Hierarchical block multi-color ordering: a new parallel ordering method for vectorization and parallelization of the sparse triangular solver in the ICCG method

Takeshi Iwashita¹ · Senxi Li² · Takeshi Fukaya¹

Received: 15 November 2019 / Accepted: 15 April 2020 / Published online: 11 May 2020
© The Author(s) 2020

Abstract

In this paper, we propose a new parallel ordering method to vectorize and parallelize the sparse triangular solver, which is called hierarchical block multi-color ordering. In this method, the parallel forward and backward substitutions can be vectorized while preserving the advantages of block multi-color ordering, that is, fast convergence and fewer thread synchronizations. To evaluate the proposed method in a parallel ICCG (Incomplete Cholesky Conjugate Gradient) solver, numerical tests were conducted using seven test matrices on three types of computational nodes. The numerical results indicate that the proposed method outperforms the conventional block and nodal multi-color ordering methods in 18 out of 21 test cases, which confirms the effectiveness of the method.

Keywords Block multi-coloring · Iterative method · Parallel processing · SIMD · Sparse triangular solver

1 Introduction

A sparse triangular solver is an important computational kernel for an iterative linear solver in various numerical simulations. It is the main component of the Gauss–Seidel (GS) smoother, SOR method and IC/ILU preconditioning, which are used as building blocks in various computational science or engineering analyses (Beauwens 2004; Saad 2003; Meurant 1999). Therefore, the development of a fast multithreaded sparse triangular solver is essential to accelerate these analyses when conducted on not only a single computing node but also a large-scale cluster system of nodes. For example, the performance of the solver significantly influences the total simulation time of large-scale partial differential equation analysis using a multigrid solver with the GS, IC, or ILU smoother (Wallin et al. 2006; Buckeridge and Scheichl 2010). However, it is well known that the sparse

triangular solver, which consists of forward and backward substitutions, cannot be straightforwardly parallelized (Dongarra et al. 1990, 1998). Thus, in this paper, we discuss an effective approach to developing a high-performance multithreaded sparse triangular solver.

There are various methods for parallelizing a sparse triangular solver or its related techniques, and we focus on the parallel ordering (reordering) method, which is one of the most common methods for parallelization of a sparse triangular solver. There are several well-known orderings, such as dissection and domain decomposition orderings, but multi-color ordering is the most commonly used technique. It has been used in various applications to parallelize, for example, the ICCG method. However, it is well known that the multi-color ordering entails a trade-off problem between convergence and the number of synchronizations (Doi and Washio 1999). An increase in the number of colors typically results in better convergence, but it also leads to an increase in the number of synchronizations, which is proportional to the number of colors. The trade-off problem between convergence and parallelism is a common issue for parallel ordering techniques (Duff and Meurant 1989).

One of the solutions for the above trade-off problem is *block* multi-coloring. In this method, multi-color ordering is applied to blocks of unknowns. The technique has been investigated in several contexts. The concept of block

✉ Takeshi Iwashita
iwashita@iic.hokudai.ac.jp

¹ Information Initiative Center, Hokkaido University, N 11 W 5 Kita-ku, Sapporo, Hokkaido, Japan

² Graduate School of Information Science and Technology, The University of Tokyo, 1-1 Yayoi Bunkyo-ku, Tokyo, Japan

coloring or block independent sets can be seen in (Saad 2003). In an early work (Block et al. 1990), it is discussed for the parallel SOR method. For parallelization of the IC/ILU preconditioned iterative solver, it was first investigated in a finite difference method, that is, structured grid analysis (Iwashita and Shimasaki 2003; Iwashita et al. 2005). In this research, block coloring proved to be effective for improving convergence without increasing thread synchronization. Following on from these research activities, the algebraic block multi-color ordering method was introduced for a general sparse linear system (Iwashita et al. 2012). Although there are various options for coloring or blocking methods (Jones and Plassmann 1994; Iwashita and Shimasaki 2002), this technique has been used in various applications because of its advantages in terms of convergence, data locality, and the number of synchronizations (Semba et al. 2013; Tsuburaya et al. 2015). Particularly, several high-performance implementations of the HPCG benchmark adopt the technique, which shows the effectiveness of the method in a fast multigrid solver with the parallel GS smoother (Park et al. 2014; Kumahata et al. 2016; Yoshifuji et al. 2016; Vermij et al. 2017; Ruiz et al. 2018). However, the block multi-coloring method has a drawback in its implementation using SIMD vectorization. Since the unknowns in each block have data dependency with one another, they cannot be processed in parallel, which prevents the use of SIMD instructions for elements in different rows.

Because the sparse triangular solver is a memory-intensive kernel, its performance on previous computer was not substantially affected by the use of SIMD instructions. However, to increase the floating-point performance, recent processors enhance the SIMD instructions and their SIMD width (vector length) is becoming large. For example, Intel Xeon (Skylake) (Hammond et al. 2018), Intel Xeon Phi (Sodani et al. 2016), and Fujitsu A64FX (ARM SVE) (Feldman 2018) processors are equipped with 512-bit SIMD instructions. We note that ARM SVE supports at most a 2,048 vector length (Stephens et al. 2017). Considering this trend of processors, we aim to develop a parallel sparse triangular solver in which both multithreading and SIMD vectorization are efficiently used.

In this paper, we propose a new parallel ordering technique in which SIMD vectorization can be used and the advantages of block multi-color ordering, that is, fast convergence and fewer synchronizations, are preserved. The technique is called “hierarchical block multi-color ordering” and it has a mathematically equivalent solution process (convergence) to block multi-color ordering. Moreover, the number of synchronizations in the multithreaded substitutions is the same as that of block multi-color ordering. We conduct seven numerical tests using finite

element electromagnetic field analysis code and matrix data obtained from a matrix collection, and confirm the effectiveness of the proposed method in the context of the parallel ICCG solver.

2 Sparse triangular solver

In this paper, we consider the following n -dimensional linear system of equations:

$$Ax = b. \quad (1)$$

We discuss the case in which the linear system (1) is solved using an iterative linear solver involving IC(0)/ILU(0) preconditioning, the Gauss-Seidel (GS) method (smoother), or the SOR method. When we discuss a parallel ICCG (precisely IC(0)-CG) solver for (1), we assume that coefficient matrix A is symmetric and positive or semi-positive definite. For the parallelization of the iterative solver that we consider, the most problematic part is in the sparse triangular solver kernel. For example, in an IC/ILU preconditioned Krylov subspace iterative solver, the other computational kernels consist of an inner product, matrix-vector multiplication, and vector updates, which can be parallelized straightforwardly. The sparse triangular solver kernel is given by following forward and backward substitutions:

$$y = L^{-1}r, \quad (2)$$

and

$$z = U^{-1}y, \quad (3)$$

where r , y , and z are n -dimensional vectors. Matrices L and U are, respectively, lower and upper triangular matrices with the same nonzero patterns as the lower and upper triangular parts of A . In ILU (IC) preconditioning, the preconditioning step is given by (2) and (3), and triangular matrices L and U are derived from the following incomplete factorization:

$$A \simeq LU. \quad (4)$$

The iteration steps in the GS and SOR methods or smoothers can be expressed by similar substitutions. The substitution is an inherently sequential process, and it cannot be parallelized (multithreaded) straightforwardly.

