

# A new Matrix Reordering Method for GPU Acceleration of an ILU Preconditioner

Kengo Suzuki<sup>1</sup>, Takeshi Fukaya<sup>1</sup>, and Takeshi Iwashita<sup>2</sup>  
<sup>1</sup>Hokkaido University, <sup>2</sup>Kyoto University

## Introduction

### Solve linear systems $Ax = b$ efficiently on a GPU

Highlights of this study:

- Improve the FGMRES solver preconditioned by ILU(0) on a GPU.
- Enhance the concurrency of ILU(0) by **matrix reordering**.
- Utilize our previous reordering method, **HBMC** [1], which can increase the concurrency of ILU(0) while maintaining the convergence rate.

Sparse triangular systems to be solved in ILU(0) have data dependencies, making it difficult to use GPUs effectively.

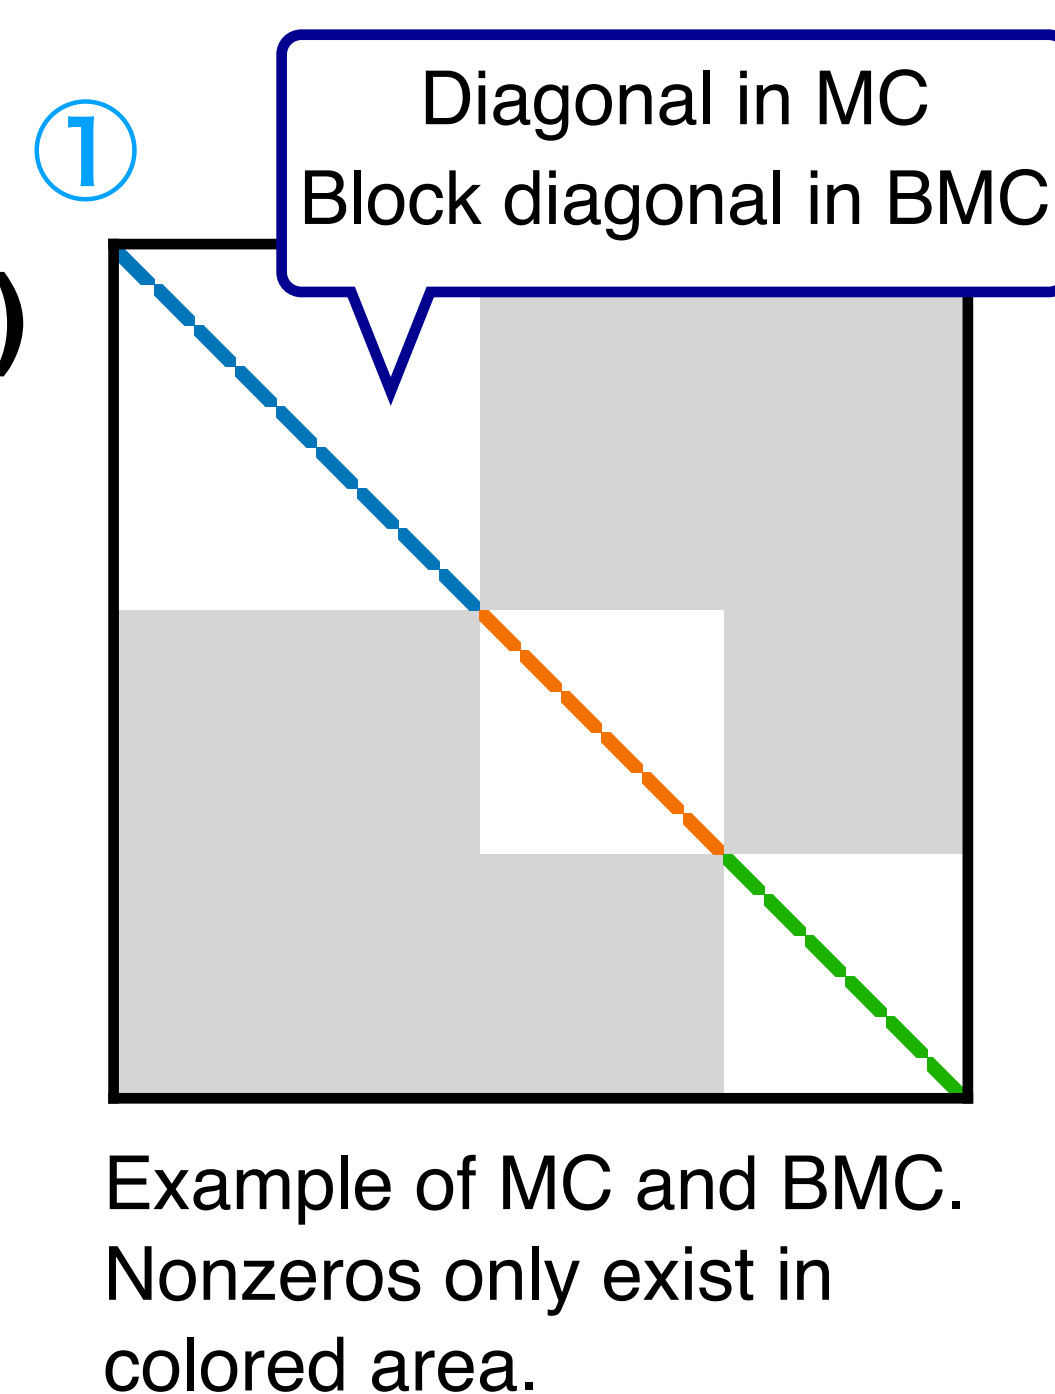
## Matrix reordering

### Enhances the concurrency of ILU(0)

However, it includes a trade-off problem between **concurrency and convergence**.

For example,

- **MC ordering** (a strong candidate on GPUs)
  - Increases the concurrency significantly.
  - Reduces the effect of ILU(0) on convergence.
- **BMC ordering** (a typical method on CPUs)
  - Maintains the effect of ILU(0) on convergence.
  - Lacks fine-grained concurrency due to the block diagonal structures.

