$ is called the stripe count. Algorithms for the construction of stripe structures for sparse matrices, in general, and for finite element matrices, in particular, are given in Melhem (1986a) and Melhem (1986b), respectively.

The elements of a striped matrix A may be stored in an $n \times \pi$ array, AE , and the position of each element may be specified in another $n \times \pi$ array, AP . More specifically, we may set

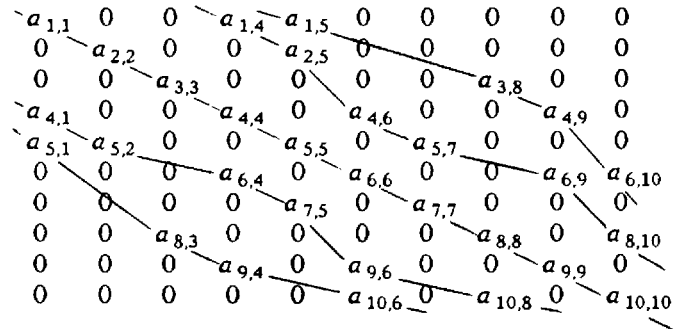
$$AE(i, k + offset) = \begin{cases} a_{i, \sigma_k(i)} & \text{if } (i, \sigma_k(i)) \in S_k \\ 0 & \text{otherwise} \end{cases} \quad (3a)$$

$$AP(i, k + offset) = \begin{cases} \sigma_k(i) & \text{if } (i, \sigma_k(i)) \in S_k \\ n + 1 & \text{otherwise} \end{cases} \quad (3b)$$

Here, $offset = \pi_1 + 1$, and $AP(i, k + offset) = n + 1$ is used to indicate that S_k does not contain a position in row i . In Figure 1, we give an example illustrating the concept of stripe data structures. Clearly, this stripe storage scheme is a generalization of the diagonal scheme used in Young et al. (1985) and Madsen, Rodrigue, and Karush (1976). It is also a more restricted form of the data structure used in Rice and Boisvert (1984), in which neither the strictly increasing property (1), nor the nonintersecting property (2) holds. The former property is crucial for the efficient manipulation of symmetric matrices on vector computers, while the latter is useful only for the manipulation of sparse matrices on linear computational networks (Melhem, 1986a).

If A is a symmetric matrix with nonzero diagonal elements, then it is possible to construct a stripe structure $\Sigma_A = \{S_{-\pi_s}, \dots, S_{\pi_s}\}$, such that the lower stripes, S_{-k} , $k = 1, \dots, \pi_s$, are mirror images of the upper stripes, S_k , $k = 1, \dots, \pi_s$. In this case, only the elements in S_0, \dots, S_{π_s} need to be stored in AE and AP , with $offset = 1$ in Eq(3a) and Eq(3b).

Given a symmetric matrix A in stripe storage form, the computation of the product vector $\mathbf{y} = \mathbf{A}\mathbf{x}$, for any vector \mathbf{x} , may be accomplished by the multiplication of each element $a_{i, \sigma_k(i)}$ in an upper stripe S_k with both x_i and



(a) The Matrix

0	0	$a_{1,1}$	$a_{1,4}$	$a_{1,5}$	11	11	1	4	5
0	0	$a_{2,2}$	$a_{2,5}$	0	11	11	2	5	11
0	0	$a_{3,3}$	0	$a_{3,8}$	11	11	3	11	8
0	$a_{4,1}$	$a_{4,4}$	$a_{4,6}$	$a_{4,9}$	11	1	4	6	9
$a_{5,1}$	$a_{5,2}$	$a_{5,5}$	$a_{5,7}$	0	1	2	5	7	11
0	$a_{6,4}$	$a_{6,6}$	$a_{6,9}$	$a_{6,10}$	11	4	6	9	10
0	$a_{7,5}$	$a_{7,7}$	0	0	11	5	7	11	11
$a_{8,3}$	0	$a_{8,8}$	$a_{8,10}$	0	3	11	8	10	11
$a_{9,4}$	$a_{9,6}$	$a_{9,9}$	0	0	4	6	9	11	11
$a_{10,6}$	$a_{10,8}$	$a_{10,10}$	0	0	6	8	10	11	11

(b) AE

(c) AP

Fig. 1 A stripe structure of a sparse matrix

“Besides matrix/vector multiplication, the stripe structure may be used in the solution of triangular linear systems.”

$x_{\sigma_k(i)}$, and then the accumulation of the results in $y_{\sigma_k(i)}$ and y_i , respectively. Using a pseudo-CRAY-FORTRAN language, the matrix/vector multiplication algorithm may be written as follows:

ALG2: Symmetric Matrix/Vector Multiplication

1. $x(n + 1) = 0$
2. $y(i) = AE(i, offset) * x(i)$ $i = 1, \dots, n$
3. DO $k = 1, \dots, \pi_s$
 - 3.1 call gather($n, w, x, AP(1, k + offset)$)
 - 3.2 $y(i) = y(i) + AE(i, k + offset) * w(i)$ $i = 1, \dots, n$
 - 3.3 $z(i) = AE(i, k + offset) * x(i)$ $i = 1, \dots, n$
 - 3.4 $w(i) = 0$ $i = 1, \dots, n$
 - 3.5 call scatter($n, w, AP(1, k + offset), z$)
 - 3.6 $y(i) = y(i) + w(i)$ $i = 1, \dots, n$

Step 2 in ALG2 accumulates in y the contribution of S_0 , steps 3.1 and 3.2 accumulate the contribution of S_k , and steps 3.3–3.6 accumulate the contribution of S_{-k} . The strictly increasing property (1) is essential for the correctness of the above algorithm because it ensures that if $i \neq j$, then $\sigma_k(i) \neq \sigma_k(j)$. If this is not satisfied, then $z(\sigma_k(i))$ and $z(\sigma_k(j))$ will be scattered to the same location in w , and thus only one of them will be accumulated in y .

Besides matrix/vector multiplication, the stripe structure may be used in the solution of triangular linear systems. For example, consider the upper triangular system $(D + U)x = b$, where D is a diagonal matrix, and U is a strictly upper triangular matrix with a stripe structure $\Sigma_U = \{S_1, \dots, S_{\pi_2}\}$. The solution of this system is a recursive process in which the calculation of x_i proceeds in the order $i = n, \dots, 1$, with x_i depending on the previously calculated values $x_j, j > i$. The minimum recursion span—that is, the minimum integer d such that x_i does not depend on x_{i+1}, \dots, x_{i+d} —is equal to the upper zero stretch of the matrix defined as follows:

DEFINITION. The upper zero stretch of a matrix A is the largest integer Δ_2 such that $a_{ij} = 0$ for $i = 1, \dots, n$ and $i < j < i + \Delta_2$. In other words, Δ_2 is the size of the maximum band above (and including) the main diagonal which contains zero elements.

For the upper triangular matrix U , the upper zero stretch is found, by properties (1) and (2), to be

$$\Delta_2 = \min\{\sigma_1(i) - i; (i, \sigma_1(i)) \in S_1\}.$$

Given Σ_U and Δ_2 , the triangular system $(D + U)x = b$ may be solved using operations on vectors of length

Δ_2 . Similarly, given a strictly lower triangular matrix L and a stripe structure $\Sigma_L = \{S_{-\pi_s}, \dots, S_{-1}\}$, a lower triangular system of the form $(L + D)\mathbf{x} = \mathbf{b}$ may be solved using operations on vectors of length Δ_1 , where Δ_1 is the lower zero stretch of L given by

$$\Delta_1 = \min\{i - \sigma_{-1}(i); (i, \sigma_{-1}(i)) \in S_{-1}\}.$$

When the IC/PCCG method is applied to symmetric matrices, it is necessary to solve in each iteration two triangular systems of the forms $(L + D)\mathbf{x} = \mathbf{b}$ and $(D + U)\mathbf{y} = \mathbf{x}$, with $L = U^T$. If $\Sigma_U = \{S_1, \dots, S_{\pi_s}\}$ is a stripe structure for U , then it is possible to construct a stripe structure $\Sigma_L = \{S_{-\pi_s}, \dots, S_{-1}\}$ for L such that each stripe S_{-k} , $1 \leq k \leq \pi_s$ is the mirror image of S_k . In this case, only U needs to be stored. Let UE and UP be the arrays used to store U in a way analogous to Eq(3a) and Eq(3b), with *offset* = 0, and let Δ_s be the lower (upper) zero stretch of L (U).

