

A highly portable heterogeneous implementation of symmetry-preserving methods for magnetohydrodynamics

J.A. Hopman, F.X. Trias and J.Rigola

Heat and Mass Transfer Technological Center (CTTC)

Technical University of Catalonia (UPC), Terrassa, Spain


Outline

Symmetry preserving methods in MHD


HPC² framework

Exploiting symmetries in geometry

MHD introduction




MHD introduction




Liquid metals in magnetic field



MHD introduction






MHD introduction



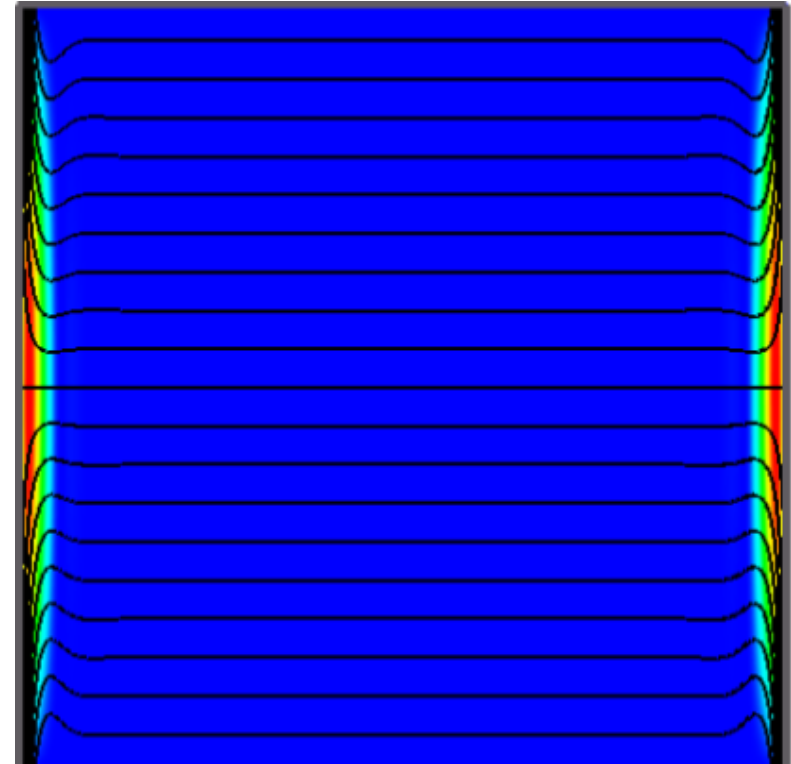


MHD introduction






MHD introduction



MHD introduction



MHD introduction



Lorentz force implementation

Following method of Ni et al.^{1, 2}

Collocated + staggered current densities, j

Conserves current density


Scheme basics

Update U using projection method

Solve 2nd Poisson equation for φ : $\nabla \cdot (\nabla \varphi) = \nabla \cdot (\mathbf{u} \times \mathbf{B})$

Update j

Interpolate to cell center to calculate F_{Lor}



¹Ni, M. J., Munipalli, R., Morley, N. B., Huang, P., & Abdou, M. A. (2007). A current density conservative scheme for incompressible MHD flows at a low magnetic Reynolds number. Part I: On a rectangular collocated grid system. *Journal of Computational Physics*, 227(1), 174-204.

²Ni, M. J., Munipalli, R., Huang, P., Morley, N. B., & Abdou, M. A. (2007). A current density conservative scheme for incompressible MHD flows at a low magnetic Reynolds number. Part II: On an arbitrary collocated mesh. *Journal of Computational Physics*, 227(1), 205-228.

Preserving symmetries

Following method of Trias et al.¹


Conserve physical properties by mimicking continuous operators

- Use projected distances in gradients
- Use midpoint interpolation
- Use face-volume weighted interpolation

Consequences for MHD


Avoid iterative correction schemes

Conserve total momentum from Lorentz force: $\int_{\Omega} \nabla \cdot (\mathbf{J}(\mathbf{B} \times \mathbf{r})) d\Omega$




¹Trias, F. X., Lehmkuhl, O., Oliva, A., Pérez-Segarra, C. D., & Verstappen, R. W. C. P. (2014). Symmetry-preserving discretization of Navier–Stokes equations on collocated unstructured grids. *Journal of Computational Physics*, 258, 246–267.

Case 1: 2D Taylor-Green vortex in transverse magnetic field




HydroDynamic TGV

Case 1: 2D Taylor-Green vortex in transverse magnetic field




HydroDynamic TGV




Transverse B -field

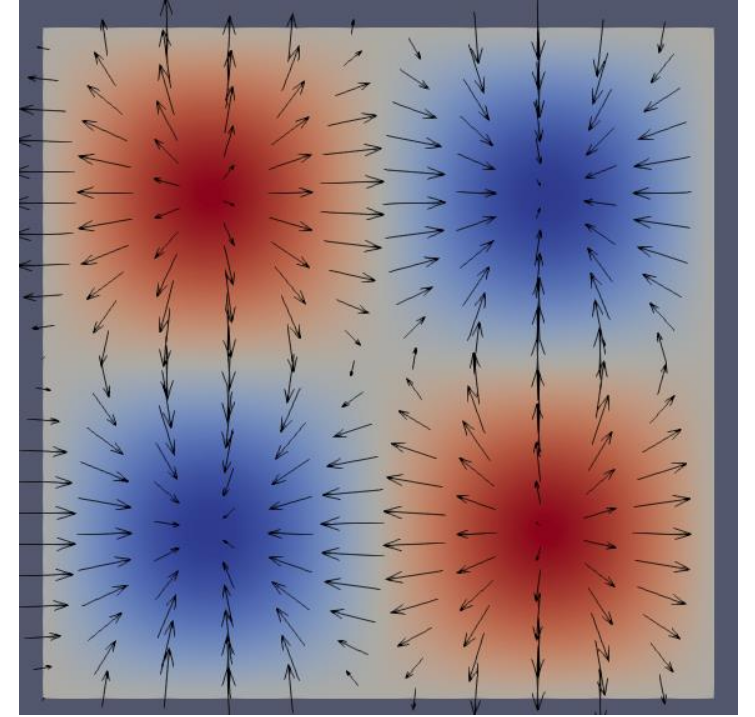
Case 1: 2D Taylor-Green vortex in transverse magnetic field



