

LA-UR 97-1605

CONF-970876--1

Approved for public release;  
distribution is unlimited

Title:

The Delayed Coupling Method: An Algorithm for Solving Banded Diagonal Matrix Problems in Parallel

Author(s):

N. Mattor  
T. J. Williams  
D. W. Hewett  
A. M. Dimitis

RECEIVED

AUG 13 1997

OSTI

Submitted to:

IMACS '97  
August 24-29, 1997  
Berlin, Germany

MASTER

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

jcp

**Los Alamos**  
National Laboratory

Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by the University of California for the U.S. Department of Energy under contract W-7405-ENG-36. By acceptance of this article, the publisher recognizes that the U.S. Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or to allow others to do so, for U.S. Government purposes. Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy. The Los Alamos National Laboratory strongly supports academic freedom and a researcher's right to publish; as an institution, however, the Laboratory does not endorse the viewpoint of a publication or guarantee its technical correctness.

## **DISCLAIMER**

**This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, make any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.**

**DISCLAIMER**

**Portions of this document may be illegible  
in electronic image products. Images are  
produced from the best available original  
document.**

# The Delayed Coupling Method: An Algorithm for Solving Banded Diagonal Matrix Problems in Parallel

N. Mattor, T. J. Williams<sup>1</sup>, D. W. Hewett, A. M. Dimits  
Lawrence Livermore National Laboratory  
Livermore, California 94550 USA  
e-mail: mattor@m5.llnl.gov

## ABSTRACT

We present a new algorithm for solving banded diagonal matrix problems efficiently on distributed-memory parallel computers, designed originally for use in dynamic alternating-direction implicit (ADI) partial differential equation solvers. The algorithm optimizes efficiency with respect to the number of numerical operations, and with respect to the amount of interprocessor communication. We refer to our approach as the “delayed coupling method” because the communication is deferred until needed. We focus here on tridiagonal and periodic tridiagonal systems.

## 1. INTRODUCTION

We discuss a new approach to parallel solution of banded linear systems, the “delayed coupling method.” The method is analogous to the solution of an inhomogeneous linear differential equation, where the solution is a “particular” solution added to an arbitrary linear combination of “homogeneous” solutions. The coefficients of the homogeneous solutions are later determined by boundary conditions. In our parallel method, each processor is given a contiguous subsection of a tridiagonal system. With no information about the neighboring subsystems, each processor obtains the solution up to two constants. Then the global solution can be found by matching endpoints.

Our earlier paper[1] has a more detailed description of the method and its application to tridiagonal systems. The algorithm is designed with the following objectives, listed in order of priority. The first objective is to minimize the number of interprocessor communications opened, since this is the most time consuming process. Second, the algorithm allows flexibility of the specific solution method of the tridiagonal submatrices. Here, we employ a variant of LU decomposition, but this is easily replaced with cyclic reduction or other. Third, we wish to minimize storage needs.

## 2. BASIC ALGORITHM

We consider the  $N \times N$  tridiagonal linear system

$$\Lambda X = R, \tag{1}$$

---

<sup>1</sup>Now at Los Alamos National Laboratory





- **Send (ToPid, data, n)**: When invoked by processor FromPid, the array data of length  $n$  is sent to processor ToPid. Send () is *nonblocking*.
- **Receive (FromPid, data, n)**: To complete data transmission, processor ToPid invokes Receive (). Upon execution, the array sent by processor FromPid is stored in the array data array of length  $n$ . Receive () is *blocking* (the processor waits for the data to be received before continuing).

Opening interprocessor communications is generally the most time-consuming step in the entire tridiagonal solution process, so it is important to minimize this. The following algorithm consumes a time of  $T = (\log_2 P)t_c$  in opening communication channels (where  $t_c$  is the time to open one channel).

1. Each processor writes whatever data it has that is relevant to Eq. (6) in the array OutData.
2. The OutData arrays from each processor are concatenated as follows (Fig. 1):
  - (a) Each processor  $p$  sends its OutData array to processor  $p - 1 \pmod{P}$ , and receives a corresponding array from processor  $p + 1 \pmod{P}$ , as depicted in Fig. 1a. The incoming array is concatenated to the end of OutData.
  - (b) At the  $i^{th}$  step, repeat the first step, except sending to processor  $p - 2^{i-1} \pmod{P}$ , and receiving from processor  $p + 2^{i-1} \pmod{P}$  (Fig. 1b,c), for  $i = 1, 2, \dots$ . After  $\log_2 P$  iterations (or the next higher integer), each processor has the contents of the reduced matrix in the OutData array.
3. Each processor rearranges the contents of its OutData array into a local copy of the reduced tridiagonal system, and then solves. At this point, each processor has all the values in Eq. (5) stored locally.