The solution of $(L + D)\mathbf{x} = \mathbf{b}$ is affected by the absence of an explicit storage for the elements of L . More specifically, consider the usual forward substitution algorithm. In this algorithm, the solution proceeds such that after the computation of $x_1, \dots, x_{r\Delta_s}$, for some r , the next Δ_s elements of \mathbf{x} are computed by first computing the Δ_s components z_i , $i = r\Delta_s + 1, \dots, (r + 1)\Delta_s$, of the vector $\mathbf{z} = \mathbf{b} - L\mathbf{x}$, and then dividing each z_i by d_i , the i th component of D . However, the above scheme is not efficient because it is difficult to access, in UE , the elements of Σ_L which belong to rows $r\Delta_s + 1, \dots, (r + 1)\Delta_s$. To overcome this difficulty, we use a column sweep algorithm (*ji* algorithm according to Dongarra, Gustavson, and Karp, 1984). More specifically, we let L_j be the j th column of L and we compute the vector $\mathbf{z} = \mathbf{b} - L\mathbf{x} = \mathbf{b} - \sum_{l=1}^n L_l x_l$ progressively by accumulating $\sum_{l=r\Delta_s+1}^{(r+1)\Delta_s} L_l x_l$ into \mathbf{z} as soon as the appropriate values of x are computed. This requires the access of the elements of Σ_L which belong to columns $r\Delta_s + 1, \dots, (r + 1)\Delta_s$; that is, the elements of Σ_U which belong to rows $r\Delta_s + 1, \dots, (r + 1)\Delta_s$. These elements may be easily accessed in UE .

More precisely, the solution of $(L + D)\mathbf{x} = \mathbf{b}$ may be described by the following algorithm, in which we assume, for simplicity, that $n = \Delta_s m$ for some integer m :

ALG3: Forward Substitution

1. $z(i) = b(i)$ $i = 1, \dots, n$
2. $x(i) = z(i)/d(i)$ $i = 1, \dots, \Delta_s$
3. FOR $r = 0, \dots, m - 1$ DO



3.0 Let $i_s = r\Delta_s + 1$ and $i_e = (r + 1)\Delta_s$
 3.1 FOR $k = 1, \dots, \pi_s$ DO
 3.1.1 $w(i) = UE(i,k) * x(i)$ $i = i_s, \dots, i_e$
 3.1.2 $u(i) = 0$ $i = i_s, \dots, n$
 3.1.3 call scatter($\Delta_s, u, UP(i_s,k), w(i_s)$)
 3.1.4 $z(i) = z(i) - u(i)$ $i = i_s, \dots, n$
 3.2 $x(i) = z(i)/d(i)$ $i = i_e + 1, \dots, i_e + \Delta_s$

Note that in 3.1.3, the Δ_s elements of w may be scattered anywhere in u . For this reason, the summation in 3.1.4 runs up to $i = n$. An alternative way for writing the inner loop in step 3.1 is as follows:

3.1.1 call gather($\Delta_s, u(i_s), z(i_s), UP(i_s, k)$)
 3.1.2 $u(i) = u(i) - UE(i, k) * x(i)$ $i = i_s, \dots, i_e$
 3.1.3 call scatter($\Delta_s, z(i_s), UP(i_s, k), u(i_s)$)

Clearly, this second alternative reduces the number of operations. Moreover, the subtraction and the multiplication are performed in the same step, which allows for the chaining of the two operations. The choice between the two alternatives should depend on the relative execution time of the different vector instructions on the specific computer used. For example, actual measurements on the CRAY X-MP show that the slow down due to the additional call to "gather" in the second alternative more than offsets the gain obtained from the chaining of the operations and the lower operation count.

2. INCOMPLETE FACTORIZATIONS WITH DIAGONAL UPDATE SETS

Given a symmetric positive definite matrix A of order n , the incomplete factorization of A is a splitting of the form

$$A = M + R \quad (4a)$$

where

$$M = (I + U)^T D (I + U) \quad (4b)$$

and I is the identity matrix, D is a diagonal matrix, and U is a strictly upper triangular matrix. The matrices D and U are determined from A by applying a factorization procedure in which only specific elements of A may be updated during the factorization. More specifically, if $P = \{(i,j); i,j \in [1,n]\}$ is any subset of positions in A , then the corresponding incomplete factors of A are obtained as follows:

ALG4: Incomplete Cholesky Factorization

FOR $i = 1, \dots, n$ DO

$$d_{i,i} = a_{i,i}$$


```

FOR  $j = i + 1, \dots, n$  DO
   $u_{ij} = a_{ij}/d_{ii}$ 
FOR  $k = i + 1, \dots, n$  DO
  FOR  $j = k, \dots, n$  DO
    IF  $(k,j) \in P$  THEN  $a_{kj} = a_{kj} - u_{ik} * u_{ij} * d_{ii}$ 

```

We call P the update set, and we call the corresponding factorization Eq (4b) and splitting Eq (4a) the $IC(P)$ factorization and splitting of A . A possible choice of P is the set $P_a = \{(i,j): a_{ij} \neq 0\}$ of positions that contain nonzero elements in A .

Manteuffel (1980) suggested a shifted incomplete Cholesky splitting (SIC) in which the off-diagonal elements of A are scaled by some factor ω before the factorization. In order to be more specific, let F and Λ be the upper triangular and the diagonal parts of A , respectively. That is,

$$A = F^T + \Lambda + F. \quad (5)$$

With this, consider the shifted matrix

$$\bar{A} = \Lambda + \omega(F^T + F)$$

and let \bar{U} and \bar{D} be the factors produced by the application of ALG4 to \bar{A} . The $SIC(P,\omega)$ splitting of A , corresponding to the update set P and the shift factor ω , is then given by

$$A = (I + \bar{U})^T \bar{D} (I + \bar{U}) + \bar{R}.$$

Note that $SIC(P,\omega)$ reduces to $IC(P)$ for $\omega = 1$.

Factorizations with update sets larger than or equal to P_a , have been considered in the literature (for example, Meijerink and van der Vorst, 1981) for matrices generated from rectangular grids. For these matrices, ALG4 may be greatly simplified. However, for matrices generated from irregular grids, the choice of $P \supseteq P_a$ does not seem useful because, even though the execution of ALG4 is costly, only few of the terms $u_{ik} * u_{ij}$ in the inner loop are nonzeros, thus causing only few updates in the entries of A . In this case, the SSOR splitting of A may be used.

2.1 SSOR SPLITTING AS INCOMPLETE FACTORIZATION

If the update set P is taken to be the empty set \emptyset , then it is easy to see that the factors U_\emptyset and D_\emptyset resulting from ALG4 are given by



Fig. 2 The spectrum of incomplete Cholesky factorizations

$$\begin{aligned} D_{\emptyset} &= \Lambda. \\ U_{\emptyset} &= \Lambda^{-1}F. \end{aligned}$$

Also, if the off-diagonal elements of A are shifted by ω before factorization, then the resulting factors are given by

$$\bar{D}_{\emptyset} = \Lambda. \quad (6a)$$

$$\bar{U}_{\emptyset} = \omega\Lambda^{-1}F. \quad (6b)$$

Hence, the $SIC(\emptyset, \omega)$ preconditioning matrix is

$$\begin{aligned} \bar{M}_{\emptyset} &= (I + \omega\Lambda^{-1}F)^T \Lambda (I + \omega\Lambda^{-1}F) \\ &= (\Lambda + \omega F)^T \Lambda^{-1} (\Lambda + \omega F), \end{aligned} \quad (7)$$

which is equal to the SSOR preconditioning matrix, up to a scalar factor of $\omega(2 - \omega)$. This scalar factor does not have any effect from the point of view of the PCCG method.

2.2 DIAGONAL UPDATE CHOLESKY FACTORIZATION

Consider nonshifted incomplete factorization. We may envision a linear scale on which the preconditioners obtained from ALG4 are laid according to the size of the update set P (see Fig. 2). Clearly, the exact factorization $IC(P_e)$, corresponding to $P_e = \{(i,j); i,j = 1, \dots, n\}$, and the SSOR splitting $IC(P_{\emptyset})$ are laid at the two ends of the scale, with $IC(P_a)$ somewhere in between. The results of Manteuffel (1980) for M matrices show that if $P_1 \subseteq P_2$, then the $IC(P_1)/PCCG$ may not outperform $IC(P_2)/PCCG$. In other words, the performance of the $IC/PCCG$ improves when we move to the right on the scale of Figure 2.

