Incomplete cyclic reduction of banded and strictly diagonally dominant linear systems

Carl Christian Kjelgaard Mikkelsen and Bo Kågström

Department of Computing Science and HPC2N
Umeå University, Sweden {spock,bokg}@cs.umu.se

Abstract. The ScaLAPACK library contains a pair of routines for solving banded linear systems which are strictly diagonally dominant by rows. Mathematically, the algorithm is complete block cyclic reduction corresponding to a particular block partitioning of the system. In this paper we extend Heller's analysis of incomplete cyclic reduction for block tridiagonal systems to the ScaLAPACK case.

Keywords: Banded or block tridiagonal linear systems, strict diagonal dominance, incomplete cyclic reduction, ScaLAPACK

1 Introduction

Let A be an N by N block tridiagonal matrix, i.e.

$$A = \begin{bmatrix} D_1 & F_1 & & \\ E_2 & \ddots & \ddots & \\ & \ddots & \ddots & F_{N-1} \\ & & E_N & D_N \end{bmatrix}$$

and consider the solution of the linear system

$$Ax = f \tag{1}$$

on a parallel machine. If $A = [a_{ij}]$ is strictly diagonally dominant by rows, i.e

$$\forall\, i \; : \; \sum_{j \neq i} |a_{ij}| < |a_{ii}|,$$

then A is nonsingular and we can use (complete) cyclic reduction to solve the linear system (1). The basic algorithm is due to R.W. Hockney [5] who worked closely with G. H. Golub on the solution of certain tridiagonal linear systems. Heller [4] showed that block cyclic reduction is well defined for matrices which are strictly diagonally dominant by rows. Moreover, Heller introduced the incomplete cyclic reduction algorithm and estimated the truncation error.

Every banded matrix A can be partitioned as a block tridiagonal matrix; only the dimension of the diagonal blocks may not be too small. Specifically, if

A has k superdiagonals and k subdiagonals, then the dimension of each of the diagonal blocks must be at least k.

ScaLAPACK [2] contains a pair of routines (PDDBTRF/PDDBTRS) which can be used to solve narrow banded linear systems which are strictly diagonally dominant by rows. The algorithm is complete block cyclic reduction corresponding to a particular partitioning of the system. The odd numbered diagonal blocks are as large as possible, while the even numbered diagonal blocks are k by k.

In this paper we extend Heller's analysis to the ScaLAPACK case. Central to the analysis is the dominance factor ϵ given by

$$\epsilon = \max_i \left\{ \frac{1}{|a_{ii}|} \sum_{j \neq i} |a_{ij}| \right\}.$$

The strict diagonal dominance of A implies that $\epsilon < 1$.

We review the incomplete cyclic reduction algorithm and some related results in Section 2. The new results are derived in Section 3. We have already proved Theorem 1 for tridiagonal matrices ($k = 1$) in a previous paper [7]. In this paper, we derive Theorem 2 and use it to prove Theorem 1 in the general case.

This presentation owes much to Heller's original paper [4] which we highly recommend. The ScaLAPACK implementation of block cyclic reduction of narrow banded linear systems is described in detail in the paper by Arbenz, Cleary, Dongarra and Hegland [1].

2 The Algorithm

Complete block cyclic reduction of the linear system (1) is equivalent to Gaussian elimination with no pivoting on a permuted system

$$(PAP^T)Px = Pf,$$

where P is a block permutation matrix, which reorders the vector $(1, 2, \dots, N)$, so the odd multiples of 2^0 come first, followed by odd multiples of 2^1, etc.

The incomplete block cyclic reduction algorithm is stated as Algorithm 1. Given an integer m we obtain an approximation y of the true solution as follows. First, we execute m steps of cyclic reduction in order to eliminate the variables which correspond to blocks numbered with an odd multiple of $2^0, 2^1, \dots, 2^{m-1}$ (lines 1 through 3). The m Schur complements can be represented by the block tridiagonal linear systems

$$A^{(j)} x^{(j)} = v^{(j)}, \quad j = 1, 2 \dots, m,$$

generated by Algorithm 1. We approximate the solution of the mth Schur complement by dropping the off diagonal blocks (line 4). The remaining components of y are obtained by backward substitution (lines 5 through 8).