3 Parallel ordering method

A parallel ordering (reordering) method is one of the most popular parallelization methods for a sparse triangular solver, that is, forward and backward substitutions. It transforms the coefficient matrix into an appropriate form for

parallel processing by reordering the unknowns or their indices. Let the reordered linear system of (1) be denoted by

$$\bar{A}\bar{x} = \bar{b}. \quad (5)$$

Then, the *reordering* is given by the transformation:

$$\bar{x} = P_{\pi}x, \quad (6)$$

where P_{π} is a permutation matrix. When we consider index set $I = \{1, 2, \dots, n\}$ that corresponds to the index of each unknown, the reordering is the permutation of the elements of I . In the present paper, the reordering function of the index is denoted by π ; that is, the i -th unknown of the original system is moved to the $\pi(i)$ -th unknown of the reordered system. In the reordering technique, the coefficient matrix and right-hand side are given as follows:

$$\bar{A} = P_{\pi}AP_{\pi}^T, \quad \bar{b} = P_{\pi}b. \quad (7)$$

3.1 Equivalence of orderings

We consider the case in which two linear systems, (1) and (5), are solved using an identical iterative method. The approximate solution vector at the j -th iteration for (1) and that for (5) are denoted by $x^{(j)}$, and $\bar{x}^{(j)}$, respectively. If it holds that

$$\bar{x}^{(j)} = P_{\pi}x^{(j)} \quad (8)$$

at every j -th step under the setting $\bar{x}^{(0)} = P_{\pi}x^{(0)}$ for initial guesses, then we can say that these two solution processes are *equivalent*. For example, in the Jacobi method and most Krylov subspace methods, reordering does not affect convergence; that is, the solution process for any reordered system is (mathematically) equivalent to that for the original system. However, in the case of the iterative solver that we consider in this paper, such as the IC/ILU preconditioned iterative solver, the solution processes are typically inequivalent. This is because of the sequentiality involved in the triangular solver (substitutions). However, there are special cases in which the reordered system has an equivalent solution process to the original system. In these cases, we say that two (original and new) orderings are equivalent or π is an equivalent reordering.

We define the equivalence of two orderings as follows: In the GS and SOR methods, equivalence is given by (8) under the proper setting of the initial guess. In IC(0)/ILU(0) preconditioning, equivalence is given as follows: We denote the incomplete factorization matrices of $\bar{A}$ by $\bar{L}$ and $\bar{U}$. Moreover, the preconditioning step of the reordered linear system is given by $\bar{z} = (\bar{L}\bar{U})^{-1}\bar{r}$. If $\bar{z} = P_{\pi}z$ is satisfied under $\bar{r} = P_{\pi}r$, then we say that the orderings are equivalent. For

example, the ICCG (IC(0)-CG) method exhibits an equivalent solution process for both original and reordered linear systems when the orderings are equivalent.

The condition for equivalent reordering is given as follows: When the following ER condition is satisfied, π is the equivalent reordering.

ER (Equivalent Reordering) Condition —

$$\forall i_1, i_2 \in I \text{ such that } a_{i_1, i_2} \neq 0 \vee a_{i_2, i_1} \neq 0, \quad (9)$$

$$\text{sgn}(i_1 - i_2) = \text{sgn}(\pi(i_1) - \pi(i_2)),$$

where a_{i_1, i_2} denotes the i_1 -th row i_2 -th column element of A . For a further explanation, we introduce an ordering graph, which is the directed graph that corresponds to the coefficient matrix. Each node of the graph corresponds to an unknown or its index. An edge between two nodes i_1 and i_2 exists only when the i_1 -th row i_2 -th column element or i_2 -th row i_1 -th column element is nonzero. The direction of the edge (arrow) shows the order of two nodes. Figure 1 shows an example of the ordering graph. Using the ordering graph, (9) can be rewritten as the statement that the new and original orderings have the same ordering graph. In (Doi and Lichniewsky 1991), the authors stated that the ordering graph provides a unique class of mutually equivalent orderings. In the appendix, we provide a proof sketch of the relationship between (9) and the equivalence of orderings.

4 Hierarchical block multi-color ordering method

In this paper, we propose a new parallel ordering method for the vectorization and parallelization of a sparse triangular solver. Additionally, the proposed ordering is intended to inherit the advantages of convergence, number of synchronizations, and data locality from block multi-color ordering (BMC). The proposed parallel ordering is called hierarchical

Fig. 1 Example of an ordering graph

Fig. 2 Coefficient matrix based on BMC

Fig. 3 Blocks of BMC and secondary reordering for HBMC

block multi-color ordering (HBMC), which is equivalent to BMC in terms of convergence.

In the technique, we first order the unknowns by using BMC, and then reorder them again. We focus on the explanation of the secondary reordering because we use the conventional algorithm shown in (Iwashita et al. 2012) for the application of BMC. Therefore, the original linear system based on BMC is written as (1) and secondary reordering is denoted by π . Thus, the final reordered linear system based on HBMC is given by (5).

4.1 Block multi-color ordering (BMC)

In this subsection, we briefly introduce BMC and some notation required for the explanation of HBMC. In BMC, all unknowns are divided into blocks of the same size, and the multi-color ordering is applied to the blocks. Because blocks that have an identical color are mutually independent, the forward and backward substitutions are parallelized based on the blocks in each color. The number of (thread)

synchronizations of parallelized (multithreaded) substitution is given by $n_c - 1$, where n_c is the number of colors. Figure 2 shows the coefficient matrix that results from BMC.

In the present paper, the block size and k -th block in color c are denoted by b_s and $b_k^{(c)}$, respectively. In BMC, each unknown (or its index) is assigned to a certain block, as shown in Fig. 3, where $n(c)$ is the number of blocks in color c .

4.2 Hierarchical block multi-color ordering (HBMC)

In the proposed HBMC, a new (hierarchical) block structure is introduced. First, we define a level-1 block (or multithreaded block) as follows. The block consists of w consecutive blocks of BMC in each color. When the k' -th level-1 block in color c is written as $\bar{b}_{k'}^{(c)}$,

$$\bar{b}_{k'}^{(c)} = \bigcup_{k=k_s+1}^{k_s+w} b_k^{(c)}, \quad (10)$$

where $k_s = (k' - 1) \times w$. We note that parameter w is determined by the length of the SIMD vector instruction (SIMD width) of the targeted processor. It is typically 4 or 8, and will be larger in the future.

In our technique, secondary reordering is performed on each level-1 block as shown in Fig. 3. Without loss of generality, we describe the reordering process for a level-1 block, that is, the blocks from $b_{k_s+1}^{(c)}$ to $b_{k_s+w}^{(c)}$ of BMC. In the first step, we pick up the first (top) unknown of each block, $b_{k_s+1}^{(c)}, b_{k_s+2}^{(c)}, \dots, b_{k_s+w}^{(c)}$, and order the picked unknowns. These w unknowns are mutually independent because the blocks in each color are independent in BMC. In the next step, we pick up the second unknown of each block, which are mutually independent, and order them after the unknowns previously ordered. We repeat the process until no unknowns remain. In total, the pick-up process is performed b_s times. Figure 4 shows the secondary reordering process in level-1 blocks when $n_c = 2$, $b_s = 2$, and $w = 4$, where each unknown is associated with the diagonal element of the coefficient matrix. In the figure, the colored elements represent nonzero elements. Although only one level-1 block exists in each color in Fig. 4, there are typically much more level-1 blocks in a color. After the reordering process is complete, we encounter the second-level block structure in the reordered coefficient matrix, which is given by the $w \times w$ (small) diagonal matrices. The level-2 block structure is used for SIMD vectorization of the substitutions. Figure 5 shows the matrix form of the coefficient matrix based on HBMC.

The summary of the reordering process for HBMC from BMC is as follows: In the first step, we generate level-1 blocks in each color, each of which consists of w blocks of

Fig. 4 Secondary reordering in a level-1 block

Fig. 5 Coefficient matrix based on HBMC

BMC. A level-1 block involves $b_s w$ unknowns in total. In the second step, we reorder the unknowns in each level-1 block. For the explanation, we introduce local labels for the unknowns in the level-1 block, which are from 1 to $b_s w$. For the application of HBMC, we reorder the unknowns in accordance with the value of $\text{mod}(i_l - 1, w)$ in the level-1 block, where i_l is the local label. Finally, we obtain the coefficient matrix form shown in Fig. 5 using the above procedure.