The diagonal update factorization is the factorization obtained from ALG4 with the update set given by $P_d = \{(i,i); i = 1, \dots, n\}$. That is, only the diagonals of A are updated in ALG4. If we shift the matrix A by ω before applying ALG4, then it is easy to see that the factors \bar{D}_d and \bar{U}_d resulting from the factorization are given by

$$\begin{aligned} \bar{D}_d &= \bar{\Lambda} \\ \bar{U}_d &= \omega(\bar{\Lambda})^{-1}F \end{aligned}$$

where, by ALG4, the elements $\bar{\lambda}_i$, $i = 1, \dots, n$ of the diagonal matrix $\bar{\Lambda}$ are given by

$$\bar{\lambda}_i = a_{ii} - \sum_{j=1}^{i-1} (u_{ji})^2 \bar{\lambda}_j.$$

The equations that express the solution of upper triangular linear systems are similar to the above equations, with the division replaced by a multiplication. As-

suming that the execution of a vector divide is at most three times slower than a vector multiply (on the CRAY the ratio is found to be 2.5 for vectors of length 6,400), then it may be shown that the cost of computing $\bar{\Lambda}$ is at most equal to the cost of one PCCG iteration. Given $\bar{\Lambda}$, the $SIC(P_d, \omega)$ preconditioning matrix is

$$\bar{M}_d = (I + \omega \bar{\Lambda}^{-1} F)^T \bar{\Lambda} (I + \omega \bar{\Lambda} F). \quad (8)$$

In addition to being easy to compute, both SSOR and $SIC(P_d, \omega)$ preconditioners have an advantage over factorizations with update sets larger than P_d . Namely, they may lead to reduced computational work in each PCCG iteration. More specifically, each iteration in ALG1 requires, in addition to $O(n)$ operations, both the matrix/vector multiplication A_p and the solution of two triangular systems in step 3. However, for SSOR and $SIC(P_d, \omega)$ preconditioners, ALG1 may be rewritten such that it requires only the solution of two triangular systems, thus reducing the work per iteration by a factor of almost 2. For example, if the SSOR matrix of Eq(7) is used in ALG1, and the following substitution is made

$$H = (\Lambda + \omega F)^{-T} A (\Lambda + \omega F)^{-1}, \quad (9)$$

then it is possible to rewrite the algorithm such that it only involves, besides $O(n)$ operations, the multiplication of a vector p by H (see Eisenstat, 1981, for details). From Eq(9) and Eq(5), this multiplication may be computed as follows:

$$H p = \omega^2 \left(v + \left(\frac{1}{\omega} \Lambda + F \right)^{-T} \left(p + \frac{\omega - 2}{\omega} \Lambda v \right) \right) \quad (10)$$

where $v = [(1/\omega)\Lambda + F]^{-1} p$. That is, the multiplication $H p$ may be performed by solving two triangular systems.

The same argument applies to the $SIC(P_d, \omega)$ matrix Eq(8), which has the same form as Eq(7), except that Λ is replaced by $\bar{\Lambda}$. In this case, the multiplication $H p$ may be obtained from Eq(10) with Λ replaced by $\bar{\Lambda}$.

Given that $P_d \supseteq \emptyset$, then according to Manteuffel (1980), $IC(P_d)/PCCG$ is expected to outperform $SSOR/PCCG$ for M matrices. The examples considered in section 4 suggest that, in general, $IC(P_d)/PCCG$ is at least as efficient as the $SSOR/PCCG$.

3. MULTICOLOR NUMBERING SCHEMES

The major difficulty in the vectorization of the IC/PCCG method concerns the recursive nature of the solution of triangular linear systems. In order to overcome this difficulty, many multicolor numbering schemes have been

"The major difficulty in the vectorization of the IC/PCCG method concerns the recursive nature of the solution of triangular linear systems."

suggested (Kincaid et al., 1984; Poole and Ortega, 1986; Saad and Schultz, 1985) for rectangular grids. Their goal is to obtain a coefficient matrix A which may be partitioned into submatrices A_{ij} , $i, j = 1, \dots, p$, for some p , such that A_{ii} , $i = 1, \dots, p$, are diagonal matrices. This goal is achieved by applying an algorithm with the following form:

ALG5: A Global Multicolor Numbering Scheme

1. Assign to each node in the grid a color from a set of p colors such that neighboring nodes have different colors. Choose p to be as small as possible.
2. For each color $c = 1, \dots, p$, number all the nodes which have color c in a column-wise sequential order.

The above numbering scheme is global in the sense that all the nodes which have the same color are numbered consecutively. Although global numbering schemes produce matrices with large zero stretch (approximately n/p , where n is the number of nodes in the grid), experience shows that these schemes worsen the condition of the matrix (Kincaid et al., 1984; Poole and Ortega, 1986) thus slowing down the convergence of the IC/PCCG method. In this section, we suggest a numbering scheme that compromises between the zero stretch of the matrix and the convergence properties of the PCCG. We first introduce the scheme for rectangular grids.

3.1 COLUMN-WISE MULTICOLOR NUMBERING OF RECTANGULAR GRIDS

Let T be an upper (or a lower) triangular matrix given in striped form. It was shown in section 1 that the larger the zero stretch of T , the more efficient is the solution of $Tx = b$ on vector computers. However, the advantage of having large zero stretches when the IC/PCCG is applied to the solution of linear systems should be put in its proper perspective. More specifically, it is important to note the following:

1. The zero stretch does not affect the matrix/vector multiplication, the inner products computation, or the vector addition operations in ALG1.
2. The solution of a triangular linear system includes data movement operations (e.g., scatter and gather in ALG3). These operations are slower than vector arithmetic operations on existing vector computers,

hence they dominate the execution time of the solution.

3. The advantage of having long vectors is relatively limited in data movement operations compared with arithmetic operations. More specifically, the execution time of a vector arithmetic operation on a pipeline computer is usually specified by $\tau_s + \delta\tau_p$, where τ_s is a vector setup time, τ_p is pipe unit time, and δ is the vector length. Usually, τ_s is much larger than τ_p , which makes it very advantageous to use long vectors. For example, a multiply/add vector operation on the CRAY X-MP consumes 70 and 170 μ -seconds for $\delta = 640$ and 6,400, respectively. On the other hand, using long vectors in data movement operations is less advantageous because the pipe unit time for vector data movement operations is relatively large and depends on the distribution of data in memory. For example, a specific gather operation on the CRAY X-MP consumes 40 and 370 μ -seconds for $\delta = 640$ and 6,400, respectively. The corresponding times for a scatter operation are 90 and 410 μ -seconds, respectively.

Given the above facts, our main goal in the vectorization of the *IC/PCCG* method should be to increase the zero stretch in the matrices to a point that prevents the recursive solution of triangular systems from dominating the entire process. However, once the execution time of ALG1 is not dominated by the time for step 3, any further increase of the zero stretches will have a relatively limited effect on the execution time of the *IC/PCCG* iteration. For instance, when step 3 consumes only one-fourth of the execution time of ALG1, the doubling of the speed of step 3 speeds up the entire algorithm by less than 1.15.

The above discussion suggests that a numbering scheme which produces "reasonably conditioned" matrices with "reasonable" zero stretches may be better than a numbering scheme which produces "ill conditioned" matrices with very large zero stretches. The column-wise multicoloring scheme belongs to the former class. It is described, for an $r \times m$ rectangular grid, by the following algorithm:

ALG6: A Column-wise Multicolor Numbering Scheme

1. Assign to each node in the grid a color from a set of q colors such that neighboring nodes *that are in the same column* have different colors. Choose q to be as small as possible.

"The main goal in the vectorization of the IC/PCCG method should be to increase the zero stretch in the matrices to a point that prevents the recursive solution of triangular systems from dominating the entire process."

Table 1
Minimal Number of Colors in Multicoloring Schemes

	Global Schemes (p)	Column-wise Schemes (q)
FD_5	2	2
FE_3	3	2
FE_4	4	2
FE_6	6	3
FE_9	9	3

FD_5 denotes the 5-point finite difference discretization; and FE_3 , FE_4 , FE_6 , and FE_9 denote finite element discretizations with 3-node triangles, 4-node rectangles, 6-node triangles, and 9-node rectangles, respectively.

2. FOR $j = 1, \dots, m$ DO
 FOR each color $c = 1, \dots, q$, number, sequentially, the nodes in column j which have color c .

In other words, ALG6 uses a multicolor technique to number all the nodes in column j before numbering the nodes in column $j + 1$. It may be shown that ALG6 produces matrices with zero stretches equal to r/q , while ALG5 produces matrices with zero stretches rm/p . The minimal values of q and p for different discretization stencils are given in Table 1.

3.2 NUMBERING PIERCED RECTANGULAR GRIDS

Pierced rectangular grids (see Fig. 3) are defined in Melhem (1986b) to be rectangular grids from which some subrectangles are removed. This type of grid is useful because any irregular domain may be covered by a pierced rectangular grid, or by a grid that is isomorphic to a pierced rectangular grid.

Let Q be a rectangular grid which contains n_Q nodes numbered by the integers $1, \dots, n_Q$ according to some numbering scheme, and let Ω be a grid which is obtained by piercing Q and which contains n_Ω nodes. A renumbering of the nodes that are in Ω may then be defined by a function ν that assigns to each node i , which is in both Q and Ω , a unique number $\nu(i)$, $1 \leq \nu(i) \leq n_\Omega$. This renumbering is said to be deduced from the numbering of Q if ν is derived as follows:

ALG7: Numbering Pierced Rectangular Grids

- ```

 $l = 0$
FOR $i = 1, \dots, n_Q$ DO
 IF node i is in Ω THEN $\{l = l + 1; \nu(i) = l\}$
 ELSE $\{\nu(i) = \uparrow(\text{undefined})\}$.
```

The upper (or lower) zero stretch of the symmetric matrix generated from the pierced rectangular grid  $\Omega$  is then given by

$$\Delta = \min\{|\nu(i) - \nu(j)|; i \text{ and } j \text{ are neighboring nodes in } \Omega\}. \quad (11)$$

The specific value of  $\Delta$  in Eq(11) depends on the shape of  $\Omega$  and on the numbering scheme. However, because of the regularity of global and column-wise multicolor numberings, it is possible, in both schemes, to establish upper and lower bounds for  $\Delta$ . In the remainder of this section, we will find such bounds for the  $FE_4$  and  $FE_9$  discretizations. Similar bounds may be de-

rived for other discretizations. We start by deriving bounds on  $\Delta$  for column-wise multicolor numberings.

