

**Parallel Methods for Solving Nonlinear
Block Bordered Systems of Equations¹**

Xiaodong Zhang², Richard H. Byrd³, Robert B. Schnabel³

CU-CS-454-89

December 1989

Department of Computer Science
Campus Box 430
University of Colorado at Boulder
Boulder, Colorado 80309-0430 USA

¹This research is partially supported by the Air Force Office of Scientific Research under AFOSR grant AFOSR-85-0251, NSF Cooperative Agreement CDA-8420944, and ARO grant DAAL-88-K-0086.

²Division of Mathematics and Computer Science, The University of Texas at San Antonio, San Antonio, Texas 78285-0664

³Department of Computer Science, Campus Box 430, University of Colorado, Boulder, Colorado 80309

Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of the National Science Foundation.

The findings in this report are not to be construed as an official Department of the Army position, unless so designated by other authorized documents.

Abstract

We discuss a group of parallel algorithms, and their implementation, for solving a special class of large sparse nonlinear equations. The type of sparsity occurring in these problems, which arise in VLSI design, structural engineering and many other areas, is called a *block bordered* structure. We describe an explicit method and several implicit methods for solving block bordered nonlinear problems, and make a mathematical analysis and computational comparisons of the two types of methods. Several variations and globally convergent modifications of the implicit method are also presented. We describe parallel algorithms for solving block bordered nonlinear equations, and present experimental results on the Intel hypercube that show the effectiveness of the parallel implicit algorithms. These experiments include a fairly large circuit simulation that leads to a multi-level block bordered system of nonlinear equations.

1 Introduction

1.1 Definition of block bordered nonlinear problems

In this paper we present a group of parallel algorithms, and their implementations, for solving a special class of large sparse nonlinear equations, instances of which occur in VLSI design, structural engineering and many other areas. The class of sparsity occurring in these problems is called a block bordered structure. In such a problem the general system of n nonlinear equations in n unknowns may be grouped into $q+1$ subvectors, x_1, \dots, x_{q+1} and f_1, \dots, f_{q+1} such that the nonlinear system of equations has the form

$$\begin{cases} f_i(x_i, x_{q+1}) = 0; & i = 1, \dots, q \\ f_{q+1}(x_1, \dots, x_{q+1}) = 0 \end{cases} \quad (1.1)$$

where

$$x_i \in R^{n_i},$$

$$f_i \in R^{n_i}, \quad i = 1, \dots, q+1,$$

$$\sum_{i=1}^{q+1} n_i = n.$$

The block bordered Jacobian matrix of (1.1) is

$$\begin{pmatrix} A_1 & & & B_1 \\ & A_2 & & B_2 \\ & & \cdot & \cdot \\ & & & \cdot \\ C_1 & C_2 & \cdot & \cdot & A_q & B_q \\ & & & & C_q & P \end{pmatrix} \quad (1.2)$$

where

$$A_i = \frac{\partial f_i}{\partial x_i} \in R^{n_i \times n_i}, \quad i = 1, \dots, q,$$

$$B_i = \frac{\partial f_i}{\partial x_{q+1}} \in R^{n_i \times n_{q+1}}, \quad i = 1, \dots, q,$$

$$C_i = \frac{\partial f_{q+1}}{\partial x_i} \in R^{n_{q+1} \times n_i}, \quad i = 1, \dots, q,$$

$$P = \frac{\partial f_{q+1}}{\partial x_{q+1}} \in R^{n_{q+1} \times n_{q+1}}.$$

Newton's method is the fundamental approach for solving a general nonlinear system of equations. Several parallel algorithms for solving general systems of nonlinear equations that are based upon Newton's method have been developed and implemented on various parallel computers. These parallel algorithms consist mainly of solving the linear Jacobian system in parallel. Many parallel algorithms have been developed for solving linear systems, such as parallel factorizations, parallel SOR methods, parallel red-black methods, parallel multicolor and others (see e.g. Ortega and Voigt [1985], O'Leary and White [1985], White [1986], Fontecilla [1987], Coleman and Li [1987]).

In the case of large sparse nonlinear problems, one cannot expect a single parallel algorithm to handle all the instances of the system of nonlinear equations problem efficiently, but rather the algorithm must take into account the sparsity structure and other special characteristics of the problem. Parallel algorithms taking advantage of this special structure may be much more efficient than algorithms ignoring the structure. This paper is an instance of developing special algorithms for a special, important structure.

1.2 Background on block bordered problems

Block bordered problems of the form (1.1) arise in many areas of science and engineering, and a few algorithms have been developed to solve linear block bordered systems of equations efficiently. In applications such as structural engineering, large spatial models may be divided into q regions such that each region only interacts directly with neighboring regions. The variables, x_i , for each region, are chosen so that the model can determine their values, given the values of the linking variables (the x_{q+1}) at the boundaries of the regions.

The linking variables are tied together by a $q + 1$ st set of equations representing the interactions between the regions. Thus the equilibrium equations for such a model will be of the form (1.1). In addition, the Jacobian matrix is symmetric. These problems, and parallel algorithms for solving the linear block bordered systems that arise from them, are discussed in Farhat and Wilson [1986], Nour-Omid and Park [1986].

Mu and Rice [1989] study parallel Gaussian elimination for the block bordered matrices arising from the discretization of partial differential equations (PDEs). Christara and Houstis [1988][1989] implement a domain decomposition spline collocation method, and a preconditioned conjugate gradient (PCG) method, for this linear block bordered system on both NCUBE/7 and Sequent multiprocessors.

All the work described above concerns parallel methods for solving linear block bordered equations. Our research is to develop, implement, and analyze parallel methods for solving nonlinear block bordered problems. To our knowledge, no one has done similar work. Chan[1985] considers methods for a class of nonlinear problems that includes (1.1), but his focus is on approximate solutions of the linearizations of these equations, and he does not specifically consider block bordered structure.

Block bordered equations also arise in VLSI circuit design, where parts of the circuits may be divided into regions. The concept of macromodeling the circuit is to decompose the circuit into subcircuits and to analyze the subcircuits separately (see Rabbat et al [1979], [1980]). Macromodeling of the circuit results in a system of nonlinear equations of form (1.1). A_i and B_i ($i = 1, \dots, q$) in the Jacobian matrix are usually used to represent internal and input-output variables in each of the q independent subcircuits. The variables represent voltages and currents. The bottom block f_{q+1} represents the voltages or currents between subcircuits. Since each voltage or current is used only in one block of equations f_i plus possibly the bottom block f_{q+1} , the nonzero columns of the B_i 's (and A_i) are disjoint, meaning that the matrices B_1, \dots, B_q in (1.2) also follow a block diagonal pattern. In addition, since f_{q+1} describes the point-to-point connections of voltage and current, it is a linear function. The size of the function f_{q+1} depends on the number of connections among the subcircuits.

We have studied parallel methods for solving block bordered nonlinear equations extensively from both theoretical and practical viewpoints. Section 2 presents an explicit method and several implicit methods for solving block bordered nonlinear equations, and gives some mathematical analysis and computational comparison of these two types of methods. Section 3 briefly discusses techniques used to make these methods globally convergent. In Section 4, we give a group of parallel algorithms for solving block bordered nonlinear systems of form (1.1) which may be implemented on both shared and distributed memory multiprocessors. The implementations and experimental results of these algorithms on the Intel hypercube, a distributed memory multiprocessor, are presented in Section 5. Finally, our conclusions and some future research directions are summarized in Section 6.

2 Explicit and Implicit Methods

2.1 Introduction

There are two basic ways in which Newton's method can be applied to the nonlinear block bordered system of equations (1.1). The explicit approach is to simply apply Newton's method to (1.1). This involves iteratively solving the linear system

$$J(X^k)\Delta X^k = -F(X^k) \quad i = 1, \dots, q \quad (2.1)$$

for ΔX^k , where $J(X^k)$ is the Jacobian of F , which has the block bordered structure (1.2).

