Data Distributions for Symmetric Linear Algebra Kernels

Olivier Beaumont, Jean-Alexandre Collin, Philippe Duchon,
Lionel Eyraud-Dubois, Julien Langou, Mathieu Vérité

LaBRI, Inria Center of the University of Bordeaux
University of Colorado, Denver

1st ACDA Workshop in Aussois
4th to 9th September 2022

Table of Contents

Context

Data placement for Linear Algebra operations

- Linear Algebra is everywhere in Scientific Computing
  - Solving Partial Differential Equations becomes $Ax = b$ after discretization
- Very computationally intensive: **distributed** execution necessary
- Tightly coupled: importance of minimizing communications
- Objective: reduce the **total volume of communications**

In this talk

- Focus on **symmetric** operations: SYRK ($C \mathrel{+}= A \cdot A^\mathsf{T}$), Cholesky
- Analyze different placement strategies and their effect on communications
- Obtain **optimal** placement and experimental validation

Note: previous works show *asymptotic* results ($O(\cdot)$); we focus on *explicit* leading coefficient

Matrix Multiplication, 2DBC, communication volume

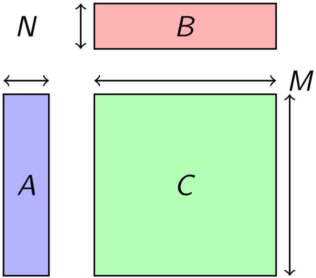

GEneral **M**atrix **M**ultiplication: $C \mathrel{+}= A \cdot B$
on P nodes

Matrix Multiplication, 2DBC, communication volume

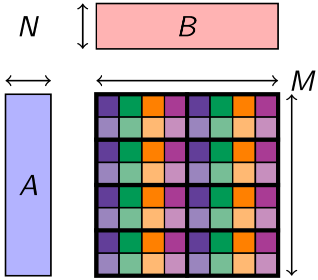

2D Block Cyclic 2×4, $P = 8$ nodes

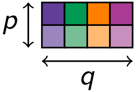

Matrix Multiplication, 2DBC, communication volume

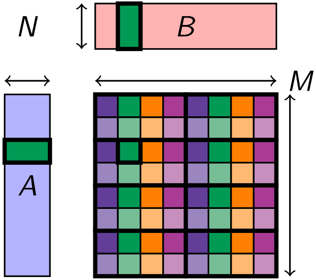

2D Block Cyclic 2×4, $P = 8$ nodes

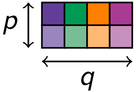

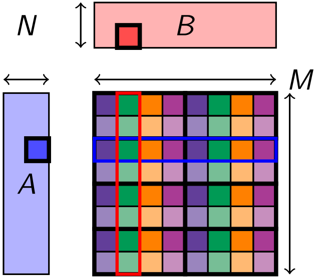

2D Block Cyclic 2×4, $P = 8$ nodes

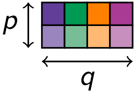

Matrix Multiplication, 2DBC, communication volume

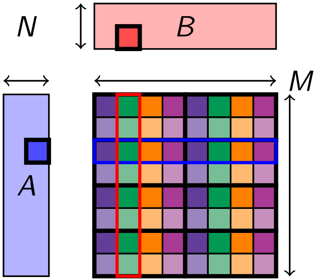

2D Block Cyclic 2×4, $P = 8$ nodes

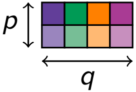

Communication volume: number of values communicated

Each tile of A is used by q nodes, each tile of B by p nodes.

$$V = MN(q-1) + MN(p-1) = MN(p+q-2)$$

2DBC, Arithmetic Intensity

Arithmetic Intensity: $\rho = \frac{\text{number of computations}}{\text{communication volume}}$

- Total computations: $2M^2N$ (N products and N additions per element of C)

$$\rho = \frac{2M^2N}{MN(p+q-2)}$$

2DBC, Arithmetic Intensity

Arithmetic Intensity: $\rho = \frac{\text{number of computations}}{\text{communication volume}}$