Algorithm 1 The incomplete block cyclic reduction algorithm

Input: An N by N block tridiagonal linear system

$$A^{(0)}x^{(0)} = v^{(0)}$$

where $A^{(0)}$ is strictly diagonally dominant by rows and an integer $m \geq 0$.

Output: An approximation y of the exact solution.

1: **for** $j = 1 : 1 : m$ **do**

2: Assemble the block tridiagonal linear system

$$A^{(j)}x^{(j)} = v^{(j)}$$

where

$$A^{(j)} := \left(E_i^{(j)}, D_i^{(j)}, F_i^{(j)}\right)_{\{i\in\mathbb{N}\,:\,1\leq i\cdot 2^j\leq N\}}, \quad \text{and} \quad \left(v_i^{(j)}\right)_{\{i\in\mathbb{N}\,:\,1\leq i\cdot 2^j\leq N\}}$$

are given by

$$\begin{aligned}
D_i^{(j)} &:= D_{2i}^{(j-1)} - E_{2i}^{(j-1)}\left({D_{2i-1}^{(j-1)}}^{-1} F_{2i-1}^{(j-1)}\right) - F_{2i}^{(j-1)}\left({D_{2i+1}^{(j-1)}}^{-1} E_{2i+1}^{(j-1)}\right)\\
E_i^{(j)} &:= -E_{2i}^{(j-1)}\left({D_{2i-1}^{(j-1)}}^{-1} E_{2i-1}^{(j-1)}\right)\\
F_i^{(j)} &:= -F_{2i}^{(j-1)}\left({D_{2i+1}^{(j-1)}}^{-1} F_{2i+1}^{(j-1)}\right)\\
v_i^{(j)} &:= v_{2i}^{(j-1)} - E_{2i}^{(j-1)}\left({D_{2i-1}^{(j-1)}}^{-1} v_{2i-1}^{(j-1)}\right) - F_{2i}^{(j-1)}\left({D_{2i+1}^{(j-1)}}^{-1} v_{2i+1}^{(j-1)}\right)
\end{aligned}$$

3: **end for**

4: Define the block diagonal matrix

$$D^{(m)} := \mathrm{diag}(D_i^{(m)})_{\{i\in\mathbb{N}\,:\,1\leq i\cdot 2^m\leq N\}}$$

and solve the system

$$D^{(m)}y^{(m)} = v^{(m)}$$

with respect to $y^{(m)}$.

5: **for** $j = m : -1 : 1$ **do**

6: Set $y_{2i}^{(j-1)} = y_i^{(j)}$

7: Solve

$$D_{2i-1}^{(j-1)} y_{2i-1}^{(j-1)} = v_{2i-1}^{(j-1)} - E_{2i-1}^{(j-1)} y_{2i-2}^{(j-1)} - F_{2i-1}^{(j-1)} y_{2i+2}^{(j-1)}$$

with respect to $y_{2i-1}^{(j-1)}$.

8: **end for**

9: Return y given by

$$y = ({y_1^{(0)}}^T, {y_2^{(0)}}^T, \ldots, {y_N^{(0)}}^T)^T.$$

Heller [4] showed that if A is strictly diagonally dominant by rows, then the incomplete cyclic reduction algorithm is well defined. Moreover, the significance of the off diagonal blocks can be measured using the auxiliary matrices $B^{(j)}$ given by

$$B^{(j)} = {D^{(j)}}^{-1}(A^{(j)} - D^{(j)}).$$

Heller [4] showed that

$$\|x-y\|_\infty \leq \|B^{(m)}\|_\infty \|x\|_\infty.$$

and he derived the central estimate

$$\|B^{(j+1)}\|_\infty \leq \|{B^{(j)}}^2\|_\infty \leq \|B^{(j)}\|_\infty^2 < 1, \quad j = 0,1,2,\dots,$$

which shows that the error decays quadratically with increasing j. Recently, we have shown that

$$\|B^{(j)}\|_\infty \leq \epsilon^{2^j}$$

and this estimate is tight [7]. This general estimate carries to the ScaLAPACK case, but it does not reflect the structure of the odd numbered diagonal blocks. Theorem 1 shows how to integrate this information into the analysis.