HydroDynamic TGV




Transverse B -field




Imposed φ -field

Case 1: 2D Taylor-Green vortex in transverse magnetic field




Magneto-
HydroDynamic TGV

=





Transverse B -field

+



Imposed φ -field

Case 2: MHD duct flow



Accuracy results

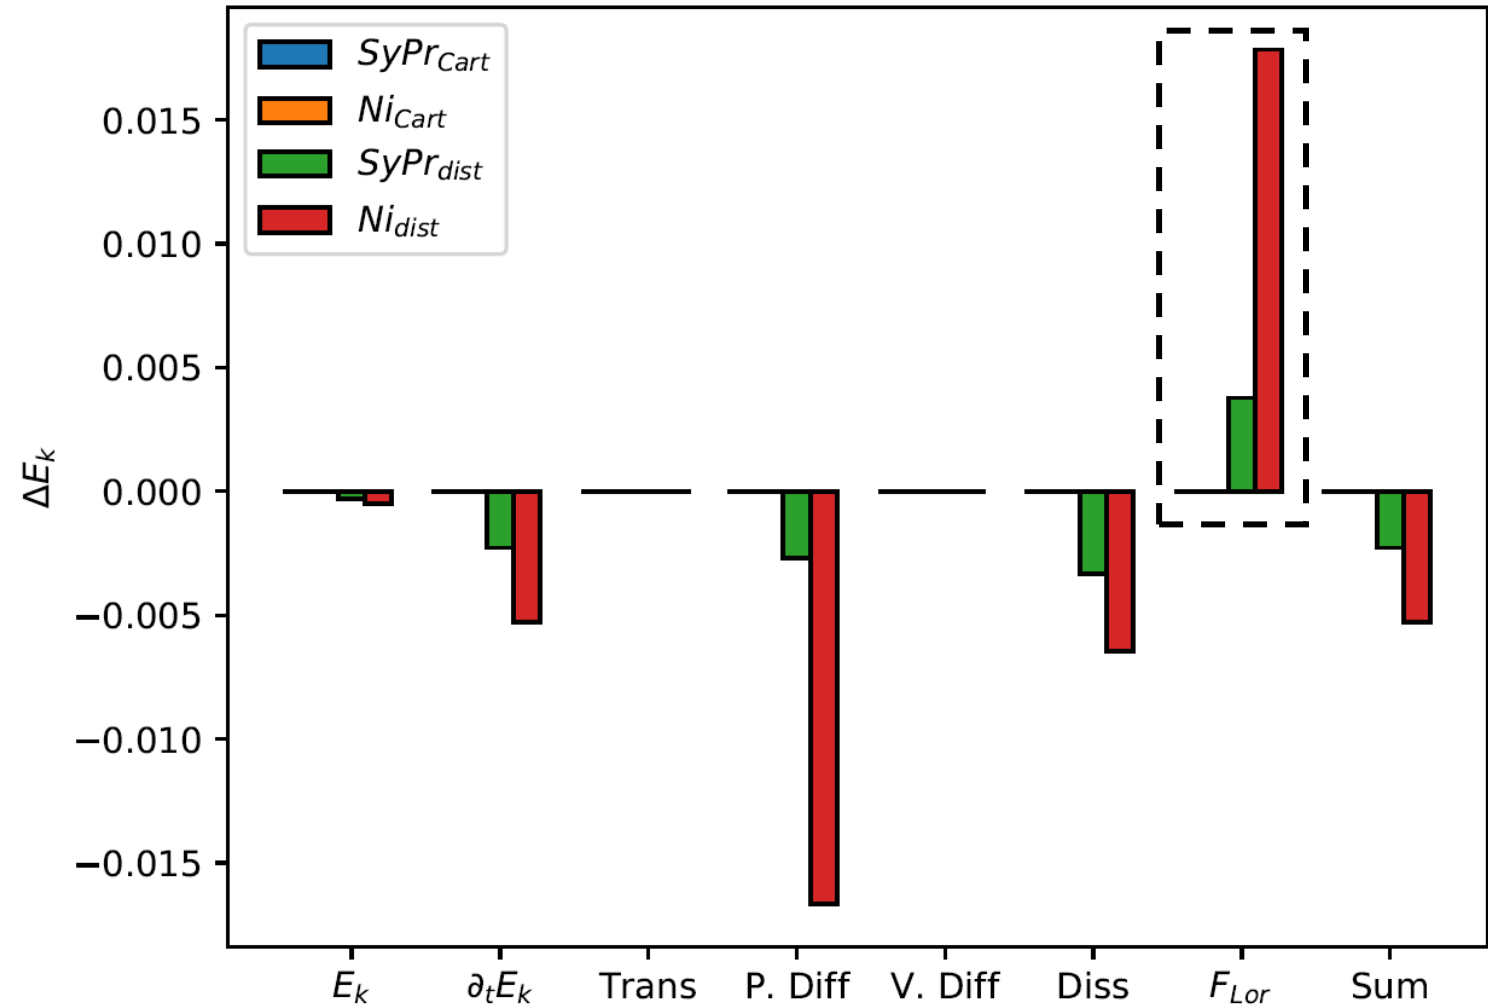
2D Taylor-Green vortex

Should not generate Lorentz Force
Should not dissipate energy

Results


Symmetry Preserving method
outperforms Ni method
Especially on distorted grids

Error of energy budgets

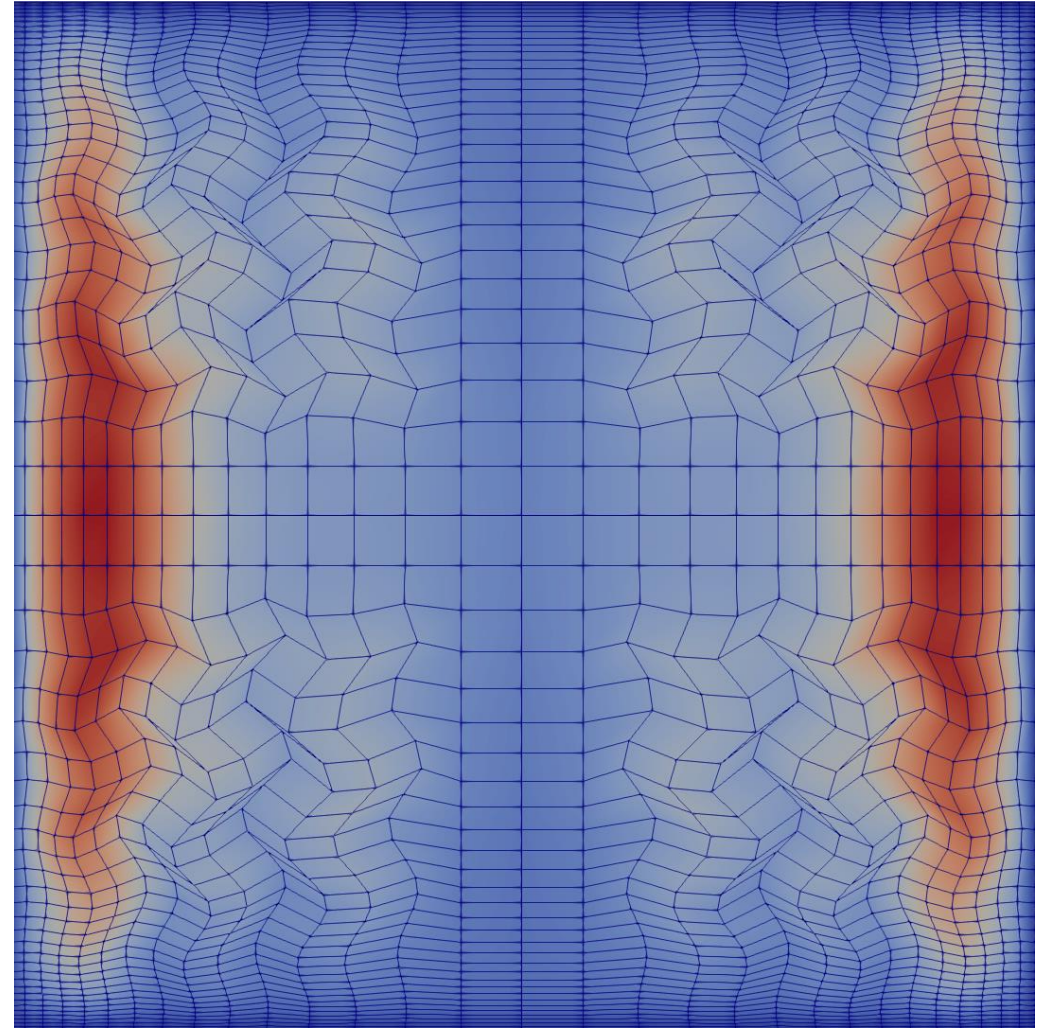


Stability for highly distorted meshes

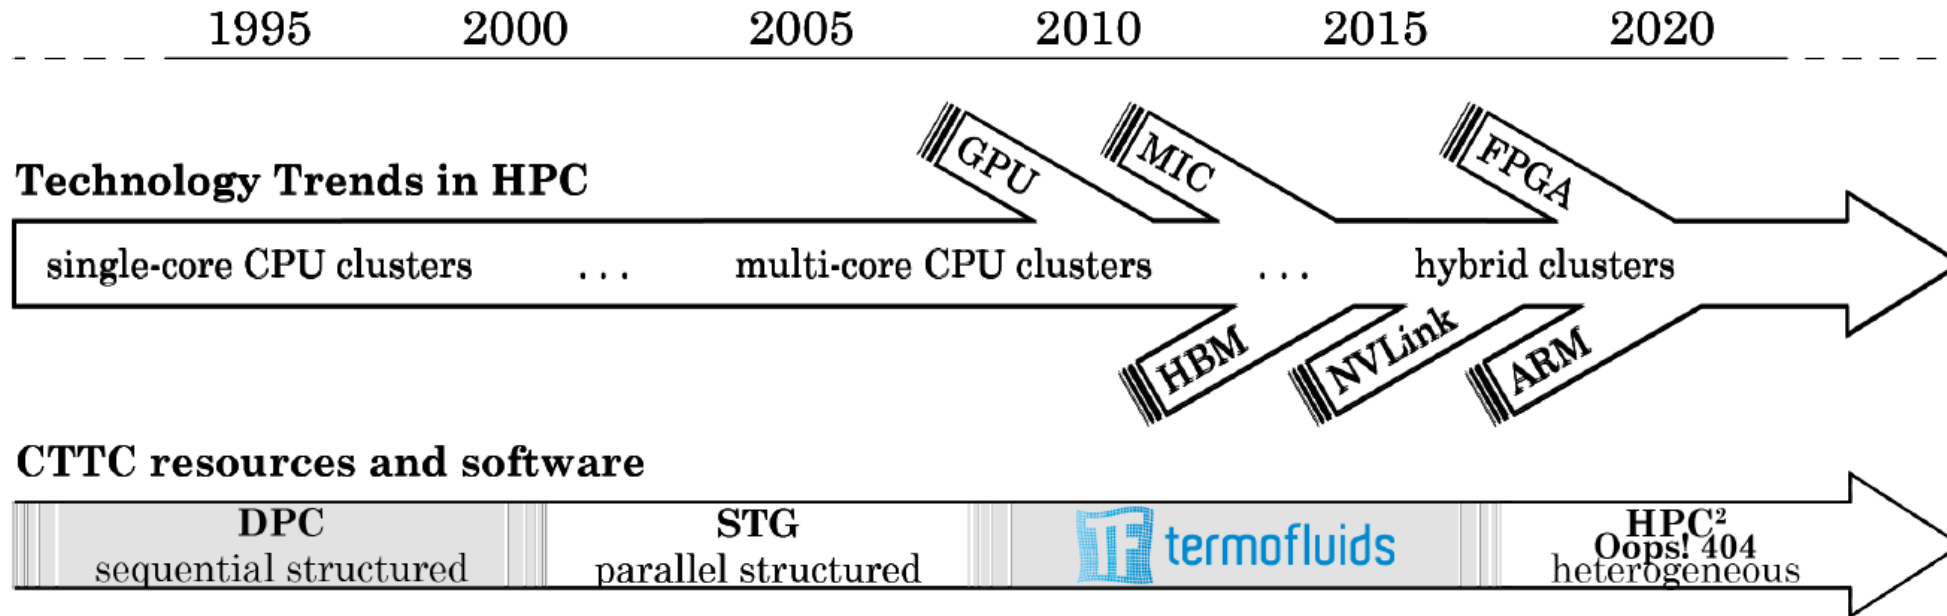
Taylor green vortex



M-profile in duct



HPC²



Highly-portable code for HPC