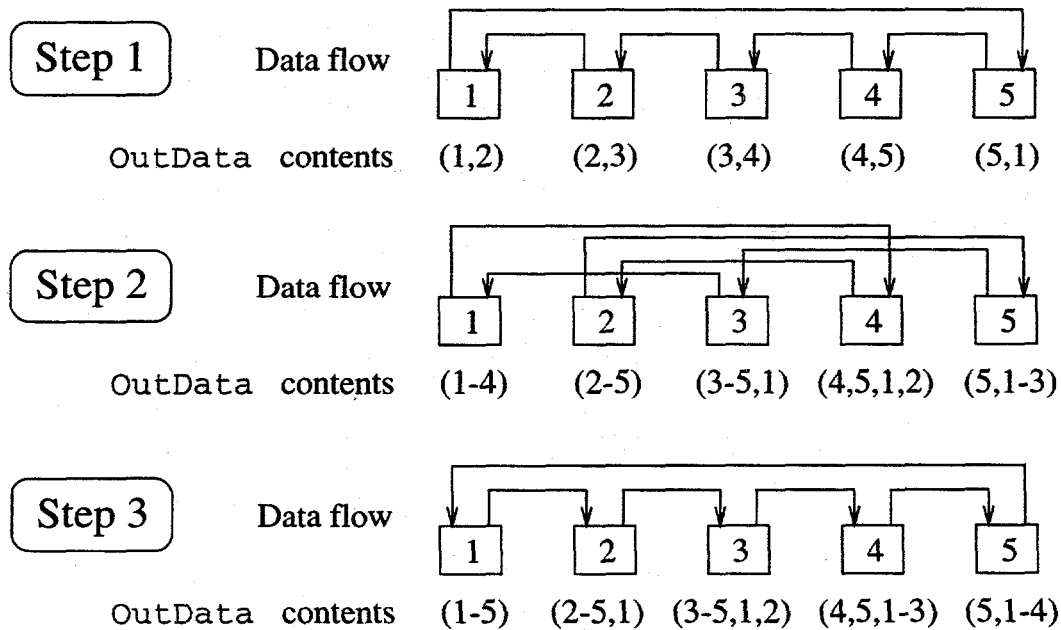


Figure 1: Illustration of the method to pass reduced matrix data between processors, shown for  $P = 5$ .

## 5. PERFORMANCE

The time consumption for this routine is as follows:

1. To calculate the three roots  $x^R$ ,  $x^{UH}$ , and  $x^{LH}$  requires  $13M$  binary floating point operations by each processor, done in parallel.
2. To assemble the reduced matrix in each processor requires  $\log_2 P$  steps where interprocessor communications are opened, and the  $i^{th}$  opening passes  $8 \times 2^{i-1}$  real numbers.
3. Solution of the reduced system through LU decomposition requires  $8(2P - 2)$  binary floating point operations by each processor, done in parallel.
4. Calculation of the final solution requires  $4M$  binary floating point operations by each processor, done in parallel.

If  $t_b$  is the time of one binary floating point operation,  $t_c$  is the time required to open a communication channel (latency), and  $t_p$  is the time to pass one real number once communication is opened, then the time to execute this parallel routine is given by (optimally)

$$\begin{aligned} T_P &\simeq 13Mt_b + (\log_2 P)t_c + 8(P-1)t_p + 8(2P-2)t_b + 4Mt_b \\ &\simeq (17M + 16P)t_b + (\log_2 P)t_c + 8Pt_p, \end{aligned} \quad (7)$$

for  $P \gg 1$ . For cases of present interest,  $T_P$  is dominated by  $(\log_2 P)t_c$  and  $17Mt_b$ . The *parallel efficiency* is defined by  $\epsilon_P \equiv \frac{T_S}{PT_P}$ , where  $T_S$  is the execution time of a serial code which solves by LU decomposition. Serial LU decomposition solves an  $N \times N$  system in a time  $T_S = 8Nt_b$ , so

$$\epsilon_P = \frac{8}{17 + 16P^2/N + (\log_2 P)Pt_c/Nt_b + 8P^2t_p/Nt_b}. \quad (8)$$