- Total computations: $2M^2N$ (N products and N additions per element of C)

$$\rho = \frac{2M^2N}{MN(p+q-2)} \simeq \frac{2M^2N}{2MN\sqrt{P}} = \frac{M}{\sqrt{P}} \quad \text{if } p \simeq q \simeq \sqrt{P}$$

2DBC, Arithmetic Intensity

Arithmetic Intensity: $\rho = \frac{\text{number of computations}}{\text{communication volume}}$

- Total computations: $2M^2N$ (N products and N additions per element of C)

$$\rho = \frac{2M^2N}{MN(p+q-2)} \simeq \frac{2M^2N}{2MN\sqrt{P}} = \frac{M}{\sqrt{P}} \quad \text{if } p \simeq q \simeq \sqrt{P}$$

- $S =$ number of elements of C per node $= \frac{M^2}{P}$

$$\boxed{\rho = \sqrt{S}}$$

2DBC, Arithmetic Intensity

Arithmetic Intensity: $\rho = \frac{\text{number of computations}}{\text{communication volume}}$

- Total computations: $2M^2N$ (N products and N additions per element of C)

$$\rho = \frac{2M^2N}{MN(p+q-2)} \simeq \frac{2M^2N}{2MN\sqrt{P}} = \frac{M}{\sqrt{P}} \quad \text{if } p \simeq q \simeq \sqrt{P}$$

- S = number of elements of C per node $= \frac{M^2}{P}$

$$\boxed{\rho = \sqrt{S}}$$

- This is **optimal**

Table of Contents

Symmetric multiplication is SYRK: $C \mathrel{+}= A \cdot A^{\mathsf{T}}$

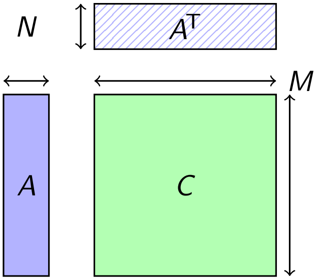

SYRK: $C \mathrel{+}= A \cdot A^{\mathsf{T}}$ (**SY**mmetric **R**ank-**K** update)

dominant part of Cholesky factorization

(solve $A = L \cdot L^{\mathsf{T}}$ for symmetric positive definite matrix A)

Symmetric multiplication is SYRK: $C \mathrel{+}= A \cdot A^{\mathsf{T}}$

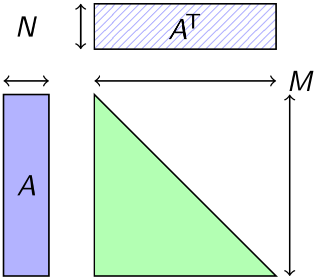

SYRK: $C \mathrel{+}= A \cdot A^{\mathsf{T}}$ (**SY**mmetric **R**ank-**K** update)

dominant part of Cholesky factorization

(solve $A = L \cdot L^{\mathsf{T}}$ for symmetric positive definite matrix A)

