Solving Systems of Nonlinear Equations on a Message-Passing Multiprocessor

Thomas F. Coleman^{*}
Guangye Li[†]

87-887
November 1987

Department of Computer Science
Cornell University
Ithaca, New York 14853-7501

^{*}Computer Science Department and Center for Applied Mathematics, Cornell University, Ithaca, NY 14853. Research partially supported by the Applied Mathematical Sciences Research Program (KC-04-02) of the Office of Energy Research of the U.S. Department of Energy under grant DE-FG02-86ER25013.A000.

[†]Department of Combinatorics and Optimization, University of Waterloo, Waterloo, Ontario, Canada. Research partially supported by the U.S. Army Research Office through the Mathematical Sciences Institute, Cornell University. Permanent address: Computer Center, Jilin University, People's Republic of China.

SOLVING SYSTEMS OF NONLINEAR EQUATIONS ON A MESSAGE-PASSING MULTIPROCESSOR*

THOMAS F. COLEMAN† AND GUANGYE LI‡

Abstract. We develop parallel algorithms for the solution of dense systems of nonlinear equations on a message-passing multiprocessor computer. Specifically, we propose a distributed finite-difference Newton method, a multiple secant method, and a rank-1 secant method. Experimental results, obtained on an Intel hypercube, indicate that our methods exhibit good parallelism.

Key Words. systems of nonlinear equations, hypercube computer, message-passing multiprocessor, secant method, finite-difference Newton method, parallel algorithms.

Subject Classification. AMS/MOS: 65H10, 65K05, 65K10.

Abbreviated Title. Nonlinear Equations on a Multiprocessor

* Written November 1987.

† Computer Science Department and Center for Applied Mathematics, Cornell University, Ithaca, NY 14853. Research partially supported by the Applied Mathematical Sciences Research Program (KC-04-02) of the Office of Energy Research of the U.S. Department of Energy under grant DE-FG02-86ER25013.A000.

‡ Department of Combinatorics and Optimization, University of Waterloo, Waterloo, Ontario, Canada. N2L 3G1. Research partially supported by the U.S. Army Research Office through the Mathematical Sciences Institute, Cornell University. Permanent Address: Computer Center, Jilin University, People's Republic of China.

1. Introduction. In this paper we investigate parallel algorithms, tailored to the hypercube multiprocessor context, for the solution of systems of nonlinear equations

$$(1) \quad \text{solve } F(x) = 0$$

where $F : R^n \rightarrow R^n$. Component i of F is denoted by f_i . We assume that F is differentiable; let $J(x)$ denote the Jacobian matrix evaluated at point x .

Our implementations are specific to a hypercube multiprocessor; however, the algorithmic ideas are applicable more generally. In particular, the parallel algorithms presented here can be tailored to any multiprocessor computer provided that the communication topology allows for efficient “fan-in” and “fan-out” operations and the processors themselves have significant local memory. Further, some of our proposed algorithms - multi-secant update, triangular solve - are most meaningful when a “ring” communication pattern is used; hence, the topology of the multiprocessor should allow for a ring embedding. Finally, we remark that we always assume that the dimension of the problem, n , is greater than the number of processors p ; indeed, the algorithms uniformly become more efficient as n/p increases.

Our ultimate interest is in large sparse problems; however, in this paper we restrict our attention to the case in which the Jacobian matrix is assumed to be dense.

In a nutshell, this paper represents our attempt to parallelize the popular globalized Newton-like approaches to (1), such as secant and finite-difference Newton methods with a dogleg step or linesearch procedure. Consequently, the algorithms

under consideration actually solve the structured minimization problem

$$(2) \quad \text{minimize } \{f(x) : f(x) = \frac{1}{2}F(x)^T F(x)\}.$$

Obviously a solution to (1) is also a solution to (2); unfortunately, the converse is not always true. Nevertheless, such algorithms often are used successfully to obtain solutions to (1).

We divide the world into two classes of functions: functions F that are most conveniently evaluated as a single entity (i.e. a single subroutine evaluates the entire vector function F , sequentially) and functions F that can be evaluated in a distributed, parallel fashion. In this paper, to be concrete, we restrict the latter category to functions that conveniently separate into component sequential subroutines for each f_i , $i = 1 : n$. We call such functions *row-separable*.¹

If F is to be evaluated as a single entity then our approach is to distribute copies of the F -evaluation subroutine to all the processors. We then assume that any node (processor) can evaluate $F(x)$, given x , with no other communication necessary. Notice that the usual (rank-1) secant method cannot be parallelized in any obvious way since the evaluation of $F(x)$ is not distributed. If the evaluation of F is cheap relative to the other computations (e.g. matrix updating, triangular solves,...) then this poses no problem - let one node evaluate F while the others remain idle; however, if F is

¹ We restrict our attention to row-separable functions in order to be provide specific explicit algorithms, and for purposes of implementation and experimentation. However, the ideas and algorithms developed for row-separable functions are easily adapted to the general situation: i.e. functions that can be evaluated in a distributed, parallel manner.

relatively expensive, then it is not clear how to effectively parallelize the rank-1 secant method. For this reason we have developed the *multi-secant method* in which each processor evaluates F at a different point (or perhaps several different points). The result is a rank- q update, where q is a multiple of the number of processors available. This approach is discussed in Section 5. In the extreme case, when each node is evaluating F at many different points, the multiple secant method resembles a parallel finite-difference Newton method. The latter approach is explored in Section 4.

When F is row-separable the evaluation of $F(x)$ can be done in parallel. This allows for an efficient parallel version of the (globalized) rank-1 secant method. The crucial problem here is the design of an effective parallel QR -updating scheme. We discuss this in Section 3 .

Next we briefly describe the salient features of a hypercube computer. See Wiley [13], for example, for more information.

A message-passing multiprocessor consists of several independent processors connected by communication links. Each processor has significant local memory. (For example, the Intel iPSC with extra memory boards has approximately 4 megabytes of available memory, per node.) There is also a host computer, connected to one or several of the nodes (processors), whose purpose is to load programs and data onto the nodes of the cube, as well as collect the answer; we take the view that the host does not participate in intermediate computations.

Each processor supports two message-passing primitives: *send* and *receive*. When a node *sends* an array, it is transported through a sequence of nodes and communi-

cation links - the sequence is usually determined by the operating system - until the target node is reached. Upon executing a *receive*, a node checks to see if a new message is in its buffer. If so, the message is read and execution continues; if not, execution is suspended (on the receiving node only) until a message arrives.

A hypercube computer is a particular kind of message-passing multiprocessor. Specifically, the name refers to the topology defined by the communication links. A 0-dimensional hypercube, or 0-cube, is a single processor. To construct a 1-cube (i.e. 2 nodes), join two 0-cubes with a single communication link. In general, construct an m -cube (i.e. 2^m nodes) from two $m - 1$ cubes: find a one-to-one correspondence between the nodes in each cube and add a communication link between each pair.

