1981

A New Ordering Scheme for the Hermite Bicubic Collocation Equations

Wayne R. Dyksen

John R. Rice
Purdue University, jrr@cs.purdue.edu

Report Number:
81-425

A NEW ORDERING SCHEME FOR THE HERMITE BICUBIC COLLOCATION EQUATIONS

Wayne R. Dyksen[1]

John R. Rice[1]

Division of Mathematical Sciences
Purdue University
West Lafayette, IN 47907

CSD-TR 425

ABSTRACT

We present a new ordering scheme for the linear equations obtained from the Hermite bicubic collocation approximation to the solution of second order, linear elliptic partial differential equations on rectangular domains. The resulting matrix has a non-zero diagonal and is a band matrix with band width $4n+11$ for a grid of n^2 rectangles.

We report on an experimental study which demonstrates that the resulting linear system may be solved safely by using Gauss elimination **without** pivoting. We show that a standard finite element order is reordered by scaled partial pivoting to give essentially the new ordering we describe. We discuss the theoretical computational savings expected from not pivoting and show that the computing time for not pivoting is approximately the same as for scaled partial pivoting. This is contrary to the expectation from the usual operations count analysis. The new ordering, as one expects, does reduce the memory use by about 40%.

A condensed version of this paper will appear in *Elliptic Problem Solvers II* edited by Birkhoff and Schoenstat, Academic Press, 1983.

A New Ordering Scheme for the Hermite Bicubic Collocation Equations

Wayne R. Dyksen[1]
John R. Rice[1]

I. INTRODUCTION

We consider a second order, linear elliptic problem on a rectangle R in the following form

$$\begin{aligned} Lu &= au_{xx} + bu_{xy} + cu_{yy} + du_x + eu_y + fu = g \quad (x,y) \in R \\ Mu &= \alpha u + \beta u_x + \gamma u_y = \delta \qquad\qquad (x,y) \in \partial R. \end{aligned}$$

Here, a, b, c, d, e, f, g, α, β, γ and δ are functions of x and y. A Hermite cubic is a C^1 piecewise cubic polynomial and a Hermite bicubic is the tensor product of Hermite cubics in x and in y. We consider a collocation method which places a rectangular grid on the rectangle R and chooses as basis functions Hermite bicubics with breakpoints along the grid lines to approximate $u(x,y)$ by

$$U(x,y) = \sum \alpha_{ij}\, h_i(x)\, h_j(y) \approx u(x,y).$$

The coefficients α_{ij} of U are determined by collocating at four interior points of each grid square plus at points along the boundary. That is, the α_{ij} satisfy the following *collocation equations*:

$$\begin{aligned} LU(x_k, y_k) &= f(x_k, y_k) \quad k = 1,\ldots,N \\ MU(x_k, y_k) &= g(x_k, y_k) \quad k = N+1,\ldots,M \end{aligned}$$

where $N = 4n^2$ and $M-N = 4(2n+1)$ for a grid of n^2 rectangles. Refer to (Rice, 1983) for more details of this method.

The unknowns and equations of the collocation equation are usually numbered so as to give a coefficient matrix of the linear system which is somewhat bi-diagonal (i.e., with the finite element ordering). With $n=3$ and a Poisson problem with Dirichlet boundary conditions, the pattern of non-zero elements is as shown in Figure 1. There are variations on this ordering which produce similar patterns.

[1] This work supported in part by Department of Energy contract DE-AC02-81ER10997

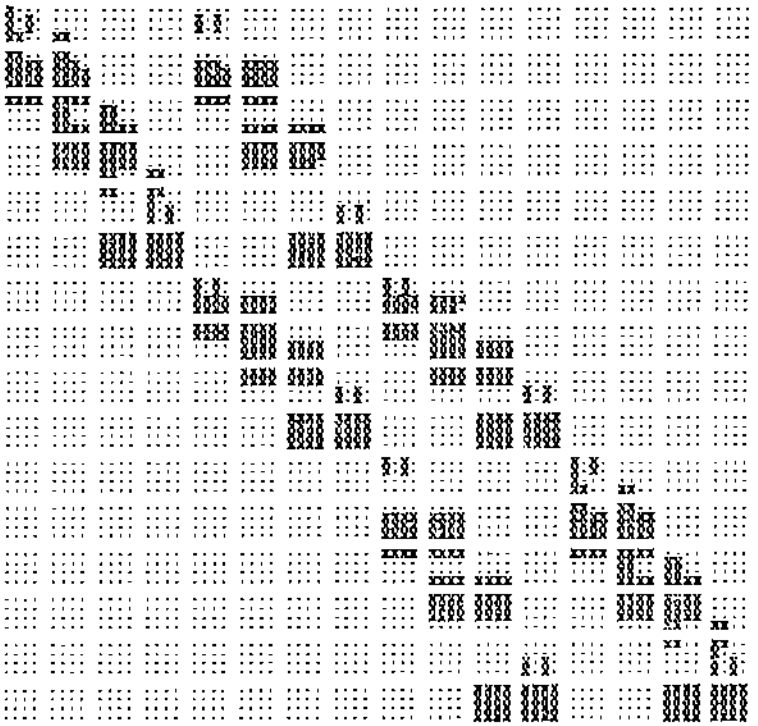

FIGURE 1. The pattern of non-zero elements in the collocation equations with the usual orderings of the equations and unknowns. This is for a 3×3 grid and a Poisson problem with Dirichlet boundary conditions.