Symmetric multiplication is SYRK: $C \mathrel{+}= A \cdot A^T$

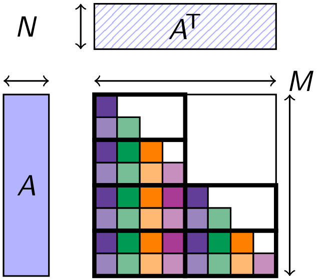

2D Block Cyclic 2×4, $P = 8$ nodes

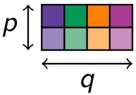

Symmetric multiplication is SYRK: $C += A \cdot A^{\mathsf{T}}$

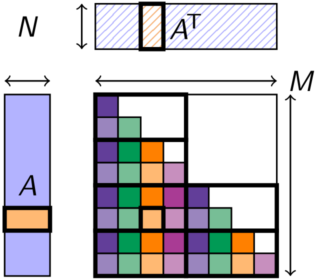

2D Block Cyclic 2×4, $P = 8$ nodes

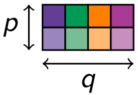

Symmetric multiplication is SYRK: $C \mathrel{+}= A \cdot A^{\mathsf{T}}$

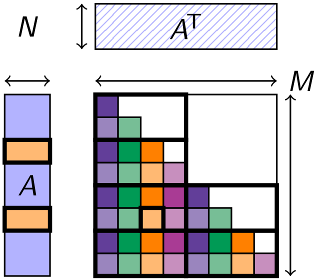

2D Block Cyclic 2×4, $P = 8$ nodes

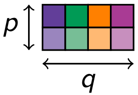

Symmetric multiplication is SYRK: $C += A \cdot A^{\mathsf{T}}$

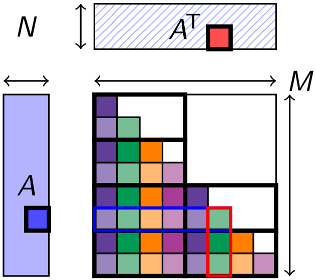

2D Block Cyclic 2×4, $P = 8$ nodes

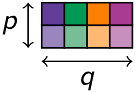

Symmetric multiplication is SYRK: $C \mathrel{+}= A \cdot A^{\mathsf{T}}$

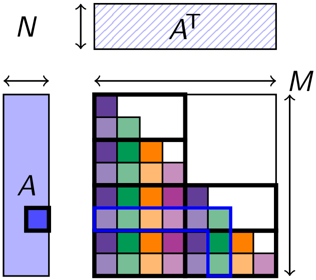

2D Block Cyclic 2×4, $P = 8$ nodes

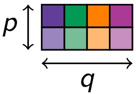

Symmetric multiplication is SYRK: $C \mathrel{+}= A \cdot A^T$

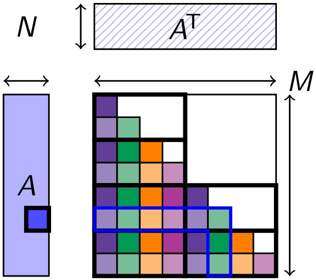

2D Block Cyclic 2×4, $P = 8$ nodes

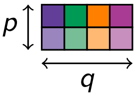

Communication volume

Each tile of A is used by $p + q - 1$ nodes: $V = MN(p + q - 2)$

Symmetric multiplication is SYRK: $C \mathrel{+}= A \cdot A^{\mathsf{T}}$

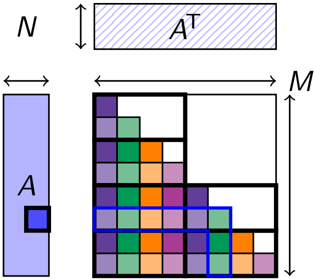

2D Block Cyclic 2×4, $P = 8$ nodes

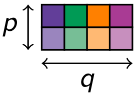

Communication volume

Each tile of A is used by $p + q - 1$ nodes: $V = MN(p + q - 2)$

Arithmetic Intensity: $\rho = \dfrac{M^2 N}{2MN\sqrt{P}} = \dfrac{M}{2\sqrt{P}}$, and since $S = \dfrac{M^2}{2P}$ now: $\boxed{\rho = \dfrac{\sqrt{S}}{\sqrt{2}}}$

Symmetric multiplication is SYRK: $C \mathrel{+}= A \cdot A^{\mathsf{T}}$

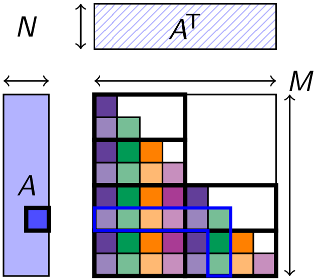

2D Block Cyclic 2×4, $P = 8$ nodes

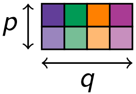

Communication volume

Each tile of A is used by $p + q - 1$ nodes: $V = MN(p + q - 2)$

Arithmetic Intensity: $\rho = \dfrac{M^2 N}{2MN\sqrt{P}} = \dfrac{M}{2\sqrt{P}}$, and since $S = \dfrac{M^2}{2P}$ now: $\boxed{\rho = \dfrac{\sqrt{S}}{\sqrt{2}}}$