Stencil based → Algebra based
Only a few algebraic kernels are needed

Algebraic kernels

From continuous NS equations:

$$\nabla \cdot \mathbf{u} = 0, \quad \partial_t \mathbf{u} + (\mathbf{u} \cdot \nabla) \mathbf{u} - \frac{1}{Re} \Delta \mathbf{u} + \nabla p = 0$$

To discrete algebraic equations:

$$M \mathbf{u}_s = \mathbf{0}_c, \quad \Omega \partial_t \mathbf{u}_c + C(\mathbf{u}_s) \mathbf{u}_c + D \mathbf{u}_c + \Omega G p_c = \mathbf{0}_c$$

Using three kernels only:

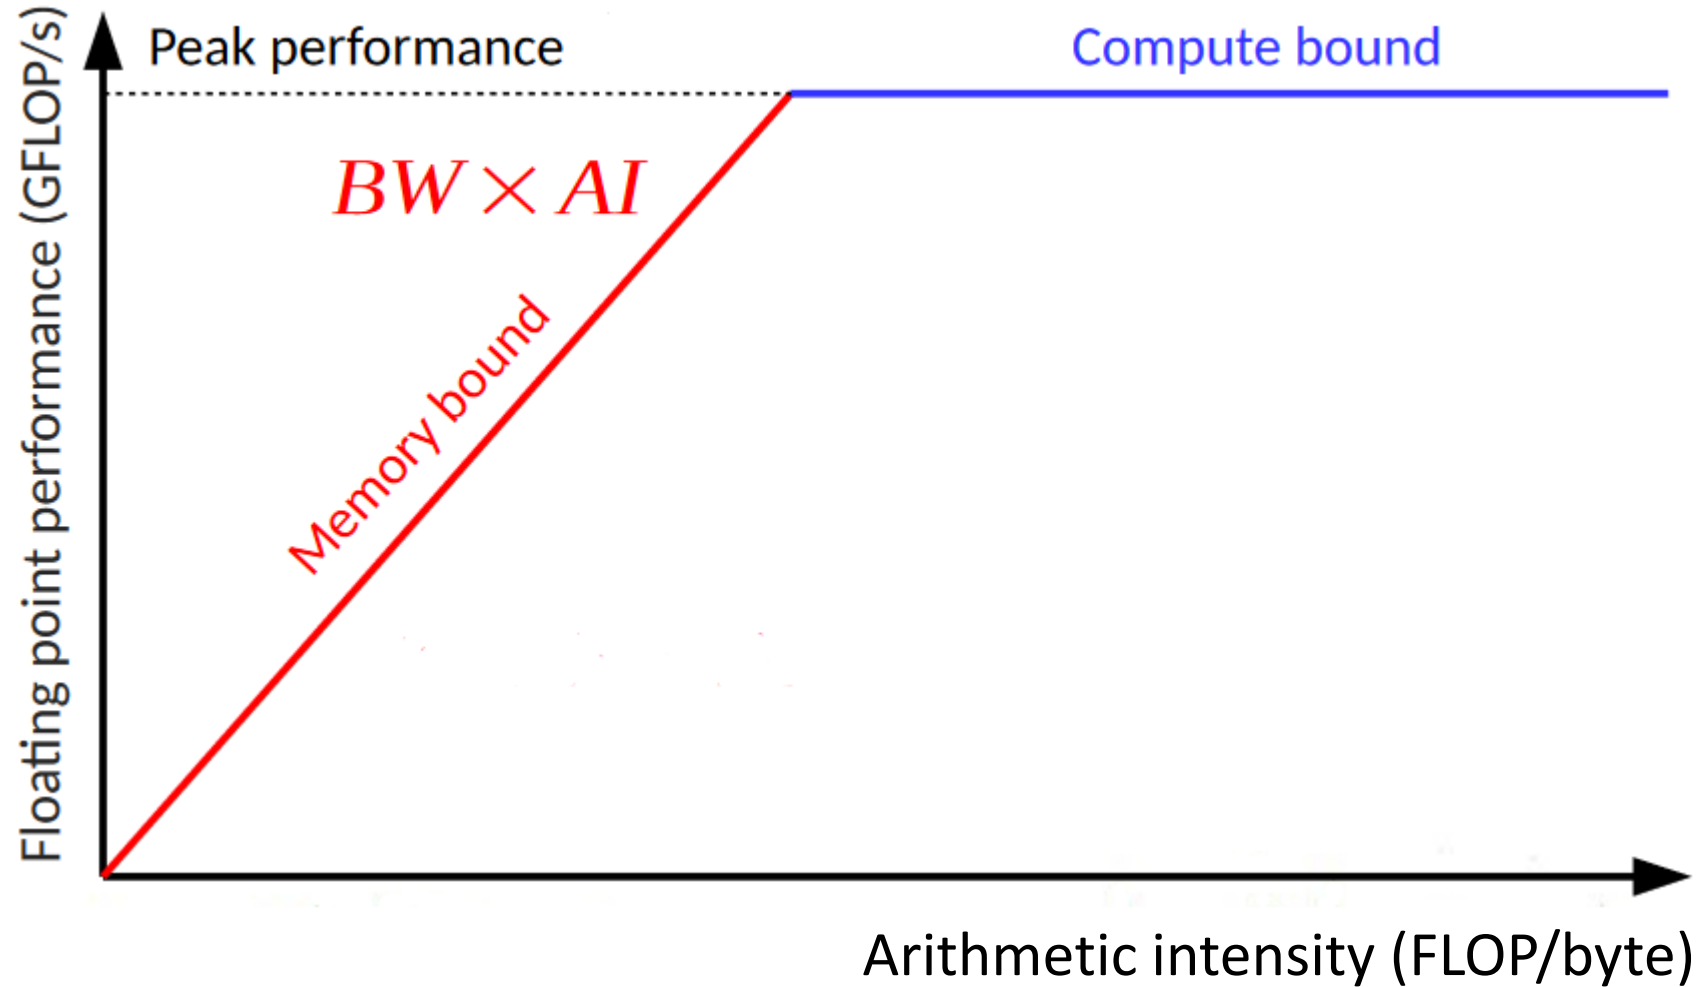
$$\mathbf{y} \leftarrow A \mathbf{x}, \quad \mathbf{z} \leftarrow a \mathbf{x} + b \mathbf{y}, \quad r \leftarrow \mathbf{x} \cdot \mathbf{y}$$