3 The Main Result

Theorem 1. *If the odd numbered diagonal blocks can be partitioned as block tridiagonal matrices with q diagonal blocks, then*

$$\|B^{(1)}\|_\infty \leq \epsilon^{1+q}. \tag{2}$$

Moreover, this estimate is tight.

It is straight forward to verify that equality is achieved for block tridiagonal matrices for which $E_i = 0$, $D_i = I_k$ and $F_i = \epsilon I_k$, where I_k denotes the k by k identity matrix.

We now reduce the proof of the central inequality (2) to a single application of Theorem 2. Normally, we would illustrate the cyclic reduction algorithm using explicit block permutations, but here we use “in-place” computations in order to estimate $B^{(1)}$. Now, let

$$G = \begin{bmatrix} E_{-1} & D_{-1} & F_{-1} & & \\ & E_0 & D_0 & F_0 & \\ & & E_1 & D_1 & F_1 \end{bmatrix}$$

be a compressed representation of three consecutive block rows drawn for a block tridiagonal linear system. Similarly, let

$$A_0^{(1)} = \begin{pmatrix} E_0^{(1)} & D_0^{(1)} & F_0^{(1)} \end{pmatrix}$$

be a representation of the corresponding block row of the Schur complement. Specifically, we have

$$D_0^{(1)} = D_0 - E_0 D_{-1}^{-1} F_{-1} - F_0 D_1^{-1} F_1$$

and

$$E_0^{(1)} = -E_0 D_{-1}^{-1} E_{-1}, \quad F_0^{(1)} = -F_0 D_1^{-1} F_1.$$

Our goal is to estimate the infinity norm of the matrix $B_0^{(1)}$ given by

$$B_0^{(1)} = \left[U_0^{(1)} \middle| V_0^{(1)} \right]$$

where

$$D_0^{(1)} \left[U_0^{(1)} \middle| V_0^{(1)} \right] = \left[E_0^{(1)} \middle| F_0^{(1)} \right].$$

To this end we will now identify $B_0^{(1)}$ with a certain submatrix of a matrix G' which is row equivalent to G. Let $\left[U_i, V_i\right]$ be the solution of the linear system

$$D_i \left[U_i, V_i\right] = \left[E_i, F_i\right], \quad i \in \{-1, 0, 1\}.$$

Then

$$G \sim \left[\begin{array}{c|c|c|c|c} U_{-1} & I & V_{-1} & & \\ \hline & U_0 & I & V_0 & \\ \hline & & U_1 & I & V_1 \end{array}\right]$$

$$\sim \left[\begin{array}{c|c|c|c|c} U_{-1} & I & V_{-1} & & \\ \hline -U_0 U_{-1} & & (I - U_0 V_{-1} - V_0 U_1) & & -V_0 V_1 \\ \hline & & U_1 & I & V_1 \end{array}\right].$$

Moreover, it is straight forward to verify that

$$G \sim \left[\begin{array}{c|c|c|c|c} U_{-1} & I & V_{-1} & & \\ \hline U_0^{(1)} & & I & & V_0^{(1)} \\ \hline & & U_1 & I & V_1 \end{array}\right]$$

$$\sim \left[\begin{array}{c|c|c|c|c} U_{-1} - V_{-1} U_0^{(1)} & I & & & -V_{-1} V_0^{(1)} \\ \hline U_0^{(1)} & & I & & V_0^{(1)} \\ \hline -U_1 U_0^{(1)} & & & I & V_1 - U_1 V_0^{(1)} \end{array}\right] =: G'.$$

Theorem 2 details the structure of the matrix G' and Theorem 1 is an immediate consequence.

Theorem 2. *Let $q \in \mathbb{N}$ and let G_q be a representation of $2q+1$ consecutive block rows of a block tridiagonal matrix A which is strictly diagonally dominant by rows with dominance factor ϵ, i.e.*

$$G_q = \left[\begin{array}{c|ccc|c|ccc|c} E_{-q} & D_{-q} & F_{-q} & & & & & & \\ & \ddots & \ddots & \ddots & & & & & \\ & & \ddots & \ddots & \ddots & & & & \\ & & & E_{-1} & D_{-1} & F_{-1} & & & \\ \hline & & & & E_0 & D_0 & F_0 & & \\ \hline & & & & & E_1 & D_1 & F_1 & \\ & & & & & & \ddots & \ddots & \ddots \\ & & & & & & & \ddots\ \ddots & \ddots \\ & & & & & & & E_q\ D_q & F_q \end{array}\right].$$