Best known bound [Olivry et al. 2020]: $2\sqrt{S}$ 

SBC: Symmetric Block Cyclic [SC'2022]

Goal: same nodes on rows and columns

SBC: Symmetric Block Cyclic [SC'2022]

Goal: same nodes on rows and columns

Symmetric Block Cyclic $P = 8$

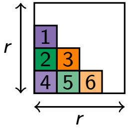

$$P = \frac{r^2}{2} \quad \Leftrightarrow \quad r = \sqrt{2P}$$

SBC: Symmetric Block Cyclic [SC'2022]

Goal: same nodes on rows and columns

Symmetric Block Cyclic $P = 8$

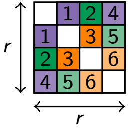

$$P = \frac{r^2}{2} \quad \Leftrightarrow \quad r = \sqrt{2P}$$

Goal: same nodes on rows and columns

Symmetric Block Cyclic $P = 8$

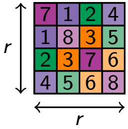

$$P = \frac{r^2}{2} \quad \Leftrightarrow \quad r = \sqrt{2P}$$

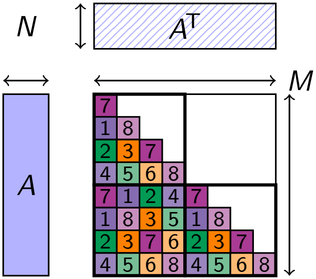

Goal: same nodes on rows and columns

Symmetric Block Cyclic $P = 8$

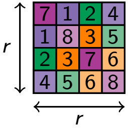

$$P = \frac{r^2}{2} \quad \Leftrightarrow \quad r = \sqrt{2P}$$

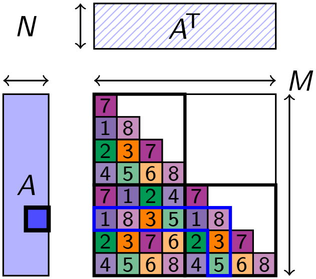

Goal: same nodes on rows and columns

Symmetric Block Cyclic $P = 8$

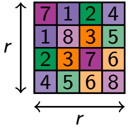

$$P = \frac{r^2}{2} \quad \Leftrightarrow \quad r = \sqrt{2P}$$

SBC: Symmetric Block Cyclic [SC'2022]

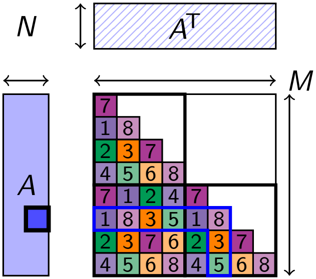

Goal: same nodes on rows and columns

Symmetric Block Cyclic $P = 8$

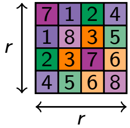

$P = \frac{r^2}{2} \quad \Leftrightarrow \quad r = \sqrt{2P}$

Communication volume

One tile of A is needed by r nodes: $V = MN(r-1)$

Arithmetic intensity: $\rho = \dfrac{M^2N}{MN(r-1)} = \dfrac{M}{\sqrt{2P}} = \boxed{\sqrt{S}}$

Experimental results: Cholesky factorization on $P \sim 28$ nodes

CHAMELEON+STARPU on `bora` cluster (36 *Intel Xeon Skylake Gold 6240* cores per node)

Experimental results: Cholesky factorization on $P \sim 28$ nodes

CHAMELEON+STARPU on `bora` cluster (36 *Intel Xeon Skylake Gold 6240* cores per node)