SpMV

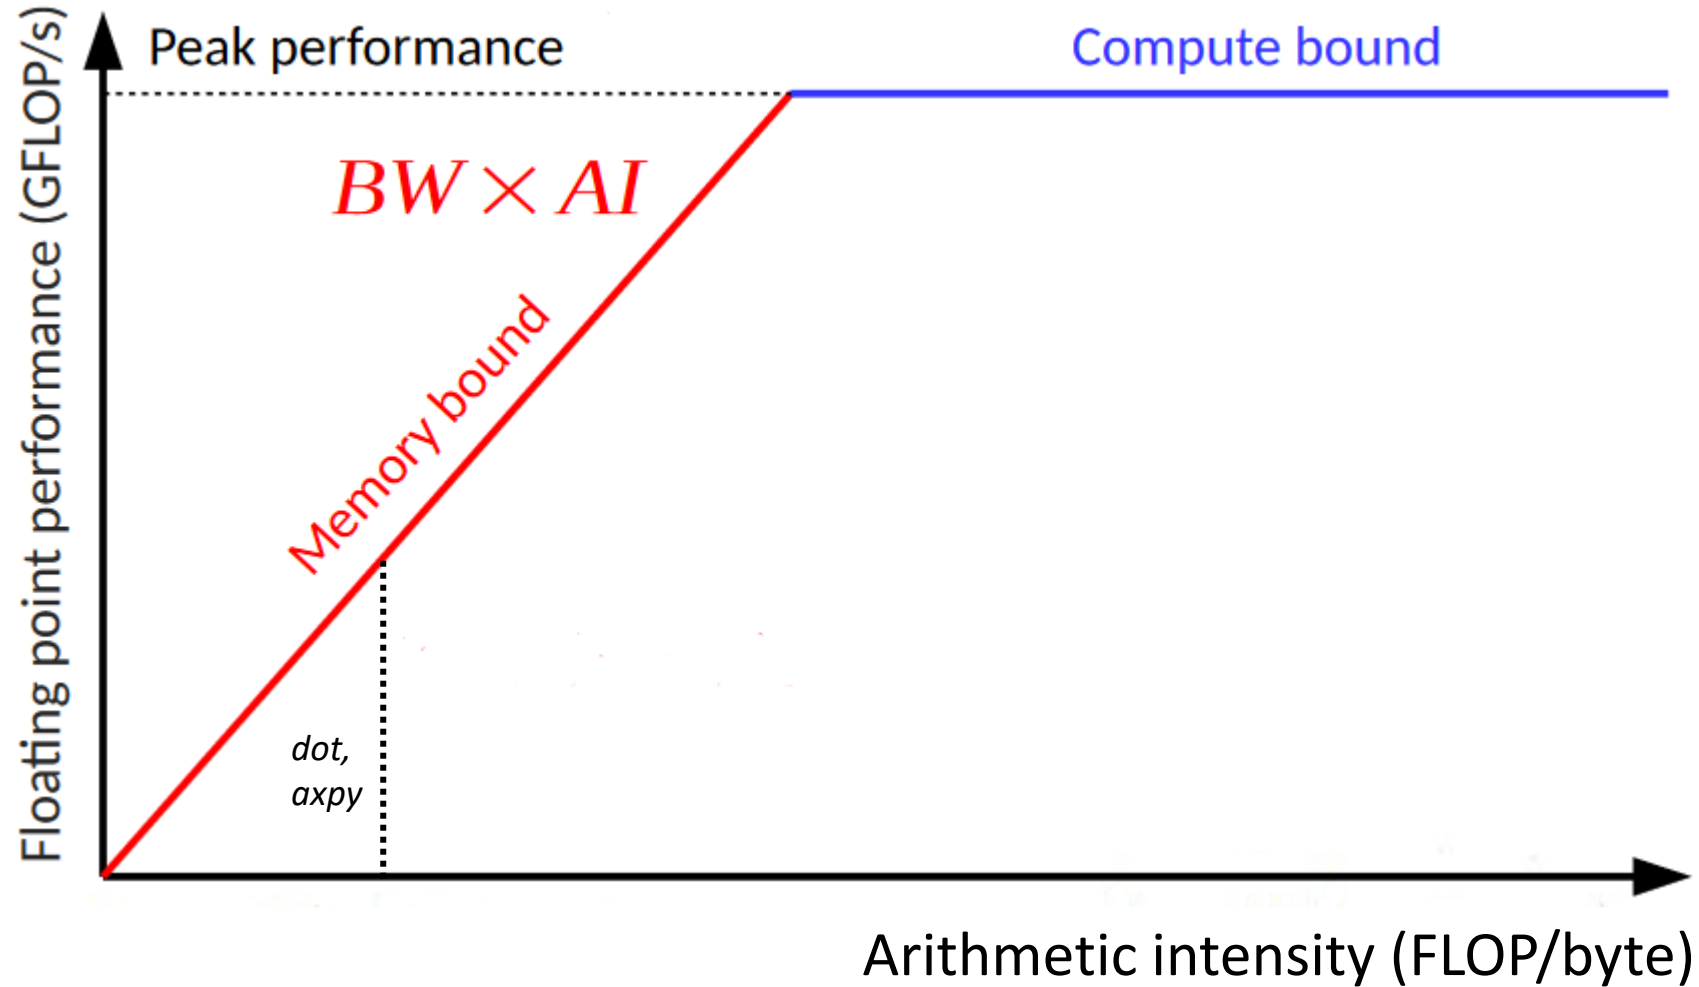
axpy

dot

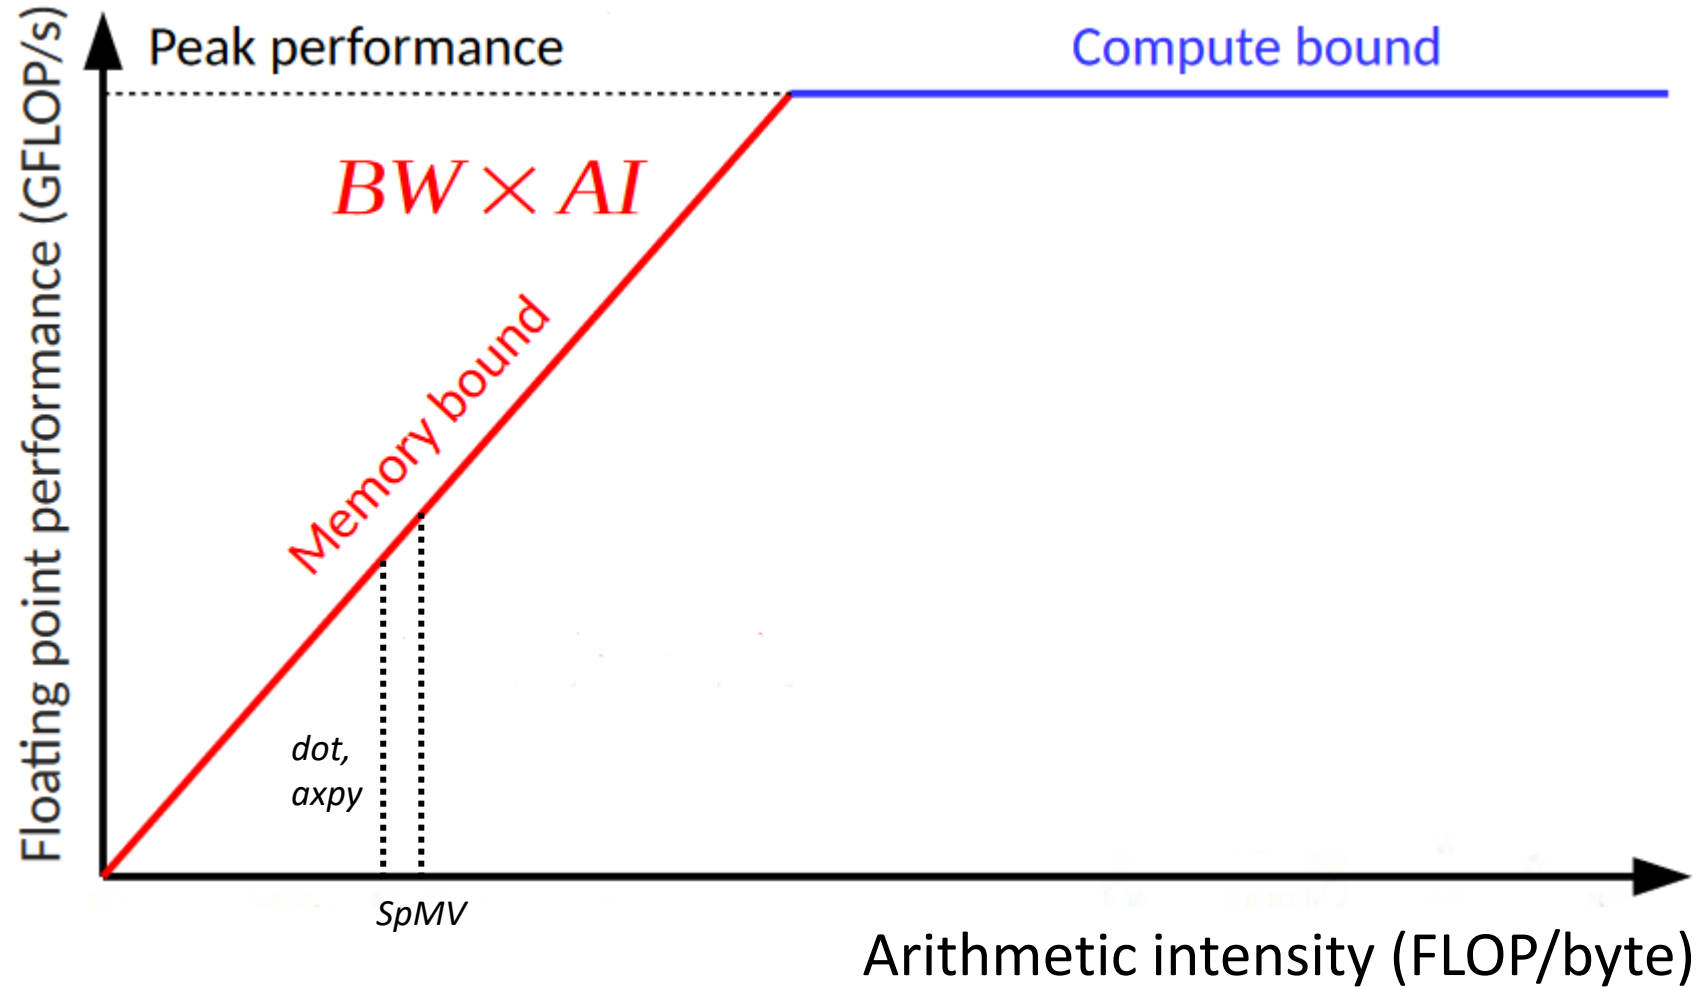
Memory boundedness




Memory boundedness