Then G_q is row equivalent to a unique matrix K_q of the form

$$K_q = \left[\begin{array}{c|ccc|c|ccc|c} \mathcal{U}_{-q}^{(q)} & I & & & & & & & \mathcal{V}_{-q}^{(q)} \\ \vdots & & \ddots & & & & & & \vdots \\ \vdots & & & \ddots & & & & & \vdots \\ \mathcal{U}_{-1}^{(q)} & & & I & & & & & \mathcal{V}_{-1}^{(q)} \\ \hline \mathcal{U}_{0}^{(q)} & & & & I & & & & \mathcal{V}_{0}^{(q)} \\ \hline \mathcal{U}_{1}^{(q)} & & & & & I & & & \mathcal{V}_{1}^{(q)} \\ \vdots & & & & & & \ddots & & \vdots \\ \vdots & & & & & & & \ddots & \vdots \\ \mathcal{U}_{q}^{(q)} & & & & & & & I & \mathcal{V}_{q}^{(q)} \end{array}\right] \tag{3}$$

where the spikes decay exponentially as we move towards the main block row. Specifically, if we define

$$Z_i^{(q)} = \left[\begin{array}{c|c} \mathcal{U}_{-i}^{(q)} & \mathcal{V}_{-i}^{(q)} \\ \hline U_i^{(q)} & V_i^{(q)} \end{array}\right], \quad 0 < i \le q,$$

and

$$Z_0^{(q)} = \left[\begin{array}{c|c} \mathcal{U}_0^{(q)} & \mathcal{V}_0^{(q)} \end{array}\right],$$

then

$$\|Z_i^{(q)}\|_\infty \le \epsilon^{1+q-i}, \quad 0 \le i \le q.$$

Proof. The existence and uniqueness of the matrix K_q is immediate. Moreover, Corollary 3.2 [6] implies that

$$\|Z_i^{(q)}\|_\infty \le \epsilon < 1, \quad 0 \le i \le q.$$

In particular, we already know that

$$\|Z_q^{(q)}\| \le \epsilon = \epsilon^{(1+q)-q}$$

and we are free to concentrate on the “interior” case of $0 \le i < q$.

We shall prove the decay of the spikes using the principle of mathematical induction. Let Ω be given by

$$\Omega = \{q \in \mathbb{N} \mid \forall i \le q \,:\, \|Z_i^{(q)}\|_\infty \le \epsilon^{1+q-i}\}.$$

Our goal is to show that $\Omega = \mathbb{N}$. We begin by showing that $1 \in \Omega$. We have

$$G_1 = \begin{bmatrix} E_{-1} & D_{-1} & F_{-1} & & \\ & E_0 & D_0 & F_0 & \\ & & E_1 & D_1 & F_1 \end{bmatrix} \sim \begin{bmatrix} U_{-1} & I & V_{-1} & & \\ & U_0 & I & V_0 & \\ & & U_1 & I & V_1 \end{bmatrix}$$

$$\sim \begin{bmatrix} U_{-1} & I & V_{-1} & & \\ \mathcal{U}_0^{(1)} & & I & & \mathcal{V}_0^{(1)} \\ & & U_1 & I & V_1 \end{bmatrix} \sim \begin{bmatrix} \mathcal{U}_{-1}^{(1)} & I & & & \mathcal{V}_{-1}^{(1)} \\ \mathcal{U}_0^{(1)} & & I & & \mathcal{V}_0^{(1)} \\ \mathcal{U}_1^{(1)} & & & I & \mathcal{V}_0^{(1)} \end{bmatrix},$$

where

$$D_i \left[U_i \middle| V_i \right] = \left[E_i \middle| F_i \right]$$

and

$$Z_0^{(1)} = \left[\mathcal{U}_0^{(1)} \middle| \mathcal{V}_0^{(1)} \right] = -(I - U_0 V_{-1} - V_0 U_1)^{-1} \left[U_0 U_{-1} \middle| V_0 V_1 \right].$$