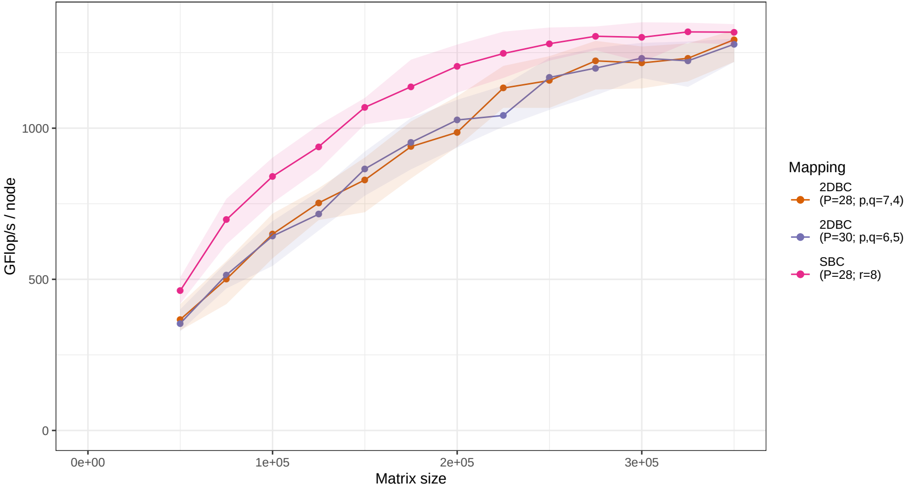

Table of Contents

Can we do better?

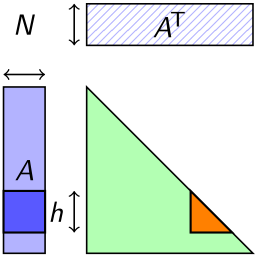

If one node computes a triangle on the diagonal

$S = \frac{h^2}{2} \qquad V = N \cdot h$

$\Rightarrow \rho = \frac{h^2 N}{Nh} = h = \sqrt{2S}$

Can we do better?

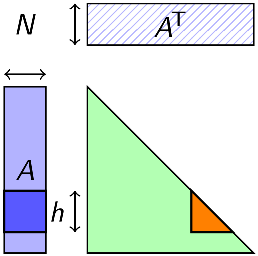

If one node computes a triangle on the diagonal

$S = \frac{h^2}{2} \qquad V = N \cdot h$

$\Rightarrow \rho = \frac{h^2 N}{Nh} = h = \sqrt{2S}$

Theorem [SPAA'2022]

- The communication volume for a SYRK operation with memory S is at least $V \geqslant \frac{M^2 N}{\sqrt{2S}}$

  Equivalent to $\boxed{\rho \leqslant \sqrt{2S}}$ ([Olivry et al 2020]: $\rho \leqslant 2\sqrt{S}$)

- There exists a **sequential, out-of-core** algorithm with $V = \frac{M^2 N}{\sqrt{2S}} + \frac{N^2}{2} + O(NM \log N)$

TBC distribution: **T**riangle **B**lock **C**yclic

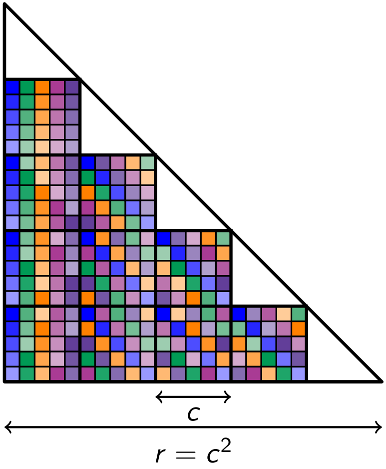

- Consider a **prime** integer c
- Build a pattern of size $c^2 \times c^2$
- c^2 nodes own $\frac{c(c-1)}{2}$ tiles  (from [SPAA'22])