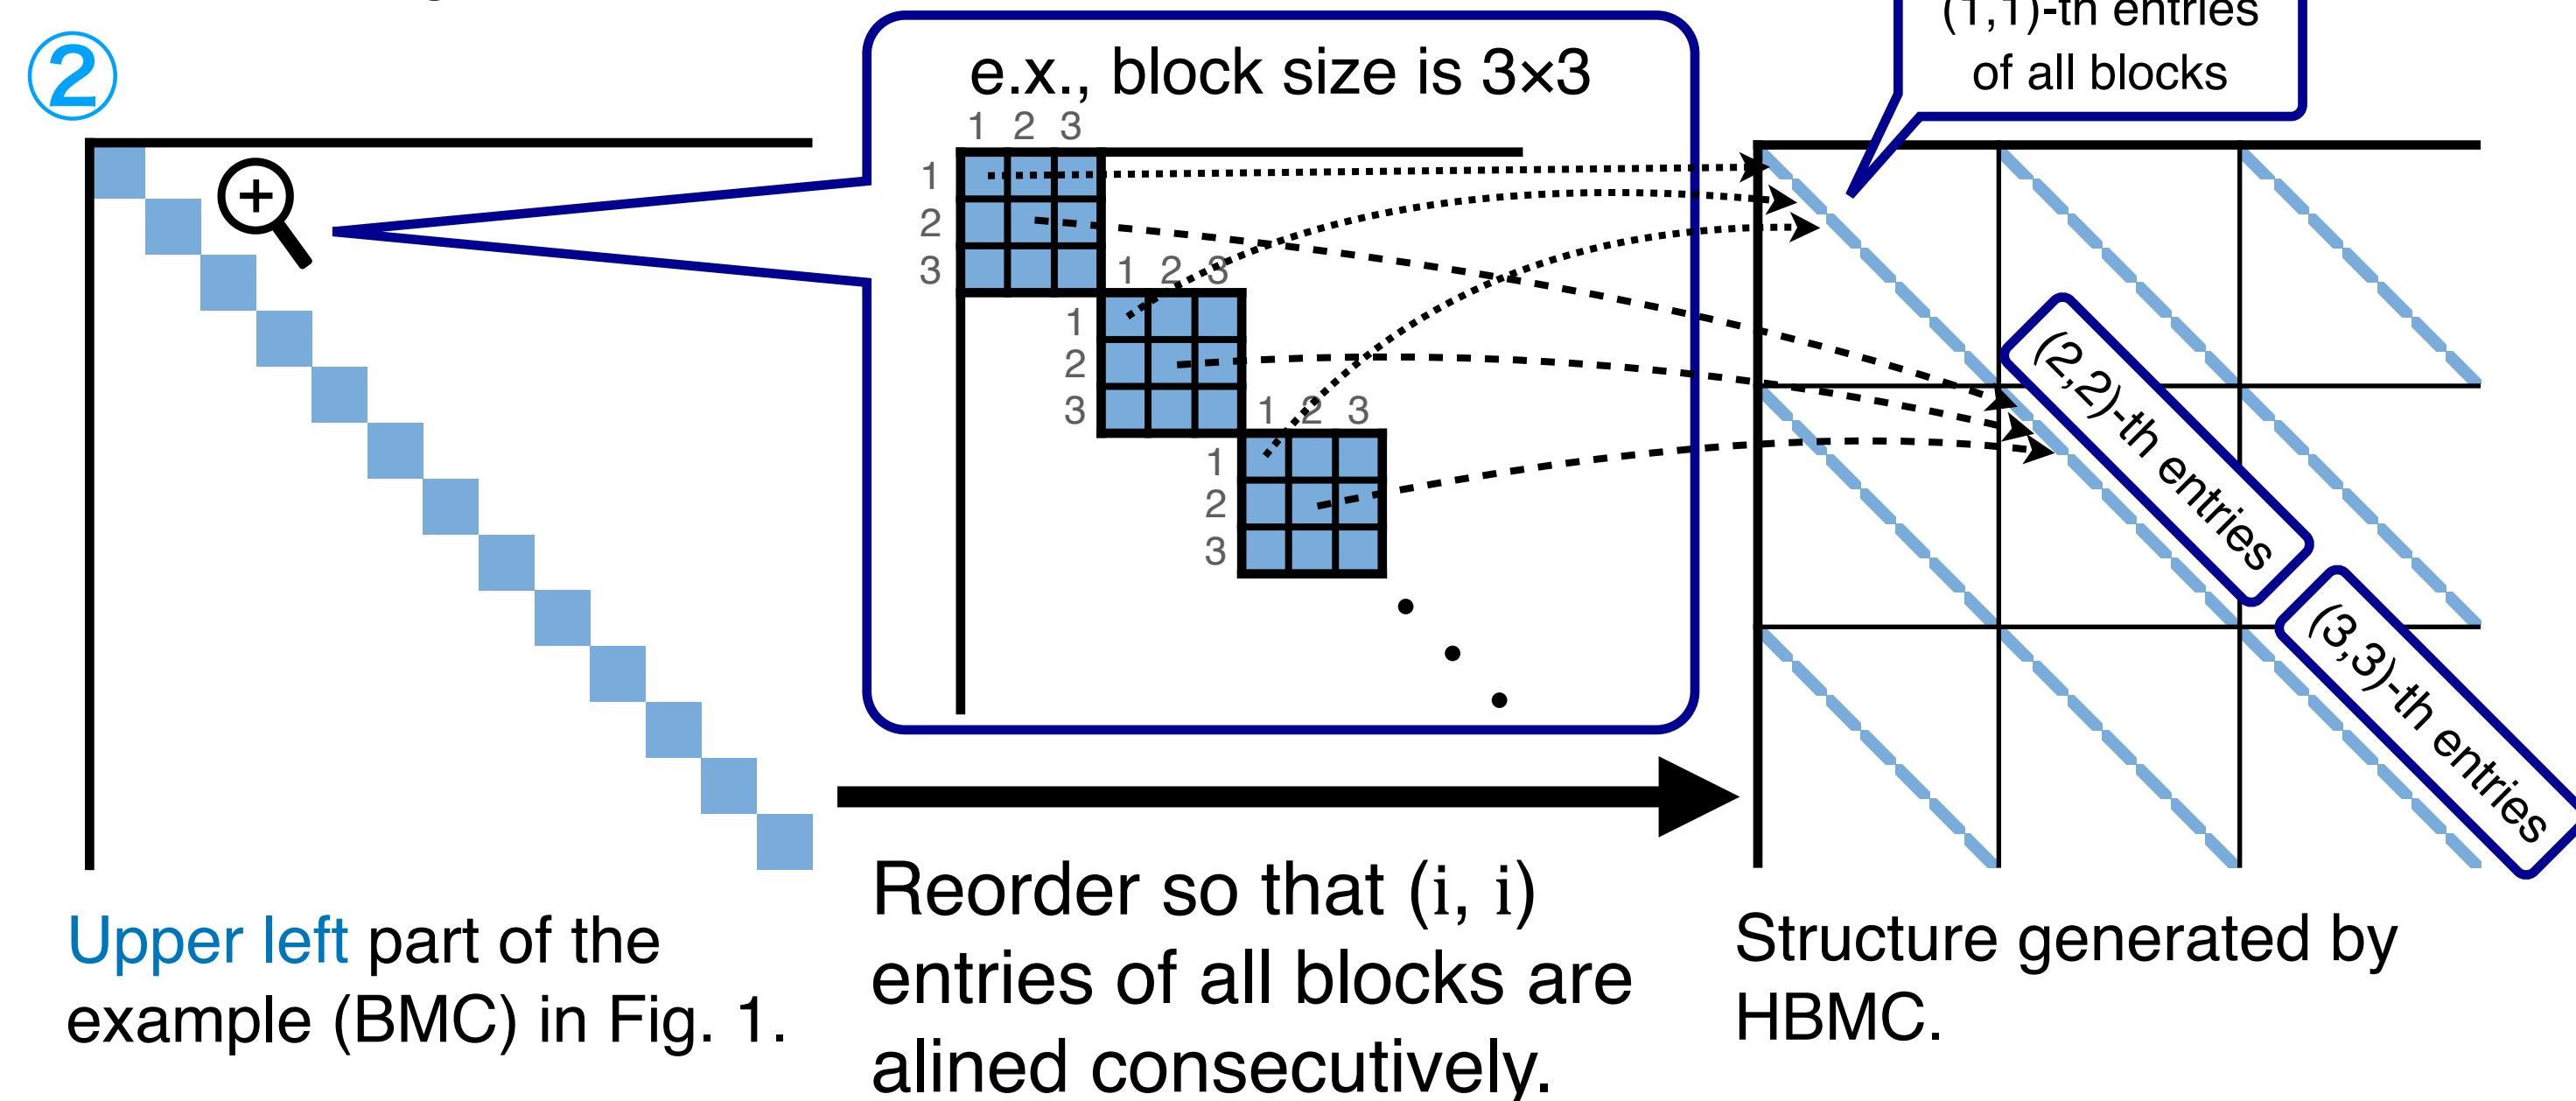


ILU(0) has the same sparse pattern as  $A$ ; transforming  $A$  reorders ILU(0).

## HBMC ordering

### Enables the vectorization of BMC-ILU(0)

HBMC expands the blocks (of size  $m \times m$ ) of BMC into a main diagonal and  $m-1$  sub-diagonals.