The exact formula for

$$\left[\mathcal{U}_i^{(1)} \middle| \mathcal{V}_i^{(1)} \right], \quad i = \pm 1$$

is irrelevant, because we already know that

$$\left\| \left[\mathcal{U}_i^{(1)} \middle| \mathcal{V}_i^{(1)} \right] \right\|_\infty \leq \epsilon, \quad i = -1, 0, 1.$$

It remains to be seen that $\|Z_0^{(1)}\|_\infty \leq \epsilon^2$. By definition

$$\begin{aligned} Z_0^{(1)} &= \left[U_0 \middle| V_0 \right] \left[\frac{V_{-1}}{U_1} \right] Z_0^{(1)} - \left[U_0 U_{-1} \middle| V_0 V_1 \right] \\ &= \left[U_0 \middle| V_0 \right] \left\{ \left[\frac{V_{-1}}{U_1} \right] \left[\mathcal{U}_0^{(1)} \middle| \mathcal{V}_0^{(1)} \right] - \left[\begin{array}{c|c} U_{-1} & 0 \\ \hline 0 & V_1 \end{array} \right] \right\} \\ &= \left[U_0 \middle| V_0 \right] \left[\begin{array}{c|c|c} V_{-1} & U_{-1} & 0 \\ \hline U_1 & 0 & V_1 \end{array} \right] \left[\begin{array}{c|c} \mathcal{U}_0^{(1)} & \mathcal{V}_0^{(1)} \\ \hline -I & 0 \\ \hline 0 & -I \end{array} \right] \end{aligned}$$

and since we already know that $\|Z_0^{(1)}\|_\infty < 1$ we can conclude.

$$\|Z_0^{(1)}\|_\infty \leq \epsilon^2 \max\{\|Z_0^{(1)}\|_\infty, 1\} = \epsilon^2,$$

and we have shown that $1 \in \Omega$. Now assume that $1 < q$ and $q - 1 \in \Omega$. We must show that $q \in \Omega$. By assumption

$$G_q \sim \left[\begin{array}{c|c|cccccc|c|c} U_{-q} & I & V_{-q} & & & & & & & \\ \hline & \mathcal{U}_{1-q}^{(q-1)} & I & & & & & & \mathcal{V}_{1-q}^{(q-1)} & \\ & \vdots & & \ddots & & & & & \vdots & \\ & \mathcal{U}_{-1}^{(q-1)} & & & I & & & & \mathcal{V}_{-1}^{(q-1)} & \\ & \mathcal{U}_{0}^{(q-1)} & & & & I & & & \mathcal{V}_{0}^{(q-1)} & \\ & \mathcal{U}_{1}^{(q-1)} & & & & & I & & \mathcal{V}_{1}^{(q-1)} & \\ & & & & & & & \ddots & & \\ & \mathcal{U}_{q-1}^{(q-1)} & & & & & & \ddots & \mathcal{V}_{1}^{(q-1)} & \\ \hline & & & & & & & U_q & I & V_q \end{array} \right]$$

$$=: \left[\begin{array}{c|c|c|c|c} U_{-q} & I & V'_{-q} & & \\ \hline & \mathcal{U}^{(q-1)} & I & \mathcal{V}^{(q-1)} & \\ \hline & & U'_q & I & V_q \end{array} \right],$$

where we have made the last definition to emphasize the similarity between our current situation and the problem of showing that $1 \in \Omega$.

We now continue to reduce G_q using elementary block row operations. We have

$$G_q \sim \left[\begin{array}{c|c|c|c|c} U_{1-q} & I & V'_{1-q} & & \\ \hline \xi^{(q)} & & I & & \nu^{(q)} \\ \hline & & U'_{q-1} & I & V_{q-1} \end{array}\right]$$
$$\sim \left[\begin{array}{c|c|c|c|c} U_{1-q} - V'_{1-q}\xi^{(q)} & I & & & -V'_{1-q}\nu^{(q)} \\ \hline \xi^{(q)} & & I & & \nu^{(q)} \\ \hline -U'_{q-1}\xi^{(q)} & & & I & V_{q-1} - U'_{q-1}\nu^{(q)} \end{array}\right] = K_q,$$

where the last equality is critical, but follows immediately from the uniqueness of the matrix K_q. Now, it is straightforward to verify that

