L. GIANINAZZI, A. ZIOGAS, P. LUCZYNSKI, L. HUANG, S. ASHKBOOS, F. SCHEIDL, A. CARIGIET, C. GE, N. ABUBAKER, M. BESTA, T. BEN-NUN, T. HOEFLER

Arrow Matrix Decomposition:

A Novel Approach for Communication-Efficient Sparse Matrix Multiplication

Iterated SpMM

dense `tall`

$$X^{t+1} = f(A\, X^{t})$$

sparse

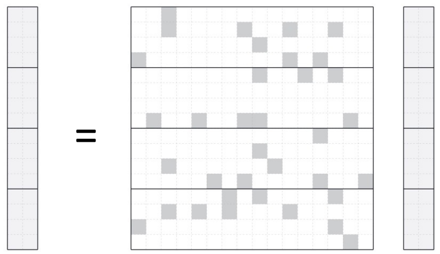

Systems of Equations	Graph Neural Networks	Eigenvectors
PDEs	Physics Simulations Network Analysis	PageRank

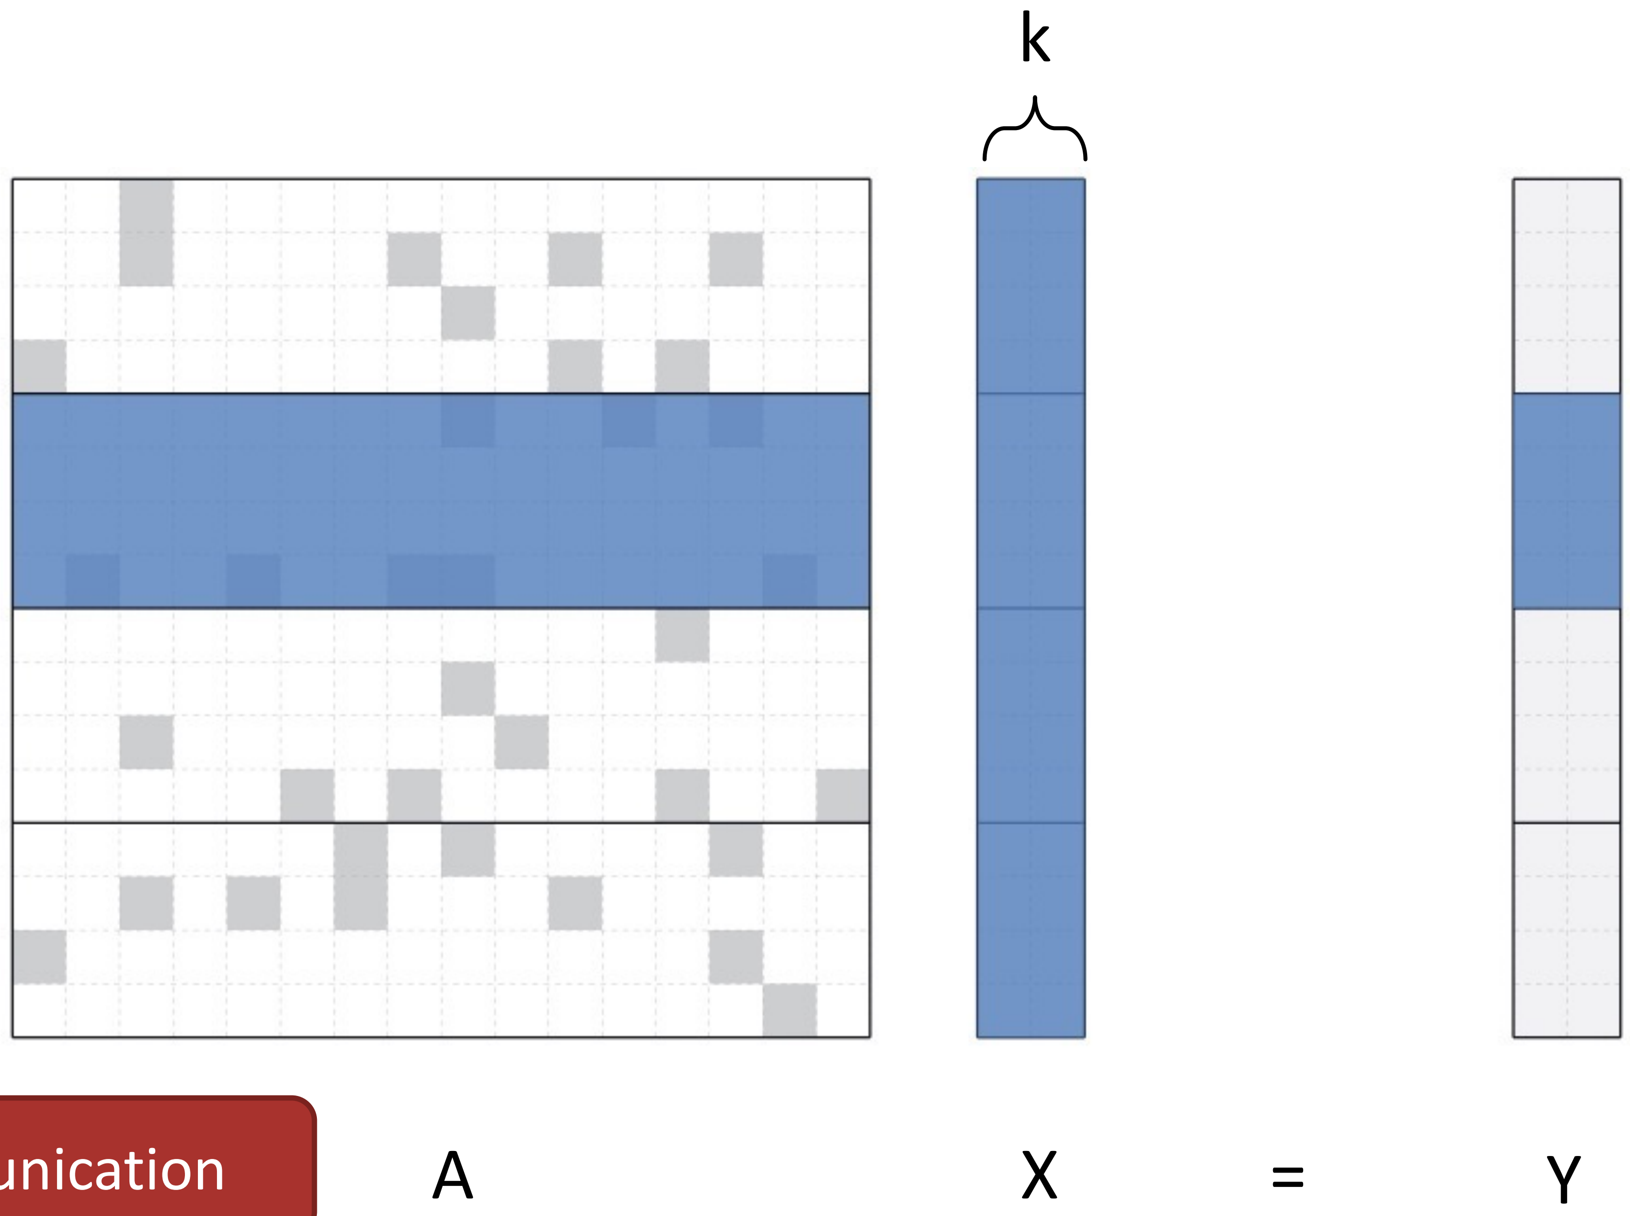

n rows

$\theta(nk)$ communication

Background – 1.5D

