<http://www.diva-portal.org>

Postprint

This is the accepted version of a paper presented at *9th International Conference on Parallel Processing and Applied Mathematics, SEP 11-14, 2011, Torun, POLAND*.

Citation for the original published paper:

Kjelgaard Mikkelsen, C C., Kågström, B. (2012)

Incomplete cyclic reduction of banded and strictly diagonally dominant linear systems.

In: Wyrzykowski, R. (ed.), *PARALLEL PROCESSING AND APPLIED MATHEMATICS, PT I* (pp. 80-91). Springer

Lecture Notes in Computer Science

https://doi.org/10.1007/978-3-642-31464-3_9

N.B. When citing this work, cite the original published paper.

Permanent link to this version:

<http://urn.kb.se/resolve?urn=urn:nbn:se:umu:diva-50786>

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. We obtain a tight estimate on the significance of the off diagonal blocks of the tridiagonal linear systems generated by the cyclic reduction algorithm. Numerical experiments illustrate the advantage of omitting all but the first reduction step for a class of matrices related to high order approximations of the Laplace operator.

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

1 Introduction

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

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

and consider the solution of the linear system

$$Ax = f \quad (2)$$

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 we can use cyclic reduction to solve the linear system (2). The basic algorithm is due to R.W. Hockney [4] who worked closely with G. H. Golub on the solution of certain tridiagonal linear systems. Heller [3] showed that if A is

strictly diagonally dominant, then the cyclic reduction algorithm runs to completion. 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 (1) only the dimension of the diagonal blocks D_i may not be too small. Specifically, if A has k superdiagonals and k subdiagonals, then the dimension of each D_i 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

$$\epsilon = \max_i \left\{ \frac{1}{|a_{ii}|} \sum_{j \neq i} |a_{ij}| \right\} \geq 0.$$

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. In Section 4 we demonstrate the existence of a class of physically relevant linear systems for which it is numerically sound to neglect all but the first reduction step, an observation which greatly improves the strong scalability of the algorithm.

This presentation owes much to Heller's original paper [3] 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 (2) 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 (lines 1 through 3) in order to eliminate the variables which correspond to blocks numbered with an odd

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

$$D_i^{(j)} := D_{2i}^{(j-1)} - E_{2i}^{(j-1)} \left(D_{2i-1}^{(j-1)} \right)^{-1} F_{2i-1}^{(j-1)} - F_{2i}^{(j-1)} \left(D_{2i+1}^{(j-1)} \right)^{-1} E_{2i+1}^{(j-1)}$$

$$E_i^{(j)} := -E_{2i}^{(j-1)} \left(D_{2i-1}^{(j-1)} \right)^{-1} E_{2i-1}^{(j-1)}$$

$$F_i^{(j)} := -F_{2i}^{(j-1)} \left(D_{2i+1}^{(j-1)} \right)^{-1} F_{2i+1}^{(j-1)}$$

$$v_i^{(j)} := v_{2i}^{(j-1)} - E_{2i}^{(j-1)} \left(D_{2i-1}^{(j-1)} \right)^{-1} v_{2i-1}^{(j-1)} - F_{2i}^{(j-1)} \left(D_{2i+1}^{(j-1)} \right)^{-1} v_{2i+1}^{(j-1)}$$

3: **end for**

4: Define the block diagonal matrix

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

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

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

multiple of $2^0, 2^1, \dots, 2^{m-1}$. The remaining components can be found by solving the block tridiagonal linear system

$$A^{(m)}x^{(m)} = v^{(m)}. \quad (3)$$

We approximate the solution of (3) by dropping the off diagonal blocks (line 4). The remaining components of y are obtained by backward substitution (lines 5 through 8).

Heller [3] measured the significance of the off diagonal blocks using the auxiliary matrices $B^{(j)}$ given by

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

Heller [3] showed that

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

and derived the central estimate

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

which shows that the error decays quadratically. 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 banded 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*

$$\|B^{(1)}\|_\infty \leq \epsilon^{1+q}. \quad (4)$$

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 (4) 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

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

where

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

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

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

Then

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

Moreover, it is straight forward to verify that

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

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.