$$\left[\xi^{(q)} \middle| \nu^{(q)}\right] = -\left(I - \mathcal{U}^{(q-1)}V'_{1-q} - \mathcal{V}^{(q-1)}U'_{q-1}\right)^{-1} \left[\mathcal{U}^{(q-1)}U_{1-q} \middle| \mathcal{V}^{(q-1)}V_{q-1}\right],$$

or equivalently

$$\left[\xi^{(q)} \middle| \nu^{(q)}\right] = \left[\mathcal{U}^{(q-1)} \middle| \mathcal{V}^{(q-1)}\right] \left\{ \left[\begin{array}{c} V'_{1-q} \\ \hline U'_{q-1} \end{array}\right] \left[\xi^{(q)} \middle| \nu^{(q)}\right] - \left[\begin{array}{c|c} U_{1-q} & 0 \\ \hline 0 & V_{q-1} \end{array}\right] \right\}$$
$$= \left[\mathcal{U}^{(q-1)} \middle| \mathcal{V}^{(q-1)}\right] \left[\begin{array}{c|c|c} V'_{1-q} & -U_{1-q} & 0 \\ \hline U'_{q-1} & 0 & -V_{q-1} \end{array}\right] \left[\begin{array}{c|c} \xi^{(q)} & \nu^{(q)} \\ \hline I & 0 \\ \hline 0 & I \end{array}\right]. \quad (4)$$

It follows that equation (4) is just a compact way of expressing the fact that

$$Z_i^{(q)} = Z_i^{(q-1)} \left[\begin{array}{c|c|c} V'_{1-q} & -U_{1-q} & 0 \\ \hline U'_{q-1} & 0 & -V_{q-1} \end{array}\right] \left[\begin{array}{c|c} \xi^{(q)} & \nu^{(q)} \\ \hline I & 0 \\ \hline 0 & I \end{array}\right], \quad i < q,$$

and this relation will allow us to estimate $\|Z_i^{(q)}\|_\infty$. By assumption, $q-1 \in \Omega$ or equivalently

$$\|Z_i^{(q-1)}\|_\infty \leq \epsilon^{q-i}, \quad i \leq q-1.$$

It follows, that

$$\|Z_i^{(q)}\|_\infty \leq \epsilon^{q-i}\epsilon \max\{\| \left[\xi^{(q)} \ \nu^{(q)}\right] \|_\infty, 1\} = \epsilon^{1+q-i}, \quad i < q,$$

because $\| \left[\xi^{(q)} \ \nu^{(q)}\right] \|_\infty < 1$. The extreme case of $i = q$ was settled at the beginning of the proof. This shows that $q \in \Omega$ and the proof is complete. □

.

4 Conclusion and Future Work

Incomplete cyclic reduction applies to systems which are block tridiagonal and strictly diagonally dominant by rows. The ScaLAPACK library contains two subroutines which can be used to factor and solve narrow banded linear systems which are strictly diagonally dominant by rows. Mathematically, the algorithm is complete cyclic reduction corresponding to a particular partitioning of the system.

We have extended Heller's analysis of incomplete cyclic reduction to the ScaLAPACK case. Our analysis can be used to determine when it is to omit all but the first few reduction steps. The argument proceeds as follows. By design the ScaLAPACK implementation applies to systems which are narrow banded ($k \ll n$) and very large compared with the number of processors ($1 \ll n/p$). Suppose we want an approximation so that the relative forward error is at most $10u$, where u is the unit round off error! When is this even possible? A row scaling will ensure that the diagonal entries are all ones and it is straightforward to verify that the condition number satisfies

$$\kappa_\infty(A) \leq \frac{1+\epsilon}{1-\epsilon}$$

and that equality is possible for matrices of the type

$$D_i = I, \quad E_i = F_i = \epsilon I.$$

In general, the coefficients are not floating point numbers, and storing the system in memory induces a relative error which satisfies

$$\frac{\|\Delta x\|_\infty}{\|x\|_\infty} \leq \kappa_\infty(A) u.$$

Therefore, it is not unreasonable to demand $\frac{1+\epsilon}{1-\epsilon} \leq 10$ or equivalently $\epsilon \leq \frac{9}{11}$. Now, a single step of incomplete cyclic reduction will reduce the truncation error below the unit round off error, provided $\epsilon^q \leq u$ or equivalently $\frac{\log u}{\log \epsilon} \leq q$. With $\epsilon = \frac{9}{11}$, we see that $184 \leq q$ will suffice. If, say, $92 \leq q$, then we can not be certain that the first step will suffice, but two steps will.