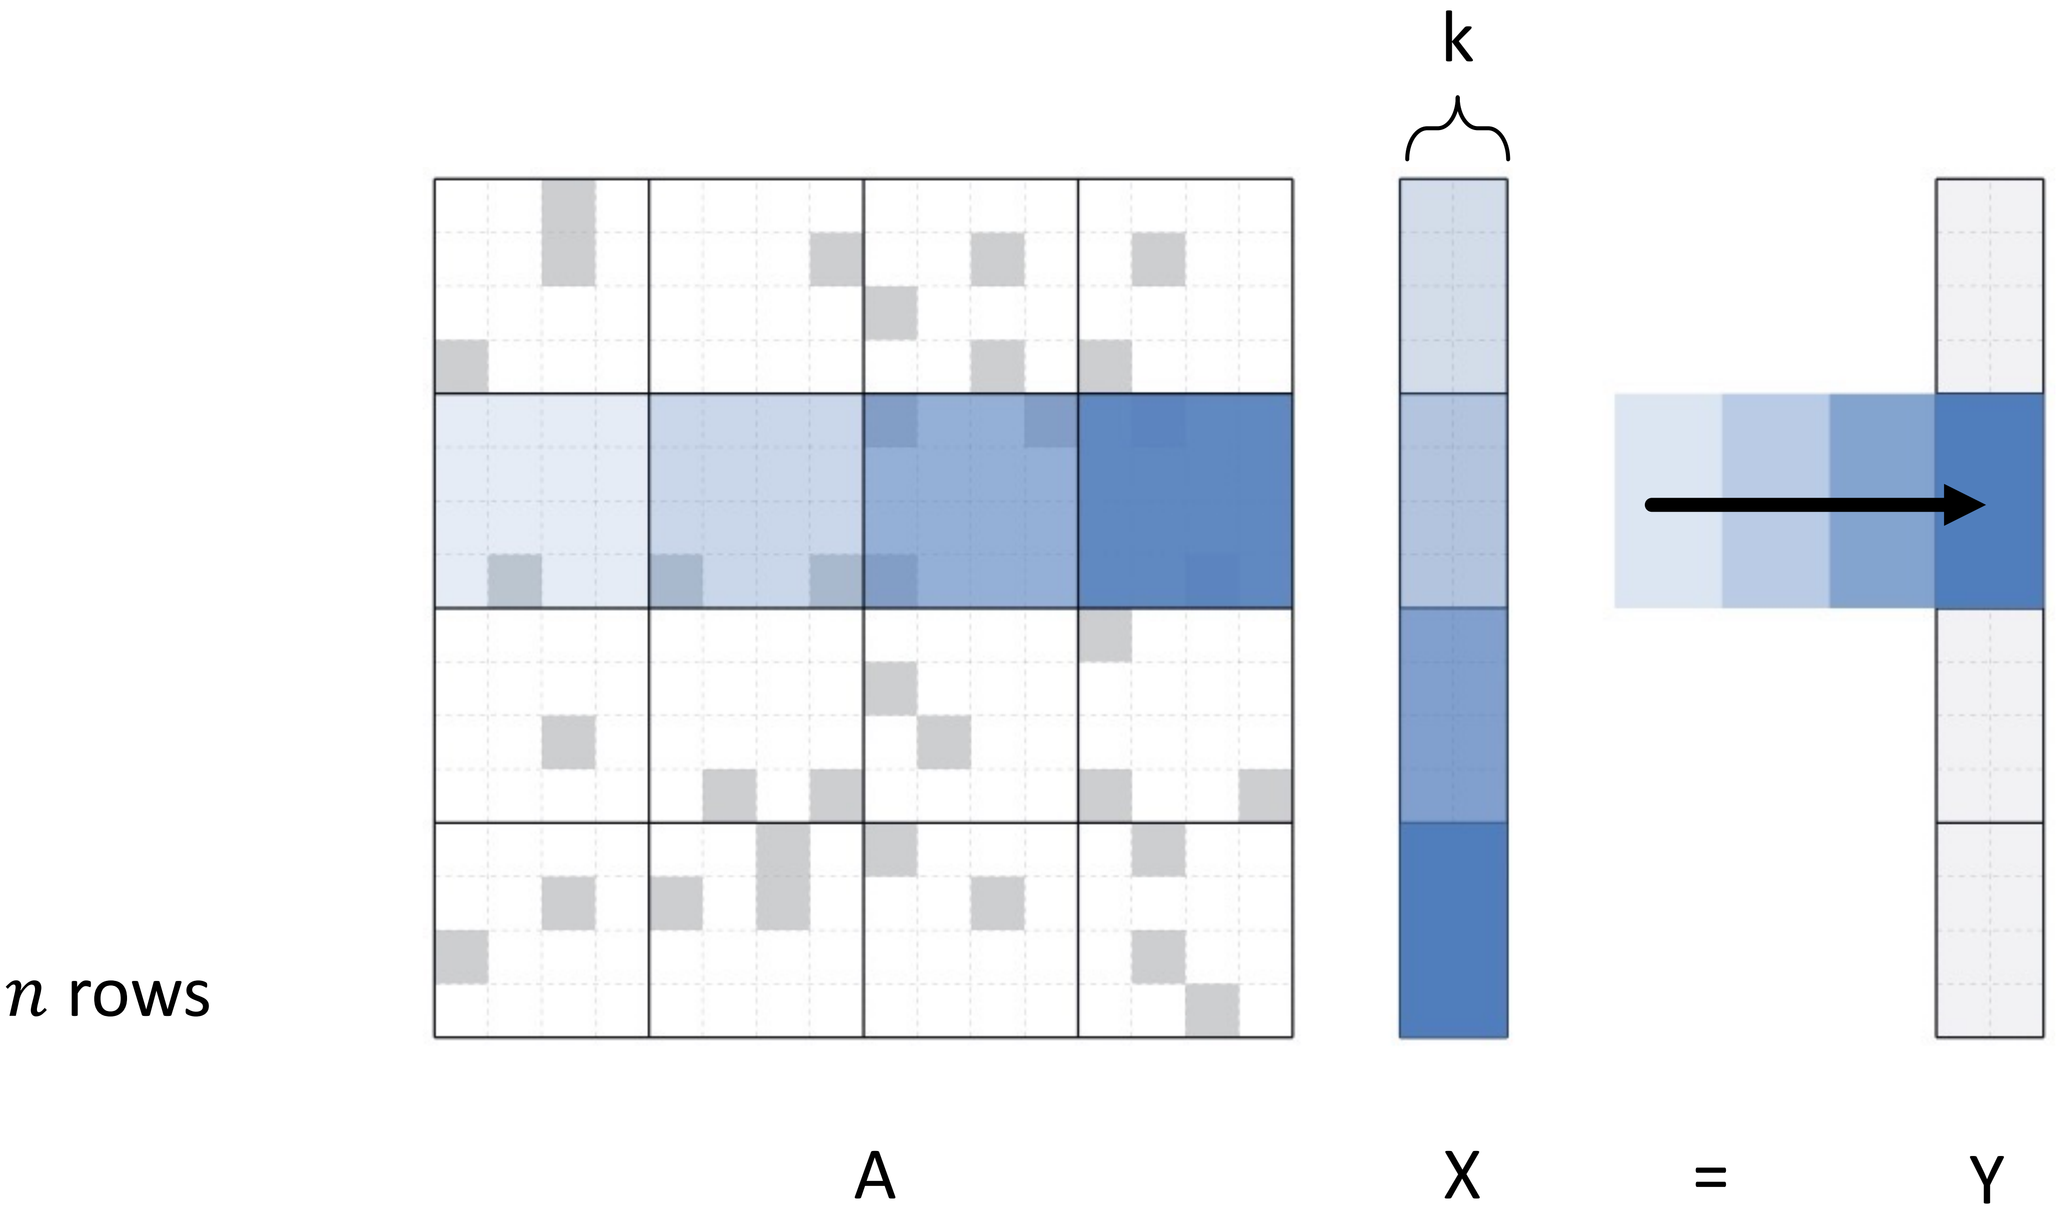

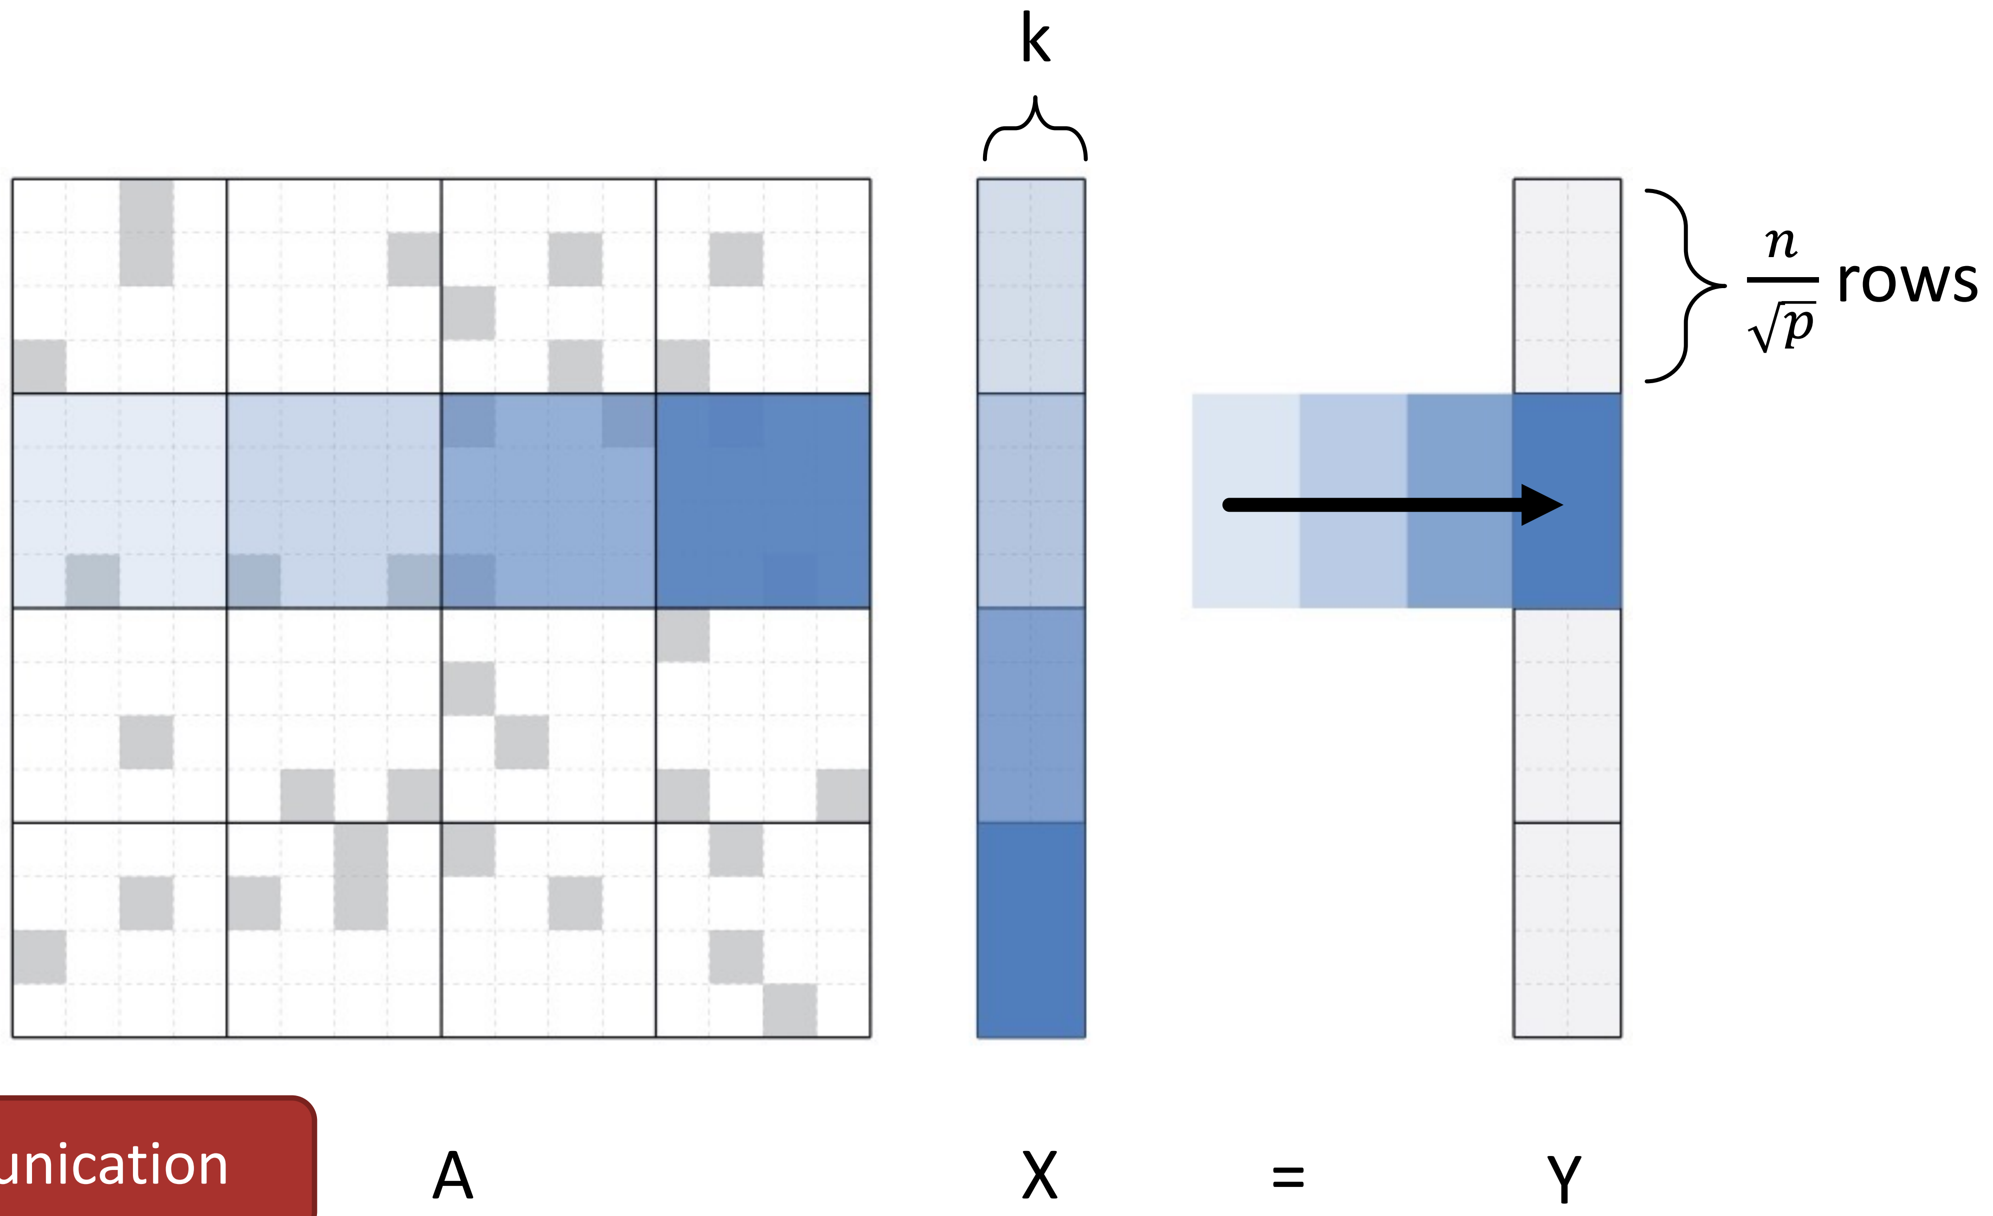

p processors

n rows

$\theta\left(\frac{nk}{\sqrt{p}}\right)$ communication

Background – Bandwidth