Memory boundedness



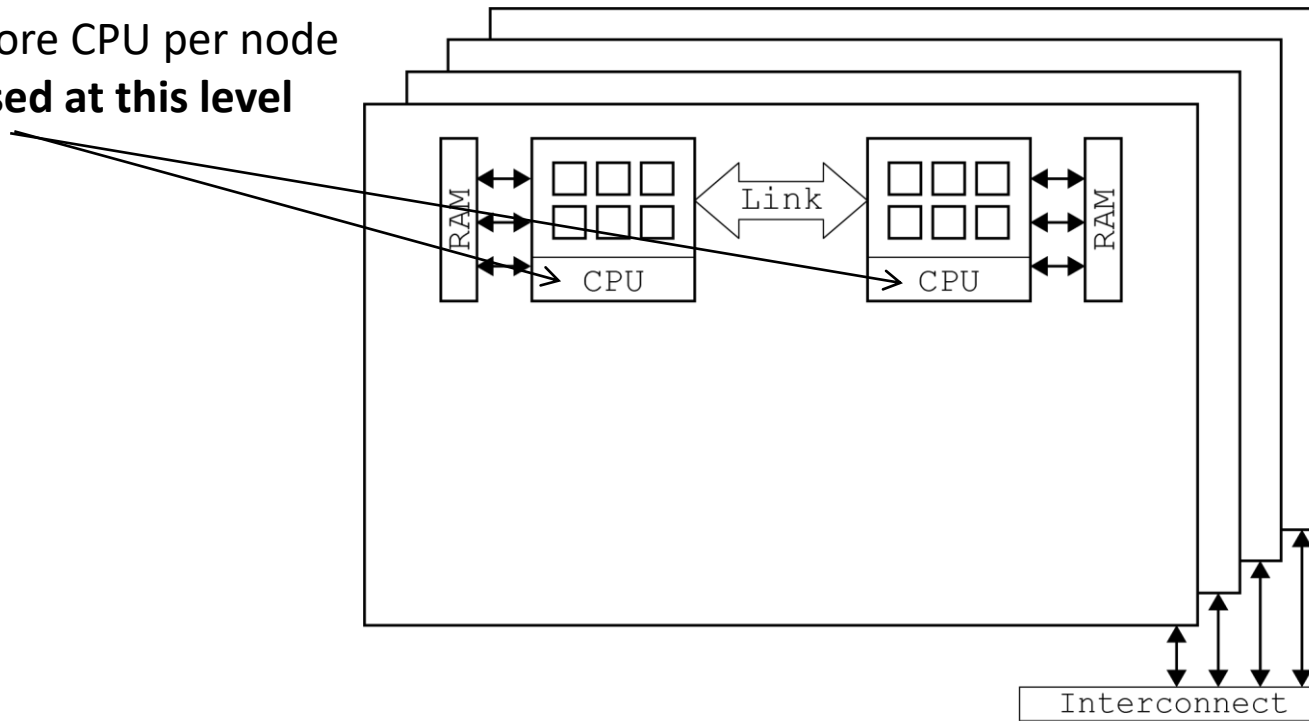
HPC²: hierarchy



multiple nodes interconnected via
high-bandwidth network
MPI is used at this level