Originally, we sought to determine if it was possible to accelerate the solution by eliminating all but the first few steps. Theorem 1 can be used to determine the smallest number of steps necessary to reduce the relative forward error to a predetermined tolerance. However, unless the interconnect is very slow, the savings are negligible, because the vast majority of the time is consumed during the very first step.

As a result we have reconsidered the entire approach. The factorization phase of the ScaLAPACK implementation has an arithmetic redundancy which can be 4 or higher [1]. In contrast, the truncated SPIKE algorithm by Polizzi and Sameh uses an UL/LU factorization scheme [9] to achieve a redundancy of about 2. Mathematically, the algorithms are similar to the point where the analysis reduces to an application of Theorem 2, see [8]. The UL/LU factorization technique applies equally to the ScaLAPACK setting and this is certainly one way of improving on the ScaLAPACK code.

However, the goal is derive an algorithm which does even better. This has been accomplished for tridiagonal matrices by M. Hegland in the context of wrap-around partitions [3]. The arithmetic redundancy for Hegland's scheme was asymptotically equal to 1 and a nearly perfect speedup was recorded.

These observations have lead us to concentrate on two separate problems. Firstly, the application of Hegland's ideas to the SPIKE/ScaLAPACK setting. Secondly, the extension of Theorem 2 to matrices which are symmetric as well as strictly diagonally dominant by rows. This is an common situation for which the general estimate of Theorem 2 can be very pessimistic.

Acknowledgments. The work is supported by eSSENCE, a collaborative e-Science programme funded by the Swedish Research Council within the framework of the strategic research areas designated by the Swedish Government. In addition, support has been provided by the Swedish Foundation for Strategic Research under the frame program A3 02:128 and the EU Mål 2 project UMIT.

References

1. Arbenz, P., Cleary, A., Dongarra, J., Hegland, M.: *A comparison of parallel solvers for diagonally dominant and general narrow-banded linear systems.* In: EuroPar '99 Parallel Processing, P. Amestoy, Ph. Berger, M. Dayd, I. Duff., V. Frayss, L. Giraud, D. Ruiz, eds., pp. 1078-1087, Springer, Berlin (1999)
2. Blackford, L.S., Choi, J., Cleary, A., D'Azevedo, E., Demmel, J., Dhillon, I., Dongarra, J., Hammarling, S., Henry, G., Petitet, A., Stanley, K., Walker, D., Whaley, R. C.: *ScaLAPACK User's Guide*, SIAM, USA (1997)
3. Hegland, M.: *On the parallel solution of tridiagonal linear systems by wrap-around partitioning and incomplete LU factorization.* Numer. Math., vol. 59 (1991), pp. 453-472 .
4. Heller, D.: *Some aspects of the cyclic reduction algorithm for block tridiagonal linear systems.* SIAM J. Numer. Anal., vol. 13 (1976), no. 4, pp. 484-496.
5. Hockney, R.W.: *A fast direct solution of Poisson's equation using Fourier analysis. J. Assoc. Comput. Mach.*, vol. 12 (1965), no. 1, pp. 95-113.
6. Mikkelsen, C.C.K., Manguoglu, M.: Analysis of the truncated SPIKE algorithm. SIAM J. Matrix Anal. Appl., vol. 30 (2008), no. 4, pp. 1500-1519.
7. Mikkelsen, C.C.K, Kågström, B.: *Parallel solution of narrow banded diagonally dominant linear systems.* Accepted for publication in the proceedings of PARA-2010: State of the Art in Scientific and Parallel Computing, Reykjavik, Iceland, June 6-9th 2010.
8. Mikkelsen, C.C.K.: The explicit SPIKE algorithm: Iterative solution of the reduced system. Accepted for publication in the proceeding of the Conference on High Performance Scientific Computing: Architectures, Algorithms and Applications, October 10-12th, 2010, Purdue University, West Lafayette, Indiana, USA.
9. Polizzi, E., Sameh, A.: A parallel hybrid banded system solver: The SPIKE algorithm. Parallel Comput., vol. 32 (2006), no. 2, pp. 177-194.