The pure implicit approach is to use each of the q systems of nonlinear equations

$$f_i(x_i, x_{q+1}) = 0, \quad i = 1, \dots, q \quad (2.2)$$

to solve for x_i , given a fixed value of x_{q+1} . This means that each of the x_i is implicitly given by a function of x_{q+1} . The whole problem (2.2) is then equivalent to solving

$$f_{q+1}(x_1(x_{q+1}), \dots, x_q(x_{q+1}), x_{q+1}) = 0. \quad (2.3)$$

The Jacobian of this system is given by

$$\hat{J} = \frac{\partial f_{q+1}}{\partial x_{q+1}} - \sum_{i=1}^q \frac{\partial f_{q+1}}{\partial x_i} \left(\frac{\partial f_i}{\partial x_i} \right)^{-1} \frac{\partial f_i}{\partial x_{q+1}} \quad i = 1, \dots, q \quad (2.4)$$

or

$$\hat{J} = P - \sum_{i=1}^q C_i A_i^{-1} B_i \quad i = 1, \dots, q \quad (2.5)$$

and we may solve (2.3) by Newton's method. We will be considering practical variants of the implicit method that do not calculate $x_i(x_{q+1})$ exactly.

In this section we describe the explicit and implicit methods and their relations to each other, and give some simple experimental results using them on a sequential computer. These results give a preliminary indication of why the implicit method will be advantageous on parallel computers.

2.2 Algorithms and analysis for the explicit and implicit methods

Newton's method applied to (1.1) in the explicit method consists of solving the following linear equations at iteration k ($k = 0, 1, \dots$): from $f_i(x) = 0$, $i = 1, \dots, q$,

$$A_i \Delta x_i^k + B_i \Delta x_{q+1}^k + f_i(x_i^k, x_{q+1}^k) = 0 \quad (2.6)$$

and from $f_{q+1}(x_1^k, \dots, x_q^k, x_{q+1}^k) = 0$

$$\sum_{i=1}^q C_i \Delta x_i^k + P \Delta x_{q+1}^k + f_{q+1}(x_1^k, \dots, x_q^k, x_{q+1}^k) = 0. \quad (2.7)$$

(For simplicity we are omitting superscripts k on A_i, B_i, C_i and P , but note that at least the A_i 's and B_i 's can change at each iteration.) Solving for Δx_i^k in (2.6) and substituting into (2.7), we obtain

$$\hat{J} \Delta x_{q+1}^k = -f_{q+1}(x_1^k, \dots, x_q^k, x_{q+1}^k) + \sum_{i=1}^q C_i A_i^{-1} f_i(x_i^k, x_{q+1}^k) \quad (2.8)$$

where \hat{J} is given by (2.5). (We assume for now that \hat{J} and each A_i is non-singular.) So

$$x_{q+1}^{k+1} = x_{q+1}^k + \Delta x_{q+1}^k \quad (2.9)$$

can be determined from (2.8), and

$$x_i^{k+1} = x_i^k + \Delta x_i^k \quad i = 1, \dots, q \quad (2.10)$$

can be determined from (2.6).

In the pure implicit method, Newton's method is applied to (2.3), and gives

$$\hat{J} \Delta x_{q+1}^k + f_{q+1}(x_1(x_{q+1}^k), \dots, x_q(x_{q+1}^k), x_{q+1}^k) = 0 \quad (2.11)$$

where $x_i(x_{q+1}^k)$ ($i = 1, \dots, q$) is implicitly determined by solving the nonlinear system

$$f_i(x_i, x_{q+1}^k) = 0 \quad (2.12)$$

for x_i . To turn this into a practical computational procedure, we use a second (or inner) iterative Newton process on (2.12) to calculate an $x_i(x_{q+1}^k)$ which solves (2.12) approximately. For each $i = 1, \dots, q$, this yields the inner iterations

$$\hat{A}_i \Delta x_i^{k,j-1} + f_i(x_i^{k,j-1}, x_{q+1}^k) = 0 \quad i = 1, \dots, q, \quad j = 1, 2, \dots, I_{in}. \quad (2.13)$$

Here $x_i^{k,0} = x_i^k$, I_{in} is the number of inner iterations, and $\hat{A}_i = A_i$ if it is only evaluated once at the beginning of each outer iteration, else it may be evaluated up to I_{in} times. At the end of each inner iteration, we set

$$x_i^{k,j} = x_i^{k,j-1} + \Delta x_i^{k,j-1}, \quad i = 1, \dots, q. \quad (2.14)$$

When we exit the inner iterations, we set

$$x_i(x_{q+1}^k) = x_i^{k+1} = x_i^{k,j}. \quad (2.15)$$

Then, x_{q+1} is determined from (2.11), and

$$x_{q+1}^{k+1} = x_{q+1}^k + \Delta x_{q+1}^k. \quad (2.16)$$

The following theorems show that the explicit method and the implicit methods just described are very closely related.

Theorem 1 If f_{q+1} is linear and only one inner Newton iteration is applied to solve for x_i^{k+1} ($i = 1, \dots, q$) in the implicit method, i.e. $I_{in} = 1$, then for a fixed k , the steps Δx_{q+1}^k are identical in both methods.

Proof In this case $x_i^{k+1} = x_i^k - A_i^{-1} f_i(x_i^k, x_{q+1}^k)$, and $f_{q+1}(x_1^{k+1}, \dots, x_q^{k+1}, x_{q+1}^k) = f_{q+1}(x_1^k, \dots, x_q^k, x_{q+1}^k) - \sum_{i=1}^q C_i A_i^{-1} f_i(x_i^k, x_{q+1}^k)$. Substituting this into (2.11) gives

$$\hat{J} \Delta x_{q+1}^k = -f_{q+1}(x_1^k, \dots, x_q^k, x_{q+1}^k) + \sum_{i=0}^q C_i A_i^{-1} f_i(x_i^k, x_{q+1}^k) \quad (2.17)$$

which is identical to the explicit formula (2.8). \square

In the implicit method described above, the steps Δx_i for $i \leq q$ do not involve any information about f_{q+1} or Δx_{q+1} , and are not the same as in the explicit method. For this reason, the method is not one-step quadratically convergent. If the value of x_i^{k+1} ($i \leq q$) calculated by the implicit method is corrected after each iteration to account for the change in x_{q+1} , however, the implicit method can be made closer to the explicit method and quadratically convergent. The problem may be defined to find a correction term δ such that

$$f_i(x_i^{k+1} + \delta, x_{q+1}^{k+1}) \approx 0 \quad (2.18)$$

or

$$f_i(x_i^{k+1} + \delta, x_{q+1}^k + \Delta x_{q+1}^k) \approx 0. \quad (2.19)$$

Making a linear approximation to f_i in (2.19) yields the condition

$$f_i(x_i^{k+1}, x_{q+1}^k) + A_i \delta + B_i \Delta x_{q+1}^k = 0. \quad (2.20)$$

The correction term δ obtained from (2.20) would then be

$$\delta = -A_i^{-1} [f_i(x_i^{k+1}, x_{q+1}^k) + B_i \Delta x_{q+1}^k]. \quad (2.21)$$

However, after I_{in} inner iterations of solving for x_i^{k+1} , $f_i(x_i^{k+1}, x_{q+1}^k) \approx 0$. Thus we make a further approximation giving the correction term

$$\delta_i = -A_i^{-1} B_i \Delta x_{q+1}^k. \quad (2.22)$$

We name the implicit method with this correction term added to each x_i^{k+1} after Δx_{q+1}^k is calculated the *corrected implicit method*. The cost of the correction term (2.22) is small since the matrices $A_i^{-1} B_i$ ($i = 1, \dots, q$) have been calculated already and in a parallel implementation, the matrix-vector products can be parallelized fully.