4.2.1 Equivalence between BMC and HBMC

We prove that HBMC is equivalent to BMC; that is, the convergence rates of the linear solvers based on the two orderings are the same. Because the secondary reordering for HBMC is *locally* performed in each level-1 block, the order between two unknowns that belong to two different level-1 blocks are preserved in the final order. Consequently, it holds that

$$\forall i_1 \in \bar{b}_{k_1}^{(c_1)}, i_2 \in \bar{b}_{k_2}^{(c_2)} \text{ such that } c_1 \neq c_2 \vee k_1 \neq k_2, \quad (11)$$

$$\text{sgn}(i_1 - i_2) = \text{sgn}(\pi(i_1) - \pi(i_2)).$$

From (11), if the *local* ordering subgraphs of BMC and HBMC that correspond to each level-1 block are identical, then the two orderings are equivalent. Next, we examine the reordering process in a level-1 block. In the secondary reordering process of HBMC, the order of unknowns that belong to different BMC blocks changes. However, the reordering process for these unknowns does not affect the ordering graph, that is, the convergence. In BMC, the unknowns that belong to two different blocks in the same color have no data relationship with one another; that is, there are no edges between them in the ordering graph. Therefore, even if we change the order of unknowns that belong to different BMC blocks, this does not affect the ordering graph. Consequently, we now pay attention to the influence of reordering inside a BMC block. When we analyze the above picking process, we can confirm that the order of the unknowns that belong to the same BMC block is preserved in the final order. In each block, we pick up the first unknown, and then

Fig. 6 Ordering graphs in a five-point stencil problem and coefficient matrix based on HBMC

the second, and continue this process. Therefore, the order does not change for these unknowns:

$$\forall i_1, i_2 \in b_k^{(c)} \quad \text{sgn}(i_1 - i_2) = \text{sgn}(\pi(i_1) - \pi(i_2)). \quad (12)$$

When we consider the mutual independence among the BMC blocks in each color, (11) and (12), we can demonstrate that secondary reordering π does not change the form of the ordering graph. This proves that HBMC is equivalent to BMC.

As an example that shows the relationship between BMC and HBMC, Fig. 6 demonstrates the ordering of nodes (unknowns) in a five-point finite difference analysis. Fig. 6a and b show that BMC and HBMC have identical ordering graphs. Consequently, the two orderings are equivalent in terms of convergence. Figure 6c shows the coefficient matrix based on HBMC, which involves the hierarchical block structures.

4.3 Parallelization and vectorization of forward and backward substitutions

Corresponding to the colors of the unknowns, solution vector $\bar{\mathbf{x}}$ and coefficient matrix $\bar{\mathbf{A}}$ are split as

$$\bar{\mathbf{x}} = \begin{pmatrix} \bar{\mathbf{x}}_1 \\ \bar{\mathbf{x}}_2 \\ \vdots \\ \bar{\mathbf{x}}_{n_c} \end{pmatrix}, \quad (13)$$

and

$$\bar{\mathbf{A}} = \begin{pmatrix} \bar{\mathbf{C}}_{1,1} & \bar{\mathbf{C}}_{1,2} & \cdots & \bar{\mathbf{C}}_{1,n_c} \\ \bar{\mathbf{C}}_{2,1} & \bar{\mathbf{C}}_{2,2} & \cdots & \bar{\mathbf{C}}_{2,n_c} \\ \vdots & \vdots & \ddots & \vdots \\ \bar{\mathbf{C}}_{n_c,1} & \bar{\mathbf{C}}_{n_c,2} & \cdots & \bar{\mathbf{C}}_{n_c,n_c} \end{pmatrix}, \quad (14)$$

where $\bar{\mathbf{x}}_c$ corresponds to the unknowns with color c , and $\bar{\mathbf{C}}_{c,d}$ represents the relationship between $\bar{\mathbf{x}}_c$ and $\bar{\mathbf{x}}_d$. Hereafter, we assume that the size of $\bar{\mathbf{x}}_c$ is a multiple of $b_s w$. In the analysis program, the assumption is satisfied using some dummy unknowns. Let the number of level-1 blocks assigned to color c be denoted by $\bar{n}(c)$, then diagonal block $\bar{\mathbf{C}}_{c,c}$ of $\bar{\mathbf{A}}$ is given by the following block diagonal matrix:

$$\bar{\mathbf{C}}_{c,c} = \begin{pmatrix} \bar{\mathbf{B}}_1^{(c)} & & & \mathbf{0} \\ & \bar{\mathbf{B}}_2^{(c)} & & \\ & & \ddots & \\ \mathbf{0} & & & \bar{\mathbf{B}}_{\bar{n}(c)}^{(c)} \end{pmatrix}, \quad (15)$$

where $\bar{\mathbf{B}}_k^{(c)}$ is the $b_s w \times b_s w$ matrix that corresponds to the unknowns in the k -th level-1 block with color c , which we denote by $\bar{b}_k^{(c)}$. Moreover, matrix $\bar{\mathbf{B}}_k^{(c)}$ is written as

$$\bar{B}_k^{(c)} = \begin{pmatrix} \bar{D}_1^{(k,c)} & \bar{E}_{1,2}^{(k,c)} & \cdots & \bar{E}_{1,b_s}^{(k,c)} \\ \bar{E}_{2,1}^{(k,c)} & \bar{D}_2^{(k,c)} & \cdots & \bar{E}_{2,b_s}^{(k,c)} \\ \vdots & \vdots & \ddots & \vdots \\ \bar{E}_{b_s,1}^{(k,c)} & \bar{E}_{b_s,2}^{(k,c)} & \cdots & \bar{D}_{b_s}^{(k,c)} \end{pmatrix}, \quad (16)$$

where $\bar{D}_l^{(k,c)}$, ($l = 1, 2, \dots, b_s$) are $w \times w$ diagonal matrices.

The forward substitution included in ILU(0)/IC(0) preconditioners or GS and SOR methods uses a lower triangular matrix with the same nonzero element pattern as the lower triangular part of $\bar{A}$. From (14) and (15), lower triangular matrix $\bar{L}$ is written as

$$\bar{L} = \begin{pmatrix} \bar{L}_{1,1} & & & \\ \bar{L}_{2,1} & \bar{L}_{2,2} & & \mathbf{0} \\ \vdots & \ddots & \ddots & \\ \bar{L}_{n_c,1} & \bar{L}_{n_c,2} & \cdots & \bar{L}_{n_c,n_c} \end{pmatrix}, \quad (17)$$

and diagonal block $\bar{L}_{c,c}$ is given by

$$\bar{L}_{c,c} = \begin{pmatrix} \bar{L}_1^{(c)} & & \mathbf{0} \\ & \bar{L}_2^{(c)} & \\ & & \ddots \\ \mathbf{0} & & & \bar{L}_{\bar{n}(c)}^{(c)} \end{pmatrix}, \quad (18)$$

where $\bar{L}_k^{(c)}$ is the $b_s w \times b_s w$ lower triangular matrix that corresponds to block $\bar{b}_k^{(c)}$. The forward substitution for the reordered linear system is given by

$$\bar{L}\bar{y} = \bar{r}, \quad (19)$$

where $\bar{r}$ is the residual vector in the case of ILU (IC) preconditioning. Let $\bar{y}_c$ and $\bar{r}_c$ represent, respectively, the segments of $\bar{y}$ and $\bar{r}$ that correspond to color c , and $\bar{y}_k^{(c)}$ and $\bar{r}_k^{(c)}$ represent the subsegments in the segments that correspond to block $\bar{b}_k^{(c)}$; that is,

$$\bar{y}_c = \begin{pmatrix} \bar{y}_1^{(c)} \\ \bar{y}_2^{(c)} \\ \vdots \\ \bar{y}_{\bar{n}(c)}^{(c)} \end{pmatrix}, \quad \text{and} \quad \bar{r}_c = \begin{pmatrix} \bar{r}_1^{(c)} \\ \bar{r}_2^{(c)} \\ \vdots \\ \bar{r}_{\bar{n}(c)}^{(c)} \end{pmatrix}. \quad (20)$$