**THEOREM 1.** *If either  $FE_4$  or  $FE_9$  discretization is used, and a column-wise multicoloring scheme is applied with  $q = 2$  or  $3$ , respectively, then  $(ce_{\min} - 1)/2 \leq \Delta \leq ce_{\min}$  where  $ce_{\min}$  is the minimum number of elements in any column of elements in  $\Omega$ .*

*Proof.* See the Appendix.

For global coloring schemes, the bounds on  $\Delta$  may be given in terms of the total number of nodes  $n_\Omega$  in  $\Omega$ . For this, we first estimate the number of nodes in  $\Omega$  which are colored by any particular color.

**LEMMA 1.** *Let the  $n_\Omega$  nodes in a pierced rectangular domain  $\Omega$  be numbered by first using a global multicolor scheme to number the smallest rectangular domain  $Q$  which encloses  $\Omega$ , and then deducing the node numbers in  $\Omega$  using ALG7. If  $FE_4$  or  $FE_9$  discretizations are used with  $p = 4$  and  $9$ , respectively, then for any particular  $l$ ,  $1 \leq l \leq p$ , the number,  $n_\Omega^l$ , of nodes given the color  $l$  is bound by*

$$\left[ \frac{1}{2\sqrt{p} - 1} \right]^2 n_\Omega \leq n_\Omega^l \leq \left[ \frac{2}{\sqrt{p} + 1} \right]^2 n_\Omega. \quad (12)$$

*Proof.* See the Appendix.

In Figure 4, we color some specific rectangular grids to show that the bounds in the Lemma are tight (the numbers in Fig. 4 refer to colors). More specifically, for  $FE_4$ , the grid in Figure 4a gives  $n_\Omega^1 = (1/9)n_\Omega$  and  $n_\Omega^4 = (4/9)n_\Omega$ . Similarly, for  $FE_9$ ,  $n_\Omega^9 = (1/25)n_\Omega$  in Figure 4b, and  $n_\Omega^1 = (1/4)n_\Omega$  in Figure 4c. Clearly, the grids in Figure 4 are rectangular grids, which are special cases of pierced rectangular grids.

The following theorem uses Lemma 1 to bound the zero stretch of the matrix resulting from the finite element discretization.

**THEOREM 2.** *Given the hypothesis of Lemma 1, the zero stretch of the matrix resulting from the finite element discretization is bounded by*

$$\left[ \frac{1}{2\sqrt{p} - 1} \right]^2 n_\Omega - 1 \leq \Delta \leq \left[ \frac{2}{\sqrt{p} + 1} \right]^2 [n_\Omega + p]. \quad (13)$$

*Proof.* See the Appendix.

Given a pierced rectangular grid  $\Omega$ , the above bounds may be used for the a priori estimation of the zero stretch which result from specific numbering schemes. The experiments of section 4 show that the convergence rate of the PCCG method is inversely re-

|   |    |    |    |    |    |    |    |    |
|---|----|----|----|----|----|----|----|----|
| 5 | 14 | 22 | 30 | 39 | 46 | 52 | 58 | 64 |
| 9 | 18 | 25 | 34 | 43 | 49 | 55 | 61 | 67 |
| 4 | 13 | 21 | 29 | 38 | 45 | 51 | 57 | 63 |
| 8 | 17 | 24 | 33 | 42 | 48 | 54 | 60 | 66 |
| 3 | 12 |    | 28 | 37 | 44 | 50 | 56 | 62 |
| 7 | 16 |    | 32 | 41 | 47 | 53 | 59 | 65 |
| 2 | 11 | 20 | 27 | 36 |    |    |    |    |
| 6 | 15 | 23 | 31 | 40 |    |    |    |    |
| 1 | 10 | 19 | 26 | 35 |    |    |    |    |

**Fig. 3** A pierced rectangular grid numbered with a two-color, column-wise scheme

|   |   |   |  |  |
|---|---|---|--|--|
| 1 | 3 | 1 |  |  |
| 2 | 4 | 2 |  |  |
| 1 | 3 | 1 |  |  |

|   |   |   |   |   |
|---|---|---|---|---|
| 2 | 5 | 8 | 2 | 5 |
| 1 | 4 | 7 | 1 | 4 |
| 3 | 6 | 9 | 3 | 6 |
| 2 | 5 | 8 | 2 | 5 |
| 1 | 4 | 7 | 1 | 4 |

|   |   |   |   |
|---|---|---|---|
| 1 | 4 | 7 | 1 |
| 3 | 6 | 9 | 3 |
| 2 | 5 | 8 | 2 |
| 1 | 4 | 7 | 1 |

**Fig. 4** Tightening the bounds of Lemma 1

*"In finite element analysis, it is often useful to change the density of the discretization points among the different regions of the domain. This is especially advantageous if a rough estimate of the solution is known or if an adaptive solution technique is applied."*

lated to the zero stretch,  $\Delta$ , of the matrix, and that beyond a specific value of  $\Delta$ , the reduction in the execution time of each PCCG iteration does not offset the increase in the number of iterations. In this context, a priori estimates of  $\Delta$  may help in choosing the best numbering scheme for  $\Omega$ .

### 3.3 A MULTICOLOR NUMBERING FOR IRREGULAR GRIDS

In finite element analysis, it is often useful to change the density of the discretization points among the different regions of the domain. This is especially advantageous if a rough estimate of the solution is known or if an adaptive solution technique is applied. In both cases, it is preferable to increase the density of the discretization points in the regions where the solution is known to (or is found to) change rapidly. This desirable variation in density may not, usually, be accomplished without the use of irregular grids.

All previous schemes for numbering irregular grids were aimed at reducing the band-width and the profile of the resulting matrices. This is particularly useful if direct methods are used for the solution of the linear systems resulting from the grids. However, as we explained in section 1, if iterative solution methods are applied and vector computers are used, then the efficiency of the solution process is determined by the zero stretch of the matrix. In this section, we describe a multicoloring technique that aims at increasing the zero stretch of the matrix.

Let  $G$  be the graph corresponding to a given irregular grid. That is, the nodes in  $G$  correspond to the nodes in the grid, and any two neighboring nodes in the grid are connected by an edge in  $G$ . Guided by the global multicoloring scheme for rectangular grids, our goal is to partition  $G$  into disjoint subsets of nodes  $G^l$ ,  $l = 1, \dots, p$ , such that (1)  $p$  is as small as possible, (2) no two nodes in any partition  $G^l$  are neighbors, and (3) the variation in the size of the partitions  $G^l$ ,  $l = 1, \dots, p$  is minimal. Given such a partitioning, it is possible to number the nodes in each subset consecutively. However, this does not guarantee a large zero stretch if the numbering of the nodes within each partition  $G^l$  is done arbitrarily. For example, if the node numbered first in  $G^l$  is connected to the node numbered last in  $G^{l-1}$ , then the zero stretch is equal to unity. In other words, the



relation between the nodes in  $G^l$  and  $G^{l-1}$  should be taken into consideration in the numbering process.

The numbering scheme which we suggest consists of the following four steps:

1. The generation of a level structure  $\{V_1, V_2, \dots\}$ , such that the nodes in each level  $V_u$  are not connected to the nodes in levels  $V_w$ ,  $w > u + 1$  or  $w < u - 1$ . In other words, nodes in  $V_u$  may only be connected to nodes in  $V_{u-1}$  and  $V_{u+1}$ .
2. For each level  $V_u$ , the partitioning of the nodes in  $V_u$  into the minimum number of independent sets  $V_u^1, V_u^2, \dots$  such that the nodes in each  $V_u^\lambda$  are independent. That is, no two nodes in  $V_u^\lambda$  are connected by an edge in  $G$ .
3. The formation of the partitions  $G^l$ .
4. The actual numbering of the nodes.

In order to generate the level structure, we start from an arbitrary initial level  $V_1$ , and we proceed in a way very similar to the Cuthill-McKee algorithm (Cuthill and McKee, 1969). More specifically, given a level  $V_u$ , we construct the next level  $V_{u+1}$  by including in it any node that is connected to a node in  $V_u$  and is not in any previous level. A criterion that should guide the choice of  $V_1$  is the desire to have  $\sum_{u=\text{even}} |V_u| \approx \sum_{u=\text{odd}} |V_u|$ , where  $|V_u|$  is the number of nodes in  $V_u$ . For example, if we denote by  $CN_u$  the  $u$ th column of nodes in a pierced rectangular grid, then for  $FE_4$  discretization,  $\{CN_1, CN_2, \dots\}$  is a level structure in which the ratio of  $\sum_{u=\text{odd}} |CN_u|$  to  $\sum_{u=\text{even}} |CN_u|$  may be proved to be larger than 0.5 and smaller than 2 (see proof of Lemma 1 in the Appendix).