The pattern seen in Figure 1 is unexpected. One expects a somewhat tridiagonal pattern of non-zeros and, potentially more serious, one cannot use any of the standard iteration methods on the equations because of the large number of zero elements on the diagonal. Thus we were led to search for a new ordering which has non-zeros on the diagonals and which is more "natural". Such an ordering, called collorder, is presented in the next section. A second candidate for such an ordering is the tensor product ordering.

Unfortunately, standard iteration methods do not converge for these equations in the collorder ordering (see (Balart, 1982), (Dyksen, 1982) for iteration methods that do converge). We thus turned our attention to the behavior of Gauss elimination applied to the linear equations with this ordering. We discovered that the standard LINPACK routines (Dongarra, 1979) give unacceptable results because of scaling problems. The difficulties of scaling these equations are analyzed in (Dyksen and Rice, 1983); we merely report here that *a priori* scaling is difficult (at best) and we believe that one should use scaled partial pivoting which we call BAND GE here. Once we switched to BAND GE, we found round-off error effects to be remarkably low. Even more surprising, we found that pivoting is **not** needed with the collorder ordering, and that the collorder ordering is essentially the ordering one obtains when applying scaled partial pivoting to the equations in the usual orderings such as shown in Figure 1.

In Section III we present the results of an experimental study which strongly supports the hypothesis: **Gauss elimination without pivoting is safe for the collorder ordering of the collocation equations.** Further discussion is given in Section IV.

Even more surprising to us is the fact: **not pivoting in the Gauss elimination does not reduce the computational time in solving these equations.** The usual operations count approach suggests that the arithmetic should be halved; this does not happen. There is about a 40% savings in the total memory required for the computation. The computational efficiency using the collorder ordering is analyzed in Section V.

Our conclusions are given in Section VI along with two conjectures; one is about a variant of collorder and the other is, for elliptic problems in general, about the relationship between round-off effects and the orderings of unknowns in Gauss elimination.

II. THE COLLORDER ORDERING

Once the grid is placed over the rectangle R, the grid intersection points are numbered in a natural way from west to east, south to north. The collocation points (equations) are then associated with the nearest grid point and are numbered in groups of four in the order of their corresponding grid point. For example, with a 3×3 grid, there are 64 collocation points whose ordering is given in Figure 2.

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

FIGURE 2. The collorder ordering of the collocation points for a 3×3 grid with 64 unknowns.

The important idea in the above ordering is that the grid points are first numbered in a natural way and then the collocation points are numbered in the order of their corresponding grid point. We could number the grid points south to north, west to east as long as we suitably modify the rest of the ordering scheme described in the following sections. Similarly, the four collocation points associated with any grid point may be locally ordered in any manner.

To the $i^{\text{th}}[j^{\text{th}}]\, x[y]$ grid coordinate there are associated two Hermite cubic functions: $\varphi_i^x[\varphi_j^y]$ which takes on function values and $\psi_i^x[\psi_j^y]$ which takes on derivative values. The Hermite bicubic basis functions (unknowns) are then ordered corresponding to their support in a natural way from west to east, south to north. For example, with 9 sub-rectangles there are 64 bicubic basis

functions whose ordering is given in Figure 3.

49 $\varphi_1^x\varphi_4^y$	53 $\varphi_2^x\varphi_4^y$	57 $\varphi_3^x\varphi_4^y$	61 $\varphi_4^x\varphi_4^y$
50 $\varphi_1^x\psi_4^y$	54 $\varphi_2^x\psi_4^y$	58 $\varphi_3^x\psi_4^y$	62 $\varphi_4^x\psi_4^y$
51 $\psi_1^x\varphi_4^y$	55 $\psi_2^x\varphi_4^y$	59 $\psi_3^x\varphi_4^y$	63 $\psi_4^x\varphi_4^y$
52 $\psi_1^x\psi_4^y$	56 $\psi_2^x\psi_4^y$	60 $\psi_3^x\psi_4^y$	64 $\psi_4^x\psi_4^y$
33 $\varphi_1^x\varphi_3^y$	37 $\varphi_2^x\varphi_3^y$	41 $\varphi_3^x\varphi_3^y$	45 $\varphi_4^x\varphi_3^y$
34 $\varphi_1^x\psi_3^y$	38 $\varphi_2^x\psi_3^y$	42 $\varphi_3^x\psi_3^y$	46 $\varphi_4^x\psi_3^y$
35 $\psi_1^x\varphi_3^y$	39 $\psi_2^x\varphi_3^y$	43 $\psi_3^x\varphi_3^y$	47 $\psi_4^x\varphi_3^y$
36 $\psi_1^x\psi_3^y$	40 $\psi_2^x\psi_3^y$	44 $\psi_3^x\psi_3^y$	48 $\psi_4^x\psi_3^y$
17 $\varphi_1^x\varphi_2^y$	21 $\varphi_2^x\varphi_2^y$	25 $\varphi_3^x\varphi_2^y$	29 $\varphi_4^x\varphi_2^y$
18 $\varphi_1^x\psi_2^y$	22 $\varphi_2^x\psi_2^y$	26 $\varphi_3^x\psi_2^y$	30 $\varphi_4^x\psi_2^y$
19 $\psi_1^x\varphi_2^y$	23 $\psi_2^x\varphi_2^y$	27 $\psi_3^x\varphi_2^y$	31 $\psi_4^x\varphi_2^y$
20 $\psi_1^x\psi_2^y$	24 $\psi_2^x\psi_2^y$	28 $\psi_3^x\psi_2^y$	32 $\psi_4^x\psi_2^y$
1 $\varphi_1^x\varphi_1^y$	5 $\varphi_2^x\varphi_1^y$	9 $\varphi_3^x\varphi_1^y$	13 $\varphi_4^x\varphi_1^y$
2 $\varphi_1^x\psi_1^y$	6 $\varphi_2^x\psi_1^y$	10 $\varphi_3^x\psi_1^y$	14 $\varphi_4^x\psi_1^y$
3 $\psi_1^x\varphi_1^y$	7 $\psi_2^x\varphi_1^y$	11 $\psi_3^x\varphi_1^y$	15 $\psi_4^x\varphi_1^y$
4 $\psi_1^x\psi_1^y$	8 $\psi_2^x\psi_1^y$	12 $\psi_3^x\psi_1^y$	16 $\psi_4^x\psi_1^y$