The hypercube topology allows for a number of interesting properties (e.g. [12]). A particularly important one, for our purposes, is that a spanning tree of depth m can be embedded in an m -cube, *rooted at any node*. This allows for the efficient implementation of a number of global operations. For example, a node can send information to every other node in m “time steps”. This is usually called a *broadcast* or *fan-out* operation. Alternatively, a vector distributed over the cube can be collected onto any single node (the target node) in m “time-steps” using a spanning tree rooted at the target node. This is often called a *fan-in* operation.

Each node has a unique name, *myid*, which is a number in the range $[0, p - 1]$ where p is the number of processors; each node is aware of its own name. We use two labellings, or assignment of node names. The first is the *natural* ordering which is an assignment of integers in $[0, p - 1]$ such that each neighbor of node i (a neighbor

of node i is a node connected to i by a single communication link) differs by a single bit in its binary representation of its name. For example, the neighbors of node 5, in a 4-cube, are nodes 4, 7, 1, and 13. A *ring* ordering is also used. In this case node i is connected by single communication links to node $(i-1) \bmod p$ and node $(i+1) \bmod p$. An m -cube always allows for an embedding of a ring on 2^m nodes.

Experimental results reported in this paper were obtained using the Cornell Theory Center 16-node Intel iPSC hypercube under Xenix 286 release 3.4 of the host operating system and iPSC release 3.0 of the node operating system. The nodes were equipped with extra memory boards yielding approximately 4 megabytes of available memory per node. All our programs were written in Fortran.

2. The Sequential Secant Algorithm. We begin by summarizing a simplified version of the Minpack [9] sequential secant algorithm. This algorithm, in turn, is based on the work of Powell [10].

Suppose x_c is the current approximation to a solution to (2) and define $F_c \stackrel{\text{def}}{=} F(x_c)$. Let B_c be the current Jacobian approximation and let $B_c = Q_c R_c$ be the QR -factorization of B_c . Assume B_c to be nonsingular.

A *trial step* s_c is computed by approximately solving the trust region problem

$$(3) \quad \text{minimize } \{\|F_c + B_c s\|_2^2 : \|s\|_2 \leq \Delta_c\}$$

where Δ_c is the current radius of the trust region. The approximate solution s_c is obtained by further restricting (3) : Specifically, s_c solves the problem

$$(4) \quad \text{minimize } \{\|F_c + B_c s\|_2^2 : \|s\|_2 \leq \Delta_c, s \in P_c\}$$

where P_c is a piecewise linear path defined as follows: First, connect x_c to the Cauchy point, $x_c + s_c^{\text{Cauchy}}$, where

$$(5) \quad s_c^{\text{Cauchy}} \stackrel{\text{def}}{=} - \frac{\|B_c^T F_c\|_2^2}{\|B_c B_c^T F_c\|_2^2} B_c^T F_c,$$

and then connect $x_c + s_c^{\text{Cauchy}}$ to the Newton point, $x_c + s_c^{\text{Newton}}$, where

$$(6) \quad s_c^{\text{Newton}} \stackrel{\text{def}}{=} -B_c^{-1} F_c.$$

The computation determining the trial step s_c boils down to the algorithm in Figure 1. (Note: *newx* is a logical input parameter. If *newx* = *false* then all quantities in Step 1. have not changed since the previous call; otherwise, *newx* = *true* and all quantities have changed values.)

```

{1: Compute Cauchy and Newton steps }
If {newx } then
    {Compute  $s_c^{Cauchy}$  (given  $Q_c, R_c$ ) }
     $u \leftarrow Q_c^T F_c$  ;
     $g \leftarrow R_c^T u$  ;
     $w \leftarrow Q_c R_c g$ ;
     $s_c^{Cauchy} \leftarrow -\frac{\|g\|_2^2}{\|w\|_2^2} g$  ;

    {Compute  $s_c^{Newton}$  }
    Solve  $R_c s_c^{Newton} = -u$ ;
Endif

{2: Solve (4) }
If {  $\|s_c^{Newton}\|_2^2 \leq \Delta_c$  } then  $s_c \leftarrow s_c^{Newton}$ 

Elseif {  $\|s_c^{Cauchy}\| \geq \Delta_c$  } then  $s_c \leftarrow \left( \frac{\Delta_c}{\|s_c^{Cauchy}\|_2} \right) s_c^{Cauchy}$ 

Else  $s_c \leftarrow s_c^{Cauchy} + \alpha(s_c^{Newton} - s_c^{Cauchy})$ 
    where  $\alpha$  is the positive root of the quadratic equation
     $\|s_c^{Cauchy} + \alpha(s_c^{Newton} - s_c^{Cauchy})\|_2^2 = \Delta_c^2$ 
Endif

```

FIG. 1. *Algorithm Dogleg*

The possible correction s_c determined by algorithm *Dogleg* is accepted (i.e. $x_+ \leftarrow x_c + s_c$) only if $\|F(x_c + s_c)\|_2 < \|F(x_c)\|_2$. If s_c is not accepted then Δ_c is reduced and Step 2. of algorithm *Dogleg* is repeated.

There must also be a mechanism for increasing Δ_c so that progress is not impeded by unnecessarily small steps. This is accomplished by comparing the predicted reduction to the actual reduction. If this ratio, *ratio*, is sufficiently large then Δ_c is increased.

Besides updating x and Δ , it is necessary to evaluate $F(x_c + s_c)$ and update the QR -factorization of B to reflect the rank-1 secant update (due to Broyden [1]). We will not go into the QR -updating details here; however, orthogonal rotations can be used to stably perform this update using a total of $26n^2$ floating point operations (e.g. [3]). The algorithm in Figure 2 is a formalization of the ideas expressed above.

The algorithm described in Figure 2 represents a simplified version of the Minpack implementation. For example, Minpack will refresh B by finite-differences when it appears that convergence is not proceeding rapidly enough. In addition, Minpack will modify R if singularity is detected. Furthermore, the Minpack 'ratio test' and subsequent adjustment of Δ is somewhat more complicated. We will not spell these details in this description since they do not bear significantly on questions of parallelization.

```

{0: Initialize }
Choose  $x_c$ , evaluate  $F(x_c)$ , determine  $B(x_c)$  by finite-differences;
Compute  $B(x_c) = Q_c R_c$ 
 $newx \leftarrow true$ ;

{1: Attempt to find a zero of  $F(x)$  }
Repeat
    Determine  $s_c$  by Algorithm Dogleg( $newx$ ) ;
    Evaluate  $F(x_c + s_c)$  ;

    { Compute ratio }
     $actred \leftarrow 1 - \frac{\|F(x_c + s_c)\|_2^2}{\|F_c\|_2^2}$ ;
     $prered \leftarrow 1 - \frac{\|F_c + Q_c R_c s_c\|_2^2}{\|F_c\|_2^2}$ ;
     $ratio \leftarrow \frac{actred}{prered}$ ;

    { Update  $x$  }
    If {  $ratio \leq .0001$  } then  $x_+ \leftarrow x_c$ 
    Else  $x_+ \leftarrow x_c + s_c$  Endif

    {Update  $\Delta$  }
    If {  $ratio \leq \frac{1}{4}$  } then  $\Delta_c \leftarrow \frac{1}{2} \Delta_c$ 
    Elseif {  $ratio \leq \frac{3}{4}$  } then  $\Delta_+ \leftarrow \Delta_c$ 
    Else  $\Delta_+ \leftarrow 2\Delta_c$  Endif

    { Update  $B, newx$  }
    If {  $x_+ \neq x_c$  } then
         $newx \leftarrow true$  ;
        Update  $Q_c R_c \rightarrow Q_+ R_+$  to reflect the rank-1 change:
        
$$B_+ \leftarrow B_c + \frac{([F_+ - F_c] - B_c s_c) s_c^T}{s_c^T s_c} ;$$


    Else  $newx \leftarrow false$ 

Endif

Until {convergence }

```

FIG. 2. *Algorithm Secant*

3. Parallel Secant Method for Row-Separable Functions.

3.1. The Algorithm. As mentioned in Section 1, a row-separable function F is defined to be one in which it is convenient to have available a separate subroutine to evaluate each $f_i(x), i = 1 : n$. Assuming row-separability ², we now develop a parallel *Secant/Dogleg* secant method for (1) based upon the sequential secant method described in Section 2.