Given a specific level,  $V_u$ , the optimal generation of the independent sets  $V_u^1, V_u^2, \dots$  is known to be NP complete (requires nonpolynomial time to execute). However, a simple algorithm that can be used for this generation proceeds by considering the nodes in  $V_u$  in any given order. For each node  $i$ ,  $i$  is added to the first set  $V_u^\lambda$  that does not contain a node that is a neighbor to  $i$ . If this condition is not satisfied for any of the sets constructed so far, then a new set that contains only node  $i$  is created. Let  $s_u$  be the number of independent sets constructed for  $V_u$ , and let  $|V_u^\lambda|$ ,  $1 \leq \lambda \leq s_u$  be the number of nodes in the set  $V_u^\lambda$ .

For pierced rectangular domains and  $FE_4$  discretizations, the number of independent subsets is constant for any level  $u$ . Namely,  $s_u = 2$ . However, for general grids,  $s_u$  may not be constant. In this case, we define  $s_{\max} =$

$\max_u\{s_u\}$ . Note that for  $s_u < \lambda \leq s_{\max}$ , the sets  $V_u^\lambda$  are empty.

Given the independent sets, the partitions  $G^1, G^2, \dots$  may be constructed by the combination of the appropriate independent sets such that the variation in the size of the different partitions is as small as possible. For this, the following algorithm may be applied.

**ALG8: Creation of Independent Partitions**

```

FOR $l = 1, \dots, 2s_{\max}$, initialize G^l to the empty set,
FOR $u = 1, 3, 5, \dots$ DO
 FOR $\lambda = 1, \dots, s_{\max}$ DO
 Choose the partition G^l , $1 \leq l \leq s_{\max}$ which
 contains, so far, the minimum number of nodes
 (break ties arbitrarily), and set $G^l = G^l \cup V_u^\lambda$.
 FOR $u = 2, 4, 6, \dots$ DO
 FOR $\lambda = 1, \dots, s_{\max}$ DO
 Choose the partition G^l , $s_{\max} < l \leq 2s_{\max}$ which
 contains, so far, the minimum number of nodes
 (break ties arbitrarily), and set $G^l = G^l \cup V_u^\lambda$.

```

After obtaining the  $2s_{\max}$  partitions, we may then start the actual numbering of the nodes. For this, we consider the partitions in the order  $G^1, G^2, \dots$ , and within each partition  $G^l$ , we number the nodes sequentially such that if  $V_u^k$  and  $V_v^\lambda$  are two subsets in  $G^l$  with  $v > u$ , then the nodes in  $V_u^k$  are numbered before the nodes in  $V_v^\lambda$ .

Clearly, ALG8 guarantees that the sizes of the partitions  $G^l$ ,  $l = 1, \dots, s_{\max}$  do not differ by more than the size of the largest subset among all subsets  $V_u^\lambda$ ,  $\lambda = 1, \dots, s_{\max}$ ,  $u = 1, 2, \dots$ . The same applies to the partitions  $G^l$ ,  $l = s_{\max} + 1, \dots, 2s_{\max}$ . However, no relation may be established between the size of two partitions  $G^l$  and  $G^\lambda$ ,  $l \leq s_{\max}$  and  $\lambda > s_{\max}$ , without an explicit assumption about the size of each level  $V_u$  with respect to the following level  $V_{u+1}$ . If an assumption of this kind is made, then it may be possible, following a reasoning similar to that of Theorem 2, to find a lower bound on the value of the zero stretch of the matrix resulting from this multicolor numbering scheme. We will not pursue this issue any further.


We illustrate our numbering scheme by applying it to the grid of Figure 5. In this figure, a level structure  $\{V_1, \dots, V_9\}$  is indicated by bold dashed lines, and within each level  $V_u$ , the nodes which are assigned to the subsets  $V_u^1, V_u^2$ , and  $V_u^3$  are marked by the symbols  $\bullet$ ,  $\square$ , and  $\triangle$ , respectively. The application of ALG8, then,

results in the following partitions:

$$\begin{aligned}
 G^1 &= V_1^1 \cup V_3^3 \cup V_5^2 \cup V_9^1 \\
 G^2 &= V_1^2 \cup V_3^2 \cup V_7^2 \cup V_9^2 \\
 G^3 &= V_3^1 \cup V_5^1 \cup V_7^1 \cup V_9^3 \\
 G^4 &= V_2^1 \cup V_4^2 \cup V_8^1 \\
 G^5 &= V_2^2 \cup V_6^1 \cup V_8^2 \\
 G^6 &= V_2^3 \cup V_4^1 \cup V_6^2
 \end{aligned}$$

which leads to the numbering shown. It is easy to check that the zero stretch resulting from this numbering is  $\Delta = 8$ . How good is this result compared with the maximum possible zero stretch for this specific grid? And is there any numbering scheme which is not NP complete, and which will produce the maximum zero stretch? We do not know the answers to these questions, but some preliminary results lead us to believe that the answer to the second question is negative.

Finally, we should note that the above numbering scheme is global in the sense that all the nodes which are given a particular color are numbered consecutively. A level-wise numbering scheme analogous to the column-wise scheme of section 3.1 may also be applied. However, the construction of the subsets  $V_u^1, V_u^2, \dots$  in this case should be done such that the sizes of these subsets are as uniform as possible.



**Fig. 5 Multicolor numbering of an irregular grid**

**Table 2**

**The Number of PCCG Iterations and the Execution Time (in seconds) for Problem 1 with  $FD_6$  Discretization ( $\omega = 1$ )**

| Grid Size | Precond. Method | One Color | Global 2-Colors | Column-wise 2-Colors | No Precond. 1-Color |
|-----------|-----------------|-----------|-----------------|----------------------|---------------------|
| 64 × 64   | SSOR            | 70/       | 97/             | 84/                  | 191/                |
|           | IC ( $P_d$ )    | 59/       | 97/             | 83/                  |                     |
| 100 × 100 | SSOR            | 110/5.38  | 153/2.54        | 131/2.72             | 301/2.32            |
|           | IC ( $P_d$ )    | 90/4.42   | 153/2.57        | 129/2.69             |                     |
| 150 × 150 | SSOR            | 163/17.6  | 231/8.6         | 197/7.9              | 439/5.56            |
|           | IC ( $P_d$ )    | 136/14.0  | 231/8.7         | 193/7.8              |                     |
| 200 × 200 | SSOR            | 218/41.3  | 308/20.3        | 262/18.4             | 586/13.1            |
|           | IC ( $P_d$ )    | 181/34.5  | 308/20.3        | 258/18.3             |                     |

**Table 3**

**The Number of PCCG Iterations (I) and the Execution Time (T) for Problem 1 with  $FD_6$  and the Optimal Choice of  $\omega$**

| Grid Size | Precond. Method        | One Color |                | Global 2-Colors |                | Column-wise 2-colors |                |
|-----------|------------------------|-----------|----------------|-----------------|----------------|----------------------|----------------|
|           |                        | I/T       | $\omega_{opt}$ | I/T             | $\omega_{opt}$ | I/T                  | $\omega_{opt}$ |
| 200 × 200 | SSOR                   | 51/10.1   | 1.95           | 308/20.3        | 1.0            | 253/17.9             | 1.15           |
|           | SIC ( $P_{d,\omega}$ ) | 53/10.2   | 1.23           | 308/20.3        | 1.0            | 253/17.9             | 1.13           |

**Table 4**

**Results for Problem 1 with  $FE_4$  on a 150 × 150 Grid**

| Precond. Method        | One Color      |                | Global 4-Colors |                | Column-wise 2-Colors |                |
|------------------------|----------------|----------------|-----------------|----------------|----------------------|----------------|
|                        | I/T            | $\omega$       | I/T             | $\omega$       | I/T                  | $\omega$       |
| SSOR                   | 120/19.6       | 1.0            | 144/11.6        | 1.0            | 134/7.8              | 1.0            |
| IC ( $P_d$ )           | 112/18.4       | 1.0            | 143/11.7        | 1.0            | 127/7.6              | 1.0            |
|                        | I/T            | $\omega_{opt}$ | I/T             | $\omega_{opt}$ | I/T                  | $\omega_{opt}$ |
| SSOR                   | 37/6.3         | 1.95           | 144/11.6        | 1.0            | 94/5.5               | 1.62           |
| SIC ( $P_{d,\omega}$ ) | 38/6.5         | 1.55           | 144/11.7        | 1.0            | 85/5.3               | 1.43           |
| No precondition.       | I/T = 273/4.89 |                |                 |                |                      |                |

## 4. NUMERICAL RESULTS

In order to evaluate the diagonal update technique of section 2, and to test the column-wise multicolor numbering scheme of section 3, we implemented the PCCG method on the CRAY X-MP using the stripe structure approach of section 1. The implementation allows for the solution of general symmetric sparse systems and provides for no preconditioning, SSOR,  $SIC(P_a, \omega)$ , or  $SIC(P_a, \omega)$  preconditioning. The PCCG iteration is stopped when the 2-norm of both the residual and the step size fall below  $10^{-5}$ .