FIGURE 3. The collorder ordering of the basis functions for a 3×3 grid with 64 unknowns.

This ordering of the collocation points and the basis functions yields an intermediate collorder matrix with a nice structure. It is a band matrix with band width $4n+11$. The matrix is block symmetric in that it consists of 4×4 blocks which if $B_{ij}\neq 0$ then $B_{ji}\neq 0$. There are at most 16 non-zero entries per row occurring in 4 blocks. All of the symmetric pairs of off-diagonal blocks can be stored in their natural order within one 4×4 block.

For example, in the case of a Poisson equation with Dirichlet boundary conditions and a 3×3 grid, the pattern of non-zeros in the intermediate collorder matrix is shown in Figure 4.

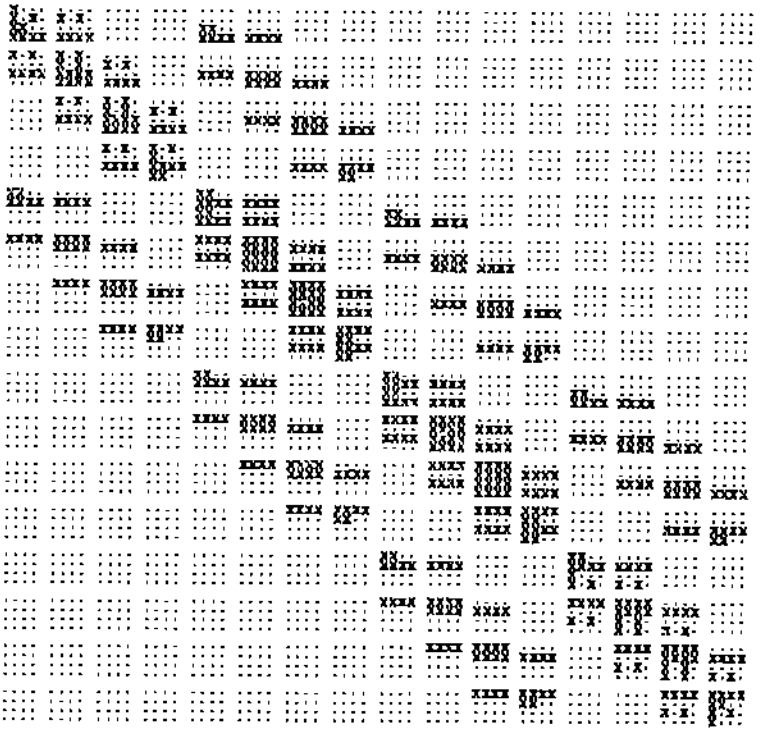

FIGURE 4. The intermediate collorder matrix for a 3×3 grid with 64 unknowns.

We see from Figure 4 that, unlike the usual collocation matrix structure of Figure 1, this matrix is tri-banded with the center band along the diagonal. However, the intermediate collorder matrix still has some zero diagonal elements corresponding to boundary conditions (e.g., rows 2 and 3).

We now outline how to further reorder this matrix to remove all zeros from the diagonal. There are eight types of boundary points which we label as follows:

Northwest	North	Northeast
West		East
Southwest	South	Southeast

From Figure 4 we see for Dirichlet boundary conditions that the 4×4 block in the upper left is associated with the southwest grid point and has the following format:

X 0 0 0
X 0 X 0
X X 0 0
X X X X

If we permute the second and third basis functions (columns 2 and 3), we obtain a non-zero diagonal in this block:

```
X  0  0  0
X  X  0  0
X  0  X  0
X  X  X  X
```

By considering all nine possible boundary conditions which may occur at the southwest corner, one can show that it is always possible to reorder the columns to obtain a non-zero diagonal in this diagonal block. The possible boundary conditions at the southwest corner with the corresponding permutations are given in Table 1.

TABLE 1. Column permutations at the southwest corner.