HPC²: hierarchy

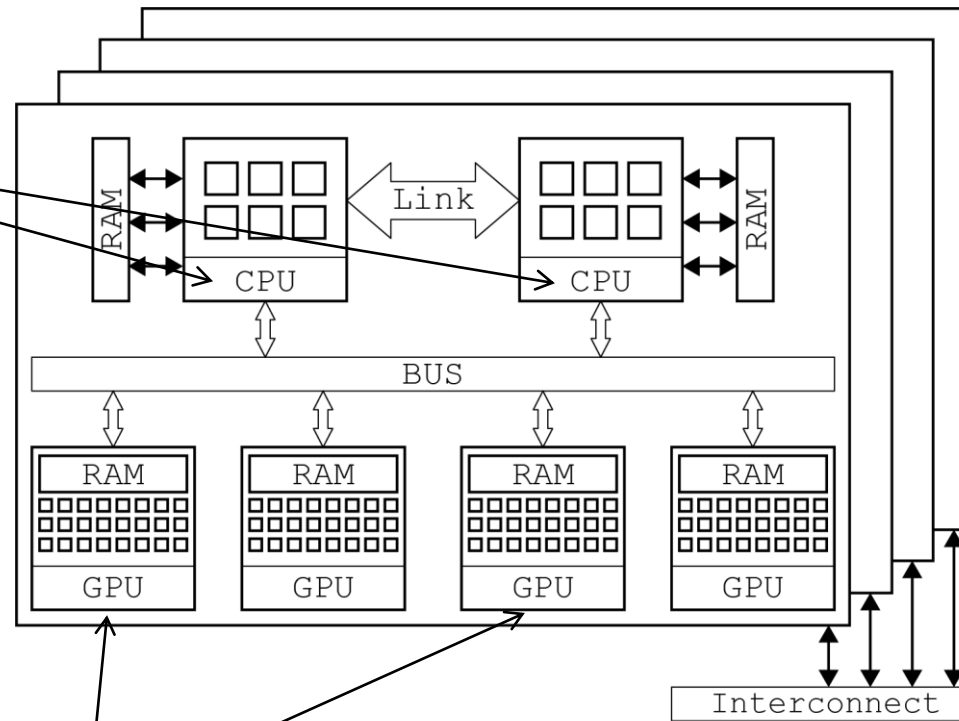
multiple multi-core CPU per node
OpenMP is used at this level



multiple nodes interconnected
via high-bandwidth network
MPI is used at this level

HPC²: hierarchy


multiple multi-core CPU per node
OpenMP is used at this level



multiple accelerators per node
OpenCL/CUDA is used at this level


multiple nodes interconnected via high-bandwidth network
MPI is used at this level

HPC²: hierarchy




HPC²: hierarchy

1. The MPI process




HPC²: hierarchy

1. The MPI process
2. The host and co-processors




HPC²: hierarchy

1. The MPI process
2. The host and co-processors
3. Multiple NUMA nodes in a manycore CPU



HPC²: tested architectures

MareNostrum 4



rank #42

3456 nodes with:

2× Intel Xeon 8160

1× Intel Omni-Path

Lomonosov-2



rank #156

1696 nodes with:

2× Intel Xeon E5-2697 v3

1× NVIDIA Tesla K40M

1× InfiniBand FDR

TSUBAME3.0



rank #31


540 nodes with:

2× Intel Xeon E5-2680 v4


4× NVIDIA Tesla P100

4× Intel Omni-Path

Exploiting symmetries




1 Symmetry




2 Symmetries

Symmetry-aware ordering

Exploiting symmetries




1 Symmetry




2 Symmetries

Symmetry-aware ordering

Exploiting symmetries




1 Symmetry




2 Symmetries

Symmetry-aware ordering

Exploiting symmetries




1 Symmetry




2 Symmetries

Symmetry-aware ordering

Exploiting symmetries




1 Symmetry




2 Symmetries

Symmetry-aware ordering

Exploiting symmetries



1 Symmetry



2 Symmetries


Symmetry-aware ordering

Block-diagonalising L

$$L = \begin{array}{c|c} L_{\text{inn}} & L_{\text{out}} \\ \hline L_{\text{out}} & L_{\text{inn}} \end{array} \in R^{N \times N}$$

$$S = \frac{1}{\sqrt{2}} \begin{array}{c|c} I_{N/2} & I_{N/2} \\ \hline I_{N/2} & -I_{N/2} \end{array} \in R^{N \times N}$$

$$\hat{L} = SLS^{-1} = \begin{array}{c|c} L_{\text{inn}} + L_{\text{out}} & 0 \\ \hline 0 & L_{\text{inn}} - L_{\text{out}} \end{array}$$

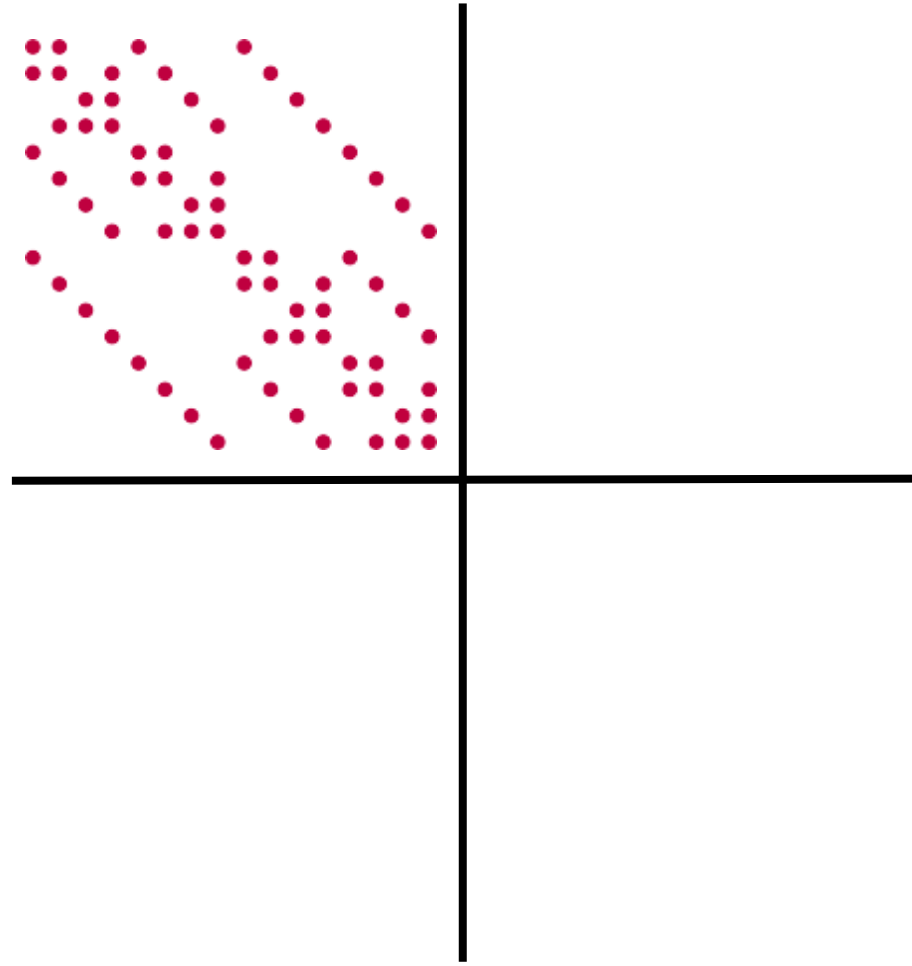


Block-diagonalising L

$$L = \begin{array}{c|c} L_{\text{inn}} & L_{\text{out}} \\ \hline L_{\text{out}} & L_{\text{inn}} \end{array} \in R^{N \times N}$$

$$S = \sqrt{\frac{1}{2}} \begin{array}{c|c} I_{N/2} & I_{N/2} \\ \hline I_{N/2} & -I_{N/2} \end{array} \in R^{N \times N}$$

$$\hat{L} = SLS^{-1} = \begin{array}{c|c} L_{\text{inn}} + L_{\text{out}} & 0 \\ \hline 0 & L_{\text{inn}} - L_{\text{out}} \end{array}$$




Block-diagonalising L

$$L = \begin{array}{c|c} L_{\text{inn}} & L_{\text{out}} \\ \hline L_{\text{out}} & L_{\text{inn}} \end{array} \in R^{N \times N}$$

$$S = \sqrt{\frac{1}{2}} \begin{array}{c|c} I_{N/2} & I_{N/2} \\ \hline I_{N/2} & -I_{N/2} \end{array} \in R^{N \times N}$$

$$\hat{L} = SLS^{-1} = \begin{array}{c|c} L_{\text{inn}} + L_{\text{out}} & 0 \\ \hline 0 & L_{\text{inn}} - L_{\text{out}} \end{array}$$