We can now see that when f_{q+1} is linear the explicit method is a special case of the corrected implicit method.

Theorem 2 If f_{q+1} is linear and only one Newton iteration is applied to solve for $x_i^{k,j}$ ($i = 1, \dots, q$) in the implicit method, i.e. $I_{in} = 1$, and the system is corrected by adding $-A_i^{-1}B_i\Delta x_{q+1}$ to x_i^{k+1} ($i \leq q$) after each iteration, then the explicit method and implicit method are identical.

Proof From Theorem 1, Δx_{q+1}^k is identical for the two methods. Combining (2.14) with $j = 1$ and (2.22) gives:

$$x_i^{k+1} = x_i^k - A_i^{-1}[f_i(x_i^k, x_{q+1}^k) - B_i\Delta x_{q+1}^k] \quad (2.23)$$

which is identical to (2.6) in the explicit method. \square

Corollary If f_{q+1} is linear, the corrected implicit method with $I_{in} = 1$ inner iterations per outer iteration is locally quadratically convergent to the solution.

Quadratic convergence can also be shown for $I_{in} > 1$ and when f_{q+1} is nonlinear. The proof is given in Zhang[1989].

2.3 Some experiments on a sequential processor

The previous subsection shows that a variant of the implicit method is equivalent to the explicit method, but doesn't indicate why the implicit method might be preferred. The main reason turns out to be that, by using more than one inner iteration per outer iteration in the implicit method, the number of outer iterations can be reduced substantially, which is advantageous especially for parallel computation. In this subsection we give a first indication of the sort of computational behavior that we have found.

We initially tested the methods discussed in this section on several artificial problems. Here we report results on a simple 20×20 nonlinear block bordered system of equations which has four 4×4 blocks, A_1, \dots, A_4 , and a 4×4 bottom block P , and f_{q+1} linear. In all cases, the starting value of x

was close to the solution, and no global strategy (e.g. line search) was used. All these experiments were run on Pyramid P90 computer. The equations are given in Zhang[1989].

First, we compare the performance of the three methods when only one inner iteration ($I_{in} = 1$) is used in the uncorrected implicit and corrected implicit methods (Table 2.1). The explicit method and the corrected implicit method with $I_{in} = 1$ are identical in this case (see Theorem 2). Thus, the same number of iterations are required to converge to the solutions. The computing times are slightly different since the implementations of the two methods are different. The uncorrected implicit method converges a little bit slower than the other two methods, which is reasonable since the correction step is needed to make it one-step quadratically convergent.

Table 2.1: Experiments with the Three Methods

$(I_{in} = 1)$ outer iterations (seconds)		
explicit	implicit	corrected implicit
13 (0.44)	14 (0.40)	13 (0.40)

Next we increased the number of inner iterations in the uncorrected and corrected implicit methods. The experimental results (Tables 2.2 and 2.3) show that the number of outer iterations is sharply decreased when the number of inner iterations is 2. However, the number of outer iterations decreases more slowly as I_{in} increases further. There exists an optimal value of I_{in} for computing time in both methods, but it is problem dependent. Our experiments also show that the corrected implicit method converges a little bit faster than the uncorrected implicit method when $I_{in} > 1$, which is consistent with our convergence analysis. In Section 5 we will see that for larger problems, the improvements in time for the corrected implicit method with $I_{in} > 1$ can be considerably larger than those seen here. Also, we will see in Sections 4 and 5 that the decrease in outer iterations is advantageous for parallel computation.

Table 2.2: Experiments with the (Uncorrected) Implicit Method

$(I_{in} > 1)$ outer iterations (seconds)				
$I_{in} = 1$	$I_{in} = 2$	$I_{in} = 3$	$I_{in} = 4$	$I_{in} = 5$
14 (0.40)	8 (0.34)	7 (0.40)	6 (0.44)	6 (0.54)

Table 2.3: Experiments with the Corrected Implicit Method

$(I_{in} > 1)$ outer iterations (seconds)				
$I_{in} = 1$	$I_{in} = 2$	$I_{in} = 3$	$I_{in} = 4$	$I_{in} = 5$
13 (0.40)	8 (0.38)	6 (0.36)	6 (0.50)	5 (0.54)

3 Globally Convergent Modifications of the Explicit and Implicit Methods

The explicit method and corrected implicit method are locally quadratically convergent to the solution of (1.1). In other words, when the initial solution approximation is good enough, these methods are guaranteed to converge rapidly to a root of (1.1). However, it is often hard to find a good initial approximation for nonlinear systems of equations in practice. In addition, many practical problems, such as the circuit equations, are highly nonlinear, and if the current solution approximation is not close enough, a Newton step may easily result in an increase in the function norm. For example, a small change in some voltage difference in a nonlinear circuit equation may result in a great change in an exponential term in a diode or transistor's function evaluation. Also, many block bordered equations result in singular or nearly singular Jacobians in the process of the iterations, for example because a transistor with an exponential model is turned on at a nearly flat function curve (see Zhang, Byrd, Schnabel [1989]).

For these reasons, the explicit and implicit methods need to be modified to handle unacceptable steps and singular Jacobian matrices in order to converge to a solution from poor starting points. In this section we briefly describe the modifications we have used, which are motivated in part by their appropriateness for parallel, distributed computation. They are described in more detail in Zhang [1989]. A global convergence analysis will be given in a forthcoming paper.

We achieve convergence from poor starting points by using a line search. The explicit method is just a standard Newton's method, so we can use a standard line search. That is, we calculate the overall step direction $d^k = (\Delta x_1^k \dots \Delta x_q^k, \Delta x_{q+1}^k)$ as described in Section 2.2, and then set

$$X^{k+1} = X^k + \lambda^k d^k$$

where the steplength parameter $\lambda^k > 0$ is chosen by a line search procedure that assures sufficient descent on $\|F(x)\|_2$. Our line search is based upon Algorithm A6.3.1 in Dennis and Schnabel [1983].

The implicit method is more complicated since we have both inner and outer iterations. We need to choose the steps $\Delta x_i^{k,j} = (x_i^{k,j+1} - x_i^{k,j})$ in the inner iterations so that the overall step direction

$$d^k = (\sum_{j=0}^{I_{in}-1} \Delta x_1^{k,j} + \delta_1^k, \dots, \sum_{j=0}^{I_{in}-1} \Delta x_q^{k,j} + \delta_q^k, \Delta x_{q+1}^k), \quad (3.1)$$

where δ_i^k is the correction step (2.22), is a descent direction on $\|F(x)\|_2$. In addition, we would like the calculation of the steps $\Delta x_i^{k,j}$ for different values of i to be independent, so that the calculations of the inner iterations can be parallelized easily and efficiently.