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

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

$$K_q = \left[\begin{array}{c|c|c|c} \mathcal{U}_{-q}^{(q)} & I & & \mathcal{V}_{-q}^{(q)} \\ \vdots & \ddots & & \vdots \\ \vdots & & \ddots & \vdots \\ \mathcal{U}_{-1}^{(q)} & & I & \mathcal{V}_{-1}^{(q)} \\ \hline \mathcal{U}_0^{(q)} & & I & \mathcal{V}_0^{(q)} \\ \hline \mathcal{U}_1^{(q)} & & I & \mathcal{V}_1^{(q)} \\ \vdots & & \ddots & \vdots \\ \vdots & & & \ddots \\ \mathcal{U}_q^{(q)} & & & I & \mathcal{V}_q^{(q)} \end{array} \right] \quad (5)$$

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

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

and

$$Z_0^{(q)} = \left[\mathcal{U}_0^{(q)} \middle| \mathcal{V}_0^{(q)} \right],$$

then

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

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

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

In particular, we already know that

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

and we can focus on the “interior” case of $0 \leq i < q$.

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

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

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

$$\begin{aligned} G_1 &= \begin{bmatrix} E_{-1} & D_{-1} & F_{-1} \\ & E_0 & D_0 & F_0 \\ & & E_1 & D_1 & F_1 \end{bmatrix} \sim \begin{bmatrix} U_{-1} & I & V_{-1} \\ & U_0 & I & V_0 \\ & & U_1 & I & V_1 \end{bmatrix} \\ &\sim \begin{bmatrix} U_{-1} & I & V_{-1} \\ \mathcal{U}_0^{(1)} & I & \mathcal{V}_0^{(1)} \\ & U_1 & I & V_1 \end{bmatrix} \sim \begin{bmatrix} \mathcal{U}_{-1}^{(1)} & I & \mathcal{V}_{-1}^{(1)} \\ \mathcal{U}_0^{(1)} & I & \mathcal{V}_0^{(1)} \\ \mathcal{U}_1^{(1)} & I & \mathcal{V}_0^{(1)} \end{bmatrix}, \end{aligned}$$

where

$$D_i [U_i | V_i] = [E_i | F_i]$$

and

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

The exact formula for

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

is irrelevant, because we already know that

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

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

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

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

$$G_q \sim \left[\begin{array}{c|c|c|c|c} U_{-q} & I & V_{-q} & & \\ \hline & \mathcal{U}_{1-q}^{(q-1)} & I & & \mathcal{V}_{1-q}^{(q-1)} \\ & \vdots & \ddots & & \vdots \\ & \mathcal{U}_{-1}^{(q-1)} & & I & \mathcal{V}_{-1}^{(q-1)} \\ & \mathcal{U}_0^{(q-1)} & & & I & \mathcal{V}_0^{(q-1)} \\ & \mathcal{U}_1^{(q-1)} & & & & I & \mathcal{V}_1^{(q-1)} \\ & \vdots & & & & \ddots & \vdots \\ & \mathcal{U}_{q-1}^{(q-1)} & & & & \ddots & \mathcal{V}_{q-1}^{(q-1)} \\ \hline & & & & U_q & I & V_q \end{array} \right]$$

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

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 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

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

or equivalently

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

$$= [\mathcal{U}^{(q-1)} | \mathcal{V}^{(q-1)}] \left[\begin{array}{c|c|c} V'_{1-q} & -U_{1-q} & 0 \\ \hline U'_{q-1} & 0 & -V_{q-1} \end{array} \right] \left[\begin{array}{c|c} \xi^{(q)} & \nu^{(q)} \\ \hline I & 0 \\ \hline 0 & I \end{array} \right]. \quad (6)$$

It follows that equation (6) 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

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

because $\|\xi^{(q)} \nu^{(q)}\|_\infty < 1$. 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. $\square$

4 Numerical Experiments

An incomplete cyclic reduction algorithm was implemented in Fortran 90 using MPI. We inherited the communication pattern from the ScaLAPACK implementation, but choose the memory layout of the Arbenz-Hegland implementation [1]. In reality we rewrote the ScaLAPACK implementation from scratch inserting a parameter controlling the number of reduction steps and replacing the blocking receive instructions from BLACS with nonblocking ones. As in the ScaLAPACK implementation [1] we split the algorithm into two phases: a factorization phase which exclusively references the matrix and a solve phase which references the right hand side in order to compute the solution.

Here we report on our findings for a very specific class of matrices, namely the pentadiagonal Toeplitz matrices $A \in \mathbb{R}^{n \times n}$ given by

$$A = \text{diag}(\beta, \alpha, 1, \alpha, \beta), \quad \alpha = \frac{334}{899}, \quad \beta = \frac{43}{1798}.$$

These matrices occur naturally in connection with the construction of highly accurate approximations of the Laplace operator [5]. We have $\epsilon = 2(\alpha + \beta) \approx 0.7909$. The worst case estimate given by Theorem 1 is pessimistic. Specifically, in the case of $q = 25$, a direct calculation establishes that

$$\|B^{(1)}\|_\infty \approx 8.5652 \times 10^{-24} \ll \epsilon^{26} \approx 2.2431 \times 10^{-3}.$$

It follows that we can safely discard all but the first reduction step even if the banded segment assigned to each core contains as little as $\mu = 50$ rows.