Our general approach is to distribute data and functions around the cube so that the work in the computationally intensive steps in algorithm *Secant/Dogleg* is well-distributed; we are adverse to redistributing information if it can be avoided. We make no attempt to parallelize steps that involve relatively insignificant computational work.

From our perspective the significant steps in algorithm *Dogleg* are the matrix-vector multiplies and the upper triangular solve in Step 1; Step 2 is relatively insignificant and can be performed on a single node. Beyond the call to *Dogleg*, the significant steps in algorithm *Hybrd* are: the initial finite-difference approximation B , the initial QR -factorization, the evaluation of $F(x_c + s_c)$, matrix-vector multiplies in the computation of *prered*, and the update of the QR -factorization.

We distribute F as follows: For $i = 0 : p - 1$, node i is assigned component

² We restrict our attention to row-separable functions in order to be provide specific explicit algorithms, and for purposes of implementation and experimentation. However, the ideas and algorithms developed for row-separable functions are easily adapted to the general situation: i.e. functions that can be evaluated in a distributed, parallel manner.

subroutine f_j for each $j = i + 1(\bmod p)$, $1 \leq j \leq n$. Therefore, to evaluate $F(y)$ each node must just evaluate its resident component functions: the simple node program is illustrated in Figure 3.

```

 $k \leftarrow myid + 1;$ 
While  $\{k \leq n\}$  do
    Evaluate  $f_k(y)$ 
     $k \leftarrow k + p$ 
Endo

```

FIG. 3. *Algorithm F-Evaluate (Node Program)*

The initial finite-difference determination of B can be accomplished in parallel using algorithm *F-Evaluate* as a subroutine. However, it turns out to be convenient to have the initial B distributed by columns; therefore, given our distribution of F , some communication is required within the parallel finite-difference method. The basic idea is to use algorithm *F-Evaluate* to parallelize the usual column-oriented finite-difference scheme. This is based, in turn, on the approximation to the j 'th column of the Jacobian matrix,

$$(7) \quad J(x)e_j \cong \frac{F(x + \tau e_j) - F(x)}{\tau}$$

where e_j is the j 'th column of the identity matrix and τ is an appropriate positive scalar. The node program for the parallel finite-difference approximation is given in Figure 4. It is assumed that every node has a copy of x , the current point, and τ , the differencing scalar.

Notice that each node determines some of the components of column j ; the "Fan-in" step collects column j onto node $(j - 1)\bmod p$ where it is stored.

```

For  $j = 1 : n$  do
    { Estimate column  $j$  }
     $y \leftarrow x + \tau e_j$  ;
     $F\text{-Evaluate}(y) \rightarrow w$  ;
    Participate in Fan-in of  $w \rightarrow \text{node } (j - 1) \bmod p$  ;
Endo

```

FIG. 4. *Algorithm J-Evaluate (Node Program)*

The initial QR -factorization of B can be accomplished by a parallel column-oriented algorithm based on orthogonal Householder transformations. Moler [7] has described the framework for such an algorithm in which the column-distributed matrix B is overwritten with R and the Householder vectors that define Q . However, our rank-1 updates require an explicit representation of Q ; therefore, we have modified Moler's algorithm to produce a *row-distributed* Q -matrix, while overwriting B with the *column-distributed* matrix R). This modification is rather straightforward and we will not describe it here; however, we do provide results of numerical experiments, in Table 1, designed to measure parallel efficiency.

There are numerous matrix-vector multiplies in algorithm *Secant/Dogleg* involving matrices Q, R, Q^T , and R^T : we have built our (straightforward) routines based on the communication utility routines provided by Intel [8]. Design of an efficient parallel triangular solver turns out to be a difficult problem. Nevertheless, there has been significant recent progress (e.g.[4,5,6,11]); we use the algorithm of Li and Coleman [6] in our implementation.

It remains to consider the QR -updating step. Indeed, our decision to cross-thread Q and R - i.e. distribute Q by rows and R by columns - was arrived at with this step

in mind. Hence we assume that if $j - 1 = i \bmod p$, $1 \leq j \leq n$, then node i houses row j of Q and column j of R . Let

$$(8) \quad B_+ = QR + rs^T = Q(R + \bar{r}s^T)$$

where $\bar{r} = Q^T r$. Assume that $\bar{r}$ is stored on a single node, node 0, say; let s^T be distributed, by column, to conform to the distribution of R .

Recall that the usual sequential algorithm introduces zeroes in $\bar{r}$ by applying orthogonal rotations to its rows, from bottom to top. Each rotation is applied, in turn, to the two corresponding rows of R and columns of Q . But the distribution we have chosen for Q and R is ideal for parallelization: each node contains a segment of the rows of R (and columns of Q) being rotated. Hence, the work involved in the rotation is well-distributed.

Upon completion of the step described above we have $B_+ = \hat{Q}\hat{R}$ where $\hat{Q}$ is orthogonal and $\hat{R}$ is upper-Hessenberg. Next we must reduce $\hat{R}$ to upper triangular form. This is done using orthogonal rotations applied from top to bottom. Rotation G_i is applied to rows $i, i + 1$ of $\hat{R}$ as well as columns $i, i + 1$ of Q (for $i = 1 : n - 1$). Rotation G_i involves computations that can be done concurrently because each node has a segment of rows (columns) $i, i + 1$ of $\hat{R}$ ($\hat{Q}$).