Then, from (17) and (19), the forward substitution for $\bar{y}_c$ is given by

$$\bar{y}_c = \bar{L}_{c,c}^{-1} \bar{q}_c, \quad (21)$$

where

$$\bar{q}_c = \bar{r}_c - \sum_{d=1}^{c-1} \bar{L}_{c,d} \bar{y}_d. \quad (22)$$

Because vector segments $\bar{y}_d$ ($d = 1, \dots, c-1$) are computed prior to the substitution (21) and shared among all threads,

$\bar{q}_c$ is a given vector in (21). When the segment of $\bar{q}_c$ that corresponds to block $\bar{b}_k^{(c)}$ is denoted by $\bar{q}_k^{(c)}$ as in (20), from (18), the forward substitution (21) is expressed as $\bar{n}(c)$ independent steps:

$$\bar{y}_k^{(c)} = (\bar{L}_k^{(c)})^{-1} \bar{q}_k^{(c)} \quad (k = 1, \dots, \bar{n}(c)). \quad (23)$$

Consequently, the forward substitution (21) for color c can be multithreaded with the degree of parallelism given by the number of level-1 blocks of each color, which is approximately $n/(n_c \cdot b_s \cdot w)$. Each thread processes one or more level-1 blocks in parallel.

Next, we explain how to vectorize each step of (23). We consider the procedure for the k -th level-1 block of color c : $\bar{y}_k^{(c)} = (\bar{L}_k^{(c)})^{-1} \bar{q}_k^{(c)}$. From (16), lower triangular matrix $\bar{L}_k^{(c)}$ is written as

$$\bar{L}_k^{(c)} = \begin{pmatrix} \tilde{D}_1^{(k,c)} & & & \\ \tilde{L}_{2,1}^{(k,c)} & \tilde{D}_2^{(k,c)} & & \mathbf{0} \\ \vdots & \ddots & \ddots & \\ \tilde{L}_{b_s,1}^{(k,c)} & \tilde{L}_{b_s,2}^{(k,c)} & \cdots & \tilde{D}_{b_s}^{(k,c)} \end{pmatrix}, \quad (24)$$

where $\tilde{D}_l^{(k,c)}$ are diagonal matrices. We split $\bar{y}_k^{(c)}$ and $\bar{q}_k^{(c)}$ into b_s segments, each of size w . Let $\bar{y}_l^{(k,c)}$ and $\bar{q}_l^{(k,c)}$ represent the segments that correspond to the level-2 block of the k -th level-1 block of color c ; that is,

$$\bar{y}_k^{(c)} = \begin{pmatrix} \bar{y}_1^{(k,c)} \\ \bar{y}_2^{(k,c)} \\ \vdots \\ \bar{y}_{b_s}^{(k,c)} \end{pmatrix}, \quad \text{and} \quad \bar{q}_k^{(c)} = \begin{pmatrix} \bar{q}_1^{(k,c)} \\ \bar{q}_2^{(k,c)} \\ \vdots \\ \bar{q}_{b_s}^{(k,c)} \end{pmatrix}. \quad (25)$$

Then, from (24), the forward substitution for level-1 block $\bar{b}_k^{(c)}$ is given by the following b_s sequential steps:

$$\bar{y}_l^{(k,c)} = (\tilde{D}_l^{(k,c)})^{-1} \bar{t}_l^{(k,c)}, \quad (l = 1, 2, \dots, b_s), \quad (26)$$

where

$$\bar{t}_l^{(k,c)} = \bar{q}_l^{(k,c)} - \sum_{m=1}^{l-1} \tilde{L}_{l,m}^{(k,c)} \bar{y}_m^{(k,c)}. \quad (27)$$

In the l -th step of (26), because $\tilde{D}_l^{(k,c)}$ is a diagonal matrix and each element of $\bar{t}_l^{(k,c)}$ can be calculated in parallel, the step consists of w independent steps that can be efficiently vectorized. In other words, the l -th step of (26) consists of a simple matrix vector multiplication and element-wise vector updates that are directly vectorized.

The backward substitution is parallelized (multithreaded) and vectorized in a similar manner, although it is performed inversely from color n_c to 1.

4.4 Implementation of HBMC

4.4.1 Reordering process

In this section, we discuss the reordering process. In the technique, any type of algorithm (heuristic) for an implementation of BMC can be used. In the application of BMC, we set the number of BMC blocks assigned to each thread as a multiple of w , except for one of the threads (typically the last-numbered thread). In this circumstance, the application of HBMC, that is, the secondary reordering from BMC, is performed in each thread. Therefore, the reordering process is fully multithreaded.

4.4.2 Storage format

In the implementation of the sparse triangular solver, a storage format for sparse matrices (Barrett et al. 1994) is typically used. For example, the factorization matrices in an IC/ILU preconditioned solver are stored in memory using such a format. Although there are several standard formats, the sliced ELL (SELL) format (Liu et al. 2013) is the most efficient for exploiting the benefit of SIMD instructions, and we used it in our implementation. In the SELL format, the slice size is an important parameter. In HBMC, we naturally set the slice size as w because the forward and backward substitutions are vectorized every w rows. This leads to the natural introduction of the concept of the SELL-C- σ format

(Kreutzer et al. 2014) to the analysis, which is a sophisticated version of SELL.

4.4.3 Multithreaded and vectorized substitutions

The program for each forward and backward substitution consists of nested loops. The outer-most loop is for the color. After the computations for each color, thread synchronization is required. Therefore, the number of synchronizations in each substitution is given by $n_c - 1$, which is the same as BMC and the standard multi-color ordering (MC). The second loop is for level-1 blocks. Because the level-1 blocks in each color are mutually independent, each thread processes single or multiple level-1 blocks in parallel. In each level-1 block, the substitution can be split into b_s steps, each of which is vectorized with an SIMD width of w . For the vectorized substitution, we used the OpenMP SIMD directive or the Intel intrinsic functions for SIMD instructions. Figure 7 shows a sample C program code for multithreaded and vectorized forward substitution using OpenMP and Intel AVX-512 intrinsic functions. In our implementation, we use padding of dummy unknowns to satisfy the condition that the number of unknowns in each color results in a multiple of $b_s w$. Consequently, we use intrinsic SIMD functions for the aligned elements in the code.

Additionally, we discuss the special nonzero pattern that appears in $\bar{L}_k^{(c)}$ that corresponds to the level-1 block. In the matrix, all nonzero elements lay on $2b_s - 1$ diagonal lines. Although we can consider using a hybrid storage format that exploits this special structure, it does not typically result in better performance because of the additional cost of processing the diagonal block and other off-diagonal elements separately. We confirmed this in some preliminary tests.

Finally, we discuss the data access locality. The access pattern for the vector elements in HBMC is different from that in BMC. Therefore, the data access locality can be different between two orderings. However, because the secondary reordering for HBMC is performed inside a level-1 block, the data access locality barely changes; at least, from the viewpoint of the last-level cache, both orderings are considered to be similar.