Southwest corner				
Boundary condition		Block before	Column permutation	Block after
West	South			
Dirichlet	Dirichlet	X000 X0X0 XX00 XXXX	2 → 3 3 → 2	X000 XX00 X0X0 XXXX
Neumann	Dirichlet	X000 X0X0 00XX XXXX	2 → 4 3 → 2 4 → 3	X000 XX00 0XX0 XXXX
Mixed	Dirichlet	X000 X0X0 XXXX XXXX	2 → 3 3 → 2	X000 XX00 XXXX XXXX
Dirichlet	Neumann	0X00 0X0X XX00 XXXX	1 → 3 2 → 1 3 → 4 4 → 2	X000 XX00 X0X0 XXXX
Neumann	Neumann	0X00 0X0X 00XX XXXX	1 → 4 2 → 1 4 → 2	X000 XX00 0XX0 XXXX
Mixed	Neumann	0X00 0X0X XXXX XXXX	1 → 4 2 → 1 4 → 2	X000 XX00 XXXX XXXX
Dirichlet	Mixed	XX00 XXXX XX00 XXXX	2 → 3 3 → 2	X0X0 XXXX X0X0 XXXX
Neumann	Mixed	XX00 XXXX 00XX XXXX	2 → 4 4 → 2	X00X XXXX 0XX0 XXXX
Mixed	Mixed	XX00 XXXX XXXX XXXX	2 → 3 3 → 2	X0X0 XXXX XXXX XXXX

By similarly considering all possible boundary conditions at every type of boundary grid point, one can show that it is always possible to reorder the basis functions to obtain a non-zero diagonal in the corresponding diagonal block. A

complete set of tables describing the reordering of the boundary basis functions is contained in Appendix A. The result for the matrix in Figure 4 is shown in Figure 5. Appendix B of (Dyksen, 1981) shows a complete set of example matrices with the collorder ordering.

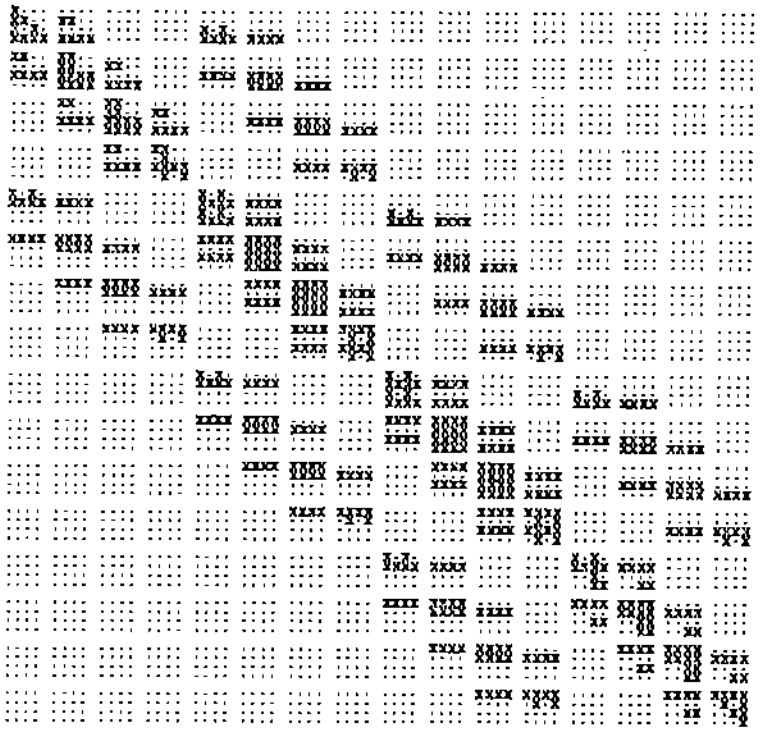

FIGURE 5. The collorder matrix for a 3 × 3 grid with 64 unknowns.

This ordering is implemented in the ELLPACK system as the indexing module HERMITE COLLORDER.

III. PIVOTING IN GAUSS ELIMINATION WITH THE COLLORDER ORDERING

We consider the following hypothesis: **Gauss elimination without pivoting is safe for the collorder ordering of the collocation equations.** We would, of course, like to establish this by analytical means but have been unable to do so. It involves an analysis of scaling in Gauss elimination, an area where little progress has been made so it is not too disappointing that we cannot analyze this particular case. We do not believe this phenomena is restricted to collocation; it is an instance of a more general phenomena in the numerical solution of elliptic problems.

We test this hypothesis in the ELLPACK environment by considering the following modules for solving elliptic problems:

HERMITE COLLOCATION
P3C1 COLLOCATION
HERMITE COLLORDER
BAND GE
BAND GE NO PIVOTING

The modules HERMITE COLLOCATION and P3C1 COLLOCATION generate the same

equations, except that HERMITE COLLOCATION scales the equations (as best we know how). BAND GE uses the LINPACK routines SGBFA and SGBSL modified to do scaled partial pivoting, while BAND GE NO PIVOTING is the same module without pivoting. HERMITE COLLORDER applies the collorder ordering to the equations.

A method is a combination of these modules and they are denoted using the following scheme:

S means use scaling = HERMITE COLLOCATION
C means use collorder ordering = HERMITE COLLORDER
P means use pivoting = BAND GE
NP means use no pivoting = BAND GE NO PIVOTING

If P3C1 COLLOCATION is used, then no letter is included in the method notation. Thus, methods SP and CNP are

HERMITE COLLOCATION + BAND GE.