Figure 5 provides the detailed algorithm. For each node k , let

$$(9) \quad I(k) = \{1 \leq i \leq n : i - 1 = k \bmod p\}$$

Remark: As we have demonstrated, the parallelization of the *Secant/Dogleg* algorithm is fairly straightforward under a row-separability assumption. The two most


```

{Reduce to upper Hessenberg form }
For  $i = n - 1 : 1(-1)$  do
    If {  $myid = 0$  } then
        Determine  $G_i$  ; { Givens rotation defined by  $\bar{r}_{i+1}, \bar{r}_i$  }
        Apply  $G_i$  to  $\bar{r}$  ;
        Broadcast  $G_i$ ;
    Endif

    For each  $k \in I(myid)$  do
        Rotate elements in rows  $i, i + 1$  of col.  $k$  of  $R$ , using  $G_i$ ;
        Rotate elements in cols  $i, i + 1$  of row  $k$  of  $Q$ , using  $G_i$ ;
    Endo
Endo

If {  $myid = 0$  } then broadcast  $\alpha \stackrel{def}{=} \bar{r}_1$  Endif

Add  $\alpha \times s$  to first row of  $R$ ;

{ Reduce  $R$  back to upper triangular form }
For  $i = 1 : n - 1$  do
    If {  $myid = (i - 1) \bmod p$  } then
        Determine  $G_i$  , based on  $R_{i+1,i}, R_{i,i}$  ; { Givens rotation }
        Broadcast  $G_i$ ;
    Endif

    For each  $k \in I(myid)$  do
        Rotate elements in rows  $i, i + 1$  of col.  $k$  of  $R$ , using  $G_i$ ;
        Rotate elements in cols  $i, i + 1$  of row  $k$  of  $Q$ , using  $G_i$ ;
    Endo
Endo

```

FIG. 5. *Algorithm QR-update (Node Program)*

challenging steps are the triangular solve and QR -update. Indeed, it is possible to avoid these two steps altogether if we recur B^{-1} instead of B . This is quite possible (e.g. [2]) though from a numerical point of view it is probably preferable to update B . There are two major reasons we do not pursue this possibility here. First, one of our goals is to determine if the Minpack algorithm - which includes updating QR -factors - can be efficiently parallelized. Second, we maintain an eye toward the sparse case in this development: in general, J^{-1} is dense when J is sparse and therefore recurring an approximation to J^{-1} is unreasonable in the large sparse situation.

3.2. Numerical Experiments. In Table 1 we present timing results reflecting the performance of the QR -updating algorithm described in Figure 5, and, for comparison, the QR -factorization routine (which we have not described in detail here) and the triangular solve. Notice that even though both the QR -update and the triangular

TABLE 1
Timing Results for the Parallel Secant Update, $p = 16$

n	QR time	QR mflps	Update time	Update mflps	Solve time	Solve mflps
100	7.2	.458	.995	.261	.70	.028
200	51.3	.52	2.6	.39	1.3	.06
300	166.0	.54	5.6	.46	2.2	.08

solve are $O(n^2)$ operations, the efficiency of the rank-1 update is much better than the triangular solve efficiency. This is due to the constant factors involved: i.e. the QR -update requires $26n^2$ arithmetic operations, whereas the triangular solve involves $1n^2$ arithmetic operations. Moreover, the efficiency of the update is not dramatically worse than for the full factorization which, in turn, exhibits close to optimal megaflop

rate. The near-maximum efficiency of the QR -factorization is due to the intensive computational work required (recall that the orthogonal matrix Q is being explicitly formed).

TABLE 2
Running Times for Parallel Secant Algorithm

Prob	p	$n = 50$	$n = 100$	$n = 300$
9	1	20.5	128.27	2494.7
9	16	5.54	15.6	162.0
10	1	60.14	418.6	10264.3
10	16	9.4	36.7	643.9
14	1	92.6	397.0	5210.3
14	16	40.19	89.2	468.7
15	1	193.14	945.1	13121.2
15	16	85.7	227.9	1445.5

TABLE 3
Speedups for Parallel Secant Method: $p = 16$

Prob	$n=50$	$n=100$	$n=300$
9	3.8	8.2	15.4
10	6.4	11.4	15.9
14	2.3	4.5	11.1
15	2.25	4.1	9.1

In Table 2 we present the running times for our parallel implementation of the secant algorithm described in Section 2. Our stopping criterion was: $\|F\| \leq 10^{-8}$; the “update Δ ” step was modified to conform with the Minpack code. Problem 15 is the well-known extended Rosenbrock function; the other problems were chosen from the Minpack collection of nonlinear equation problems.

To provide a measure of speedup, in Table 3 we have divided the numbers in Table 2 by the sequential running times of a *modified* Minpack code running on a single processor. The Minpack code was modified to force secant updates (after the

initial finite-difference Jacobian estimation). Moreover, the Minpack stopping criteria were replaced by the rule mentioned above. Hence the two codes produce exactly the same sequence of x -iterates (we verified that this claim held on the four test problems in question).

Remarks

1. Obviously speedups improve as n increases. This is due to the increase in distributed work to be performed.

2. Problems 9 and 10 require only unit steps each iteration; on the other hand, problems 14 and 15 require many non-unit steps - many trial steps are rejected. It is this fact, *in combination* with the fact that function evaluations are *extremely cheap* that accounts for the relatively poor parallel efficiency demonstrated on problems 14 and 15.

To support this claim numerically, we have artificially increased the expense of each function evaluation in problem 14 by a factor of 100. The cost of a function evaluation is then about the same order of magnitude as the cost of a function evaluation in problem 10. For $p = 16$ and $n = 100$ the running time is 177.4 compared to 1562.3 when $p = 1$; the resultant speedup is 8.9, which compares favourably to the speedups obtained on problems 9 and 10.

TABLE 4
Secant Algorithm Breakdown for $p = 16$, $n = 100$

Prob	Init. J/QR	QR-Update	F-Evaluation	Tri. Solve	Other
9	1.0	.26	.07	.22	.18
10	1.0	.14	.06	.10	.10
14+	1.0	.28	.42	.28	.06
15	.10	1.0	.27	.56	.40

The purpose of Table 4 is to provide some indication of how the computing time is distributed amongst the various tasks. The table entries represent the normalized time to do each task - for each row, each task time is divided by the time required by the most expensive task. Column "Init. J/QR " represents the time to do the initial estimation of J by finite-differences plus the time required to do the initial QR -factorization. The "F-Evaluation" column represents the total time spent on evaluations of F excluding the initial estimation of J . The column labeled "Other" reflects the time spent on all remaining tasks. This includes several parallel matrix multiplies (used in the determination of Cauchy and Newton steps) as well as some sequential work (such as the adjustment of Δ).

Problem 14+ is the Minpack problem 14 with the cost of a function evaluation increased by a factor of 100.

It is now easy to see why problem 15 experiences relatively poor speedup: too little (relative) time is spent on the highly parallel tasks such as the initialization step and "F-evaluation". Instead, the QR -update and the triangular solve tasks consume the most time: unfortunately, neither task is as parallel-efficient as the QR -factorization or the distributed evaluation of F .

Our numerical experiments indicate to us that the proposed parallel secant method is acceptably efficient (for both cheap and expensive functions) except when F is cheap *and* many iterations are required. Note that many iterations may be required due to a poor Jacobian approximation; a hybrid routine such as the Minpack algorithm would tend to refresh the approximation (e.g. by finite-differences) under such circumstances

and this, in turn, would tend to decrease the number of iterations. In particular, we note that the parallel finite-difference actually outperforms the parallel secant method on problems 14 and 15: this is due to many fewer iterations and (ridiculously) cheap function evaluations.