Block-diagonalising L

$$L = \begin{array}{c|c} L_{\text{inn}} & L_{\text{out}} \\ \hline L_{\text{out}} & L_{\text{inn}} \end{array} \in R^{N \times N}$$

$$S = \frac{1}{\sqrt{2}} \begin{array}{c|c} I_{N/2} & I_{N/2} \\ \hline I_{N/2} & -I_{N/2} \end{array} \in R^{N \times N}$$

$$\hat{L} = SLS^{-1} = \begin{array}{c|c} L_{\text{inn}} + L_{\text{out}} & 0 \\ \hline 0 & L_{\text{inn}} - L_{\text{out}} \end{array}$$

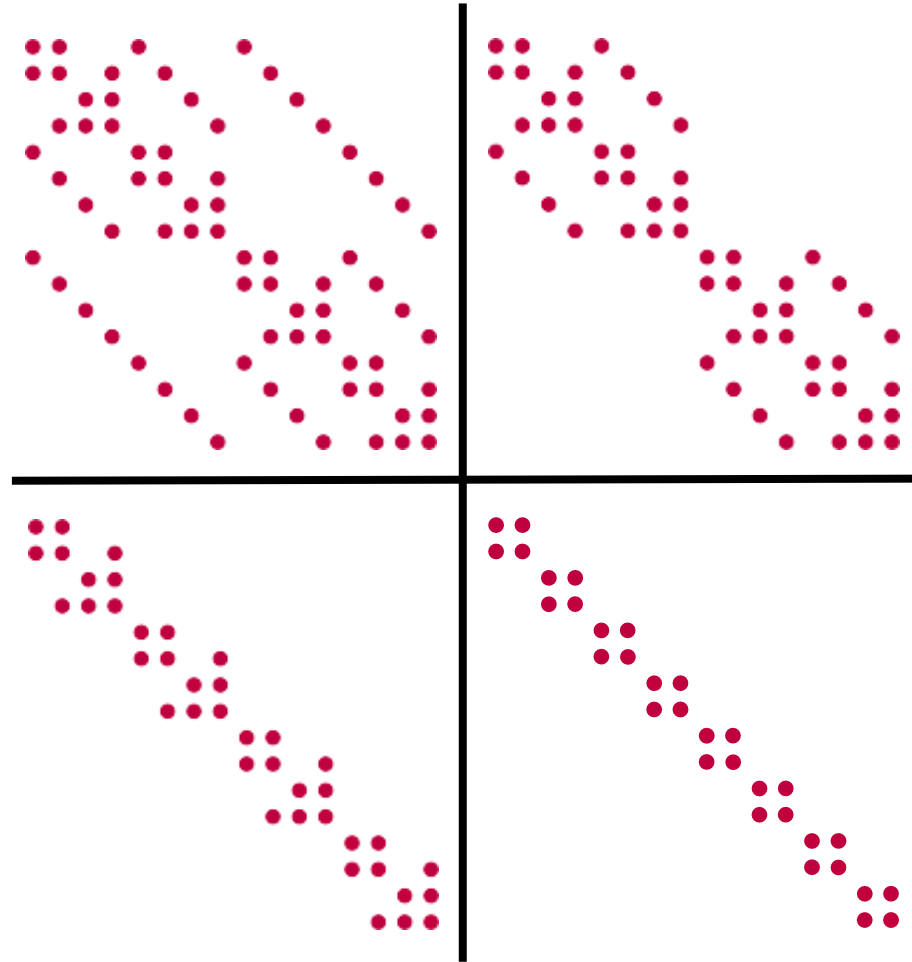


Block-diagonalising L

$$L = \begin{array}{c|c} L_{\text{inn}} & L_{\text{out}} \\ \hline L_{\text{out}} & L_{\text{inn}} \end{array} \in R^{N \times N}$$

$$S = \sqrt{\frac{1}{2}} \begin{array}{c|c} I_{N/2} & I_{N/2} \\ \hline I_{N/2} & -I_{N/2} \end{array} \in R^{N \times N}$$

$$\hat{L} = SLS^{-1} = \begin{array}{c|c} L_{\text{inn}} + L_{\text{out}} & 0 \\ \hline 0 & L_{\text{inn}} - L_{\text{out}} \end{array}$$



Sparse matrix-matrix product

Rewriting the SpMV product:

$$\hat{L}\mathbf{v} = \underbrace{\begin{pmatrix} L_{inn}^0 & & 0 \\ & \ddots & \\ 0 & & L_{inn}^p \end{pmatrix}}_{L_{inn}(\mathbf{v}^0 \mid \dots \mid \mathbf{v}^p)} \begin{pmatrix} \mathbf{v}^0 \\ \vdots \\ \mathbf{v}^p \end{pmatrix} + L_{out}\mathbf{v}$$