P3C1 COLLOCATION + HERMITE COLLORDER + BAND GE NO PIVOTING

These methods were applied to the following 20 elliptic problems from the population of (Rice et al, 1981)

2-1, 3-1, 5-1, 6-1, 9-2, 10-2, 10-3, 12-3, 17-2, 20-2, 22-1, 23-6, 33-1, 35-3, 38-1, 40-1, 50-1, 53-3, 54-2, 59-1

There is a variety of operators in this set; 15 of these problems have Dirichlet boundary conditions (5 are homogeneous).

We restate the general hypotheses considered here: **Gauss elimination without pivoting is safe for the collorder ordering of the collocation equations.** We state below more specific, testable hypotheses for accuracy in terms of "method A" = "method B" along with an interpretation (pivoting always means scaled partial pivoting).

H1: SCNP = SP
For scaled equations, collorder and no pivoting is the same as pivoting

H2: CNP = P
For unscaled equations, collorder and no pivoting is the same as pivoting

H3: SCNP = SCP
For scaled equations, pivoting is not needed for the collorder ordering

H4: SCP = SP
For scaled equations, collorder has no effect if pivoting is used.

These four hypotheses are tested with the sample described above and accuracy measured as follows:

(a) subjective evaluation of accuracy from plots for each PDE,

(b) grid size required to achieve 5%, 0.5% and 0.05% accuracy in the solution.

The second set of measures is complicated by the fact that the accuracy is dominated by round-off effects for fine grids. However, we are fortunate in that all four measures give consistent results. We give data for measure (a) in Table 2 since it is simpler.

Hypothesis	No difference	First method better	Second method better
H1: SCNP = SP	13	5	2
H2: CNP = P	13	5	2
H3: SCNP = SCP	13	5	2
H4: SCP = SP	19	1	0

TABLE 2. Accuracy results in testing the four hypotheses with 20 PDE problems.

None of the differences observed in this experiment are statistically significant using the non-parametric Friedman, Kendall and Babington-Smith test (Hollander and Wolfe, 1973).

We have also constructed a special set of eight PDE problems whose exact solutions are bicubic polynomials. Every method under consideration here provides the exact solution except for round-off effects. For this population we studied the growth of round-off effects for solving the collocation equations in the collorder ordering using pivoting. We solved these problems with n varied from 2 to 28 (36 to 3364 unknowns) and computed the ratios of the maximum to minimum round-off effects; these results are summarized below.

PDE	Maximum round-off / Minimum round-off
1	27
2	17
3	13
4	10
5	30
6	17
7	12
8	13

For 4 of the problems there is a definite linear trend in the log-log plot of error versus n; the maximum slope is 0.75 suggesting that round-off grows at worst like $n^{0.75}$ with the collorder ordering and no pivoting. In 2 of the problems there is a hint of a linear trend and in the other 2 the round-off effect is constant over this range of problem sizes.

For comparison, we solved these same eight PDE problems with the unscaled equations using the LINPACK software, the round-off growth ratios observed ranged from 500 to 15,000 with an average of 7500. Even these growth ratios are small compared to what one would expect from the condition numbers of the linear systems which are in the range 10^7 to 10^{11} when $n=28$. Detailed data for three other problems is given in Table 3. The numbers in Table 3 differ from those given above where round-off growth is measured with respect to the minimum observed round-off and not in units in the last place.

TABLE 3. Round-off effects (in units in the last place) observed using the collorder ordering with the unscaled collocation equations.

Problem	Number of equations	Condition number	Round-off No pivoting	Pivoting
2-1	1156	$1.8*10^6$	50	400
	3364	$9.8*10^6$	500,000	2,000
22-1	1156	$2.2*10^{10}$	<20	500,000
	3364	$2.4*10^{11}$	80	10,000,000
59-1	1156	$1.0*10^8$	8	40,000
	3364	$9.1*10^8$	50	300,000

We now summarize other data and mention other work that is relevant to this question. We have computed the growth factor and maximum multiplier for Gauss elimination on the scaled equations with the collorder ordering. Typical results are tabulated; they indicate that the equations do not become more difficult as the size of the system grows.

	Number of equations 84	1228	4492
Growth factor with pivoting	1.21	1.21	-
Growth factor without pivoting	2.04	1.56	-
Maximum multiplier without pivoting	22.4	22.5	22.5

In (Trottenberg and Witsch, 1981) there is an analysis that shows the following. For a fixed continuous elliptic problem, the condition of the corresponding discretized problem is independent of h, the fineness of the discretization. Note that the condition of a particular numerical computation may be much worse than that of the problem itself. This suggests that there are numerical methods for elliptic problems which are numerically stable even for discretizations with a very large number of equations. One sample point to support this conclusion is seen in Figure 5(c) of (Rice, 1981) where essentially full machine accuracy is obtained by a high order finite difference plus FFT method involving over 65,000 unknowns. Another sample point is given in Example 7.6.2 in (Dyksen, 1982) where no round-off effect is observed for a tensor product discretization with 82,368 unknowns.