4. Parallel Finite-Difference Newton Method.

4.1. The Algorithm. In the remainder of this paper we assume that the evaluation of $F(x)$ is not a distributed computation; every node has a copy of the F -evaluation subroutine. In addition, we discontinue the use of the subscript c ; e.g. x and s refer to vectors x_c and s_c , respectively. The finite-difference approximation of the Jacobian matrix can obviously be done in parallel, given $F(x)$, with each node computing its resident columns independently. Communication between nodes is not required. Since efficient parallel routines for the LU and QR factorizations exist, and since the triangular solve problem has been extensively researched, with acceptable results, the only remaining difficulty is the evaluation of $F(x + s)$ when determining if $x + s$ is an acceptable point. If F is relatively expensive to compute, then it is not reasonable to designate one distinguished node to evaluate $F(x + s)$ while the others idle.

Our solution breaks into two parts.

First, if in the course of a run previous experience suggests that the initial trial step s is likely to be accepted, then we take a chance and overlap the computation of $F(x + s)$ with the estimation of $J(x + s)$. This is at some risk because s might be determined to be unacceptable - due to the value of $\|F(x + s)\|$ - and then any work expended on the computation of $J(x_c + s)$ has been wasted. But, as indicated in Figure 6, each node will waste at most one evaluation of F before the suitability of s is determined.

Assign the task of evaluating $F(x + s)$ to node $n \bmod p$ (think of F as column

$n + 1$ of the Jacobian): $n + 1 \in I(n \bmod p)$. Figure 6 describes the algorithm.

```

If {  $myid = n \bmod p$  } then {i.e. if I am the node that evaluates  $F(x + s)$  }
    Evaluate  $F(x + s)$ ;
    Broadcast  $F(x + s)$  ;
Else
    Choose  $j \in I(myid)$ ;
    Evaluate  $F(x + s + \tau e_j)$  ;
Endif

If  $x + s$  is not acceptable then exit
Else
    Evaluate the remainder of the Jacobian columns by finite differences;
Endif

```

FIG. 6. *Algorithm JandF Evaluation (Node Program)*

On the other hand, if previous experience suggests that there is a good chance that the initial step s will not be adequate then our strategy is quite different. Specifically, each step of the parallel linesearch procedure involves p function evaluations, $F(w_{p-1}), \dots, F(w_0)$, done in parallel, where $w_{p-1}, \dots, w_0$ are points along the dogleg step P_c (see Section 2). If the first step is unsuccessful then a second parallel search is performed along a smaller segment of P_c . This process is repeated until a suitable point is found. The linesearch procedure is judged successful if the following “alpha condition” is satisfied for some $w_* \in \{w_{p-1}, \dots, w_0\}$:

$$(10) \quad f(w_*) \leq f(x) + \alpha \nabla f(x)^T [w_* - x],$$

where $0 < \alpha < \frac{1}{2}$, and $f(x) \stackrel{def}{=} \frac{1}{2} \|F(x)\|_2^2$. See Dennis and Schnabel [2] for a discussion of the “alpha condition”.

The procedure can be considered a parallel *generalized* bisection algorithm. In

each step we begin with a stepsize bound of Δ . Specifically, in the first step

$$(11) \quad \Delta \leftarrow \min\{\|s^{Newton}\|, BOUND \times XNORM\}$$

where $XNORM \stackrel{def}{=} \max\{\|x\|, TYPX\}$ and $TYPX$ is a positive user-supplied constant representing the norm of a “typical” x -iterate; $BOUND$ is a positive user-supplied constant.

In subsequent steps (if needed) Δ is defined by the (unsuccessful) evaluation point nearest to x (used in the previous step).

In each step the evaluation points are defined as follows:

$$(12) \quad w_i \in P_c, \quad \|w_i - x\| = \frac{\Delta}{\gamma^i}, \quad i = 0 : p-1$$

where γ is a positive number strictly greater than unity. The choice of γ is guided by the following two concerns.

First, γ should be chosen so that the point nearest to x , w_{p-1} , is not too close to x : i.e. we require

$$(13) \quad \frac{\Delta}{\gamma^{p-1}} \geq \mu \times XNORM$$

where μ is small positive number (usually unit roundoff). We assume $BOUND \gg \mu$.

Obviously expression (13) yields the upper bound on γ ,

$$(14) \quad \gamma \leq \left(\frac{\Delta}{\mu \times XNORM} \right)^{\frac{1}{p-1}}$$

Second, it is usually advantageous to spread the evaluation points so that at least one point (i.e. w_{p-1}) is on the Cauchy segment $(x, x + s^{Cauchy}]$. This leads to the

```

Choose  $\Delta$  as in 11
Repeat
  If {myid = 0 } then
    Determine  $\gamma > 1$  satisfying (14) and (16) ;
    Broadcast  $\gamma$  ;
  Endif
   $i \leftarrow myid$  ;
  Determine  $w_i \in P_c : \|w_i - x\| = \frac{\Delta}{\gamma^i}$  ;
  Evaluate  $z_i \stackrel{def}{=} F(w_i)$  ;
  Participate in fan-in:  $z \rightarrow z^*, w \rightarrow w^*$  on node 0 ;
  Decrease  $\Delta$  ;
Until  $w^*$  satisfies condition (10)

```

FIG. 7. *Algorithm Dogleg-Linesearch (Node Program)*

condition, if $\|s^{Cauchy}\| < \Delta$, we require

$$(15) \quad \frac{\Delta}{\gamma^{p-1}} \leq \|s^{Cauchy}\|.$$

However, this condition may lead to a γ so large that a very large segment of $s^{Newton} - s^{Cauchy}$ is devoid of function evaluations. Hence we compromise (15) and require only

$$(16) \quad \gamma \geq \min\left\{\left(\frac{\Delta}{\|s^{Cauchy}\|}\right)^{\frac{1}{p-1}}, 2\right\}.$$

Note that if we assume that

$$(17) \quad \|s^{Cauchy}\| > \mu \times XNORM$$

then (14) and (16) are consistent and $\gamma > 1$. If condition (17) does not hold then it is reasonable to stop, claiming optimality.

Figure 7 exhibits the algorithm. The "Decrease Δ " step is implemented as follows: Let $p^+ = p + 1$; determine $\gamma^+ > 1$ satisfying (14) and (16), replacing p with

p^+ . Finally, assign

$$(18) \quad \Delta \leftarrow \frac{\Delta}{[\gamma^+](p^+-1)}$$

The “Fan-in” step determines $w^* \in W = \{w_{p-1}, \dots, w_0\}$ such that $\|w^* - x\|$ is maximum and w^* satisfies (10). If no such point exists we define $w^* = x$. Determining if (10) is satisfied at point w_i is a simple computation. Specifically, we can write

$$(19) \quad w_i - x = \lambda_i^C s^{Cauchy} + \lambda_i^N s^{Newton}.$$

But, $\nabla f^T s^{Newton} = -\|F(x)\|_2^2 \stackrel{def}{=} \omega^N$ and $\nabla f^T s^{Cauchy} = -\beta \|J^T F\| \stackrel{def}{=} \omega^C$. (The constant β is computed when P_c is determined.) Therefore,

$$(20) \quad \nabla f^T [w_i - x] = \lambda_i^C \omega^C + \lambda_i^N \omega^N$$

which is a trivial expression to compute on a single node.

The overall procedure is sketched in Figure 8. Note that we have provided a high-level global view, as opposed to a node program. The parallelization of each computationally significant step has been discussed above.