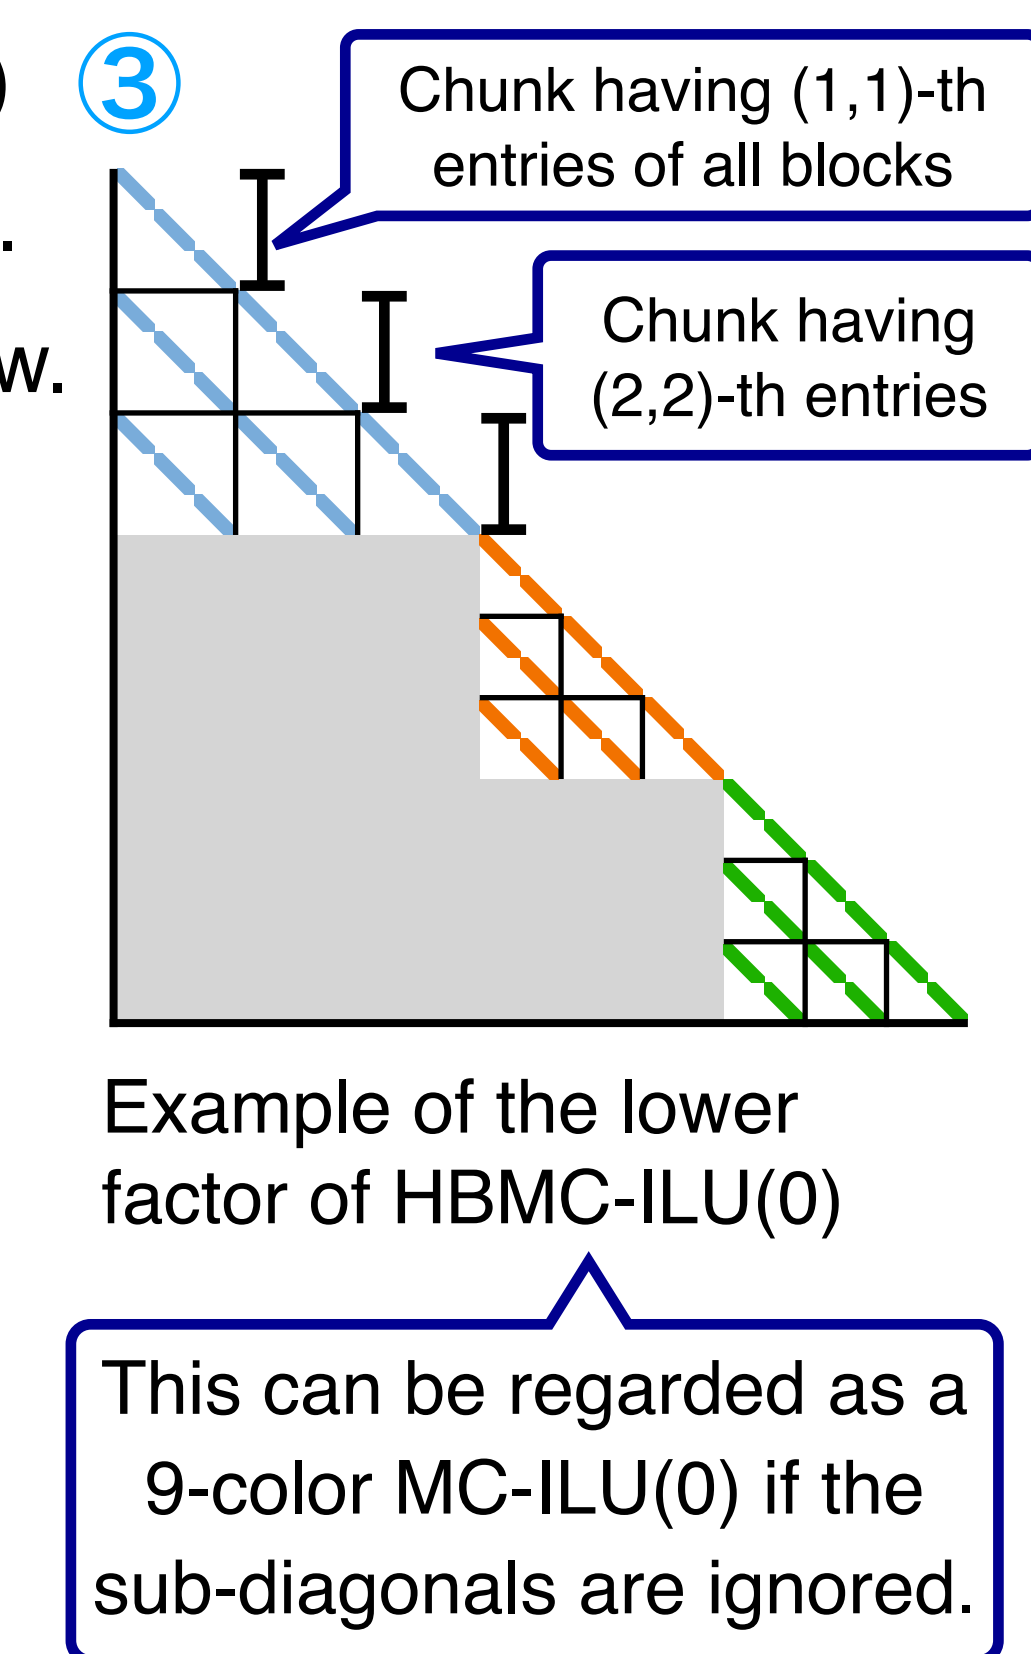


The center and lower right parts are reordered in the same way. BMC- and HBMC-ILU(0) are mathematically the same in convergence.

### Two possible GPU Implementations of HBMC-ILU(0)

- **Imp-1** (does not use the sub-diagonal structure)
  - Considers HBMC as just a  $(\#colors \times m)$ -color MC.
  - Can be programmed easily with simple data flow.
  - Requires  $(\#colors \times m) - 1$  synchronization.
- **Imp-2** (Utilizes the sub-diagonal structure)
  - Removes the redundant synchronization.
  - Requires only  $\#colors - 1$  synchronization.
  - Needs a more complex program; Per color, the thread responsible for a certain row of the chunk having (1,1)-th entries computes sequentially the same rows in the chunks including (i, i)-th ( $i = 2, \dots, m$ ) entries.


Imp-2 is not described in the extended abstract.



## MMC ordering (proposed method)

### A variation of Imp-1 of HBMC-ILU(0)

Since Imp-1 accepts a structure shown in Fig. 4, coloring can be performed more aggressively.