The linear systems used in our experiments are generated by using a modified version of the finite element code of Becker, Carey, and Oden (1981). The modification incorporates the stripe data structure and provides options for the automatic specification of pierced rectangular domains and grids, the choice of the numbering scheme (one color, column-wise multicolor, or global multicolor), and the application of finite difference discretization. The results obtained for three problems are reported next.

### 4.1 PROBLEM 1

$$\begin{aligned}\nabla^2 u - u &= 0 && \text{on } D \\ u &= 1 + xy && \text{on } \partial D\end{aligned}$$

where  $D$  is the unit square. For this problem, we consider both  $FD_5$  and  $FE_4$  discretizations. For  $FD_5$ , we consider many grid sizes, and we report in Table 2 the number of iterations  $I$  and the execution time  $T$  (in seconds) for the PCCG corresponding to the numbering schemes discussed in section 3. We also report  $I$  and  $T$  for the basic (nonpreconditioned) conjugate gradient method. The execution times are omitted for the  $64 \times 64$  grid because they are relatively small and thus inconclusive. Note that the global two-color numbering scheme is the usual red-black numbering scheme.

The results in Table 2 are obtained by fixing the shift factor,  $\omega$  (over-relaxation factor), to unity. In order to observe the effect of changing  $\omega$ , we report in Table 3 the results for the SSOR/PCCG and the  $SIC(P_a, \omega)$ /PCCG methods with  $\omega = \omega_{opt}$ , the optimal value of  $\omega$  which gave the fastest convergence. In Table 4, we combine the results for  $\omega = 1$  and  $\omega = \omega_{opt}$  in one table for the  $FE_4$  discretization.

**4.2 PROBLEM 2: (PROB 29 in Rice, Houstis, and Dyksen, 1981)**

$$\begin{aligned} \nabla y \nabla u &= 0 && \text{on } D \\ u &= x - y && \text{on } \partial D \end{aligned}$$

where  $D$  is the unit square. The results of  $FE_4$ ,  $FE_3$ , and  $FE_9$  discretizations for this problem are given in Tables 5 and 6. At this point, we would like to note that the running times reported for this and the previous problems seem high compared with the running times which are reported in other papers (for example, van der Vorst, 1985; Young et al., 1985). This should not be surprising because our implementation is designed for general sparse matrices, and hence does not assume any special form for the stripes. Clearly, an implementation that is designed specifically for matrices generated from rectangular grids would use the fact that all the stripes are parallel to the diagonal of the matrix, thus eliminating the need for the time-consuming gather and scatter operations.

**4.3 PROBLEM 3: (PROB 26 in Rice, Houstis, and Dyksen, 1981)**

$$\begin{aligned} \nabla(1 + x^2)^3 \nabla u &= 60x && \text{on } D \\ u &= 0 && \text{on } \partial D \end{aligned}$$

where  $D$  is the pierced rectangular domain of Figure 3. The result of the  $FE_4$  discretization for this problem using a  $100 \times 200$  pierced rectangular grid is given in Table 7.

**4.4 CONCLUSIONS**

By studying the results obtained for problems 1, 2, and 3, we may note the following:

1. For large grid sizes, the column-wise numbering scheme is consistently better than the global numbering scheme. In fact, it seems that the global scheme may be recommended (on the CRAY) only if the grid size is relatively small, more specifically, if the zero stretch for the column-wise scheme is smaller than 100. Clearly, for such small problems, the power of a supercomputer is not needed.
2. If no adaptive procedure is incorporated in the PCCG solver to choose the optimal  $\omega$ , then the natural choice is  $\omega = 1$ . In this case, it is clear from the results that the use of the  $IC(P_d)$  preconditioner instead of the SSOR preconditioner reduces the

**Table 5**  
**Results for Problem 2 with  $FE_4$  on a  $150 \times 150$  Grid**

| Precond.<br>Method    | One<br>Color   |                | Global<br>4-Colors |                | Column-wise<br>2-Colors |                |
|-----------------------|----------------|----------------|--------------------|----------------|-------------------------|----------------|
|                       | I/T            | $\omega$       | I/T                | $\omega$       | I/T                     | $\omega$       |
| SSOR                  | 136/22.7       | 1.0            | 171/13.5           | 1.0            | 140/8.5                 | 1.0            |
| $IC(P_d)$             | 126/21.1       | 1.0            | 170/13.5           | 1.0            | 132/8.1                 | 1.0            |
| $IC(P_s)$             | 96/17.4        | 1.0            | 166/16.7           | 1.0            | 109/8.1                 | 1.0            |
|                       | I/T            | $\omega_{opt}$ | I/T                | $\omega_{opt}$ | I/T                     | $\omega_{opt}$ |
| SSOR                  | 39/6.8         | 1.95           | 171/13.5           | 1.0            | 98/5.8                  | 1.6            |
| $SIC(P_{dr}, \omega)$ | 39/6.9         | 1.55           | 170/13.5           | 1.0            | 86/5.2                  | 1.4            |
| $SIC(P_{sr}, \omega)$ | 37/7.5         | 1.2            | 166/16.7           | 1.0            | 71/5.7                  | 1.2            |
| No precondition.      | I/T = 647/12.8 |                |                    |                |                         |                |

**Table 6**  
**Results for Problem 2 with  $FE_3$  and  $FE_9$  on a  $99 \times 99$  Grid**

| Precond.<br>Method    | $FE_3$       |                |                    | $FE_9$                  |     |                | $FE_9$                  |                |         |                |
|-----------------------|--------------|----------------|--------------------|-------------------------|-----|----------------|-------------------------|----------------|---------|----------------|
|                       | One<br>Color |                | Global<br>4-Colors | Column-wise<br>2-Colors |     | One<br>Color   | Column-wise<br>3-Colors |                |         |                |
|                       | I            | $\omega$       | I                  | $\omega$                | I   | $\omega$       | I/T                     | $\omega$       | I/T     | $\omega$       |
| SSOR                  | 108          | 1.0            | 450                | 1.0                     | 120 | 1.0            | 109/19.3                | 1.0            | 114/9.9 | 1.0            |
| $IC(P_d)$             | 92           | 1.0            | 432                | 1.0                     | 113 | 1.0            | 97/17.3                 | 1.0            | 103/9.1 | 1.0            |
|                       | I            | $\omega_{opt}$ | I                  | $\omega_{opt}$          | I   | $\omega_{opt}$ | I/T                     | $\omega_{opt}$ | I/T     | $\omega_{opt}$ |
| SSOR                  | 35           | 1.9            | 135                | 1.3                     | 109 | 1.3            | 35/6.4                  | 1.95           | 86/7.5  | 1.6            |
| $SIC(P_{dr}, \omega)$ | 32           | 1.3            | 122                | 1.2                     | 105 | 1.2            | 46/8.4                  | 1.4            | 80/7.1  | 1.3            |
| No precondition.      | I = 565      |                |                    |                         |     |                | I/T = 556/12.1          |                |         |                |

**Table 7**  
**Results for Problem 3 with  $FE_4$**

| Precond.<br>Method | One<br>Color   |                | Global<br>4-Colors |                | Column-wise<br>2-Colors |                |
|--------------------|----------------|----------------|--------------------|----------------|-------------------------|----------------|
|                    | I/T            | $\omega$       | I/T                | $\omega$       | I/T                     | $\omega$       |
| SSOR               | 90/10.8        | 1.0            | 112/6.5            | 1.0            | 106/5.0                 | 1.0            |
| $IC(P_d)$          | 69/8.4         | 1.0            | 108/6.4            | 1.0            | 100/4.8                 | 1.0            |
| $IC(P_a)$          | 69/9.0         | 1.0            | 113/8.3            | 1.0            | 104/5.8                 | 1.0            |
|                    | I/T            | $\omega_{opt}$ | I/T                | $\omega_{opt}$ | I/T                     | $\omega_{opt}$ |
| SSOR               | 29/3.6         | 1.9            | 109/6.2            | 1.2            | 102/4.7                 | 1.2            |
| $SIC(P_d, \omega)$ | 29/3.8         | 1.22           | 106/6.1            | 1.15           | 95/4.5                  | 1.2            |
| $SIC(P_a, \omega)$ | 30/4.4         | 1.2            | 107/8.1            | 1.25           | 86/4.8                  | 1.3            |
| No precondition.   | I/T = 533/7.26 |                |                    |                |                         |                |

number of PCCG iterations for both the one-color scheme and the column-wise coloring scheme. The advantage of the  $IC(P_d)$  preconditioner is clear, given that the cost of the computation of the  $IC(P_d)$  matrix is less than the cost of one PCCG iteration.