Zhang [1989] shows that if each $\Delta x_i^{k,j}$ is calculated by

$$\Delta x_i^{k,j} = -\lambda_i^{k,j} A_i^{-1} f(x_i^{k,j}, x_{q+1}^k), \quad (3.2)$$

(i.e. the step discussed in Section 2.2 multiplied by a line search parameter $\lambda_i^{k,j} > 0$), the corrections steps δ_i^k are calculated by (2.22), Δx_{q+1}^k is calculated by (2.11) as before, and f_{q+1} is linear, then d^k given by (3.1) satisfies

$$J(x^k)d^k = \left(\sum_{i=0}^{I_{in}-1} \lambda_1^{k,j} f_1(x_1^{k,j}, x_{q+1}^k), \dots, \sum_{i=0}^{I_{in}-1} \lambda_q^{k,j} f_q(x_q^{k,j}, x_{q+1}^k), f_{q+1}(x_1^k, \dots, x_q^k, x_{q+1}^k) \right).$$

From this, it is straightforward to show that a sufficient condition for d^k given by (3.1) to be a descent direction on $\|f(x)\|_2$ is that for each $i = 1, \dots, q$,

$$\sum_{j=0}^{I_{in}-1} \lambda_i^{k,j} f_i(x_i^{k,j}, x_{q+1}^k)^T f(x_i^{k,0}, x_{q+1}^k) < 0. \quad (3.3)$$

Note that (3.3) is always true for $I_{in} = 1$ (since each $\lambda_i^{k,0} > 0$), and that it is true for $I_{in} = 2$ if $\|f(x_i^{k,1}, x_{q+1}^k)\| < \|f(x_i^{k,0}, x_{q+1}^k)\|$ (which any line search will enforce) and $\lambda_i^{k,1} \leq \lambda_i^{k,0}$. Since (3.3) can be monitored independently for each i , the following strategy will guarantee that a descent direction is generated. For each j , the procedure calculates each $\Delta x_i^{k,j}$ by (3.2) using a standard line search as mentioned above, and then checks whether the corresponding partial sum of (3.3) is satisfied. If it is not, it sets $\Delta x_i^{k,l} = 0$ for $l = j, \dots, I_{in} - 1$ and exits the inner iteration for x_i . The outer line search can be performed as in the explicit method.

Our approach for dealing with (nearly) singular Jacobians is based upon the Levenberg-Marquardt approach as described in Dennis and Schnabel [1983]. For a general system of nonlinear equations, if the current Jacobian matrix J is (nearly) singular, this approach modifies the search direction to be $-(J^T J + \mu I)^{-1} J^T F$, where F is the current function value, and μ is a small positive number. This direction is a descent direction on $\|F(x)\|_2$ and is the solution to the trust region problem

$$\underset{d}{\text{minimize}} \quad \|F + Jd\|_2 \quad \text{subject to} \quad \|d\|_2 \leq \Delta$$

for some $\Delta > 0$. In the limit as $\mu \rightarrow 0$, this direction equals $-J^+ F$, where J^+ is the pseudoinverse of J .

In the explicit and implicit methods described in Section 2, we need to solve systems of linear equations using the matrices A_1, \dots, A_q and the matrix $\hat{J} = P - \sum_{i=1}^q C_i A_i^{-1} B_i$. If any of these matrices, say M , is nearly singular (i.e., either the factorization detects numerical singularity or the estimated condition number of M is greater than $\text{macheps}^{-1/2}$) we simply

replace M^{-1} by $(M^T M + \mu I)^{-1} M^T$ in the formulas as of Section 2, where μ is chosen by the strategy suggested by Dennis and Schnabel [1983] and is a function of M . These perturbations again have interpretations in terms of trust regions. Note also that the algorithms for deciding whether to perturb each A_1, \dots, A_q , and for perturbing them if necessary, are totally independent so that they can be performed in parallel.

Combining these perturbation techniques with the inner line searches in a way that assures descent at the outer iteration and global convergence is somewhat more complex, and will be addressed in Byrd, Feng, Schnabel, and Zhang [1990]. In our implementations, we have simply taken I_{in} inner iterations for each block $i, i = 1, \dots, q$. We have used a standard line search to choose each $\lambda_i^{k,j}$ (requiring sufficient descent on f_i) but have not checked a condition like (3.3) that assures global descent. To our knowledge, the algorithm has still always produced a descent direction.

The algorithm we implement is summarized below. If \hat{J} or any A_i below is (nearly) singular, \hat{J}^{-1} or A_i^{-1} is replaced by $(\hat{J}^T \hat{J} + \mu I)^{-1} \hat{J}^T$ or $(A_i^T A_i + \mu_i I)^{-1} A_i^T$ for a small positive μ or μ_i , respectively. When f_{q+1} is linear, the explicit method is just a special case with $I_{in} = 1$ and each $\lambda_i^{k,1} = 1$.

Implicit Method with Global Modification

1. For $j = 0, \dots, I_{in} - 1$, calculate $x_i^{k,j+1} = x_i^{k,j} - \lambda_i^{k,j} A_i^{-1} f(x_i^{k,j}, x_{q+1}^k)$ where $\lambda_i^{k,j} \geq 0, i = 1, \dots, q$.
2. Form and factor $\hat{J} = P - \sum_{i=1}^q C_i A_i^{-1} B_i$.
3. Calculate $\Delta x_{q+1}^k = -\hat{J}^{-1} f_{q+1}(\bar{x}_1^{k+1}, \dots, \bar{x}_q^{k+1}, x_{q+1}^k)$ where $\bar{x}_i^{k+1} = x_i^{k, I_{in}}$.
4. Calculate the corrections $\delta_i^k = -A_i^{-1} B_i \Delta x_{q+1}^k$ and set $\bar{x}_i^{k+1} = \bar{x}_i^{k+1} + \delta_i^k, i = 1, \dots, q$.
5. Calculate $X^{k+1} = X^k - \lambda^k d^k$ where $d^k = (\bar{x}_1^{k+1} - x_1^k, \dots, \bar{x}_q^{k+1} - x_q^k, \Delta x_{q+1}^k)$.

4 Parallel Explicit and Implicit Algorithms

4.1 Motivation - LU factorization of block bordered linear equations

Note that the LU factorization of the block bordered Jacobian matrix

$$\begin{pmatrix} A_1 & & & B_1 \\ & A_2 & & B_2 \\ & & \ddots & \vdots \\ & & & A_q & B_q \\ C_1 & C_2 & \dots & C_q & P \end{pmatrix}$$

is

$$\begin{pmatrix} L_1 & & & & & \\ & L_2 & & & & \\ & & \ddots & & & \\ & & & L_q & & \\ \hat{C}_1 & \hat{C}_2 & \dots & \hat{C}_q & L_{q+1} & \end{pmatrix} \begin{pmatrix} U_1 & & & & & \hat{B}_1 \\ & U_2 & & & & \hat{B}_2 \\ & & \ddots & & & \vdots \\ & & & & & \hat{B}_q \\ & & & & U_q & \hat{B}_{q+1} \\ & & & & & U_{q+1} \end{pmatrix}$$

where for $i = 1, \dots, q$,

$$\begin{aligned} A_i &= L_i U_i \\ \hat{B}_i &= L_i^{-1} B_i \\ \hat{C}_i &= C_i U_i^{-1} \end{aligned}$$

and

$$L_{q+1} U_{q+1} = \hat{J} = P - \sum_{i=1}^q C_i A_i^{-1} B_i = P - \sum_{i=1}^q \hat{C}_i \hat{B}_i.$$

(This is the same matrix \hat{J} as in Section 2.) The calculations of L_i, U_i, \hat{B}_i , and \hat{C}_i for each i are independent, and thus can be computed very efficiently in parallel. The factorization of \hat{J} must follow these calculations and will not parallelize as efficiently, especially on distributed memory multiprocessors, because it will require considerable communication.

A parallel version of the explicit method essentially consists of performing the above factorization in parallel at each iteration. The parallel version of the implicit method that we discuss next will be seen to perform closely related operations. The major difference will be that, by performing more than one inner iteration per outer iteration, it will spend a larger portion of its time on the calculations that parallelize very efficiently, those for blocks 1, ..., q , and a smaller portion of its time on the calculations that parallelize less well, the formation and factorization of \hat{J} and the outer line search. Thus the implicit method can be expected to parallelize more effectively than the explicit method, especially on distributed memory computers. If the two methods require similar amounts of time on sequential computers, as indicated in Section 2, then the implicit method can be expected to be faster on parallel computers.