TBC distribution: **T**riangle **B**lock **C**yclic

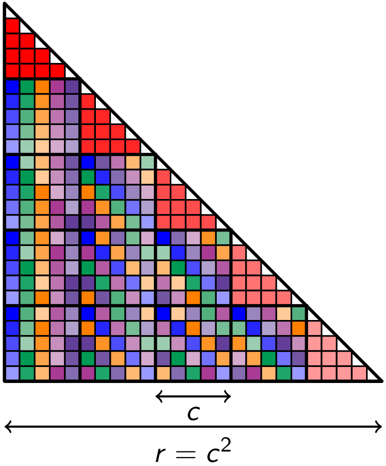

- Consider a **prime** integer c
- Build a pattern of size $c^2 \times c^2$
- c^2 nodes own $\frac{c(c-1)}{2}$ tiles  (from [SPAA'22])
- c nodes own tiles close to diagonal 

TBC distribution: **T**riangle **B**lock **C**yclic

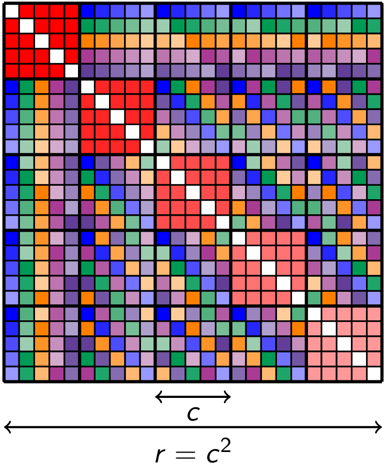

- Consider a **prime** integer c
- Build a pattern of size $c^2 \times c^2$
- c^2 nodes own $\frac{c(c-1)}{2}$ tiles  (from [SPAA'22])
- c nodes own tiles close to diagonal 
- Symmetric balanced pattern with $P = c(c+1)$
  - spans all tiles, except the diagonal

TBC distribution: **T**riangle **B**lock **C**yclic

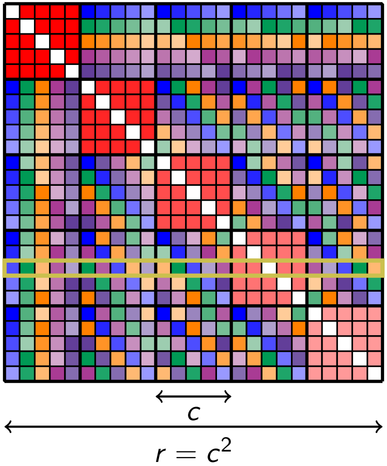

- Consider a **prime** integer c
- Build a pattern of size $c^2 \times c^2$
- c^2 nodes own $\frac{c(c-1)}{2}$ tiles  (from [SPAA'22])
- c nodes own tiles close to diagonal 
- Symmetric balanced pattern with $P = c(c+1)$
  - spans all tiles, except the diagonal
- Each row contains $c + 1$ nodes

TBC distribution: **T**riangle **B**lock **C**yclic

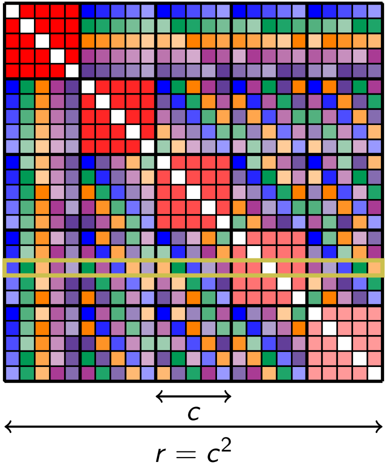

- Consider a **prime** integer c
- Build a pattern of size $c^2 \times c^2$
- c^2 nodes own $\frac{c(c-1)}{2}$ tiles  (from [SPAA'22])
- c nodes own tiles close to diagonal 
- Symmetric balanced pattern with $P = c(c+1)$
  - spans all tiles, except the diagonal
- Each row contains $c+1$ nodes
- Diagonal tiles can be assigned to any node on the row