5 Numerical results

5.1 Computers and test problems

We conducted seven numerical tests on three types of computational nodes to evaluate the proposed reordering technique in the context of the ICCG method: the computational nodes were Cray XC40, CS400 (2820XT), and Fujitsu CX2550 (M4). The two Cray systems are operated by Academic Center for Computing and Media Studies, Kyoto

```
for (c = 0; c < nc; c++){
    #pragma omp for private(p, j, num, t, index, \
                          mtmp, mval, pos, mb, mdiag)
    for (k = lev1b[c]; k < lev1b[c+1]; k++){
        num = k * 8 * bs;
        index = mat.offset[k];
        j = lev2b[k];
        for (p = j; p < j + bs; p++){
            mtmp = _mm512_load_pd( &r[num] );
            for (t = 0; t < mat.slen.lev2b[p]; t++){
                mval = _mm512_load_pd( &val[index] );
                pos = _mm512_load_epi32( &col[index] );
                mb = _mm512_i32logather_pd( pos, z, 8 );
                mtmp = _mm512_sub_pd( mtmp, \
                                   _mm512_mul_pd(mval,mb) );
                index += 8;
            }
            mdiag = _mm512_load_pd( &diaginval[num] );
            mtmp = _mm512_mul_pd( mtmp, mdiag );
            _mm512_store_pd( &z[num], mtmp );
            num = num + 8;
        }
    }
}
```

Fig. 7 C program code of multithreaded and vectorized forward substitution with OpenMP directives and Intel AVX-512 intrinsic functions

Table 1 Information about the test environments

System	Cray XC40	Cray CS400 (2820XT)	Fujitsu CX2550 (M4)
Processor type	Intel Xeon Phi (7250, KNL)	Intel Xeon (E5-2695 v4, Broadwell)	Intel Xeon (Gold6148, Skylake)
# processors / node	1	2	2
# cores / processor	68	18	20
Clock (GHz)	1.4	2.1	2.4
Compiler	icc 18.0.3	icc 17.0.6	icc 17.0.6
Compile options	-qopenmp -ipo -O3 -xMIC-AVX512	-mcmmodel=medium -shared-intel -qopenmp -xHost -ipo -O3	-mcmmodel=medium -shared-intel -qopenmp -xCORE-AVX512 -ipo -O3

Table 2 Matrix information for the test problems

Data set	Problem type	Dimension	# nonzero
Thermal2	Thermal problem	1,228,045	8,580,313
Parabolic_fem	CFD	525,825	3,674,625
G3_circuit	Circuit problem	1,585,478	7,660,826
Audikw_1	Structural problem	943,695	77,651,847
Tmt_sym	Electromagnetics problem	726,713	5,080,961
Apache2	Structural problem	715,176	4,817,870
Ieej	Eddy current analysis	1,011,920	31,468,056

Table 3 Comparison of the number of iterations

Dataset\method	MC	BMC	HBMC
Thermal2	2283	2129	2129
Parabolic_fem	1145	1052	1052
G3_circuit	1521	1227	1227
Audikw_1	1728	1714	1715
Tmt_sym	1912	1756	1755
Apache2	1615	1493	1493
Ieej	580	446	446

Univ., whereas the Fujitsu system is at Information Initiative Center, Hokkaido University. Table 1 lists information about the computational node and compiler used. In the numerical tests, we used all cores of the computational node for execution.

The program code was written in C and OpenMP for the thread parallelization. For vectorization, we used the intrinsic functions of the Intel Advanced Vector Extensions. The AVX2 (256 bits SIMD) instruction set was used for the Xeon (Broadwell) processor, whereas the AVX-512 (512 bits SIMD) instruction set was used for the Xeon Phi (KNL) and Xeon (Skylake) processors. Although we also developed a vectorized program using the OpenMP SIMD directive, its performance was slightly worse than the version with the intrinsic function in most of the test cases. Thus, in this paper, we only report the results from using the intrinsic function.

For the test problems, we used a linear system that arises from finite element electromagnetic field analysis and six linear systems picked up from the SuiteSparse Matrix Collection (Davis and Hu 2011). We selected symmetric positive definite matrices that are mainly derived from computational science or engineering problems, and have a relatively large dimension compared with other matrices in the collection.

In the electromagnetic field analysis test, the linear system that arises from the finite element discretization of the

IEEJ standard benchmark model (Nakata et al. 1991) was used. The basic equation for the problem is given as

$$\nabla \times (\nu \nabla \times \mathbf{A}_m) = \mathbf{J}_0, \quad (28)$$

where $\mathbf{A}_m$, ν , and $\mathbf{J}_0$ are the magnetic vector potential, magnetic reluctivity, and excitation current density, respectively. The analysis solved (28) using a finite edge-element method with hexahedral elements. The boundary condition is given by setting electric walls around the model. Applying the Galerkin method to (28), we obtained a linear system of equations for the test problem. Because no gauge condition is enforced in the air region of the model, the coefficient matrix of the resulting linear system is semi-positive definite (Igarashi 2001). Although the linear system has multiple solutions, the distribution for the magnetic flux density derived from them is unique. The linear system was solved using the shifted ICCG method, with the shift parameter given as 0.3. The name of the dataset of the linear system is denoted by Ieej. Table 2 lists the matrix information for the test problems. The right-hand side vector for the linear systems except for Ieej is given by the vector of ones.

In this paper, we report the performance comparison of four multithreaded ICCG solvers. The solver denoted by MC is based on the multi-color ordering which is the most popular parallel ordering method. The solver “BMC” is the solver based on the block multi-color ordering method.

Fig. 8 Convergence behavior of BMC and HBMC

The solvers denoted by “HBMC (csr_spmv)” and “HBMC (sell_spmv)” are based on the proposed HBMC, where the former solver used compressed storage row (CSR) format (Barrett et al. 1994) for the implementation of sparse matrix vector multiplication (SpMV) and the latter used the SELL format. In MC and BMC, the CSR format was used.

For the blocking method in BMC and HBMC, we used the simplest one among the heuristics introduced in (Iwashita et al. 2012), in which the unknown with the minimal number is picked up for the newly generated block. For the coloring of nodes or blocks, the greedy algorithm was used for all the solvers. The convergence criterion was set as the relative residual norm (2-norm) being less than 10^{-7} .

5.2 Numerical results

5.2.1 Equivalence of orderings in convergence and use of SIMD instructions

First, we examine the equivalence of BMC and HBMC in terms of convergence. Table 3 lists the number of iterations of the solvers tested on Cray XC40, where the block size of BMC and HBMC was set to 32. Equivalence was confirmed by the numerical results. Moreover, to examine the entire solution process, Fig. 8 shows the convergence behaviors of BMC and HBMC in the G3_circuit and Ieej tests. In the figure, the two lines of the relative residual norms overlap, which indicates that the solvers had an equivalent solution process. The equivalence of convergence was also confirmed in all test cases (seven datasets $\times$ three block sizes $\times$ three computational nodes). Furthermore, Table 3 shows the advantage in terms of convergence of BMC over MC, which coincides with the results reported in (Iwashita et al. 2012).

Next, we checked the use of SIMD instructions in the solver using the Intel Vtune Amplifier (application performance snapshot) in the G3_circuit test conducted on Fujitsu CX2550. The snapshot showed that the percentage of all packed floating point instructions in the solver based on HBMC (sell_spmv) reached 99.7%, although that in the solver using BMC was 12.7%.

5.2.2 Performance comparison

Table 4 (a) shows the performance comparison of four solvers on Cray XC40. In the tests, HBMC attained the best performance for all datasets, except Audikw_1. In the Thermal2 and G3_circuit tests, HBMC (sell_spmv) was more than two times faster than the standard MC solver. When HBMC (csr_spmv) was compared with BMC, it attained better performance in 18 out of 21 cases (seven datasets $\times$ three block sizes), which demonstrates the effectiveness of HBMC for the sparse triangular solver. In all test cases, HBMC (sell_spmv) outperformed HBMC (csr_spmv), which implies that an efficient use of SIMD instructions was important on the Xeon Phi-based system.

Table 4 (b) shows the test results on Cray CS400 (Xeon Broadwell). HBMC attained the best performance for all datasets. When HBMC (csr_spmv) was compared with BMC, it attained better performance in 16 out of 21 cases, which shows the effectiveness of HBMC. Table 4 (b) also indicates that using the SELL format for the coefficient matrix mostly led to an improvement in solver performance.