We measured the performance of our implementation on the supercomputer AKKA at HPC2N. Briefly, AKKA consists of 672 nodes using an Infiniband interconnect. Each node contains two Intel Xeon quad-core L5420 CPU with a total of 16 GB RAM per node.

We used $p \in \{1, 2, 4, \dots, 512\}$ cores, drawn from a fixed set of 64 nodes containing 512 cores. We used $n \in \{26622, 52222, 103422\}$. The dimensions were chosen so that each core would always receive a banded segment of exactly the same size regardless of the number of cores. For example $n = 26622$ corresponds to having $\mu = 50$ when using $p = 512$ cores, because $n = \mu p + k(p-1)$, where $k =$

2 is the number of super/subdiagonals. Finally, we used $r \in \{1, 10, 20, 30, 40, 50\}$. Every column of every solution was obtained with a forward normwise relative error no greater than $12u$, where $u = 2^{-53}$ is the double precision round-off error. Every experiment was programmed into a single job and submitted to the queue.

The *strong scalability efficiency* ξ_{strong} is defined as

$$\xi_{\text{strong}} = \frac{T_1/p}{T_p},$$

where p is the number of cores, T_1 is the execution time on a single core and T_p is the execution time on p cores. When $p = 1$ the parallel overhead vanishes and our implementation reduces to LAPACKs banded solver save for a few conditional statements and system calls. The measured efficiencies are displayed in Figures 1 through 3. For the sake of visual clarity we only display the solve efficiencies for $r \in \{1, 10, 50\}$. We shall now discuss the graphs in some detail.

Gaussian elimination has a very low arithmetic intensity for narrow banded linear systems. Spreading the same system across an increasing number of cores will eventually create a situation where a significant fraction of the system resides in the cache memory allowing for efficiencies greater than unity which is exactly what we have experienced.

The graphs are not monotone and are a bit erratic. By applying the law of large numbers we conclude that it is likely that the oscillations can be reduced by a factor of, say, 100 by repeating each experiment $100^2 = 10^4$ times and computing the average run-time. However, even 25 repetitions required more than 250 machine hours, so 10^4 repetitions was not feasible for this paper.

For the factorization phase we were surprised to see that the experiments are inconclusive as the efficiencies are very similar. However, for the solve phase we see that the efficiency is improved considerably by neglecting all but the first reduction step. This effect is present regardless of n and r , but it most pronounced for $n = 26666$ and $r = 1$, where the ratio of the parallel work to the communication cost is minimal.

For both algorithms we see the efficiency of the solve phase increase with the number of right hand sides. This effect is present for all n and is due to the fact that the number of messages is independent of r , while the parallel work is a monotone increasing function of r .

Acknowledgments. The work is supported by eSENCE, 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 UNIT. This research was conducted using the resources of High Performance Computing Center North (HPC2N).

Fig. 1: Strong scalability efficiency curves for $n = 26622$. Notice the difference in the solve efficiencies between the two algorithms for large values of p .

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: P. Amestoy, Ph. Berger, M. Dayd, I. Duff, V. Frayss, L. Giraud, D. Ruiz. (eds.), EuroPar '99, 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. Heller, D.: Some aspects of the cyclic reduction algorithm for block tridiagonal linear systems. SIAM J. Numer. Anal. 13, 484-496 (1976)
4. Hockney, R.W.: A fast direct solution of Poisson's equation using Fourier analysis. J. Assoc. Comput. Mach. 12, 95-113 (1965)
5. Lele, S. K.: Compact finite difference schemes with spectral-like resolution. J. Comput. Phys. 103, 16-42 (1992)
6. Mikkelsen, C.C.K., Manguoglu, M.: Analysis of the truncated SPIKE algorithm. SIAM J. Matrix Anal. Appl. 30, 1500-1519 (2008)
7. Mikkelsen, C.C.K., Kågström, B.: Parallel solution of narrow banded diagonally dominant linear systems. In: Proceedings of PARA-2010: State of the Art in Scientific and Parallel Computing, Reykjavik, Iceland, June 6-9th 2010.

Fig. 2: Strong scalability efficiency curves for $n = 52222$. Fig. 3: Strong scalability efficiency curves for $n = 103422$. Notice that the solve efficiency for the incomplete algorithm is greater than unity for all values of r and $\log_2(p) \geq 6$.