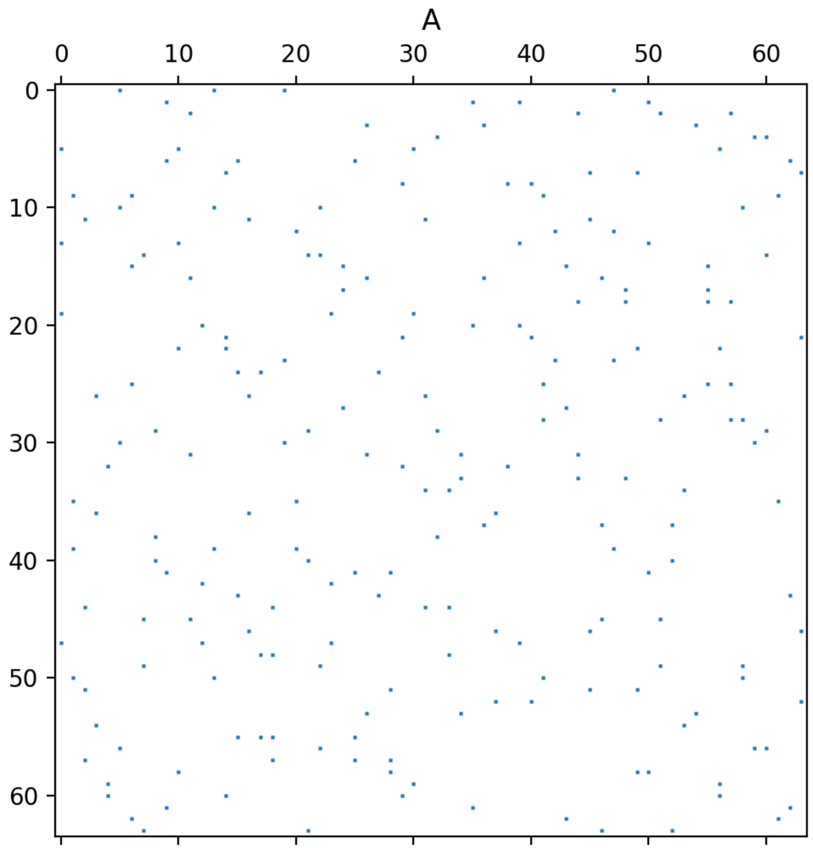

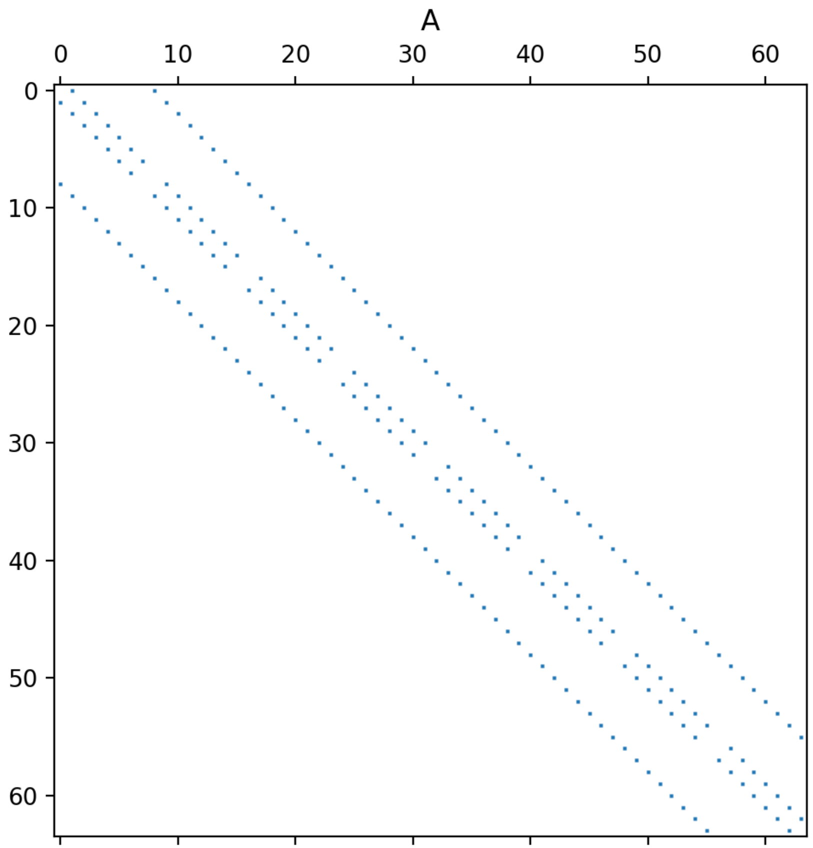

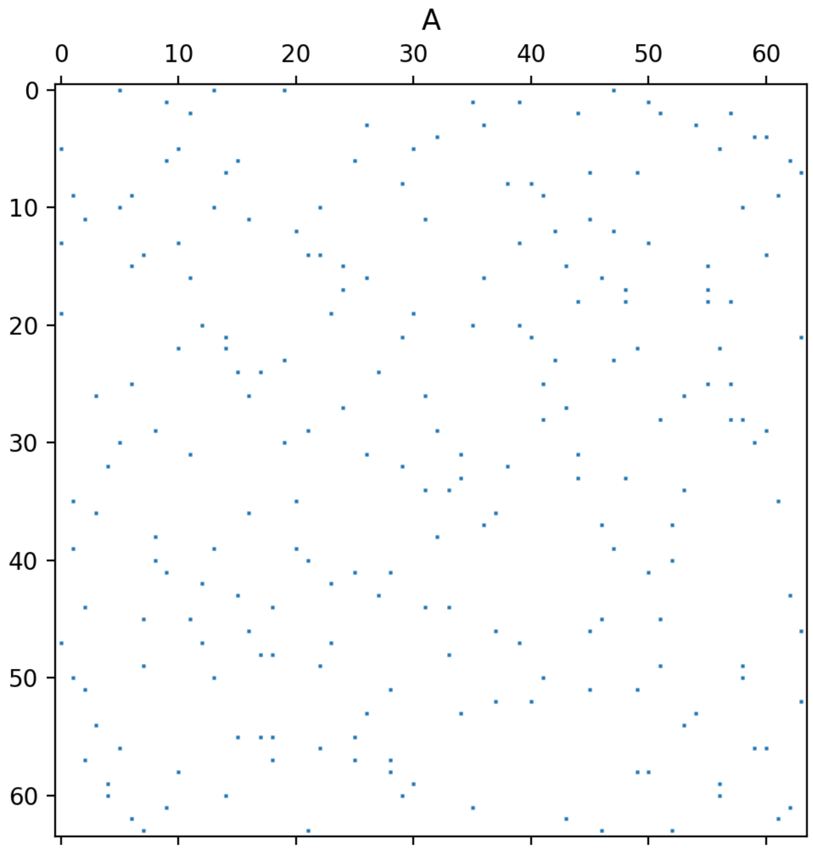

= P

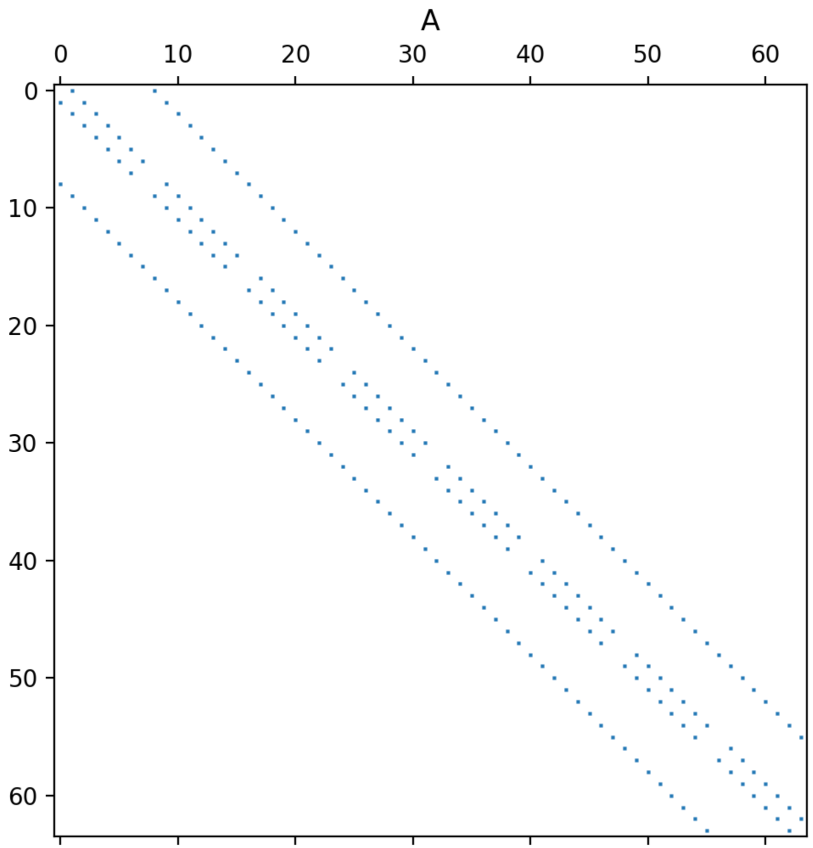

P^T