Triangle Block Cyclic: replicating the pattern

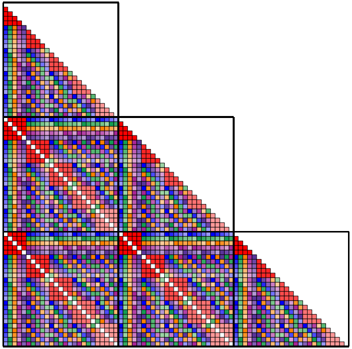

$c = 5, c^2 = 25, P = 30, M = 75$

- Replicate the pattern on the whole matrix A
- Assign remaining tiles *greedily*: for each unassigned tile,
  - among all nodes **in this row**
  - pick the one with the **lowest** number of assigned tiles
- This does not increase the communication volume
- Technically, no longer a *cyclic* distribution

Triangle Block Cyclic: replicating the pattern

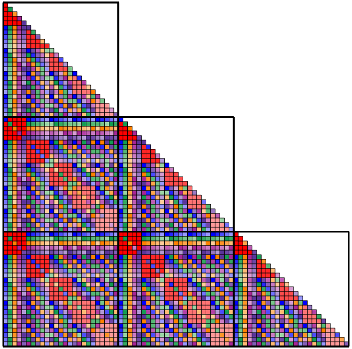

$c = 5, c^2 = 25, P = 30, M = 75$

- Replicate the pattern on the whole matrix A
- Assign remaining tiles *greedily*: for each unassigned tile,
  - among all nodes **in this row**
  - pick the one with the **lowest** number of assigned tiles
- This does not increase the communication volume
- Technically, no longer a *cyclic* distribution

Triangle Block Cyclic: replicating the pattern

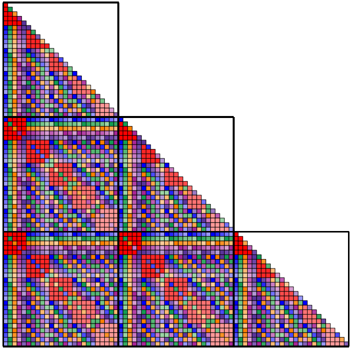

$c = 5, c^2 = 25, P = 30, M = 75$

- Replicate the pattern on the whole matrix A
- Assign remaining tiles *greedily*: for each unassigned tile,
  - among all nodes **in this row**
  - pick the one with the **lowest** number of assigned tiles
- This does not increase the communication volume
- Technically, no longer a *cyclic* distribution

Communication volume

Each tile of A is used by $c+1$ nodes: $V = MN \cdot c$
$(P = c(c+1) \Rightarrow c \simeq \sqrt{P})$

$$\rho = \frac{M^2 N}{MNc} \simeq \frac{M}{\sqrt{P}} = \boxed{\sqrt{2S}}$$

Experimental results on Cholesky factorization

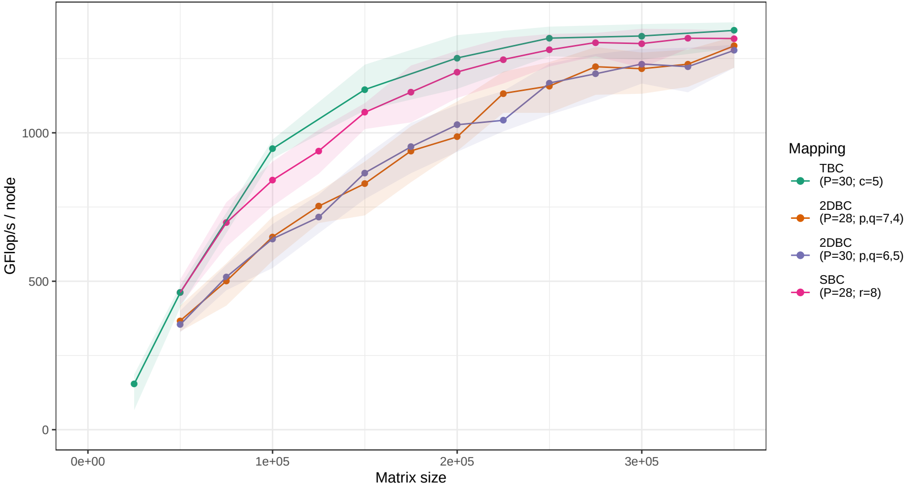

Table of Contents

Conclusions

Summary of Arithmetic Intensities

Operation	2DBC	SBC	TBC	Bound
GEMM	$1\ \sqrt{S}$	$\frac{1}{\sqrt{2}}\ \sqrt{S}$	$1\ \sqrt{S}$	$1\ \sqrt{S}$
SYRK	$\frac{1}{\sqrt{2}}\ \sqrt{S}$	$1\ \sqrt{S}$	$\sqrt{2}\ \sqrt{S}$	$\sqrt{2}\ \sqrt{S}$