4.2 Parallel Algorithms

Below we give a general description of a parallel corrected implicit method that is based upon the sequential method presented in Sections 2 and 3. The parallelism comes mainly from executing all the operations on blocks 1 through q , which have been designed to be independent, concurrently. The parallel explicit method is just the special case with $I_{in} = 1$ and no inner line search.

Inner Iterations

1. For $i = 1, q$, Do in parallel:
 - 1.1 Factor A_i and estimate its condition number $Cond(A_i)$
 - 1.2 If $Cond(A_i) \leq Tol$ then set $M_i = A_i$, $N_i = I$
 Else choose $\mu_i > 0$, form and factor $M_i = A_i^T A_i + \mu_i I$,
 set $N_i = A_i^T$
 - 1.3 For $j = 0, I_{in} - 1$, Do:
 Solve $M_i \Delta x_i^{k,j} = -N_i f_i(x_i^{k,j}, x_{q+1}^k)$ for $\Delta x_i^{k,j}$
 Inner line search: $x_i^{k,j+1} = x_i^{k,j} + \lambda_i^{k,j} \Delta x_i^{k,j}$ for some $\lambda_i^{k,j} > 0$
 - 1.4 Solve $M_i W_i = N_i B_i$ for W_i
 - 1.5 Calculate $T_i = C_i W_i$

Outer Iteration

- *2. Form $\hat{J} = P - \sum_{i=1}^q T_i$
- *3. Factor \hat{J} and estimate its condition number $Cond(\hat{J})$
- *4. If $Cond(\hat{J}) \leq Tol$ then set $M = \hat{J}$, $N = I$
Else choose $\mu > 0$, form and factor $M = \hat{J}^T \hat{J} + \mu I$, set $N = \hat{J}^T$
- *5. Solve $M \Delta x_{q+1}^k = -N f_{q+1}(x_1^{k, I_{in}}, \dots, x_q^{k, I_{in}}, x_{q+1}^k)$ for Δx_{q+1}^k .
- 6. For $i = 1, q$, Do in parallel:
Calculate corrections $\delta_i^k = -W_i \Delta x_{q+1}^k$ and set $\bar{x}_i^{k+1} = x_i^{k, I_{in}} + \delta_i^k$
- *7. Outer line search: $x^{k+1} = x^k + \lambda^k (\bar{x}_1^{k+1} - x_1^k, \dots, \bar{x}_q^{k+1} - x_q^k, \Delta x_{q+1}^k)$
for some $\lambda^k > 0$.

The steps marked with stars requires synchronization (on a shared memory multiprocessor) or communication (on a distributed memory multiprocessor). Step 2 requires synchronization if the matrices T_i are full. In the VLSI problems, however, the nonzero columns of B_i , and hence T_i , are disjoint (see Section 1.2) and hence step 2 can be performed efficiently in parallel.

On shared memory machines, steps 3-5 can be performed in parallel using standard parallel methods for solving linear equations. On a distributed memory machine, it will only be efficient to perform steps 3-5 in parallel if the dimension of \hat{J} is rather large. In our test problems, \hat{J} was fairly small, so we performed steps 3-5 on one processor, on which we kept P , \hat{J} , and x_{q+1} . The remaining data was distributed in the obvious way: A_i , B_i , C_i , and x_i were stored together on the processor that handled block i . Step 7 includes two main operations, the calculation of trial points \bar{x}^{k+1} and the evaluations of F at the points, that are performed in parallel on a shared memory machine, and may be performed sequentially or in parallel on a distributed memory machine depending on their costs relative to the cost of communication.

5 Experimental Results on a Hypercube Multiprocessor

5.1 The test problem: a nonlinear block bordered circuit equation

The nonlinear block bordered application we considered for testing our methods is the VLSI circuit simulation problem. Standard circuit simulation methods consist of stiffly stable implicit integration formulae to discretize the differential equations, Newton's method to solve the resulting nonlinear algebraic equations,

$$F(X) = 0, \quad (5.1)$$

and sparse LU decomposition to solve the linear equations that arise at each iteration of Newton's method,

$$J\Delta X = -F(X), \quad (5.2)$$

where $J \in R^{n \times n}$ is the Jacobian matrix of (5.1). Typically, less than 2 percent of the entries of J are nonzero for $n > 500$ (see e.g. Sangiovanni-Vincentelli and Webber [1986]). The Newton iteration for (5.1) is repeated until a root is found or the upper bound on the number of iterations is reached. The program then decides whether to accept the solution, based on its estimate of local truncation error and the number of iterations required.

As mentioned in Section 1.2, partitioning the circuit leads to a block bordered system of nonlinear equations of the form (1.1) (see e.g. Rabbat et al [1979]). Given a circuit network Γ , a group of partitioned subnetworks $\gamma_i, i = 1, \dots, q$, and the connecting current and voltage equations, the block bordered nonlinear system of equations is defined as follows. Currents between two subnetworks and voltages at the boundary are each represented by two variables, one in each subnetwork, which are set equal to each other by the equations of f_{q+1} . Variables $x_i (i = 1, \dots, q)$ are used to represent internal voltage and current variables in each of the q independent subnetworks. They also include the current connecting variables among the q subnetworks. The variable x_{q+1} is used to represent the voltage connecting variables among the q subnetworks. Equations f_i represent the current

equations for subnetwork γ_i and its connecting variables. These equations for voltages and currents are standard current equations involving resistors, transistors, diodes, voltage sources and other elements. Since the connecting equation f_{q+1} is linear, the coefficient matrices C_i , $i = 1, \dots, q$ for the current connecting functions are constant, and the coefficient matrix P for the voltage connecting function is also constant.

For a very large circuit, the network Γ may be divided into subnetworks recursively, which leads to a multi-level block bordered system of nonlinear equations. In such a case, the diagonal blocks A_i , ($i = 1, \dots, q$) are themselves block bordered matrices. The border elements of the multilevel system represent the connections of the highest level.

We first applied our algorithm to the circuit shown in the figure in Appendix A. It is the 741 op-amp circuit (see e.g. Sedra and Smith [1982]), which was introduced in 1966 and is currently produced by almost every analog semiconductor manufacturer. The circuit is partitioned into 4 parts with a roughly equal number of nodes in each sub-circuit. A transistor is viewed as a nonlinear three terminal device in the circuit. Thus, applying the Ebers-Moll transistor model (see Ebers and Moll [1954]), 24, 27, 23 and 27 KCL functions are defined in the 1st, 2nd, 3rd and 4th block respectively. The 7 connections among the 4 blocks result in 14 linear current and voltage connecting functions. The total number of variables is $24 + 27 + 23 + 27 + 14 = 115$. The structure of the block bordered Jacobian matrix and the linear connecting border is shown in Appendix B.

We also applied our algorithm to a large analog filter, shown in Appendix C, that is composed of three 741 op-amp circuits (see e.g. Smith [1971], Valkenburg [1982]). This circuit leads to a 2-level block-bordered nonlinear system, as follows. The analog filter is first partitioned into 3 parts, each of which contains one 741 op-amp circuit. The first level block bordered structure is thus formed with 3 diagonal blocks and 1 connecting block. Each of the diagonal blocks is a 741 op-amp circuit which is partitioned into the second level block bordered structure.

5.2 The 741 op-amp Circuit Simulation on the Intel Hypercube

The nonlinear block bordered equations of the 741 op-amp circuit were solved in parallel on an Intel iPSC1 hypercube using the algorithm of Section 4.2. The 4 blocks of the circuit were distributed among 4 nodes of the hypercube. For convenience, the steps involving the connection function f_{q+1} (steps 3-5 of the parallel algorithm) were performed on a different node which plays the control role. They could just as well have been done on one of the 4 nodes. Identical initial values were used as the inputs for all the above experiments, and the convergence tolerances were also the same for those experiments. The solutions of the experiments were verified by comparing them to the solutions computed by the program SPICE which is a general-purpose circuit simulation program for nonlinear dc, nonlinear transient, and linear ac analysis (see Newton, Pederson and Sangiovanni-Vincentelli [1988]).