To test these claims empirically, we measured the execution times of working serial and parallel codes, and calculated  $\epsilon_P$  both through its definition and through Eq. (8). Fig. 2 shows  $\epsilon_P$  as a function of  $P$  for two cases,  $N = 200$  and  $N = 50,000$ . We conclude from Fig. 2 that Eq. (8) (smooth lines) is reasonably accurate, both for the theoretical maximum efficiency (47%, achieved for small  $P$  and large  $N$ ) and for the scaling with large  $P$ .

We made these timings on the BBN TC2000 machine at Lawrence Livermore National Laboratory, using 64-bit floating point arithmetic. This machine had 128 M88100 RISC processors, connected by a butterfly-switch network. To calculate the predictions of Eq. (7) we chose  $t_c = 750\mu\text{sec}$ , based on the average time of a send/receive pair measured in our code; based on other measurements, we chose the passage time of a single 64-bit real number as  $t_p = 9\mu\text{sec}$ ; we chose  $t_b = 1.4\mu\text{sec}$ , based on our measured timing of 0.00218 sec for the serial algorithm on the  $N = 200$  case.

## 6. PERIODIC TRIDIAGONAL SYSTEM

We have generalized our algorithm to a "periodic tridiagonal system." This is a tridiagonal system with additional nonzero elements in the far upper and lower corners of the matrix, that is, Eq. (1)



This system necessitates a new solution algorithm. The most efficient we know (not shown here) uses  $LU$  decomposition, and requires  $15P$  binary operations. The interesting consequence is that the parallel efficiency nearly doubles over the nonperiodic case, since the operation count in the corresponding serial solver also rises—from  $8N$  to  $15N$ . Thus, Eq. (8) for the predicted efficiency becomes

$$\epsilon_P = \frac{15}{17 + 16P^2/N + (\log_2 P) Pt_c/Nt_b + 15P^2t_p/Nt_b}$$

## 7. DISCUSSION AND CONCLUSIONS

Stability of the parallel tridiagonal algorithm is similar to that of serial  $LU$  decomposition of a tridiagonal matrix. If the  $L_i$  are unstable to  $LU$  decomposition, then pivoting could be used. If the  $L_i$  are singular, then  $LU$  decomposition fails and some alternative should be devised. If the large matrix  $\Lambda$  is diagonally dominant, then so too are the  $L_i$ . If the reduced system is unstable to  $LU$  decomposition, this can be replaced by a different solution scheme, with little loss of overall speed (if  $P \ll M$ ).

This routine is generalizable from tridiagonal to higher systems. For example, in a 5-diagonal system, there would be four homogeneous solutions, each with an undetermined coefficient. The coefficients of the homogeneous solutions would be determined by a reduced system analogous to Eq. (6), except with  $O(4P)$  equations, not  $2P - 2$ .

In our applications of the parallel tridiagonal solver we solve a tridiagonal system along each line of grid points parallel to a given direction. In two or higher dimensions, each processor owns a segment of each of many systems, giving us a strong advantage in interprocessor communication over solving only a single system: each processor solves *all* of its triplets of independent subsystems, then packs together all of the data it needs to send to other processors—there is only one send volley for solving all of the systems. Furthermore, that number of processors each processor communicates with is the number of processors collinear in the one direction, which will generally be smaller than the total number of processors in a multidimensional domain decomposition; this improves parallel efficiency by reducing the value of  $P$  below the total number of processors.

This work was performed for the U.S. Department of Energy at Lawrence Livermore National Laboratory under contract W-7405-ENG-48 and Los Alamos National Laboratory under contract W-7405-ENG-36.

## REFERENCES

- [1] N. Mattor, T. J. Williams, D. W. Hewett, *Algorithm for solving tridiagonal matrix problems in parallel*, *Parallel Computing* 25, p. 1769 (1995).
- [2] R. W. Hockney and J. W. Eastwood, *Computer Simulation Using Particles*, Adam Hilger, Bristol, p. 185 (1988).