Contributions

- **Improved lower bound** on communication volume for SYRK & Cholesky
- Proposed **symmetric** distributions
- Significantly **improved performance** for Cholesky factorization
- Can be applied to many other symmetric computations

Conclusions

Summary of Arithmetic Intensities

Operation	2DBC	SBC	TBC	Bound
GEMM	$1\ \sqrt{S}$	$\frac{1}{\sqrt{2}}\ \sqrt{S}$	$1\ \sqrt{S}$	$1\ \sqrt{S}$
SYRK	$\frac{1}{\sqrt{2}}\ \sqrt{S}$	$1\ \sqrt{S}$	$\sqrt{2}\ \sqrt{S}$	$\sqrt{2}\ \sqrt{S}$

Contributions

- **Improved lower bound** on communication volume for SYRK & Cholesky
- Proposed **symmetric** distributions
- Significantly **improved performance** for Cholesky factorization
- Can be applied to many other symmetric computations

Conclusions

Summary of Arithmetic Intensities

Operation	2DBC	SBC	TBC	Bound
GEMM	$1\ \sqrt{S}$	$\frac{1}{\sqrt{2}}\ \sqrt{S}$	$1\ \sqrt{S}$	$1\ \sqrt{S}$
SYRK	$\frac{1}{\sqrt{2}}\ \sqrt{S}$	$1\ \sqrt{S}$	$\sqrt{2}\ \sqrt{S}$	$\sqrt{2}\ \sqrt{S}$

Contributions

- **Improved lower bound** on communication volume for SYRK & Cholesky
- Proposed **symmetric** distributions
- Significantly **improved performance** for Cholesky factorization
- Can be applied to many other symmetric computations

Limitations

Limits on the feasible values of P

- For SBC, $P = \frac{r^2}{2}$ or $\frac{r(r-1)}{2}$ for r integer
- For TBS, $P = c(c+1)$ for c **prime**

Limitations

Limits on the feasible values of P

- For SBC, $P = \frac{r^2}{2}$ or $\frac{r(r-1)}{2}$ for r integer
- For TBS, $P = c(c+1)$ for c **prime**

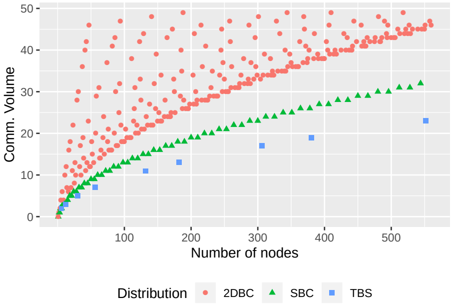

Limitations

Limits on the feasible values of P

- For SBC, $P = \frac{r^2}{2}$ or $\frac{r(r-1)}{2}$ for r integer
- For TBS, $P = c(c+1)$ for c **prime**

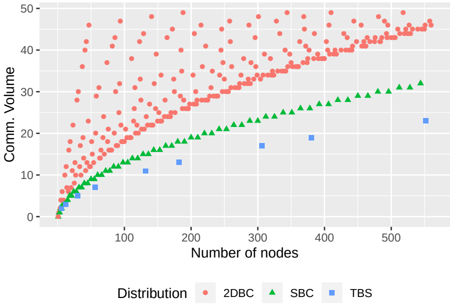

Generalize: patterns for more values of P ?

- For a given P and pattern size r
- Try to find "good" solutions
- Can be seen as an optimization problem
- For now, greedy algorithms and ILP formulations