The study of (Hammarling and Wilkinson, 1980) concentrates primarily on one dimensional problems, but they do study the usual finite difference systems arising from the Laplace operator. They test whether the relative error grows like the condition number (about n^2 or $1/h^2$) of the linear system. The experimental data given for four problems stop at n=30 or 35 and do not show conclusive trends. A similar analysis and study appears in (Birkhoff and Lynch, 1983) and they also suggest that the error due to solving the finite difference equations also grows like n^2. Implicit in both these analyses is the use of ordinary Gauss elimination; the FFT example cited in the above paragraph shows that the error growth can be much less.

The error in solving the collocation equations is much less correlated with the condition number of the linear system as is seen from the data in Table 3.

For the collorder ordering the round-off grows at worse like $n^{0.75}$ while the condition number grows like n^4. Problem 5−1 provides another data point where we used the Galerkin method instead of collocation and then solved the linear system by the LINPACK software for symmetric positive definite band systems. The condition number of the linear system is $2.3*10^9$ but only 1600 units in the last place of accuracy are lost.

IV. WHY PIVOTING IS NOT NEEDED WITH THE COLLORDER ORDERING

The reason that pivoting is not needed with the collorder ordering is that this ordering is essentially the order obtained when one does scaled partial pivoting of the scaled equations. Recall that the role of scaling is to rationalize partial pivoting, so if one knows in advance what the partial pivoting order is, then one does not need to do scaling. We have tried without success to prove that the collorder ordering is as we state, so our statement is a conjecture rather than a fact. We have checked this conjecture on a large number of PDEs and found it to hold.

A typical example is shown in Figure 6 for $n=3$. At the top of Figure 6 we show the collocation points (equations) ordered two ways. At the bottom we give the order in which the equations are eliminated using scaled partial pivoting. We see that the elimination orders are almost identical to each other and to the collorder ordering. The differences in the elimination order and the collorder order is almost always restricted to permutation of the order of the four points associated with a grid point. The subordering of collorder used at a grid point is somewhat arbitrary and we believe the choice for this subordering to be somewhat irrelevant.

V. COMPUTATIONAL EFFICIENCY

For a grid of n^2 rectangles we have a linear system with $4(n+1)^2$ unknowns and band width $4n+11$. Standard operation counts for memory and arithmetic are (neglecting lower order terms).

	Multiplies	Memory
Elimination with pivoting	$64\,n^4$	$48\,n^3$
Elimination without pivoting	$32\,n^4$	$32\,n^4$

A naive implementation of Gauss elimination with pivoting would reflect these counts directly. We were surprised to find that BAND GE NO PIVOTING actually requires **more** computing time than BAND GE for these equations. The reason is that no pivoting leads to much more fill-in in the band; this fact is reflected in the number of calls on the routine SAXPY (from the BLAS). A specific tabulation for the elimination is given below in the case of $n=28$ (3364 equations).

Module used	Calls on SAXPY	Zero returns	Difference	Total seconds to run
BAND GE	410,866	104,979	305,887	767
BAND GE NO PIVOTING	406,146	40,785	365,361	846

FINITE ELEMENT

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

HERMITE COLLORDER

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

PARTIAL PIVOTING

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

PARTIAL PIVOTING

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

FIGURE 6. (Top) Two orderings of the scaled collocation equations. (Bottom) The order of elimination of the collocation equations using scaled partial pivoting.

Recall that SAXPY computes the vector sum $ax+y$ and is in the inner loop of Gauss elimination so the total computation time should be (and is) proportional to the number of its executions. A zero return from SAXPY does not involve any significant computation.

For a sample of 45 cases, we observed that no pivoting requires an average of 8% more execution time than pivoting with a range of 4 to 14% more. Pivoting requires about 40% more total memory than not pivoting.

VI. CONCLUSIONS AND CONJECTURES

The collorder ordering has several nice properties: (1) the matrix has a non-zero diagonal and is somewhat symmetric, (2) the linear system can be solved by elimination without pivoting or scaling, and the computation is exceptionally stable, (3) there is a savings in the elimination of about 40% in memory at the cost of 5–10% longer execution time.

We conjecture that: **it is completely safe to use a "near collorder" ordering where the equations are ordered according to the grid points together with local partial pivoting used on the 4 by 4 blocks associated with the grid points.** This would greatly simplify the algorithm to obtain the collorder ordering by eliminating the reordering associated with boundary conditions. It is not likely to increase the speed and would increase the memory requirement by a factor of $1+O(h)$.

We make a second and somewhat more amorphous conjecture: **it is essential for numerical stability that the elimination order follow a continuous geometric progression through the domain of the problem.** We believe this to be a general property of discretizations of elliptic problems. We have made various experiments with both finite elements and finite differences and observed that "disrupting" the geometric pattern during elimination usually produces numerical instability, sometimes of enormous magnitude. It is as though the round-off error satisfies an initial-value problem which is parabolic and smoothing if the geometry is followed properly and is hyperbolic and unstable if the geometry is not followed.

REFERENCES

Balart, R., Houstis, E. N. and Papatheodorou, T. S. (1982), On the iteration solution of collocation method equations, *Proc. 10th IMACS World Congress*, 98-100.

Birkhoff, G. and Lynch, R. E. (1983), "Numerical Solution of Elliptic Partial Differential Equations", SIAM Publications, Philadelphia.