Tables 5.1 and 5.2 show the experimental results for the explicit method. Tables 5.3 and 5.4 list the experimental results for the corrected implicit method with one or more than one inner iterations per outer iteration and with inner line searches. Note that as long as the inner line search is applied, the corrected implicit method, even with one inner iteration per outer iteration, is not the same as the explicit method.

In Tables 5.1 and 5.3, T_i ($i = 1, \dots, 4$) is the total computing time for all computations involving the i th diagonal block on node i , T_b is the total computing time for all computations involving the bottom block on the control node, T_c is the total communication time for the computation, N_{out} is the total number of outer iterations required to converge to the solution, and I_{in} in Table 5.3 is the number of inner iterations used in the corrected implicit

Table 5.1: Times for the Explicit Method on the op-amp 741 Circuit

T_1	T_2	T_3	T_4	T_b	T_c	N_{out}
12.81	14.20	13.45	14.58	3.42	0.43	20

Table 5.2: Parallel Performance of the Explicit Method on the op-amp 741 Circuit

T_s	T_p	sp	eff
58.46	18.43	3.14	78.5%

Table 5.3: Times for the Corrected Implicit Method on the op-amp 741 Circuit

I_{in}	T_1	T_2	T_3	T_4	T_b	T_c	N_{out}
1	11.81	13.06	11.96	13.28	2.84	0.38	18
2	12.60	14.01	12.83	14.45	1.65	0.32	15
3	11.28	12.84	11.46	13.01	1.38	0.25	12
4	18.48	20.57	18.81	21.34	1.32	0.25	11
5	29.04	32.12	29.92	32.89	1.34	0.24	11

Table 5.4: Parallel Performance of the Corrected Implicit Method on the op-amp 741 Circuit

I_{in}	T_s	T_p	sp	eff
1	52.95	16.5	3.21	80.25%
2	55.54	16.42	3.38	84.50%
3	49.97	14.64	3.43	85.75%
4	80.52	22.91	3.50	87.50%
5	125.31	34.47	3.60	90.0%

method. In the performance Tables 5.2 and 5.4, T_s is the total computing time for solving the same problem sequentially on one node,

$$T_s = \sum_{i=1}^4 T_i + T_b,$$

T_p is the parallel computing time,

$$T_p = \max(T_1, \dots, T_4) + T_b + T_c,$$

sp is the speedup of the parallel computation in comparison to its sequential counterpart,

$$sp = \frac{T_s}{T_p},$$

and eff is the parallel efficiency defined by

$$eff = \frac{sp}{\text{number of processors}}.$$

Our experiments show that using the implicit method, with inner line search and multiple inner iterations per outer iteration, does indeed speed up the convergence to the solution. For example, the implicit method with one inner iteration per outer iteration and inner line search used 18 iterations to converge to the solution. The same algorithm without inner line search, i.e. the explicit method, used 20 iterations. As the number of inner iterations per outer iteration, I_{in} , was increased, the total number of outer iterations N_{out} decreased from 18 to 11, and the speedup increased because the bottom block computations constituted a smaller percentage of the overall computation. However, the sequential computing time only decreased by a small amount until $I_{in} = 3$, and then increased dramatically because the cost of the extra inner iterations swamped the small savings from the further decrease in the number of outer iterations. Both the sequential and parallel computing times were minimized when the number of inner iterations was $I_{in} = 3$, and the speedup in this case, 3.43, was good. The parallel method for this case costs 21% less than the parallel explicit method.

Our experiments also show that the bottle-neck computing time T_b of the corrected implicit method is more than 50% lower than in the explicit

method if more than one inner iteration is applied. The communication time is also lower since the total number of iterations, and hence the time to send the updated variables among the nodes, is less than for the explicit method. Consequently, the advantage of the implicit method over the explicit method should be greater in larger problems where the amount of communication and cost of solving the bottom block are larger.

5.3 Experiments for 2-level Block Bordered Circuit Equations

The 2-level nonlinear block bordered equations for the analog filter composed of 3 blocks of the 741 op-amp circuit were also solved in parallel on the Intel iPSC1 hypercube multiprocessor. The linearization of these equations at each iteration has the form

$$J\Delta X = -F \quad (5.3)$$

where J is the 2-level block bordered matrix shown in Appendix D,

$$\Delta X = (\Delta x_1^1, \dots, \Delta x_q^1, \Delta x_{q+1}^1, \Delta x_1^2, \dots, \Delta x_q^2, \Delta x_{q+1}^2, \Delta x_1^3, \dots, \Delta x_q^3, \Delta x_{q+1}^3, \Delta \hat{x}_{q+1})^T$$

and

$$F = (f_1^1, \dots, f_q^1, f_{q+1}^1, f_1^2, \dots, f_q^2, f_{q+1}^2, f_1^3, \dots, f_q^3, f_{q+1}^3, \hat{f}_{q+1})^T.$$

The system (5.3) could be solved by applying the block bordered solver to each of the 3 block bordered submatrices, and then solving the whole system by applying the block bordered solver again. Alternatively, the block bordered Jacobian matrix may be reordered to

Table 5.5: Times for the Explicit Method on the Analog Filter

T_1	T_2	T_3	T_4	T_5	T_6	T_7	T_8	T_9
13.49	14.93	14.21	15.34	13.44	14.85	14.16	15.46	13.36
T_{10}	T_{11}	T_{12}	T^1	T^2	T^3	T_b	T_c	N_{out}
14.76	14.25	15.63	3.01	3.12	3.17	0.58	1.31	21

Table 5.6: Parallel Performance of the Explicit Method on the Analog Filter

T_s	T_p	sp	eff
183.76	20.69	8.88	74.00%

In our test program, the 12 diagonal block equations of the analog filter were distributed among 12 nodes of the Intel hypercube. The first level or internal connection functions in each amplifier, f_{q+1}^j ($j = 1, \dots, 3$) were distributed to 3 of the 12 nodes, and the second level connection function among the 3 amplifiers in the analog filter, \hat{f}_{q+1} , was handled sequentially by one of the 12 nodes.

Tables 5.5 and 5.6 give the experimental results and the performance of the explicit method for solving the 2-level analog filter equation. Tables 5.7-5.11 show the experimental results and performance of the corrected implicit method for solving the 2-level block bordered analog filter equations with one to four inner iterations. The symbols in the tables have the same meanings as in Tables 5.1-5.4, with the following exception: T_i , $i = 1, \dots, 12$ is the total time for the 12 first level blocks, while T^j , $j = 1, 2, 3$ is the time for the 3 second level blocks.