maximum distance from diagonal:
bandwidth **b**

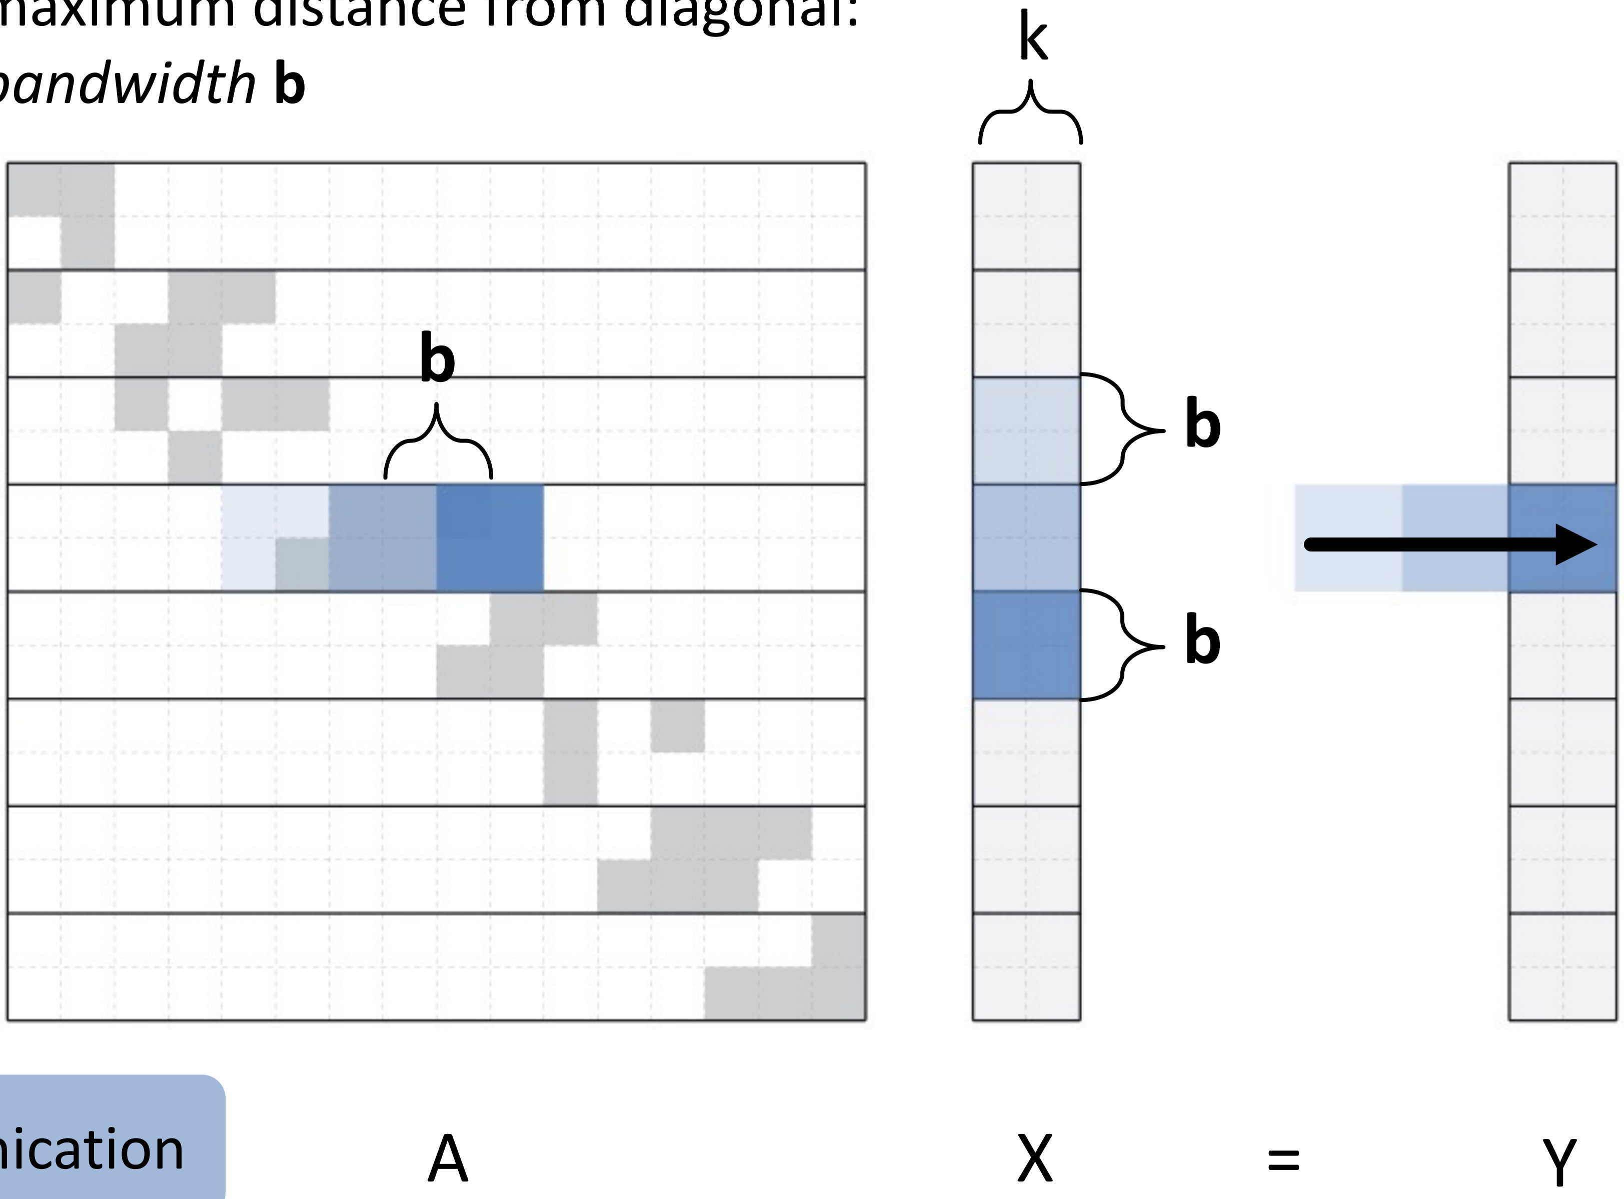

p processors
n rows

O(bk) communication

Challenge 1: *Low Diameter*

n vertices and diameter *D*