Dongarra, J. J. et al (1979), "LINPACK User's Guide", SIAM Publications, Philadelphia.

Dyksen, W. R. (1981), A new ordering scheme for the Hermite bicubic collocation equations, CSD-TR 364, Computer Science Department, Purdue University.

Dyksen, W. R. (1982), Tensor product generalized alternating direction implicit methods for solving separable second order linear elliptic partial differential equations, Ph. D. Thesis, Purdue University.

Dyksen, W. R. and Rice, J. R. (1983), The importance of scaling in the collocation method.

Hammarling, S. and Wilkinson, J. H. (1980), On linear systems arising from finite difference approximations to elliptic differential equations, Rpt. DNACS 34/80, National Physical Laboratory, England.

Hollander, M. and Wolfe, D. A. (1973), "Non parametric Statistical Methods", John Wiley, New York.

Rice, J. R. (1981), ELLPACK: Progress and plans, in *Elliptic Problem Solvers* (M. Schultz, ed.) Academic Press, New York, 135-162.

Rice, J. R., Houstis, E. N. and Dyksen, W. R. (1981), A population of linear, second order elliptic partial differential equations on rectangular domains, *Math. Comp.*, **36**, 475-484.

Rice, J. R. (1983), "Numerical Methods, Software and Analysis", Chapter 10, McGraw-Hill.

Trottenberg, U. and Witsch, K. (1981), Zur Kondition diskreter elliptischer Randwertaufgaben, Preprint No. 457, GMD-Bonn.

APPENDIX A

The tables in this appendix contain the basis function (unknown, column) permutations to be performed on the diagonal blocks which correspond to the boundary grid points. There are eight different types of boundary boundary points which we label as follows:

Northwest	North	Northeast
West		East
Southwest	South	Southeast

A table for each type of boundary grid point follows.

Note: HERMITE COLLOCATION uses the south boundary condition at the southwest and southeast corners. It uses the west and east boundary conditons at the northwest and northeast corners, respectively.

Note: The column permutations given in the tables are local in the sense that they refer to the column number within the diagonal block. Hence, switching columns 2 and 3 within a given diagonal block may actually involve switching columns 45 and 46 in the intermediate collocation matrix.

The boundary condition labels given in the tables have the following meanings:

Label	Boundary condition
Dirichlet	$u = g$
Neumann	$\frac{\partial u}{\partial n} = g$
Mixed	$\alpha u + \beta \frac{\partial u}{\partial n} = g$

TABLE A-1. Column permutations at the southwest corner.

Southwest corner				
Boundary condition		Block before	Column permutation	Block after
West	South			
Dirichlet	Dirichlet	X000 X0X0 XX00 XXXX	2 → 3 3 → 2	X000 XX00 X0X0 XXXX
Neumann	Dirichlet	X000 X0X0 00XX XXXX	2 → 4 3 → 2 4 → 3	X000 XX00 0XX0 XXXX
Mixed	Dirichlet	X000 X0X0 XXXX XXXX	2 → 3 3 → 2	X000 XX00 XXXX XXXX
Dirichlet	Neumann	0X00 0X0X XX00 XXXX	1 → 3 2 → 1 3 → 4 4 → 2	X000 XX00 X0X0 XXXX
Neumann	Neumann	0X00 0X0X 00XX XXXX	1 → 4 2 → 1 4 → 2	X000 XX00 0XX0 XXXX
Mixed	Neumann	0X00 0X0X XXXX XXXX	1 → 4 2 → 1 4 → 2	X000 XX00 XXXX XXXX
Dirichlet	Mixed	XX00 XXXX XX00 XXXX	2 → 3 3 → 2	X0X0 XXXX X0X0 XXXX
Neumann	Mixed	XX00 XXXX 00XX XXXX	2 → 4 4 → 2	X00X XXXX 0XX0 XXXX
Mixed	Mixed	XX00 XXXX XXXX XXXX	2 → 3 3 → 2	X0X0 XXXX XXXX XXXX

TABLE A-2. Column permutations at the southeast corner.

Southeast corner				
Boundary condition		Block before	Column permutation	Block after
South	East			
Dirichlet	Dirichlet	X0X0 X000 XXXX XX00	1 → 2 2 → 4 3 → 1 4 → 3	XX00 0X00 XXXX 0X0X
Dirichlet	Neumann	X0X0 X000 XXXX 00XX	1 → 2 2 → 3 3 → 1	XX00 0X00 XXXX X00X
Dirichlet	Mixed	X0X0 X000 XXXX XXXX	1 → 2 2 → 3 3 → 1	XX00 0X00 XXXX XXXX
Neumann	Dirichlet	0X0X 0X00 XXXX XX00	1 → 4 4 → 1	XX00 0X00 XXXX 0X0X
Neumann	Neumann	0X0X 0X00 XXXX 00XX	1 → 3 3 → 4 4 → 1	XX00 0X00 XXXX X00X
Neumann	Mixed	0X0X 0X00 XXXX XXXX	1 → 4 4 → 1	XX00 0X00 XXXX XXXX
Mixed	Dirichlet	XXXX XX00 XXXX XX00	1 → 4 4 → 1	XXXX 0X0X XXXX 0X0X
Mixed	Neumann	XXXX XX00 XXXX 00XX	1 → 3 3 → 1	XXXX 0XX0 XXXX X00X
Mixed	Mixed	XXXX XX00 XXXX XXXX	1 → 4 4 → 1	XXXX 0X0X XXXX XXXX

TABLE A-3. Column permutations at the northwest corner.