Our experimental results show that the corrected implicit method is also more efficient than the explicit method on this larger block bordered system of equations. The total number of iterations N_{out} is 21 for the explicit method, while for the implicit method it decreases from 18 to 11 as the num-

Table 5.7: Corrected Implicit Method Times for the Analog Filter, $I_{in} = 1$

T_1	T_2	T_3	T_4	T_5	T_6	T_7	T_8	T_9
11.54	12.78	12.13	13.13	11.52	12.69	12.10	13.21	11.42
T_{10}	T_{11}	T_{12}	T^1	T^2	T^3	T_b	T_c	N_{out}
12.61	12.20	13.40	2.56	2.67	2.73	0.5	1.12	18

Table 5.8: Corrected Implicit Method Times for the Analog Filter, $I_{in} = 2$

T_1	T_2	T_3	T_4	T_5	T_6	T_7	T_8	T_9
11.83	11.93	11.84	12.09	11.91	12.01	11.89	12.12	11.84
T_{10}	T_{11}	T_{12}	T^1	T^2	T^3	T_b	T_c	N_{out}
12.05	11.94	12.13	1.65	1.65	1.64	0.36	0.73	12

Table 5.9: Corrected Implicit Method Times for the Analog Filter, $I_{in} = 3$

T_1	T_2	T_3	T_4	T_5	T_6	T_7	T_8	T_9
21.25	23.67	23.25	24.51	21.67	23.29	24.21	24.35	21.22
T_{10}	T_{11}	T_{12}	T^1	T^2	T^3	T_b	T_c	N_{out}
23.56	23.21	24.75	1.52	1.51	1.53	0.34	0.69	11

Table 5.10: Corrected Implicit Method Times for the Analog Filter, $I_{in} = 4$

T_1	T_2	T_3	T_4	T_5	T_6	T_7	T_8	T_9
29.03	32.24	29.98	32.75	29.11	32.04	29.87	32.71	29.40
T_{10}	T_{11}	T_{12}	T^1	T^2	T^3	T_b	T_c	N_{out}
32.21	31.05	32.34	1.50	1.51	1.50	0.35	0.69	11

Table 5.11: Parallel Performance of the Corrected Implicit Method on the Analog Filter

I_{in}	T_s	T_p	sp	eff
1	157.19	17.19	8.86	73.83%
2	148.68	14.86	10.01	83.40%
3	283.84	27.31	10.39	86.58%
4	373.08	34.89	10.69	89.08%

ber of inner iterations I_{in} is increased from 1 to 4. However the sequential computing time T_b for the implicit method only decreases from 157.19 for $I_{in} = 1$ to 148.68 for $I_{in} = 2$ (for the explicit method it is 183.76), then it increases sharply due to the cost of the additional inner iterations. Thus the high speedups for the implicit method for $I_{in} = 3$ and 4 in comparison to the same sequential method are not significant since the large number of inner iterations makes the algorithm inefficient, and the sequential time is suboptimal. For the optimal number of inner iterations, $I_{in} = 2$, the speedup is 10.01 for 12 processors and the efficiency is 83.40%. The parallel computation time improvement over the parallel explicit method is 28%, as compared to 19% in the sequential case. We feel that this experiment indicates that applying the implicit method to solve large block bordered circuit equations on a distributed memory multiprocessor can result in a highly efficient method.

6 Summary and Future Research

We have introduced a corrected implicit method for solving block bordered systems of nonlinear equations. It allows multiple “inner” iterations, iterations on the variables and equations of the q diagonal blocks, to be performed per each “outer” iteration, which involves all the variables and equations including the connecting, bottom block. If only one inner iteration is performed per outer iteration, no line search is used, and the bottom, connecting equations are linear, then the corrected implicit method is identical to the

explicit method (Newton's method). When more than one inner iteration is performed per outer iteration, however, the methods are different, and in our experiments the corrected implicit method solves problems in somewhat less time than the explicit method on sequential computers. On parallel computers, the corrected implicit method has a larger advantage over the explicit method because it parallelizes more effectively, since the inner iterations constitute a larger percentage of the total computation and parallelize far more effectively than the outer iterations. On one and two level block bordered problems from VLSI circuit design that we tested, the parallel efficiency of the fastest (sequential and parallel) corrected implicit method on an Intel iPSC1 hypercube was about 85%.

The methods presented in this paper all assume that the Jacobian matrix is available at each iteration, either analytically or by finite differences, and that it isn't too expensive to evaluate. In some applications, however, the nonlinear equations are given by an expensive computational procedure, and analytic or finite difference Jacobians are very expensive to obtain. In such cases for general systems of nonlinear equations, secant approximations to the Jacobian are used that are based entirely on function values at the iterates (see e.g. Dennis and Schnabel [1983]). The development of related secant approximations to the Jacobian for block bordered nonlinear equations seems to be an attractive research topic, since it appears possible to construct approximations that retain the block bordered sparsity pattern of the Jacobian and also allow the factorization of the Jacobian approximation to be updated efficiently.

REFERENCES

- R. H. Byrd, D. Feng, R. B. Schnabel and X. Zhang [1990], Convergence of Implicit Methods for Nonlinear Equations, in preparation.
- T. Chan [1985], An Approximate Newton Method for Coupled Nonlinear Systems, *SIAM Journal on Numerical Analysis*, No. 5, pp. 904-913.
- C. Christara [1988], Spline Collocation Methods, Software and Architectures for Linear Elliptic Boundary Value Problems, *Ph.D Thesis*, Purdue University, August, 1988.
- C. Christara and E. Houstis [1989], A Domain Decomposition Spline Collocation Method for Elliptic Partial Differential Equations, *Proceedings of the 4th Conference on Hypercube Concurrent Computers and Applications*. Monterey, California, March 6-8, 1989.
- L. Chua, P. Lin [1975], *Computer-Aided Analysis of Electronic Circuits: Algorithms and Computational Techniques*, Prentice-Hall Inc., 1975.
- J. Dennis, R. Schnabel [1983], *Numerical Methods for Unconstrained Optimization and Nonlinear Equations*, Prentice-Hall, Englewood Cliffs, New Jersey.
- J. Ebers, J. Moll [1954], Large-signal Behavior of Junction Transistors, *Proc. IRE*, Vol. 42, pp. 1761-1772.
- C. Farhat, E. Wilson [1986], Concurrent Iterative Solution of Large Finite Element Systems, *Technical Report*, Civil Engineering Department, University of California at Berkeley.
- R. Fontecilla [1987], A Parallel Nonlinear Jacobi Algorithm for Solving Nonlinear Equations, *Technical Report*, Computer Science Department, University of Maryland, May, 1987.
- M. R. Garey, D. S. Johnson [1979], *Computers and Intractability, A Guide*

to the Theory of NP-Completeness, W. H. Freeman and Company, 1979.

P. Gill, W. Murray and M. Wright [1981], *Practical Optimization*, Academic Press, 1981.

D.P. O'Leary, R.E. White [1985], Multi-Splittings of Matrices and Parallel Solution of Linear Systems, *SIAM Journal of Alg. & Disc Math.*, No. 4, pp. 137-149.

M. Mu and J. Rice [1989], Solving Linear Systems with Sparse Matrices on Hypercubes, *Technical Report*, CSD-TR-870, Computer Science Department, Purdue University, February, 1989.

B. Nour-Omid, K.C. Park [1986], Solving Structural Mechanics Problems on a Caltech Hypercube Machine, *Technical Report*, Mechanical Engineering Department, University of Colorado at Boulder.

N. Rabbat, H. Hsieh [1976], A Latent Macromodular Approach to Large-Scale Sparse Networks, *IEEE Transactions on Circuits and Systems*, Vol. CAS-23, No. 12.

A. Newton, D. Pederson, A. Sangiovanni-Vincentelli [1988], *SPICE 3B1 User's Guide*, Department of Electrical Engineering and Computer Sciences, University of California at Berkeley.

N. Rabbat, A. Sangiovanni-Vincentelli, H. Hsieh [1979], A Multilevel Newton Algorithm with Macromodeling and Latency for the Analysis of Large-Scale Nonlinear Circuits in the Time Domain, *IEEE Transactions on Circuits and Systems*, Vol. CSA-26, No. 9.

N. B. Rabbat, A. L. Sangiovanni-Vincentelli [1980], Techniques of Time-Domain Analysis of LSI Circuits, *Technical Report*, IBM T. J. Watson Research Center, RC 8351 (#36320), July, 1980.