→ bandwidth $\geq \frac{n-1}{D}$

Challenge 2: *High Degree*

Maximum degree *d*

→ bandwidth $\geq \frac{d}{2}$

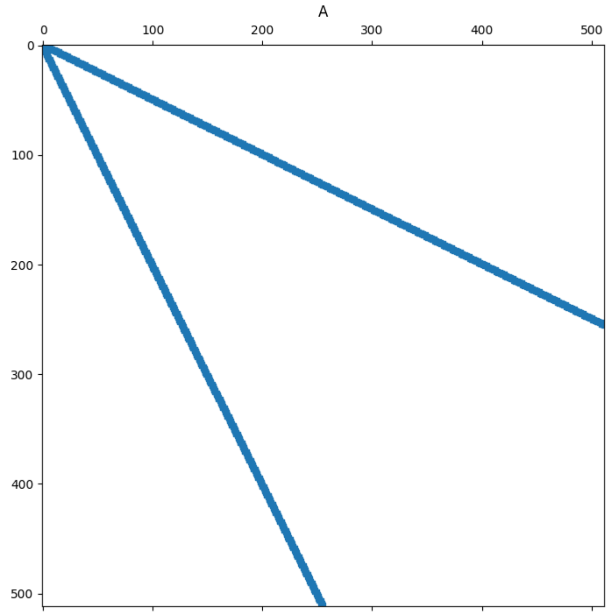

Complete Binary Tree