3. The  $IC(P_d)$  preconditioner may reduce the number of iterations over the  $IC(P_a)$  preconditioner. However, the  $IC(P_d)/PCCG$  usually runs faster than the  $IC(P_a)/PCCG$  because it costs too much to compute the  $IC(P_d)$  factorization. Here, we note that our implementation does not take advantage of the technique of Eisenstat (1981), which is described at the end of section 2, to reduce the work in  $SSOR/PCCG$  and  $IC(P_d)/PCCG$ . Clearly, the use of this technique will reduce further the execution time of these methods compared with the  $IC(P_a)/PCCG$ .
4. If a one-color numbering scheme is applied, then it is possible to reduce the number of iterations in the  $SSOR/PCCG$  drastically by using the optimum value of  $\omega$ . This drastic reduction offsets the disadvantage of having a unit zero stretch. In other words, if an adaptive strategy may be used to choose the optimum  $\omega$ , then the one-color numbering scheme seems to be the most appropriate scheme to apply.
5. For problems that are well conditioned, as for example problem 1, the fastest way to solve the problem is the nonpreconditioned CG method.



However, this is not the case for problems 2 and 3, which are ill conditioned.

6. There seems to be an inverse relation between the value of the zero stretch and the convergence rate of the PCCG. In order to be more specific, we consider the  $FD_5$  discretization of problem 1, and in addition to the numbering schemes reported in Table 2, we use a three-color, column-wise scheme, and a four-color global scheme. The results we report in Table 8 confirm this inverse relation.

Finally, we would like to note that only rectangular and pierced rectangular domains were chosen in the experiments because, first, we did not have a means for the automatic generation of large irregular grids, and second, the numbering algorithm of section 3.3 is, to our knowledge, the first algorithm which aims at increasing the zero stretch for general matrices. Hence, its performance may be compared only with an algorithm that produces the maximum zero stretch. We believe that such an algorithm is NP complete.

## APPENDIX

*Proof of Theorem 1.* Let  $i$  and  $j$  be two nodes in some element  $e$  in  $\Omega$ , and let  $l$ ,  $1 \leq l \leq q$ , be the color given to  $i$  and let  $l_+$  be the color given to  $j$ , where  $l_+ = l + 1$  if  $1 \leq l < q$ , and  $l_+ = 1$  if  $l = q$ . Note that either  $i$  and  $j$  lie in the same column of nodes,  $CN_w$  in  $\Omega$ , or  $i$  lies in  $CN_u$  and  $j$  lies in the following column of nodes, namely  $CN_{u+1}$ . Let also  $CE_w$  be a column of elements in  $\Omega$  which contains element  $e$ , and let  $ce_w$  be the number of elements in  $CE_w$ .

From ALG6 and ALG7, it is clear that the number of nodes that are numbered between nodes  $i$  and  $j$ , namely  $\nu(j) - \nu(i) + 1$ , is equal to the sum of (1) the number of nodes with color  $l$  which are above node  $i$  (including  $i$ ) in column  $CN_w$ , and (2) the number of nodes with color  $l_+$  which are below node  $j$  (including  $j$ ) in column  $CN_u$  or  $CN_{u+1}$ , whichever applies.

Given that at most two elements in the same column may share a single node, and that for any specific color, every element should contain a node of that color, we may conclude that  $\nu(j) - \nu(i) + 1$  should be larger than, or equal to, the sum of (1) half the number of elements in  $CE_w$  above, and including, element  $e$ , and (2) half the number of elements in  $CE_w$  below, and including, ele-

**Table 8**  
**The Effect of  $\Delta$  on the Solution of Problem 1**  
**with  $FD_5$  Discretization on a**  
 **$150 \times 150$  Grid ( $\omega = 1$ )**

|                                                  | One<br>Color | Col.-wise<br>3-Colors | Col.-wise<br>2-Colors | Global<br>4-Colors | Global<br>2-Colors |
|--------------------------------------------------|--------------|-----------------------|-----------------------|--------------------|--------------------|
| No. <i>SSOR/PCCG</i><br>iterations               | 163          | 187                   | 197                   | 204                | 231                |
| No. <i>IC (P<sub>d</sub>)/PCCG</i><br>iterations | 136          | 165                   | 193                   | 199                | 231                |
| The zero<br>stretch $\Delta$                     | 1            | 50                    | 74                    | 5,624              | 11,175             |

ment  $e$ . In other words,

$$v(j) - v(i) + 1 \geq \frac{ce_w + 1}{2}$$

But  $\Delta = \min_e \{v(j) - v(i)\}$ , where  $e$  may be in any column in  $\Omega$ , Hence

$$\Delta + 1 \geq \frac{ce_{\min} + 1}{2}$$

which gives the lower bound.

In order to establish the upper bound on  $\Delta$ , we note that  $v(j) - v(i) - 1$  is equal to the sum of nodes of color  $l$  above node  $i$  and the nodes of color  $l_+$  below node  $j$  (not including  $i$  and  $j$ ). But, each element in  $\Omega$  may contain at most one node of any given color. Then

$$v(j) - v(i) - 1 \leq ce_v - 1$$

and the upper bound follows directly.  $\square$

*Proof of Lemma 1.* Let  $CN_u$  be the  $u$ th column of nodes in  $\Omega$  and let  $cn_u$  be the number of nodes in this column. In the global multicolor numbering scheme, the colors  $1, \dots, \sqrt{p}$  are, first, used to color the nodes in  $CN_{1+v\sqrt{p}}$ ,  $v = 0, 1, 2, \dots$ , such that any two neighboring nodes in a column are given different colors. Then, the colors  $\sqrt{p} + 1, \dots, 2\sqrt{p}$  are used to color the nodes in  $CN_{2+v\sqrt{p}}$ ,  $v = 0, 1, 2, \dots$ , and so on. Thus, if the given color  $l$  satisfies  $(k - 1)\sqrt{p} < l \leq k\sqrt{p}$ , for some  $k$ ,  $1 \leq k \leq \sqrt{p}$ , then the number of nodes in  $\Omega$  which are given the color  $l$ , namely  $n_{\Omega}^l$ , is obtained from

$$n_{\Omega}^l = \sum_{v=0,1,\dots} cn_{k+v\sqrt{p}}^l \quad (14)$$

where  $cn_u^l$  is the number of nodes which are given the color  $l$  in column  $CN_u$ .

In order to estimate each term in Eq(14), we consider a particular column of nodes  $CN_u$ , where  $u = k + v\sqrt{p}$ , for some  $v$ , and we let  $L = (k - 1)\sqrt{p}$ . Clearly, the colors used to color the nodes in  $CN_u$  are  $\lambda = L + 1, \dots, L + \sqrt{p}$ . We also define a clique in  $CN_u$  as a subset of nodes within  $CN_u$  such that any two nodes in this subset are neighbors. Note that the nodes in  $CN_u$  which lie on a vertical side of an element form a clique. Let  $cl_u$  be the number of such cliques in  $CN_u$ . Given that any particular node in  $CN_u$  may be in at most two cliques, we obtain

$$cl_u \geq 2cn_u^{\lambda} \quad \lambda = L + 1, \dots, L + \sqrt{p}. \quad (15a)$$

Also, the coloring scheme is such that any clique may

contain at most one node of any particular color. This gives

$$cn_u^\lambda \leq cn_u^\lambda \quad \lambda = L + 1, \dots, L + \sqrt{p}. \quad (15b)$$

From Eq(15a) and Eq(15b), we get, for the given color  $l$ ,

$$cn_u^l \leq 2cn_u^\lambda \quad \lambda = L + 1, \dots, L + \sqrt{p}.$$

If we write the above inequality  $\sqrt{p} - 1$  times with  $\lambda = L + 1, \dots, L + \sqrt{p}$ ,  $\lambda \neq l$ , sum these inequalities, and then add  $2cn_u^l$  to both sides of the result, we get

$$(\sqrt{p} + 1)cn_u^l \leq 2 \sum_{\lambda=L}^{L+\sqrt{p}} cn_u^\lambda = 2cn_u. \quad (16)$$

Similarly, from Eq(15a) and Eq(15b) we may obtain

$$cn_u^l \geq \frac{1}{2}cn_u^\lambda \quad \lambda = L + 1, \dots, L + \sqrt{p},$$

which leads to

$$\left(\sqrt{p} - \frac{1}{2}\right)cn_u^l \geq \frac{1}{2}cn_u. \quad (17)$$

Eq(16) and Eq(17) may now be combined into

$$\frac{1}{2\sqrt{p} - 1}cn_u \leq cn_u^l \leq \frac{2}{\sqrt{p} + 1}cn_u$$

which, together with Eq(14) gives

$$\frac{1}{2\sqrt{p} - 1} \sum_{v=0,1,\dots} cn_{k+v\sqrt{p}} \leq n_\Omega^l \leq \frac{2}{\sqrt{p} + 1} \sum_{v=0,1,\dots} cn_{k+v\sqrt{p}}. \quad (18)$$