4.2. Numerical Results. We performed computational experiments using problems 9,10,14, and 15 referred to previously. The standard starting point (Factor = 1) was used in all cases. The stopping criterion was $\|F\| \leq 10^{-8}$.

Problem 10 is distinguished from the others: F is relatively expensive to compute. In this sense problem 10 probably represents a more realistic test function. However, problems 14 and 15 are useful for testing because non-unit steplengths are required whereas Newton iterations converge quickly, with unit steps, for problems 9 and 10.

```

Guess an initial  $x$ ;
Evaluate  $F(x)$ ,  $J(x)$ ;
Assign  $prev\text{-}step \leftarrow \text{"not-Newton"}$  ;
Repeat
    Factor  $J = LU$  ;
    Determine  $P_c$  :
        Compute  $s^{Cauchy}$ ;    { see Figure 1 }
        Compute  $s^{Dogleg}$ ;    { see Figure 1 }
        Determine  $\Delta$ ;    { use (11) }

    If {  $prev\text{-}step = \text{"Newton"}$  } then
        Try algorithm  $JandF$  ;
        If {  $JandF$  is unsuccessful } then
             $prev\text{-}step \leftarrow \text{"not-Newton"}$ 
        Endif
    Endif

    If {  $prev\text{-}step = \text{"not-Newton"}$  } then
        Perform Dogleg-Linesearch ;
        Evaluate  $J(x^+)$  ;
    Endif

Until { convergence }

```

FIG. 8. *Algorithm Newton with Dogleg-Linesearch*

In Table 5 we have recorded the running times and iteration counts obtained for the finite-difference Newton method described above; in Table 6 we divide the $p = 1$ running times by the $p = 4$ and $p = 16$ times to obtain a measure of speedup. In this case we do *not* compare our algorithm to a finite-difference version of the Minpack code because we feel such a comparison would be unfair. Specifically, the Minpack code is *QR*-based and ours is *LU*-based: such a comparison would give meaningless advantage to our method.

TABLE 5
Results for Newton Algorithm with Dogleg-Linesearch

Prob	p	$n = 50$	$n = 100$	$n = 300$
9	1	3/7.6	3/43.5	3/911.0
9	4	3/3.3	3/13.6	3/240.1
9	16	3/3.0	3/7.9	3/79.8
10	1	4/107.9	4/826.8	4/21701.4
10	4	4/30.4	4/220.0	4/5528.2
10	16	4/13.1	4/69.2	4/1447.4
14	1	9/41.1	9/218.2	9/4050.2
14	4	9/15.5	9/65.7	9/1062.6
14	16	9/12.3	9/33.4	9/344.0
15	1	22/65.0	27/496.3	31/12871.5
15	4	22/30.1	27/158.2	31/3398.8
15	16	11/13.7	12/40.6	14/496.1

Our implementation struggles with problem 15 for small p (relative to its' performance with $p = 16$). This is because our linesearch parameter γ was chosen with a moderately large p in mind. In particular, in a first iteration of the linesearch algorithm, we choose γ so that condition (16) is an equality. If a second iteration of the linesearch is needed then γ is chosen to satisfy condition (14) exactly. This strategy appears to work quite well for $p = 16$ but can lead to small steps if p is small.

In Table 6 we have normalized the execution times reported in Table 5: for each

problem divide the execution times in sub-column by the $p = 1$ execution time. Hence the entries in Table 6 reflect the speed-up factor over the running time on a single node.

TABLE 6
Speed-ups for Newton Algorithm with Dogleg-Linesearch

Prob	p	$n = 50$	$n = 100$	$n = 300$
9	4	2.3	3.2	3.8
9	16	2.5	5.5	11.4
10	4	3.5	3.8	3.9
10	16	8.2	11.9	15.0
14	4	2.7	3.3	3.8
14	16	2.5	5.5	11.4
15	4	2.2	3.1	3.8
15	16	4.7	12.2	25.9

Remarks on Table 6 :

1. For fixed $p > 1$, the speed-up improves as n increases. The primary reason for this is that as n increases the parallel factorization becomes increasingly efficient (e.g. [7]). As Table 7 indicates, the factorization accounts for a significant percentage of the total computational expense in the test problems.
2. In general, the 4-processor speeds are closer to optimality (optimal speed-up = 4) than the 16-processor speed-ups (optimal speed-up = 16). Again, the primary factor here is the increased efficiency of the parallel LU -factorization as n/p increases. There are two exceptions to this trend in Table 6.

First, a nearly optimal speed-up is obtained on problem 10, $n = 300$. Table 7 explains this: the Jacobian estimation time dominates the factorization time - parallel finite-difference is a highly parallel task.

The other exception occurs on problem 15, $n = 300$: a speedup of 25.9 is attained

(considerably better than the “optimal” speed-up of 16 !). This is possible because the sequence of points generated is a function of the number of processors used (due to the linesearch). This dependence contrasts with the parallel secant algorithm described in the previous section.

In Table 7 we break the total execution time down into the times required by the different substeps. Each row of Table 7 is normalized so that the maximum entry in each row is unity.

TABLE 7
Newton Algorithm Breakdown for $p = 16$, $n = 300$

Prob	Jac. Est.	Factor	Linesearch	Tri. Solve
9	.08	1.0	.01	.08
10	1.0	.11	.02	.01
14	.17	1.0	.00	.08
15	.04	1.0	.01	.08

Except for problem 10, the factorization represents the dominant cost. We believe this is an anomaly: in practise function evaluations are often quite expensive. However, in either case, the linesearch and triangular solve times are relatively insignificant.

We are satisfied with the performance of this parallel finite-difference algorithm which combines a dogleg step with a generalized bisection algorithm. Our experience on our test collection indicates the required number of iterations is almost always fewer than for a dogleg/trust region strategy (i.e. no linesearch). However, we admit that from an aesthetic point of view the linesearch procedure is unattractive: moreover, it is somewhat heuristic in nature and is unrelated to the quadratic model philosophy. Considering these remarks, it might be preferable, overall, to stick with the usual

dogleg/trust-region philosophy and always employ the algorithm in Figure 6 to obtain parallelism. Our experiments indicate this would be slightly less efficient.

5. Parallel Multiple Secant Method.

5.1. The Algorithm. The obvious disadvantage to the finite-difference Newton method discussed in Section 4 is that the estimation of the Jacobian matrix can be extremely time-consuming. This is less true in the parallel context since finite-differencing is a highly parallel task; nevertheless, it is often unnecessary to obtain such accuracy. The success of the sequential rank-1 secant method attests to this claim.

Section 3 presented a parallel secant method under the row-separability assumption; however, if F is to be treated as a single entity we do not know how to implement an efficient rank-1 secant method (when the evaluation of F is expensive). But, it is possible to fill the gap between rank-1 and rank- n (i.e. finite-difference approximation) with an efficient parallel rank- q secant method where q is a multiple of p , the number of nodes.

To introduce the *multi-secant method* let us consider the case when $q = p$. For the moment we also assume that our algorithm is purely local: a unit step $x^+ \leftarrow x + s$ is always taken (we consider the general situation later).