C. Romine, J. Ortega [1986], Parallel Solution of Triangular Systems of Equations, *ICASE Technical Report*, NASA.

A. Sangiovanni-Vincentelli, L. Chen, L. Chua [1977], An efficient Heuristic Cluster Algorithm for Tearing Large-Scale Networks, *IEEE Transactions on Circuits and Systems*, Vol. CAS-24, No. 12.

A. Sangiovanni-Vincentelli and D. Webber [1986], Computer Architecture Issues in Circuit Simulation, *High Speed Computing*, edited by R. Wilhelmson, University of Illinois Press.

A. Sedra, K. Smith [1982], *Microelectronic Circuits*, CBS College Publishing.

J. Smith [1971], *Modern Operational Circuit Design*, Wiley, New York.

M. Valkenburg [1982], *Analog Filter Design*, CBS College Publishing.

R. S. Varga [1973], *Matrix Iterative Analysis*, Prentice-Hall, Englewood Cliffs, N.J. (1973)

R.E. White [1986], Parallel Algorithms For Nonlinear Problems, *SIAM Journal of Alg. & Disc Math.*, No. 7, pp. 137-149.

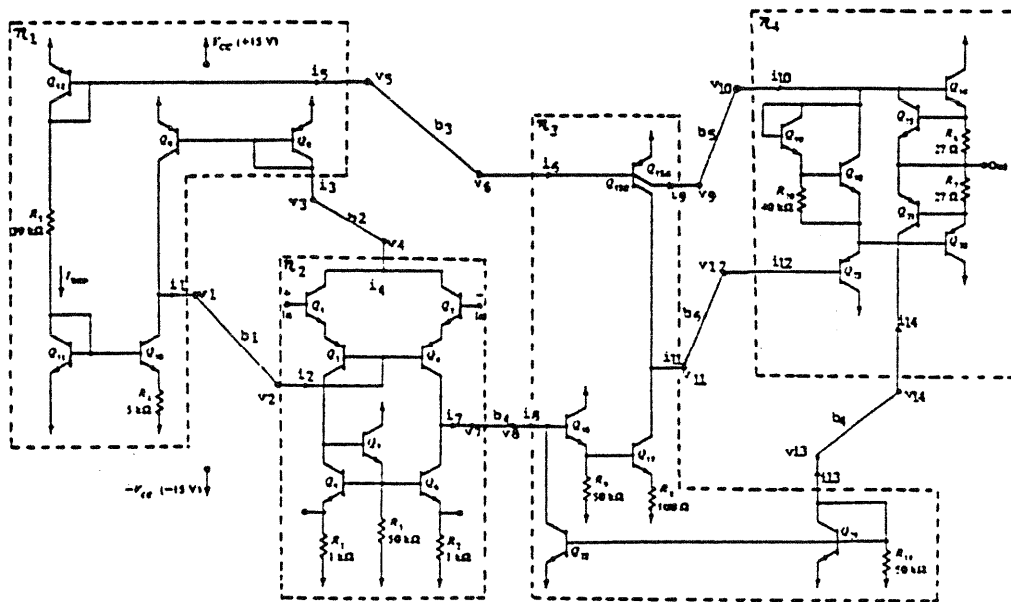
X. Zhang [1989], Parallel Computation for the Solution of Nonlinear Block Bordered Equations and Their Applications, *Ph.D Thesis*, University of Colorado at Boulder, July, 1989.

X. Zhang, R. Byrd, R. Schnabel [1989], Solving Nonlinear Block Bordered Circuit Equations on Hypercube Multiprocessors, *Proceedings of the 4th Conference on Hypercube Concurrent Computers and Applications*. Monterey, California, March 6-8, 1989.

APPENDIX A

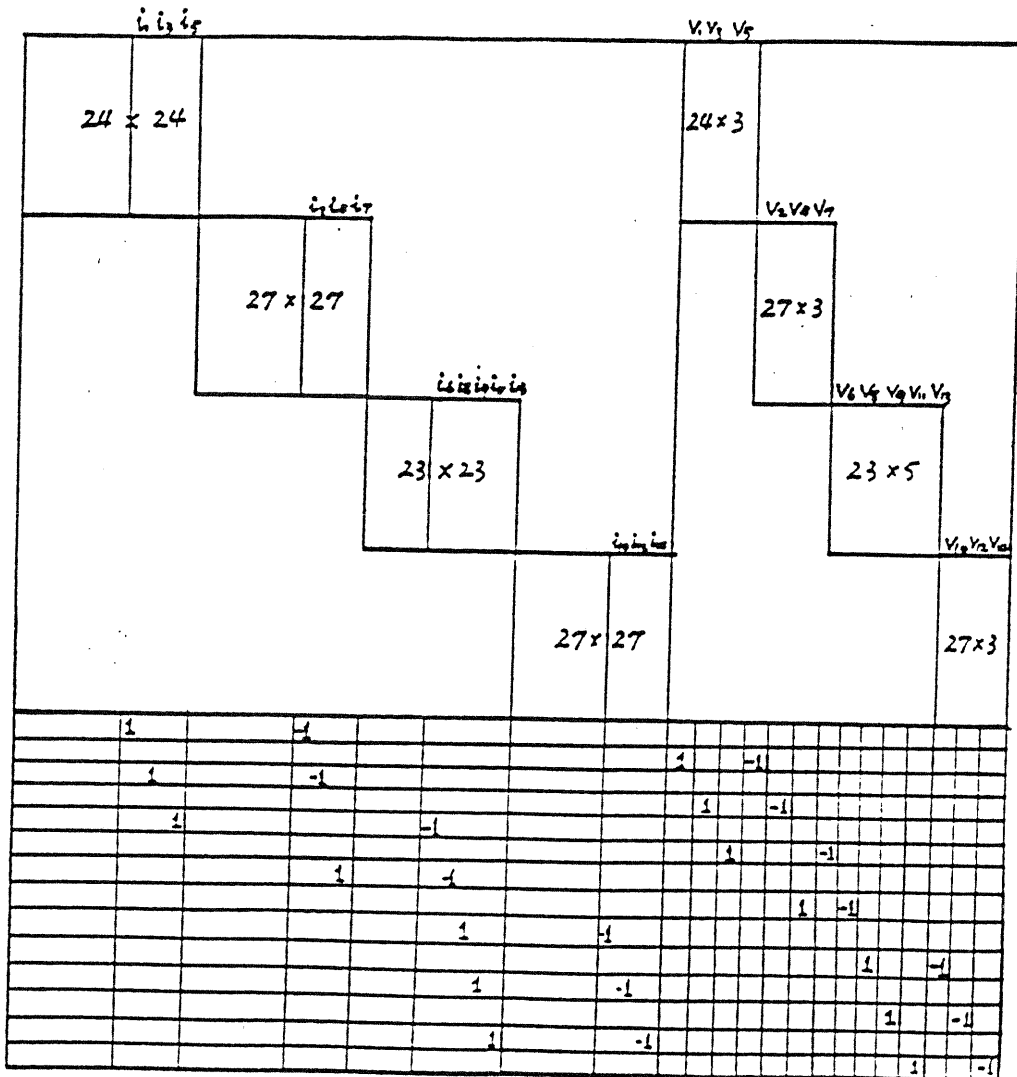
THE PARTITIONED 741 OP-AMP CIRCUIT

The 741 op-amp circuit was introduced in 1966 and is currently produced by almost every analog semiconductor manufacturer. The circuit is partitioned into 4 parts with roughly equal elements in each sub-circuit.



APPENDIX B

THE BLOCK BORDERED JACOBIAN MATRIX
OF THE PARTITIONED 741 OP-AMP CIRCUIT



APPENDIX C
AN ANALOG FILTER