Table 4 (c) shows the test results on Fujitsu CX2550 (Xeon Skylake). In the numerical tests, HBMC outperformed MC and BMC for six out of seven datasets. For the Audikw_1 dataset, HBMC did not outperform BMC on

Table 4 Numerical results: execution time (sec.)

Dataset	MC	BMC			HBMC (csr_spmv)			HBMC (sell_spmv)		
		$b_s = 8$	$b_s = 16$	$b_s = 32$	$b_s = 8$	$b_s = 16$	$b_s = 32$	$b_s = 8$	$b_s = 16$	$b_s = 32$
(a) Cray XC40 (Intel Xeon Phi)										
Thermal2	20.2	17.8	17.9	19.4	13.8	14.5	15.6	9.28	9.79	11.1
Parabolic_fem	2.64	2.77	2.79	2.72	2.35	2.33	2.48	1.72	1.72	1.79
G3_circuit	7.98	7.97	7.72	9.91	5.16	5.27	5.49	3.35	3.55	3.61
Audikw_1	109.4	73.2	69.3	63.9	72.7	73.2	73.3	68.3	68.0	69.1
Tmt_sym	6.60	5.18	5.08	4.91	3.60	3.63	3.70	2.63	2.68	2.71
Apache2	3.03	2.75	3.59	3.52	1.79	2.59	2.48	1.40	1.96	1.92
Ieej	4.58	5.35	5.07	4.12	4.94	5.02	4.55	3.60	3.49	3.18
(b) Cray CS400 (Intel Xeon Broadwell)										
Thermal2	16.3	14.8	14.8	15.3	13.5	14.2	14.8	14.4	14.2	14.7
Parabolic_fem	2.91	2.69	2.69	2.53	2.55	2.54	2.51	2.41	2.33	2.31
G3_circuit	9.39	7.70	8.15	7.92	7.62	7.90	8.20	7.23	7.44	7.46
Audikw_1	70.9	64.2	66.5	60.1	58.7	55.8	55.0	59.2	56.3	55.9
Tmt_sym	5.50	5.02	5.06	4.94	5.05	4.67	4.78	4.91	4.84	4.74
Apache2	3.91	3.18	3.78	3.84	3.12	4.22	4.40	3.31	4.22	4.08
Ieej	6.98	6.49	6.72	5.34	6.93	6.42	5.18	5.97	5.51	4.83
(c) Fujitsu CX2550 M4 (Intel Xeon Skylake)										
Thermal2	10.1	9.74	9.38	9.21	8.58	9.62	9.33	8.70	8.93	9.13
Parabolic_fem	1.81	1.70	1.75	1.64	1.47	1.57	1.78	1.52	1.49	1.49
G3_circuit	6.17	5.07	5.23	5.01	4.83	5.19	5.60	4.54	5.09	4.87
Audikw_1	44.6	34.6	36.4	36.6	35.5	37.5	37.9	37.4	37.3	39.8
Tmt_sym	3.65	3.03	2.99	3.01	2.95	3.08	3.13	2.97	3.08	3.05
Apache2	2.41	1.85	2.41	2.41	1.76	2.20	2.21	1.72	2.24	2.43
Ieej	3.76	4.71	3.72	3.56	4.20	3.69	3.37	3.81	3.51	3.15

Bold value shows the best result for each dataset on three computational nodes

Xeon Phi and Skylake, although it was better than BMC on Xeon Broadwell. This result is thought to be because of the effect of increasing the slice size that was given by the SIMD width, w . In the SELL format, some zero elements were considered as nonzero in the slice. When there was a significant imbalance among the number of nonzero elements in each row in the slice, the number of elements processed largely increased compared with the implementation with the CSR format. The Audikw_1 dataset had this property. For Audikw_1, the number of processed elements in SELL increased by 40% compared with that in CSR, although the increase was 10% for the G3_circuit dataset. For this type of dataset, when the size of the slice increased, the number of processed elements often increased. The size of the slice was set to 8 for the Xeon Phi and Skylake processors and 4 for the Xeon Broadwell processors. On the Broadwell processors, the increase in the number of elements when changing CSR to SELL was 28%, which resulted in better performance for HBMC compared with BMC. In the future, for further acceleration of the solver, we will develop an implementation in which we introduce some remedies for this SELL issue, for example, splitting the row that includes

an extremely large number of nonzero elements compared with other rows.

6 Related works

The parallelization of the sparse triangular solver for iterative solvers has been mainly investigated in the context of GS or IC/ILU smoothers for multigrid solvers, the SOR method, and IC/ILU preconditioned iterative solvers. Most of these parallelization techniques are classified into two classes: domain decomposition type methods and parallel orderings (Duff and van der Vorst 1999). A simple but commonly used technique in the former class is the additive Schwarz smoother or preconditioner. The hybrid Jacobi and GS smoother, and block Jacobi IC/ILU preconditioning are typical examples, and they are used in many applications (Vollaie and Nicolas 1998; Baker et al. 2010). However, these techniques typically suffer from a decline in convergence when the number of threads (processes) is increased. Although there are some remedies for the deterioration in convergence, for example, the overlapping technique

(Radicati di Brozolo and Robert 1989), it is generally difficult to compensate for it when many cores are used.

A parallel ordering or reordering technique is a standard method to parallelize the sparse triangular solver. We focus on the parallelization of IC/ILU preconditioners or smoothers; however, there are many studies that discuss the application of parallel ordering for GS and SOR methods, for example, (Park et al. 2015). Ref. (Duff and van der Vorst 1999) provides an overview of early works on the parallel ordering method when applied to IC/ILU preconditioned iterative solvers. Parallel orderings were mainly investigated in the context of a structured grid problem (a finite difference analysis), and the concepts of typical parallel ordering techniques, such as red-black, multi-color, zebra, domain decomposition (four or eight-corner), and nested dissection, were established in the 1990s. In these early research activities, a major issue was the trade-off problem between convergence and the degree of parallelism. After Duff and Meurant indicated the problem in (Duff and Meurant 1989), both analytical and numerical investigations were conducted in (Doi and Washio 1999; Iwashita et al. 2005; Doi and Lichniewsky 1991; Kuo and Chan 1990; Doi and Lichniewsky 1990; Eijkhout 1991; Doi 1991; Benzi et al. 1999). The concept of equivalence of orderings and some remedies for the trade-off problem, such as the use of a relatively large number of colors in multi-color ordering or block coloring, were introduced as the results of these research activities.

In practical engineering and science domains, unstructured problems are solved more frequently than structured problems. Therefore, parallel ordering techniques were enhanced for unstructured problems, and several heuristics were proposed. Typical examples are hierarchical interface decomposition (HID) (Hénon and Saad 2006) and heuristics for nodal or block multi-coloring (Iwashita et al. 2012; Jones and Plassmann 1994; Iwashita and Shimasaki 2002). These techniques and other related methods have been used in various application domains, such as CFD, computational electromagnetics, and structure analyses (Semba et al. 2013; Tsuburaya et al. 2015; Rivera et al. 2010; Nakajima 1999; Naumov 2011).

Finally, we address the recently reported research results that are related to parallel linear solvers that involve sparse triangular solvers. Gonzaga de Oliveira et al. reported their intensive numerical test results to evaluate various reordering techniques in the ICCG method in (2018). Gupta introduced a blocking framework to generate a fast and robust preconditioner based on ILU factorization in (2017). Chen et al. developed a couple of ILU-based preconditioners on GPUs in (2018). Ruiz et al. reported the evaluation results of HPCG implementations using nodal and block multi-color orderings on the ARM-based system, which confirmed the superiority of the block coloring method in (2018).