Assume that the function value $F(x)$ is known to every node. The first step is to re-label the nodes so that a *ring* is induced (i.e. node i is a neighbor of both node $(i + 1) \bmod p$ and node $(i - 1) \bmod p$, for $i = 0 : p - 1$). A gray code mapping can be used for this purpose (e.g. [8]). Assume that B_c is distributed in the usual fashion, using this labelling. Hence, node i is assigned column j provided $j - 1 = i \bmod p$. Let s be the correction to x : i.e. $x_+ \leftarrow x + s$. (We assume for the moment that s will

be accepted.)

The next step is key. Each node evaluates F at a different point. Node 0 evaluates $F(x + s)$; node i , $1 \leq i \leq p - 1$, evaluates $F(x + s^i)$ where s^i is a sparse projection of s . That is, component j of s^i will be either s_j or 0. In particular,

$$(21) \quad \text{if } \{j - 1 = i \bmod p\} \text{ or } \{s_k^i = 0, 0 \leq k < j\} \text{ then } s_j^i = 0$$

$$(22) \quad \text{otherwise } s_j^i = s_j$$

For example, if $p = q = 4$, $n = 8$,

$$(23) \quad s^0 = (s_1, s_2, s_3, s_4, s_5, s_6, s_7, s_8)$$

$$(24) \quad s^1 = (s_1, 0, s_3, s_4, s_5, 0, s_7, s_8)$$

$$(25) \quad s^2 = (s_1, 0, 0, s_4, s_5, 0, 0, s_8)$$

$$(26) \quad s^3 = (s_1, 0, 0, 0, s_5, 0, 0, 0)$$

After evaluation, each node sends a copy of its newly computed function value to its higher numbered neighbor on the ring. Hence, after this shift, node j will have the vectors $F(x)$, $F(x + s^j)$, and $F(x + s^{(j-1) \bmod p})$. Therefore, if $p = q = 4$ then in addition to $F(x)$, each node has the pair of function values listed below:

$$(27) \quad \text{node 0 : } F(x + s), F(x + s^3)$$

$$(28) \quad \text{node 1 : } F(x + s^1), F(x + s)$$

$$(29) \quad \text{node 2 : } F(x + s^2), F(x + s^1)$$

$$(30) \quad \text{node 3 : } F(x + s^3), F(x + s^2)$$

We now demand that each node satisfy its own local secant equation.

Notation: For a matrix M let $M_{I(k)}$ denote the matrix of the same dimensions which matches M in columns $I(k)$ and whose other columns are zero columns. For $k = 1 : p - 1$ the secant equation for node k is

$$(31) \quad B_{I(k)}^+ [(x + s^k) - (x + s^{k-1})] = y^k$$

where $y^k \stackrel{\text{def}}{=} F(x + s^k) - F(x + s^{k-1})$. Equation (31) is reasonable because $(s^k - s^{k-1})_i \neq 0 \Rightarrow i \in I(k)$ and

$$(32) \quad \left\{ \int_0^1 J_{I(k)}([x + s^{k-1}] + \tau[s^k - s^{k-1}]) d\tau \right\} (s^k - s^{k-1}) = y^k.$$

On node 0 we demand satisfaction of the secant equation

$$(33) \quad B_{I(0)}^+ [x - (x + s^{p-1})] = y^0$$

where $y^0 \stackrel{\text{def}}{=} F(x + s^{p-1}) - F(x)$.

This requirement is also reasonable because $s_i^{p-1} \neq 0 \Rightarrow i \in I(0)$, and

$$(34) \quad \left\{ \int_0^1 J_{I(0)}(x + \tau s^{p-1}) d\tau \right\} s^{p-1} = y^0.$$

Of course “reasonableness” does not establish that the method possesses desirable local convergence properties: we will consider this theoretical question elsewhere.

An important property of this parallel multi-secant update is that there is very little communication required: each node sends (receives) exactly one vector to (from) an *adjacent* node. Moreover, beyond satisfying p local secant equations, the updated matrix B^+ also satisfies the *global* secant equation

$$(35) \quad B^+ s = y$$

where $y \stackrel{def}{=} F(x + s) - F(x)$. To see this note that

$$(36) \quad B_{I(0)}^+ s^{p-1} + B_{I(1)}^+ (s^1 - s^0) + \cdots + B_{I(p-1)}^+ (s^{p-1} - s^{p-2}) = -B^+ s$$

and

$$(37) \quad y^0 + \cdots + y^{p-1} = F(x) - F(x + s) = -y.$$

Figure 9 summarizes the node program representing the algorithm sketched above.

Note that the product Bd^j can be computed entirely locally on node j : the vector d^j has nonzero components only in locations corresponding to columns residing on node j (i.e. the columns in $B_{I(j)}$).

```

j ← myid;
Evaluate  $F(x + s^j)$ ;
Send a copy of the vector  $F(x + s^j)$  to node  $(j + 1) \bmod p$  ;
Receive a copy of the vector  $F(x + s^{(j-1) \bmod p})$  from node  $(j - 1) \bmod p$ ;

If { myid = 0 } then
     $y^0 \leftarrow F(x + s^{p-1}) - F(x)$  ;
     $d^0 \leftarrow s^{p-1}$  ;
Else
     $y^j \leftarrow F(x + s^j) - F(x + s^{j-1})$  ;
     $d^j \leftarrow s^{j+1} - s^j$  ;
Endif

{ Update resident columns }
If {  $d^j \neq 0$  } then
     $B_{I(j)}^+ \leftarrow B_{I(j)} + \frac{(y^j - Bd^j)d^{jT}}{d^{jT}d^j}$  ;
Endif

```

FIG. 9. *Multiple Secant Update (Node Program)*

Two subtopics remain to be discussed: the generalization of the multi-secant method to the case where q is a multiple of p , and the globalization strategy.

The generalization of the multi-secant method to the case where q is a multiple of p is quite straightforward. Divide the columns on each node into q/p groups (typically with as many equal-sized groups as possible). Each group accounts for one evaluation of F ; a local secant equation is defined with respect to each group. Otherwise, the method is the same as the $p = q$ case except now the ring is viewed as having q conceptual nodes. Note that the number of transferred messages between physical nodes remains at p (neighbor-to-neighbor, each message of size n).

Globalization can be achieved in a manner similar to the parallel finite-difference algorithm. Indeed, the only change is to replace the finite-difference Jacobian calculation with the local secant update. With this exception, the algorithm described in Figure 8 can be used unaltered.³ We note that if $p = q$ and Newton steps are being successfully used, then each node is involved in exactly one F-evaluation per iteration (in each Newton iteration there are p F-evaluations). However, when the secant update follows a linesearch there is a slight redundancy. Specifically, after the linesearch procedure is completed the value $F(x^+)$ is known. But the local secant update requires only $p - 1$ additional F-evaluations beyond $F(x^+)$. Therefore, under these circumstances either one node remains idle during the parallel evaluation of F for the multiple secant update, or $F(x^+)$ is computed twice.

5.2. Numerical Results. Experimentally we compared the multi-secant method to the parallel finite-difference procedure discussed in Section 4 using problems 9,10,

³ The remarks concerning linesearch versus trust region, made at the end of the previous section, are applicable to the globalized multi-secant method as well.