Finally, we need to relate the sum in Eq(18) to the total number of nodes in  $\Omega$ . Instead of dealing with complex formulas, we treat each of the cases  $FE_4$  and  $FE_9$  separately. For  $FE_4$ , any column of nodes,  $CN_v$ , in  $\Omega$  is enclosed between two consecutive columns of elements. This implies that  $cn_u \leq cn_{u-1} + cn_{u+1}$ . Thus

$$\sum_{u=odd} cn_u \leq 2 \sum_{u=even} cn_u. \quad (19)$$

Noting that  $\sum_{u=odd} cn_u + \sum_{u=even} cn_u = n_\Omega$ , we may use Eq(19) to obtain

$$\frac{1}{3}n_\Omega \leq \sum_{u=even} cn_u \leq \frac{2}{3}n_\Omega. \quad (20)$$

The same bounds as for Eq(20) may be obtained for  $\sum_{u=odd} cn_u$ , which proves the following for  $FE_4$  and  $p = 4$ :

$$\frac{1}{2\sqrt{p} - 1} n_{\Omega} \leq \sum_{v=0,1,\dots} cn_{k+v\sqrt{p}} \leq \frac{2}{\sqrt{p} + 1} n_{\Omega}. \quad (21)$$

For  $FE_9$  discretizations, a similar procedure, starting from  $cn_u \leq cn_{u+1} + cn_{u-2}$  and  $cn_u \leq cn_{u-1} + cn_{u+2}$ , may be used to prove Eq(21) for the case  $p = 9$ . The result Eq(12), then, follows directly from Eq(21) and Eq(18).  $\square$

*Proof of Theorem 2.* From the definition of  $\Delta$ , there exist two nodes  $i$  and  $j$  in an element  $e$  such that  $v(j) - v(i) = \Delta$ . Let  $CE_u$  be the column of elements which contains  $e$ , and let  $l$  and  $l + 1$  be the colors of nodes  $i$  and  $j$ , respectively. The proof proceeds by the construction of two pierced rectangular domains  $\Omega_1$  and  $\Omega_2$  which overlap at element  $e$ . Namely,  $\Omega_1$  consists of element  $e$ , the elements in  $CN_u$  above  $e$ , and the elements in the following columns  $CN_{u+1}, CN_{u+2}, \dots$ . The grid  $\Omega_2$  consists of element  $e$ , the elements in  $CN_u$  below  $e$ , and the elements in the previous columns  $CN_1, \dots, CN_{u-1}$ . Let  $n_{\Omega_1}$  and  $n_{\Omega_2}$  be the number of nodes in  $\Omega_1$  and  $\Omega_2$ , respectively, and note that  $n_{\Omega_1} + n_{\Omega_2} = n_{\Omega} + p$ .

From the global numbering scheme, it is clear that the difference  $v(j) - v(i) + 1$  is equal to the number of nodes in  $\Omega_1$  which have the color  $l$ , namely  $n_{\Omega_1}^l$ , plus the number of nodes in  $\Omega_2$  which have the color  $l + 1$ , namely  $n_{\Omega_2}^{l+1}$ . That is,

$$v(j) - v(i) + 1 = n_{\Omega_1}^l + n_{\Omega_2}^{l+1}. \quad (23)$$

But,  $\Omega_1$  and  $\Omega_2$  are pierced rectangular domains; thus,  $n_{\Omega_1}^l$  and  $n_{\Omega_2}^{l+1}$  obey the bound of Lemma 1. The application of these bounds on Eq(23) gives

$$\left[ \frac{1}{2\sqrt{p} - 1} \right]^2 [n_{\Omega_1} + n_{\Omega_2}] \leq v(j) - v(i) + 1 \leq \left[ \frac{2}{\sqrt{p} + 1} \right]^2 [n_{\Omega_1} + n_{\Omega_2}].$$

Finally, Eq(13) results by using the definition of  $\Delta$  and  $n_{\Omega} = n_{\Omega_1} + n_{\Omega_2} - p$ .  $\square$

#### ACKNOWLEDGMENT

This work is, in part, supported under ONR contract N00014-85-K-0339 and Air Force contract AFOSR-84-0131. The experimental results were obtained on the CRAY X-MP of the Pittsburgh Supercomputing Center.

#### BIOGRAPHY

Rami Melhem was born in Cairo, Egypt, on June 30, 1954. He received his B.S. degree in electrical engineering in 1979 from Cairo University and his M.S. degree in mathematics and his Ph.D. degree in computer science,

both from the University of Pittsburgh, in 1981 and 1983, respectively.

In 1984, he joined the faculty of the Computer Science Department at Purdue University. He is currently on leave from Purdue and with the Department of Computer Science at the University of Pittsburgh. His main research interest is the application of parallel and systolic architectures to scientific computing, in general, and to finite element analysis, in particular.

## REFERENCES

- Axelsson, O. 1984. A survey of vectorizable preconditioning method for large scale finite element matrix problems. Tech. Report CNA-190, Center for Numerical Analysis, The University of Texas at Austin.
- Becker, E., Carey, G., and Oden, J. T. 1981. *Finite Elements, An Introduction, Vol. 1*. Englewood Cliffs, New Jersey: Prentice-Hall.
- Concus, P., Golub, G., and Meurant, G. 1985. Block preconditioning for the conjugate gradient methods. *SIAM J. Sci. Statist. Comput.* 6:220–252.
- Cuthill, E., and McKee, J. 1969. Reducing the bandwidth of sparse symmetric matrices. *Proc. of the ACM National Conf.* New York: Association for Computing Machinery, pp. 157–172.
- Dongarra, J., Gustavson, F., and Karp, A. 1984. Implementing linear algebra algorithms for dense matrices on a vector pipeline machine. *SIAM Rev.* 26:91–112.
- Eisenstat, S. 1981. Efficient implementation of a class of preconditioned conjugate gradient methods. *SIAM J. Sci. Statist. Comput.* 2(1):1–4.
- Gustafsson, I. 1978. A class of first order factorization methods. *BIT.* 18:142–156.
- Hageman, A., and Young, D. 1981. *Applied iterative methods*. New York: Academic Press.
- Kincaid, D., Oppe, T., and Young, D. 1984. Vector computations for sparse linear systems. Tech. Report CNA-189, Center for Numerical Analysis, The University of Texas at Austin.
- Madsen, N., Rodrigue, G., and Karush, J. 1976. Matrix multiplication by diagonals on a vector/parallel processor. *Inform. Proc. Lett.* 5(2):41–45.
- Manteuffel, T. 1980. An incomplete factorization technique for positive definite linear systems. *Math. Comp.* 34(150):673–697.
- Meijerink, J. A., and van der Vorst, H. 1977. An iterative solution method for linear systems of which the coefficient matrix is symmetric M-matrix. *Math. Comp.* 31(137):148–162.
- Meijerink, J. A., and van der Vorst, H. 1981. Guidelines for the usage of incomplete decompositions in solving sets of linear equations as they occur in practical problems. *J. Comput. Phys.* 44:134–155.
- Melhem, R. 1986a. Parallel solution of linear systems with striped sparse matrices. Tech. Report ICMA-86-91; *Parallel Comput.*, in press.
- Melhem, R. 1986b. A study of the stripe structure of finite element stiffness matrices. Tech. Report ICMA-86-92; *SIAM J. Numer. Anal.*, in press.

Meurant, G. 1984. The block preconditioned conjugate gradient method on vector computers. *BIT*. 24:623-633.

Poole, E., and Ortega, J. 1986. Multicolor ICCG methods for vector computers. Applied Math. Report RM-86-06, University of Virginia.

Rice, J., and Boisvert, R. 1984. *Solving elliptic problems with Ellpack*. New York: Springer Verlag.

Rice, J., Houstis, E., and Dyksen, W. 1981. A population of linear second order elliptic partial differential equations on rectangular domains; parts 1 and 2. *Math. Comp.* 36:475-484.

Saad, Y. 1985. Practical use of polynomial preconditionings for the conjugate gradient method. *SIAM J. Sci. Statist. Comput.* 6(4):865-881.

Saad, Y., and Schultz, M. 1985. Parallel implementations of preconditioned conjugate gradient methods. Tech. Report YALEU/DCS/RR425, Dept. of Computer Science, Yale University.

van der Vorst, H. 1982. A vectorizable variant of some ICCG methods. *SIAM J. Sci. Statist. Comput.* 3(3):350-356.

van der Vorst, H. 1986. The performance of FORTRAN implementations for preconditioned conjugate gradients on vector computers. *Parallel Comput.* 3:49-58.

Young, D., Oppe, T., Kincaid, D., and Hayes, L. 1985. On the use of vector computers for solving large sparse linear systems. Tech. Report CNA-199, Center for Numerical Analysis, The University of Texas at Austin.