In this paper, we proposed a parallel ordering that is different from the techniques described above. To the best of our knowledge, there is no parallel ordering method that vectorizes the sparse triangular solver while maintaining the same convergence and number of synchronizations as the block multi-color ordering. Since the vectorization of SpMV has been intensively investigated (Liu et al. 2013; Kreutzer et al. 2014), one of conventional approaches is the use of multi-color ordering, in which the substitution is represented as an SpMV in each color. However, the multi-color ordering suffers from the problems of convergence and data locality, which are also indicated in the latest report (Ruiz et al. 2018). When we consider the numerical results and mathematical properties of the proposed hierarchical block multi-color ordering, it can be regarded as an effective technique for multithreading and vectorizing the sparse triangular solver.

7 Conclusions

In this paper, we proposed a new parallel ordering technique, hierarchical block multi-color ordering (HBMC), for vectorizing and multithreading the sparse triangular solver. HBMC was designed to maintain the advantages of the block multi-color ordering (BMC) in terms of convergence and the number of synchronizations. In the method, the coefficient matrix was transformed into the matrix with hierarchical block structures. The level-1 blocks were mutually independent in each color, which was exploited in multithreading. Corresponding to the level-2 block, the substitution was converted into w ($=$ SIMD width) independent steps, which were efficiently processed by SIMD instructions. In this paper, we demonstrated analytically the equivalence of HBMC and BMC in convergence. Furthermore, numerical tests were conducted to examine the proposed method using seven datasets on three types of computational nodes. The numerical results also confirmed the equivalence of the convergence of HBMC and BMC. Moreover, numerical tests indicated that HBMC outperformed BMC in 18 out of 21 test cases (seven datasets $\times$ three systems), which confirmed the effectiveness of the proposed method.

In the future, we will examine our technique for other application problems, particularly a large-scale multigrid application and an HPCG benchmark. Moreover, we will examine the effect of the matrix format for preconditioners and coefficient matrices on the performance of the solver. We also intend to introduce a sophisticated storage format or its related technique to our solver. As other research issues, we will examine the effect of other coloring and blocking strategies on the performance of the solver.

Acknowledgements This work was supported by JSPS KAKENHI Grant Numbers 19H04122 and 19H05662.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

A Equivalent reordering

In this paper, we provide a sketch of the proof with respect to the ER condition. In the GS and SOR cases, the proof is given through the double use of the inductive method. The first is for the iteration step. In each iteration step, we use the inductive method again in a row-by-row manner. The proof is relatively straightforward.

In the IC/ILU preconditioning case, the proof consists of two parts. We first prove the equivalence of preconditioners $\bar{L} = P_{\pi} L P_{\pi}^T$ and $\bar{U} = P_{\pi} U P_{\pi}^T$. We consider the right-looking factorization process with overwriting, and first prove that the updating process in the factorization is equivalent. Next, we show that the application of the updating process to a_{ij} is equivalent to that to $\bar{a}_{\pi(i), \pi(j)}$ (the $\pi(i)$ -th row $\pi(j)$ -th column element of $\bar{A}$) when (9) holds. This leads to the equivalence of preconditioners. In the second part, the equivalence in the substitutions is shown using the inductive method, similar to the proof for the GS method.

References

- Baker, A.H., Falgout, R.D., Gamblin, T., Kolev, T.V., Schulz, M., Yang, U.M.: Scaling algebraic multigrid solvers: on the road to exascale. In: Bischof, C., Hegering, H.G., Nagel, W., Witum, G. (eds.) *Competence in High Performance Computing*. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-24025-6_18
- Barrett, R., Berry, M., Chan, T.F., Demmel, J., Donato, J., Dongarra, J., Eijkhout, V., Pozo, R., Romine, C., van der Vorst, H.: *Templates for the Solution of Linear Systems: Building Blocks for Iterative Methods*. SIAM, Philadelphia (1994)
- Beauwens, R.: Iterative solution methods. *Appl. Numer. Math.* **51**, 437–450 (2004)
- Benzi, M., Joubert, W., Mateescu, G.: Numerical experiments with parallel orderings for ILU preconditioners. *Electron. Trans. Numer. Anal.* **8**, 88–114 (1999)
- Block, U., Frommer, A., Mayer, G.: Block colouring schemes for the SOR method on local memory parallel computers. *Parallel Comput.* **14**, 61–75 (1990)
- Buckeridge, S., Scheichl, R.: Parallel geometric multigrid for global weather prediction. *Numer. Linear Algebra Appl.* **17**, 325–342 (2010). <https://doi.org/10.1002/nla.699>
- Chen, Y., Tian, X., Liu Zhen, B., Yang, C., Liao, W., Zhang, P., He, R., Tang, M.: Parallel ILU preconditioners in GPU computation. *Soft. Comput.* **22**, 8187–8205 (2018). <https://doi.org/10.1007/s00500-017-2764-7>
- Davis, T.A., Hu, Y.: The university of Florida sparse matrix collection. *ACM Trans. Math. Software.* **38**, 1–25 (2011). <https://doi.org/10.1145/2049662.2049663>
- Doi, S., Lichnewsky, A.: A graph-theory approach for analyzing the effects of ordering on ILU preconditioning. Res. Report No. 1452. INRIA, France, (1991)
- Doi, S.: On parallelism and convergence of incomplete LU factorizations. *Appl. Numer. Math.* **7**, 417–436 (1991)
- Doi, S., Lichnewsky, A.: Some parallel and vector implementations of preconditioned iterative methods on Cray-2. *Internat. J. High Speed Comput.* **2**, 143–179 (1990)
- Doi, S., Washio, T.: Ordering strategies and related techniques to overcome the trade-off between parallelism and convergence in incomplete factorization. *Parallel Comput.* **25**, 1995–2014 (1999)
- Dongarra, J.J., Duff, I.S., Sorensen, D.C., van der Vorst, H.A.: *Solving Linear Systems on Vector and Shared Memory Computers*. SIAM, Philadelphia (1990)
- Dongarra, J.J., Duff, I.S., Sorensen, D.C., van der Vorst, H.A.: *Numerical Linear Algebra for High-Performance Computers*. SIAM, Philadelphia (1998)
- Duff, I.S., Meurant, G.A.: The effect of ordering on preconditioned conjugate gradients. *BIT* **29**, 635–657 (1989)
- Duff, I.S., van der Vorst, H.A.: Developments and trend in the parallel solution of linear systems. *Parallel Comput.* **25**, 1931–1970 (1999)
- Eijkhout, V.: Analysis of parallel incomplete point factorizations. *Linear Algebra Appl.* **154–156**, 723–740 (1991)
- Feldman, M.: Fujitsu reveals details of processor that will power Post-K supercomputer. <https://www.top500.org/news/fujitsu-reveals-details-of-processor-that-will-powerpost-k-supercomputer> (2018). Accessed 26 March 2019
- Gonzaga de Oliveira, S.L., Bernardes, J.A.B., Chagas, G.O.: An evaluation of reordering algorithms to reduce the computational cost of the incomplete Cholesky-conjugate gradient method. *Comput. Appl. Math.* **37**, 2965–3004 (2018). <https://doi.org/10.1007/s40314-017-0490-5>
- Gupta, A.: Enhancing performance and robustness of ILU preconditioners by blocking and selective transposition. *SIAM J. Sci. Comput.* **39**, A303–A332 (2017). <https://doi.org/10.1137/15M1053256>
- Hammond, S.D., Vaughan, C.T., Hughes, C.: Evaluating the Intel Skylake Xeon Processor for HPC Workloads. In: *Proceedings 2018 International Conference High Performance Computing and Simulation, IEEE*, 342–349 (2018). <https://doi.org/10.1109/HPCS.2018.00064>
- Hénon, P., Saad, Y.: A parallel multistage ILU factorization based on a hierarchical graph decomposition. *SIAM J. Sci. Comput.* **28**, 2266–2293 (2006)
- Igarashi, H.: On the property of the curl-curl matrix in finite element analysis with edge elements. *IEEE Trans. Magn.* **37**, 3129–3132 (2001)
- Iwashita, T., Nakashima, H., Takahashi, Y.: Algebraic Block Multi-Color Ordering Method for Parallel Multi-Threaded Sparse Triangular Solver in ICCG Method. In: *Proc. 2012 IEEE 26th International Parallel and Distributed Processing Symposium (IPDPS2012)*, 474–483 (2012). <https://doi.org/10.1109/IPDPS.2012.51>
- Iwashita, T., Shimasaki, M.: Algebraic multi-color ordering for parallelized ICCG solver in finite element analyses. *IEEE Trans. Magn.* **38**, 429–432 (2002)