14 of the Minpack collection and the extended Rosenbrock function (problem 15).

Table 8 provides the results.

TABLE 8
Multi-secant method versus finite-difference Newton method: $p = 16$

Prob	Method	n=50	n=100	n=300
9	MS	3/4.6	3/12.3	3/140.9
9	FD	3/3.0	3/7.9	3/79.8
10	MS	4/12.4	4/46.6	4/733.6
10	FD	4/13.1	4/69.2	4/1447.4
14	MS	20/32.4	23/99.3	23/1175.8
14	FD	9/12.3	9/33.4	9/344.0
15	MS	21/35.2	25/101.9	41/1912.6
15	FD	11/13.7	12/40.6	14/496.1

In general, the secant method requires more iterations. Therefore, when the factorization is the dominant cost - as it is in problems 9, 14, and 15 - then the parallel finite-difference Newton method is faster. However, when F is expensive - as it is in problem 10 - then the multi-secant method is probably preferable. Indeed, we have tried problem 14+ (i.e. increase the expense of a function evaluation in problem 14 by a factor of 100) with $n = 100$: the multi-secant then requires 405.5 seconds compared to 551.7 required by the parallel finite-difference method.

In order to determine where the algorithm spends most of its time, on this test collection, Table 9 provides a breakdown of the total execution time.

TABLE 9
Multi-secant Algorithm Breakdown for $p = 16$, $n = 300$

Prob	Jac. Est.	Factor	Line search	Tri. Solve
9	.04	1.0	.00	.07
10	1.0	.38	.04	.02
14	.05	1.0	.00	.06
15	.02	1.0	.01	.07

Table 9 is fairly similar to Table 7; however, the factorization cost for problem 10 is relatively more expensive (.38 .vs. .11). This is because the dominant cost - Jacobian estimation - has decreased considerably. In a similar vein, the relative Jacobian costs have almost become insignificant in problems 9,10, and 15.

As mentioned above, the multi-secant method can allow for the independent estimation of several groups per node (as opposed to just one group per node). In Table 10 we provide results indicating the effect of varying the number of groups, q , per node.

TABLE 10
Multi-secant: Vary # groups per node. $p = 16$, $n = 300$

Prob	# groups-per-node	# iterations	Total Time
10	1	4	733.6
10	2	4	804.7
10	4	4	932.4
10	8	4	1087.2
15	1	41	1912.6
15	2	40	1679.8
15	4	29	1003.6
15	8	21	853.6

The results are as expected. For problem 10 the computational expense increases as q increases. This is because of the expensive nature of F (and the number of iterations stays constant). Problem 15 exhibits exactly the opposite behaviour: as q increases the execution time decreases. The reason for this is the number of iterations decreases as q increases: this is reasonable since the Jacobian approximations become increasingly accurate as q increases. In general then, the optimal q will depend on the particular problem: the relative costs of evaluating F , factoring the matrix, and the convergence dependence on q , play a role.

6. Conclusions. We have proposed parallel algorithms for the solution of systems of nonlinear equations, $F(x) = 0$. The algorithms are applicable on local-memory multiprocessors (such as a hypercube computer) provided that each processor has significant memory and computational power and the problem dimension is greater than the number of processors.

When F is efficiently computable in a distributed parallel manner then the globalized rank-1 secant method can be efficiently parallelized. Specifically, it is possible to parallelize the Minpack implementation, updating the QR factorization of an approximate Jacobian matrix every step. On balance we are satisfied with the parallel performance of our global parallel secant implementation. However, it should be noted that we have efficiently distributed and parallelized the F -subroutines in our experiments. In general such a task falls to the user and speedup will be strongly affected by the user's success in this task. Note that distributing F by rows, even if feasible, does not always lead to the most efficient distribution of work since it does not exploit the presence of common expensive subexpressions.

Since it is not always possible to efficiently parallelize the computation of F , we have developed parallel finite-difference and multi-secant methods. In general it is difficult to achieve good speedup, relative to the sequential rank-1 secant method, for this class of functions; however, the finite-difference algorithm is efficiently parallelized and the multi-secant method will generally improve upon this (especially for large n/p). It is perhaps possible to further improve upon the efficiency of the multi-secant method by incorporating parallel multi-rank updates to the current distributed QR -factorization.

We have not investigated this possibility yet.

Finally, we note that sparse systems are partially separable (i.e. each component function depends on only a few variables); therefore, in theory, it is usually possible to effectively evaluate $F(x)$ in a distributed parallel manner. Hence, it may be possible to efficiently solve large sparse systems of nonlinear equations using a parallel sparse secant method.

Acknowledgements: Our numerical experiments were performed with the assistance of the Advanced Computing Facility at the Cornell Center for Theory and Simulation in Science and Engineering, which is supported by the National Science Foundation and New York State. Specifically, we are grateful to Ashley Doershug for her help with the Intel iPSC.

REFERENCES

- [1] C. BROYDEN, *A class of methods for solving nonlinear simultaneous equations*, Mathematics of Computation, 19 (1965), pp. 577-593.
- [2] J. E. DENNIS AND R. B. SCHNABEL, *Numerical Methods for Unconstrained Optimization*, Prentice-Hall, 1983.
- [3] G. H. GOLUB AND C. F. VAN LOAN, *Matrix Computations*, The Johns Hopkins University Press, 1983.
- [4] M. HEATH AND C. ROMINE, *Parallel solution of triangular systems on distributed-memory multiprocessors*, Tech. Rep. TM-10384, Oak Ridge National Laboratory, 1987.
- [5] G. LI AND T. F. COLEMAN, *A parallel triangular solver for a distributed memory multiprocessor*, Tech. Rep. CS-86-787, Computer Science Department, Cornell University, 1986.
- [6] ———, *A new method for solving triangular systems on a distributed memory message-passing multiprocessor*, Tech. Rep. CS-87-812, Computer Science Department, Cornell University, 1987.
- [7] C. MOLER, *Matrix computations on distributed memory multiprocessors*, Tech. Rep., Intel Scientific Computers, 1987.
- [8] C. MOLER AND D. S. SCOTT, *Communications utilities for the iPSC*, Tech. Rep., Intel Scientific Computers, 1986.
- [9] J. J. MORÉ, B. S. GARBOW, AND K. E. HILLSTROM, *User guide for Minpack-1*, Tech. Rep. ANL-80-74, Argonne National Laboratory, 1980.
- [10] M. J. D. POWELL, *A hybrid method for nonlinear equations*, in Numerical Methods for Nonlinear Algebraic Equations, P. Rabinowitz, ed., Academic Press, 1970, pp. 87-114.
- [11] C. ROMINE AND J. ORTEGA, *Parallel solution of triangular systems of equations*, Tech. Rep. RM-86-05, Dept. of Applied Mathematics, University of Virginia, 1986.
- [12] Y. SAAD AND M. H. SCHULTZ, *Topological properties of hypercubes*, Tech. Rep. CS-RR-389, Computer Science Department, Yale University, 1985.
- [13] P. WILEY, *A parallel architecture comes of age at last*, IEEE Spectrum, (1987), pp. 46-50.