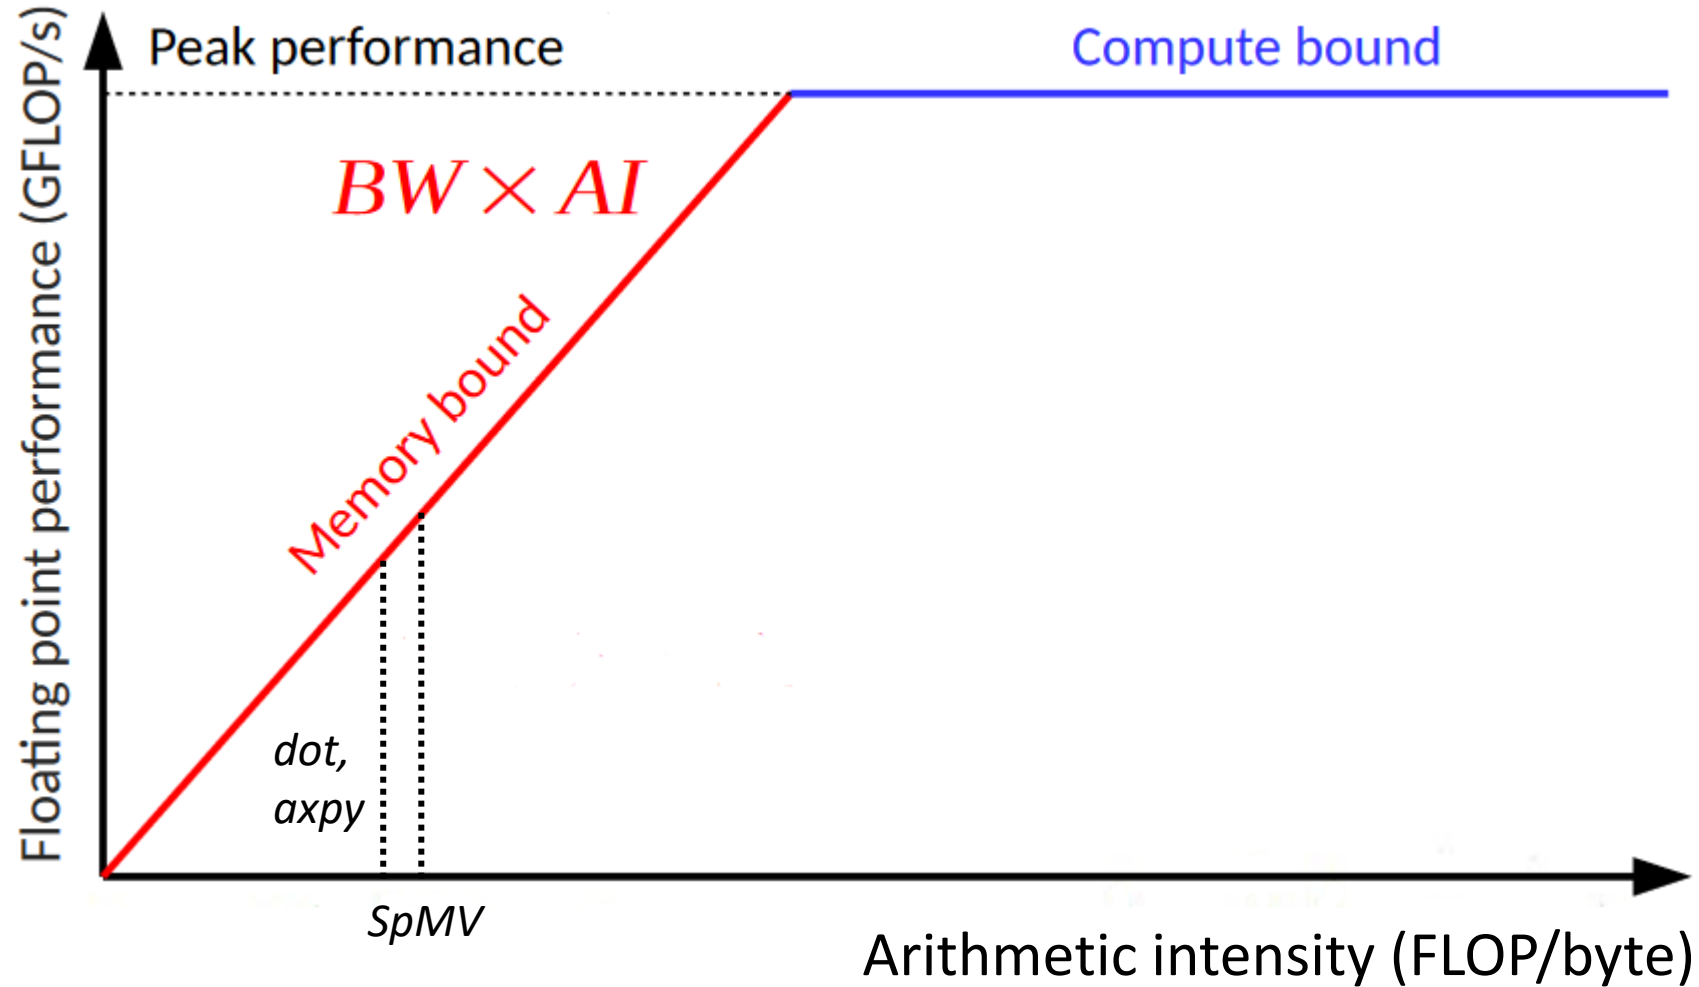
into an SpMM product

A reduction in time complexity

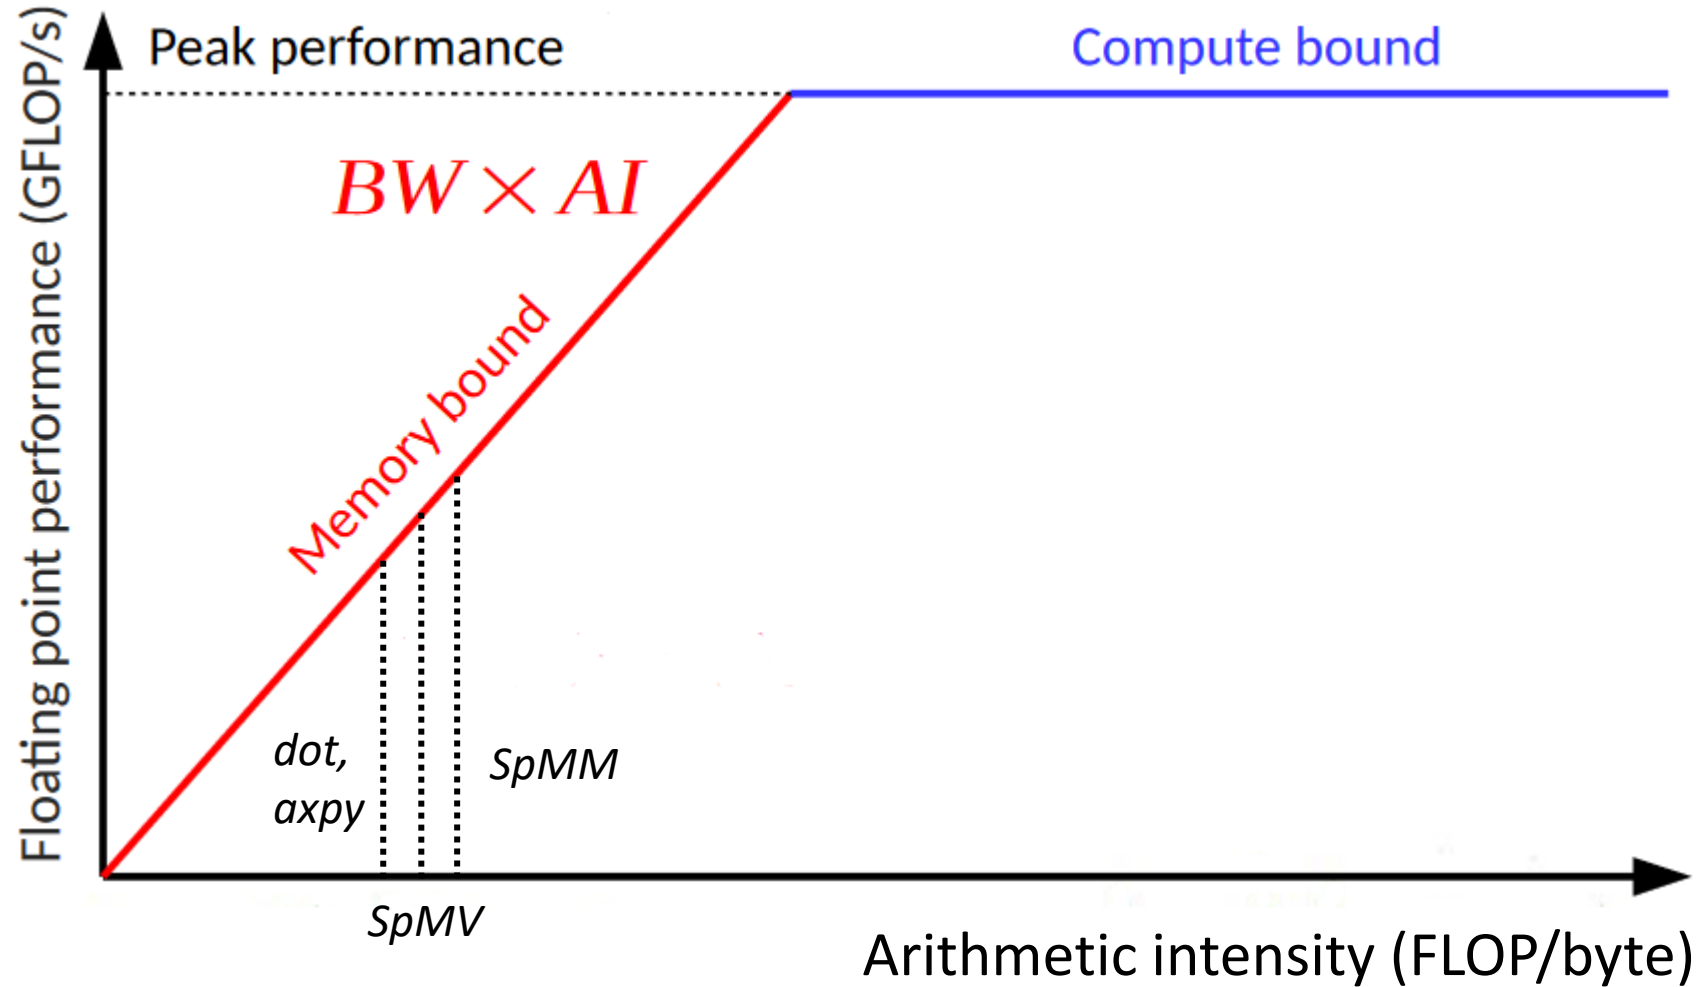
A reduction in memory footprint

An increase in arithmetic intensity

Increasing arithmetic intensity



Increasing arithmetic intensity



Summarising

Symmetry preserving methods in MHD

HPC² framework

Exploiting symmetries in geometry

Thank you for attending!