Northwest corner				
Boundary condition		Block before	Column permutation	Block after
West	North			
Dirichlet	Dirichlet	XX00 XXXX X000 X0X0	1 → 3 2 → 1 3 → 4 4 → 2	X0X0 XXXX 00X0 00XX
Neumann	Dirichlet	00XX XXXX 00X0 X0X0	1 → 4 4 → 1	X0X0 XXXX 00X0 00XX
Mixed	Dirichlet	XXXX XXXX X0X0 X0X0	1 → 4 4 → 1	XXXX XXXX 00XX 00XX
Dirichlet	Neumann	XX00 XXXX X000 0X0X	1 → 3 2 → 1 3 → 2	X0X0 XXXX 00X0 X00X
Neumann	Neumann	00XX XXXX 00X0 0X0X	1 → 2 2 → 4 4 → 1	X0X0 XXXX 00X0 X00X
Mixed	Neumann	XXXX XXXX X0X0 0X0X	none	XXXX XXXX X0X0 0X0X
Dirichlet	Mixed	XX00 XXXX X000 XXXX	1 → 3 2 → 1 3 → 2	X0X0 XXXX 00X0 XXXX
Neumann	Mixed	00XX XXXX 00X0 XXXX	1 → 4 4 → 1	X0X0 XXXX 00X0 XXXX
Mixed	Mixed	XXXX XXXX X0X0 XXXX	none	XXXX XXXX X0X0 XXXX

TABLE A-4. Column permutations at the northeast corner.

Northeast corner				
Boundary condition		Block before	Column permutation	Block after
North	East			
Dirichlet	Dirichlet	XXXX XX00 X0X0 X000	1 → 4 4 → 1	XXXX 0X0X 00XX 000X
Dirichlet	Neumann	XXXX 00XX X0X0 00X0	1 → 3 2 → 1 3 → 4 4 → 2	XXXX 0X0X 00XX 000X
Dirichlet	Mixed	XXXX XXXX X0X0 X0X0	1 → 4 4 → 1	XXXX XXXX 00XX 00XX
Neumann	Dirichlet	XXXX XX00 0X0X X000	1 → 4 3 → 1 4 → 3	XXXX 0X0X 0XX0 000X
Neumann	Neumann	XXXX 00XX 0X0X 00X0	2 → 3 3 → 4 4 → 2	XXXX 0X0X 0XX0 000X
Neumann	Mixed	XXXX XXXX 0X0X X0X0	3 → 4 4 → 3	XXXX XXXX 0XX0 X00X
Mixed	Dirichlet	XXXX XX00 XXXX X000	1 → 4 4 → 1	XXXX 0X0X XXXX 000X
Mixed	Neumann	XXXX 00XX XXXX 00X0	2 → 3 3 → 4 4 → 2	XXXX 0X0X XXXX 000X
Mixed	Mixed	XXXX XXXX XXXX X0X0	1 → 4 4 → 1	XXXX XXXX XXXX 00XX

TABLE A-6. Column permutations on the south boundary.

North boundary			
Boundary condition	Block before	Column permutation	Block after
Dirichlet	XXXX XXXX X0X0 X0X0	1 → 4 4 → 1	XXXX XXXX 00XX 00XX
Neumann	XXXX XXXX 0X0X 0X0X	2 → 3 3 → 2	XXXX XXXX 00XX 00XX
Mixed	XXXX XXXX XXXX XXXX	2 → 3 3 → 2	XXXX XXXX XXXX XXXX

TABLE A-5. Column permutations on the north boundary.

South boundary			
Boundary condition	Block before	Column permutation	Block after
Dirichlet	X0X0 X0X0 XXXX XXXX	2 → 3 3 → 2	XX00 XX00 XXXX XXXX
Neumann	0X0X 0X0X XXXX XXXX	1 → 4 4 → 1	XX00 XX00 XXXX XXXX
Mixed	XXXX XXXX XXXX XXXX	2 → 3 3 → 2	XXXX XXXX XXXX XXXX

TABLE A-7. Column permutations on the west boundary.

West boundary			
Boundary condition	Block before	Column permutation	Block after
Dirichlet	XX00 XXXX XX00 XXXX	2 → 3 3 → 2	X0X0 XXXX X0X0 XXXX
Neumann	00XX XXXX 00XX XXXX	1 → 4 4 → 1	X0X0 XXXX X0X0 XXXX
Mixed	XXXX XXXX XXXX XXXX	2 → 3 3 → 2	XXXX XXXX XXXX XXXX

TABLE A-8. Column permutations on the east boundary.

East boundary			
Boundary condition	Block before	Column permutation	Block after
Dirichlet	XXXX XX00 XXXX XX00	1 → 4 4 → 1	XXXX 0X0X XXXX 0X0X
Neumann	XXXX 00XX XXXX 00XX	2 → 3 3 → 2	XXXX 0X0X XXXX 0X0X
Mixed	XXXX XXXX XXXX XXXX	2 → 3 3 → 2	XXXX XXXX XXXX XXXX