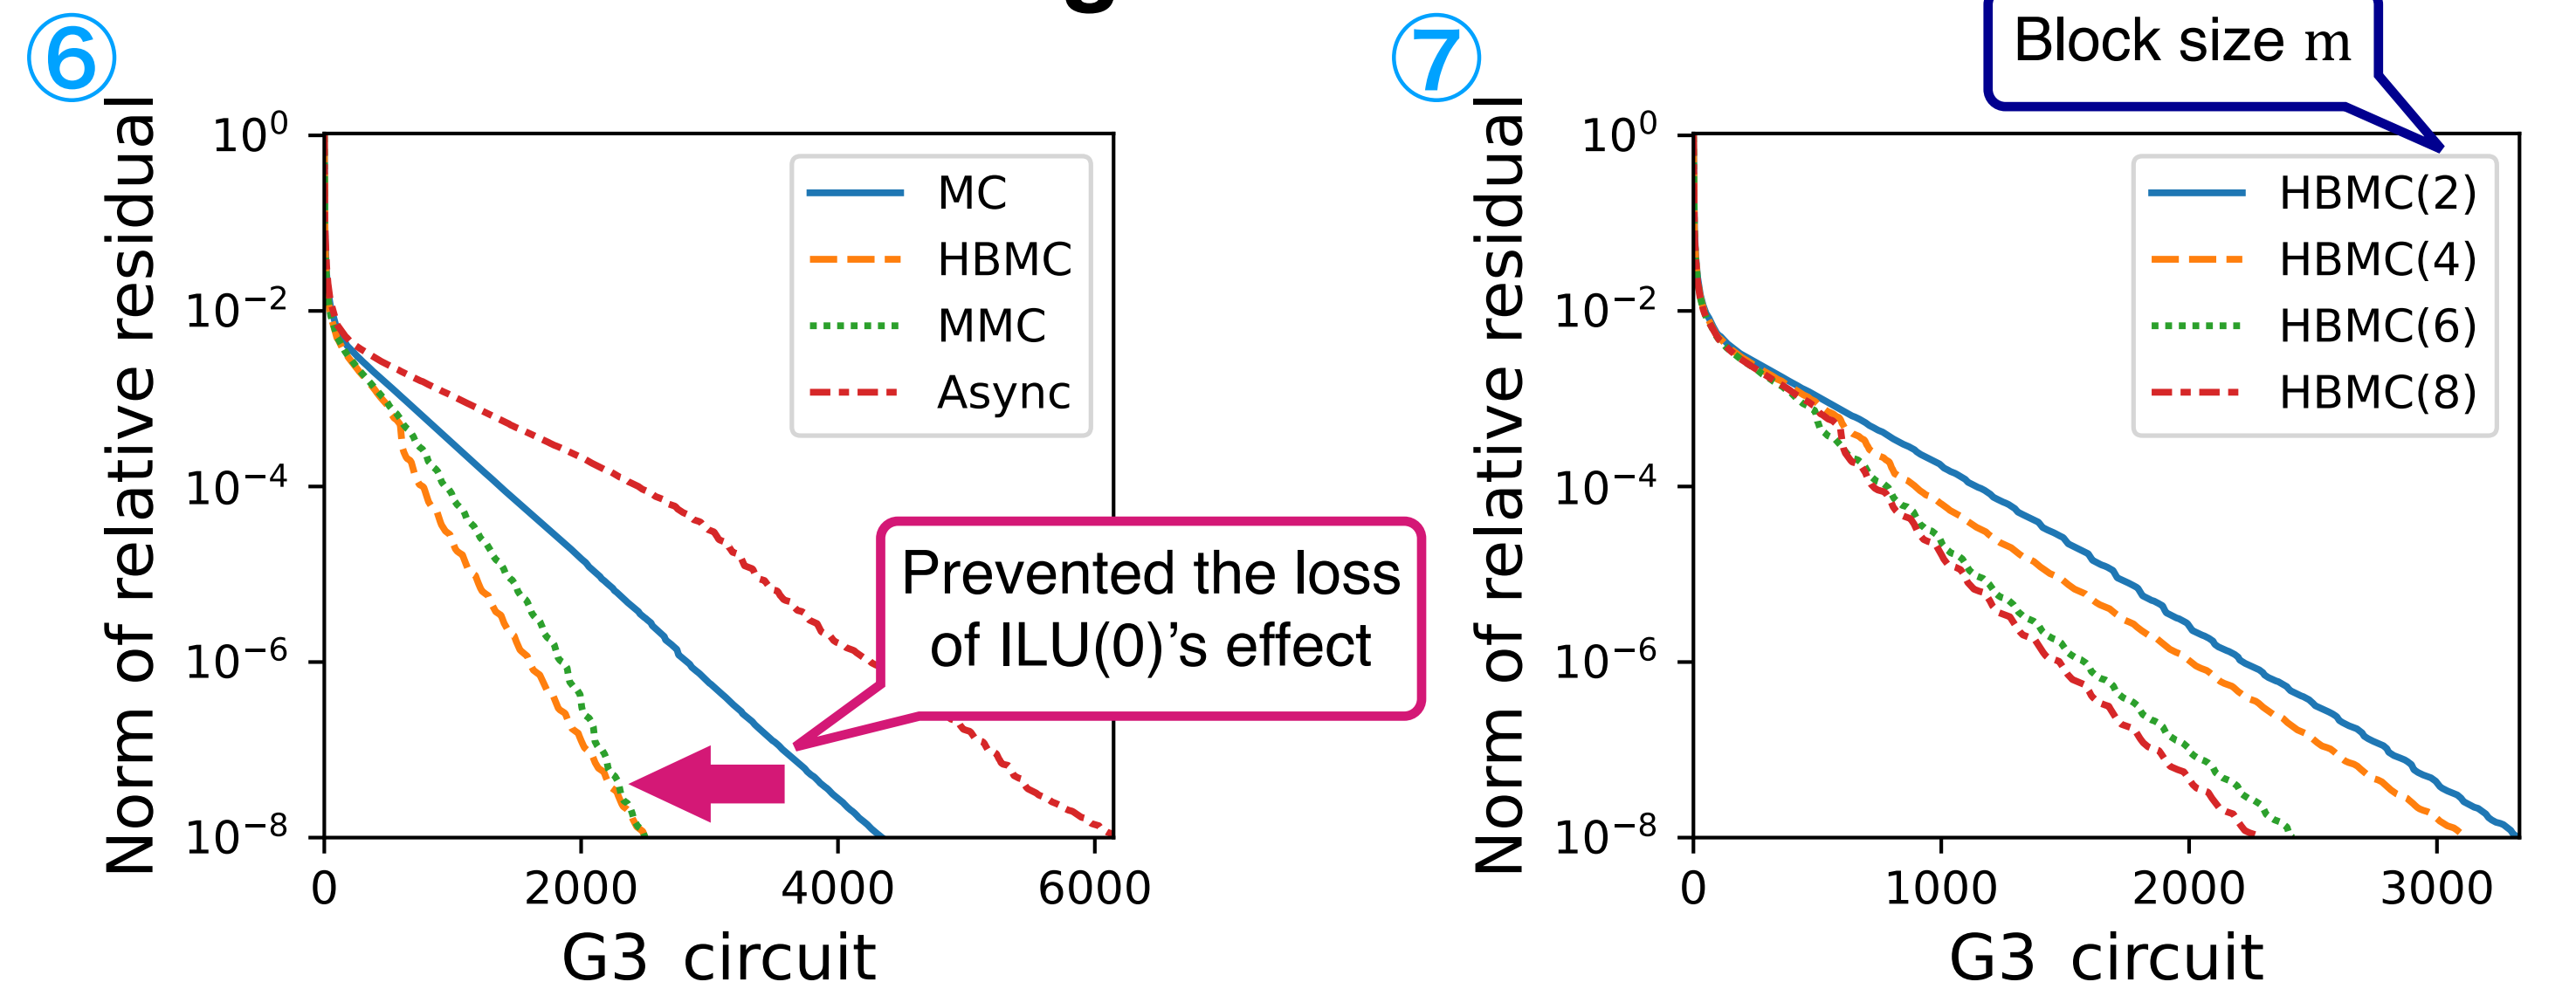
The aggressive coloring may degrade the convergence of HBMC-ILU(0).

## Numerical results

- Selected the 9 test problems from the SuiteSparse Matrix Collection.
- Solved the problems on an NVIDIA V100 controlled by an Intel Xeon.
- Compared the five preconditioners below using the FGMRES(50) solver.


Name	Description
MC (baseline)	ILU(0) with MC ordering
HBMC-1	ILU(0) with HBMC in Imp-1
HBMC-2	ILU(0) with HBMC in Imp-2
MMC	ILU(0) with MMC, a variation of HBMC-1
Async	ILU(0) with an asynchronous block Jacobi solver

### Evaluation on convergence



- As the block size increased, HBMC-ILU(0) achieved better convergence, although execution time per iteration increased.
- The reordering-based ILU(0) performed well even on problems where the iterative solver-based ILU(0) (Async) required many iterations.

### Evaluation on solution time



## Conclusions

- **Block-based matrix reordering** is effective in accelerating ILU(0) even on a GPU and can outperform the iterative solver-based approach.
- **HBMC (in imp-2)** is effective for problems with few nonzeros, while **MC and MMC** are useful for problems with many nonzeros.

[1] T. Iwashita, et. al. CCF Transaction on High Performance Computing 2 (2020), 84–97.