- Iwashita, T., Shimasaki, M.: Block red-black ordering: A new ordering strategy for parallelization of ICCG method. *Internat. J. Parallel Progr.* **31**, 55–75 (2003)
- Iwashita, T., Nakanishi, Y., Shimasaki, M.: Comparison criteria for parallel orderings in ILU preconditioning. *SIAM J. Sci. Comput.* **26**, 1234–1260 (2005)
- Jones, M.T., Plassmann, P.E.: The efficient parallel iterative solution of large sparse linear systems. In: George, A., Gilbert, J.R., Liu, J.W.H. (eds.) *IMA 56 (Graph Theory and Sparse Matrix Computations)*, pp. 229–245. Springer, Berlin (1994)
- Kreutzer, M., Hager, G., Wellein, G., Fehske, H., Bishop, A.R.: A unified sparse matrix data format for efficient general sparse matrix-vector multiplication on modern processors with wide SIMD units. *SIAM J. Sci. Comput.* **36**, C401–C423 (2014). <https://doi.org/10.1137/130930352>
- Kumahata, K., Minami, K., Maruyama, N.: High-performance conjugate gradient performance improvement on the K computer. *Int. J. High Perform. C.* **30**, 55–70 (2016). <https://doi.org/10.1177/1094342015607950>
- Kuo, C.-C.J., Chan, T.F.: Two-color Fourier analysis of iterative algorithms for elliptic problems with red/black ordering. *SIAM J. Sci. Statist. Comput.* **11**, 767–793 (1990)
- Liu, X., Smelyanskiy, M., Chow, E., Dubey, P.: Efficient Sparse Matrix-Vector Multiplication on x86-Based Many-Core Processors. In: *Proc. of the 27th international ACM conference on supercomputing (ICS '13)*, 273–282 (2013)
- Meurant, G.: *Computer Solution of Large Linear Systems*. Elsevier, Amsterdam (1999)
- Nakajima, K.: Flat MPI vs. Hybrid: Evaluation of Parallel Programming Models for Preconditioned Iterative Solvers on “T2K Open Supercomputer”. In: *Proc. 2009 Intl. Conf. Parallel Processing Workshops*, 73–80 (1999). <https://doi.org/10.1109/ICPPW.2009.68>
- Nakata, T., Takahashi, N., Imai, T., Muramatsu, K.: Comparison of various methods of analysis and finite elements in 3-D magnetic field analysis. *IEEE Trans. Magn.* **27**, 4073–4076 (1991)
- Naumov, M.: *Parallel Solution of Sparse Triangular Linear Systems in the Preconditioned Iterative Methods on the GPU*. NVIDIA Tech. Report NVR-2011-001, NVIDIA, Santa Clara, (2011)
- Park, J., Smelyanskiy, M., Vaidyanathan, K., Heinecke, A., Kalamkar, D.D., Liu, X., Patwary, Md.M.A., Lu, Y., Dubey, P.: Efficient Shared-Memory Implementation of High-Performance Conjugate Gradient Benchmark and Its Application to Unstructured Matrices. In: *Proc. Intl. Conf. High Performance Computing, Networking, Storage and Analysis (SC '14)*, 945–955 (2014). <https://doi.org/10.1109/SC.2014.82>
- Park, J., Smelyanskiy, M., Yang, U.M., Mudigere, D., Dubey, P.: High-Performance Algebraic Multigrid Solver Optimized for Multi-Core Based Distributed Parallel Systems. In: *Proc. of Intl. Conf. High Performance Computing, Networking, Storage and Analysis (SC '15)*, Article number: 54 (2015). <https://doi.org/10.1145/2807591.2807603>
- Radicati di Brozoloand, G., Robert, Y.: Parallel conjugate gradient-like algorithms for solving sparse nonsymmetric linear systems on a vector multiprocessor. *Parallel Comput.* **11**, 223–239 (1989)
- Rivera, C.A., Heniche, M., Glowinski, R., Tanguya, P.A.: Parallel finite element simulations of incompressible viscous fluid flow by domain decomposition with Lagrange multipliers. *J. Comput. Phys.* **229**, 5123–5143 (2010)
- Ruiz, D., Mantovani, F., Casas, M., Labarta, J., Spiga, F.: The HPCG benchmark: analysis, shared memory preliminary improvements and evaluation on an Arm-based platform. *UPCommons, External research report* (2018). <http://hdl.handle.net/2117/116642>
- Saad, Y.: *Iterative Methods for Sparse Linear Systems*, 2nd edn. SIAM, Philadelphia (2003)
- Semba, K., Tani, K., Yamada, T., Iwashita, T., Takahashi, Y., Nakashima, H.: Parallel performance of multithreaded ICCG solver based on algebraic block multicolor ordering in finite element electromagnetic field analyses. *IEEE Trans. Magn.* **49**, 1581–1584 (2013). <https://doi.org/10.1109/TMAG.2013.2244858>
- Sodani, A., Gramunt, R., Corbal, J., Kim, H., Vinod, K., Chinthamani, S., Hutsell, S., Agarwal, R., Liu, Y.: Knights landing: Second-generation intel xeon phi product. *IEEE Micro.* **36**, 34–46 (2016)
- Stephens, N., Biles, S., Boettcher, M., Eapen, J., Eyole, M., Gabrielli, G., Horsnell, M., Magklis, G., Martinez, A., Premillieu, N., Reid, A., Rico, A., Walker, P.: The ARM scalable vector extension. *IEEE Micro.* **37**, 26–39 (2017). <https://doi.org/10.1109/MM.2017.35>
- Tsuburaya, T., Okamoto, Y., Sato, S.: Parallelized ICCG method using block-multicolor orderings in real symmetric linear system derived from voltage-driven FEM in time domain. *COMPEL* **34**, 1433–1446 (2015). <https://doi.org/10.1108/COMPEL-02-2015-0069>
- Vermij, E., Fiorin, L., Hagleitner, C., Bertels, K.: Boosting the efficiency of HPCG and Graph500 with near-data processing. In: *Proc. 46th Intl. Conf. Parallel Processing (ICPP '17)*, 31–40 (2017). <https://doi.org/10.1109/ICPP.2017.12>
- Vollaire, C., Nicolas, L.: Preconditioning techniques for the conjugate gradient solver on a parallel distributed memory computer. *IEEE Trans. Magn.* **34**, 3347–3350 (1998)
- Wallin, D., Löf, H., Hagersten, E., Holmgren, S.: Multigrid and Gauss-Seidel Smoothers Revisited: Parallelization on Chip Multiprocessors. In: *Proc. 20th Annual Intl. Conf. Supercomputing (ICS '06)*, 145–155 (2006). <https://doi.org/10.1145/1183401.1183423>
- Yoshifuji, N., Sakamoto, R., Nitadori, K., Makino, J.: Implementation and evaluation of data-compression algorithms for irregular-grid iterative methods on the PEZY-SC processor. In: *2016 6th Workshop on Irregular Applications: Architecture and Algorithms (IA'3 '16)*, 58–61 (2016). <https://doi.org/10.1109/IA3.2016.15>