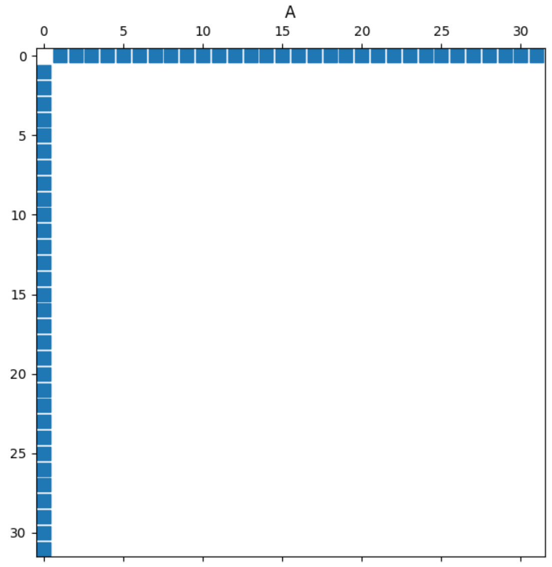

Star

$$AX = \sum_i P_i B_i (P_i^\intercal X)$$

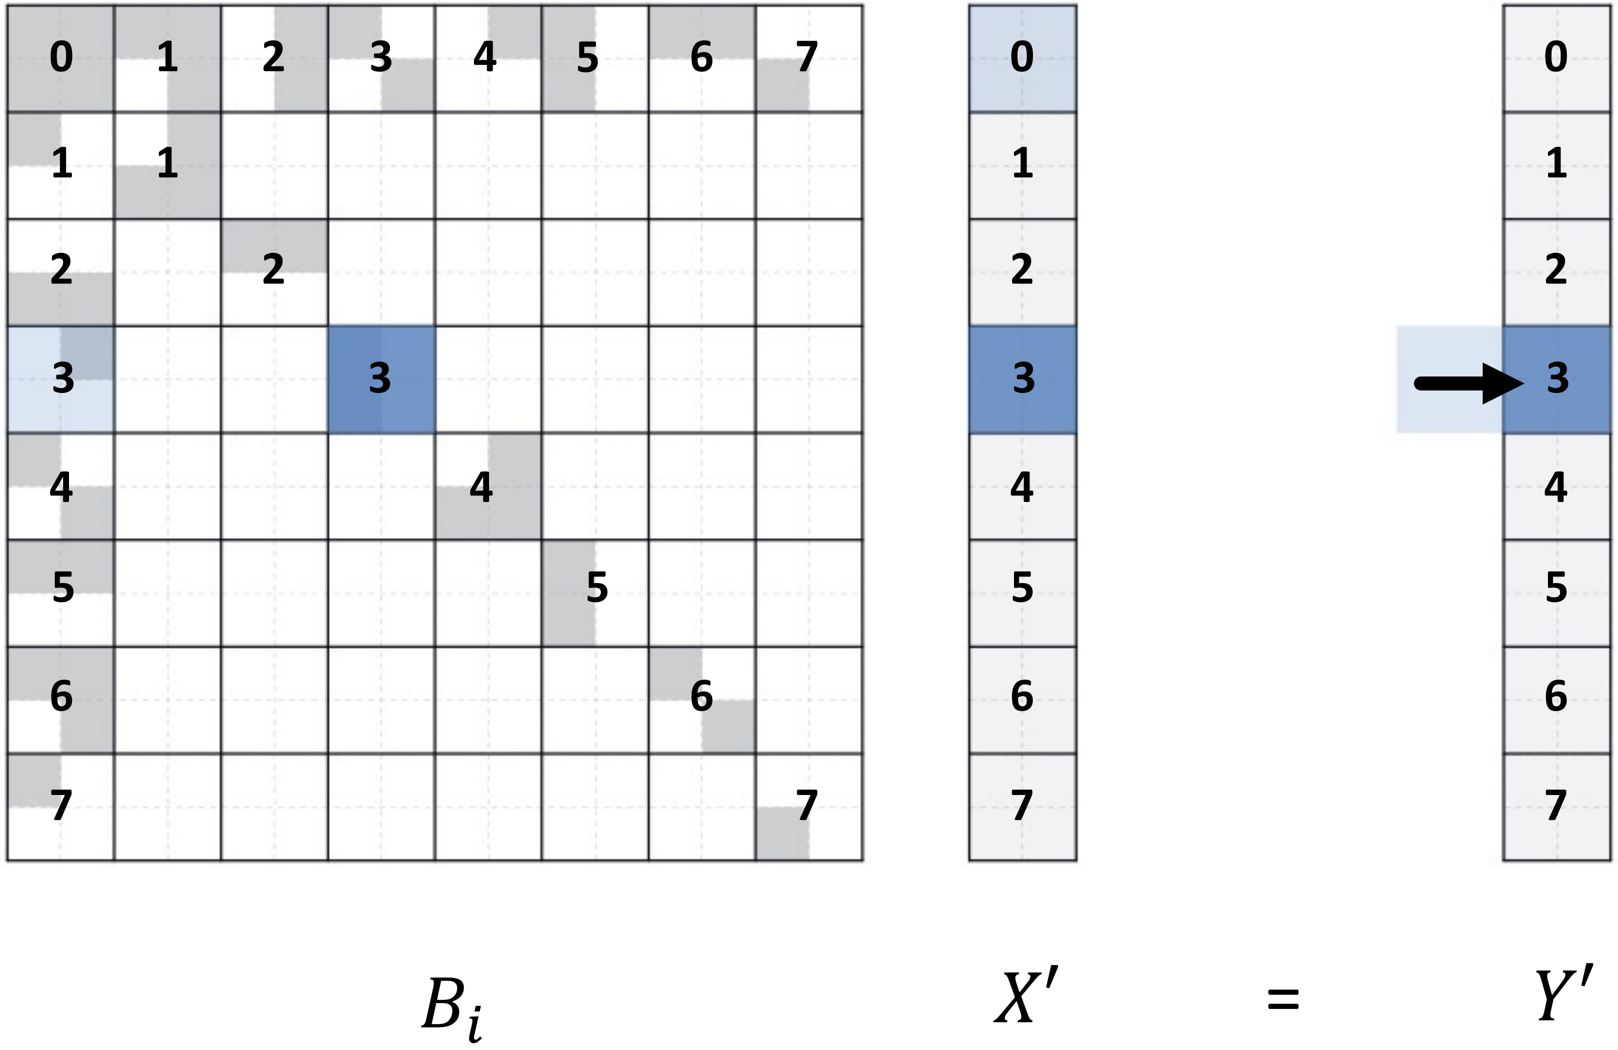

Challenge 1: *Low Diameter*

Idea: Multiple Matrices

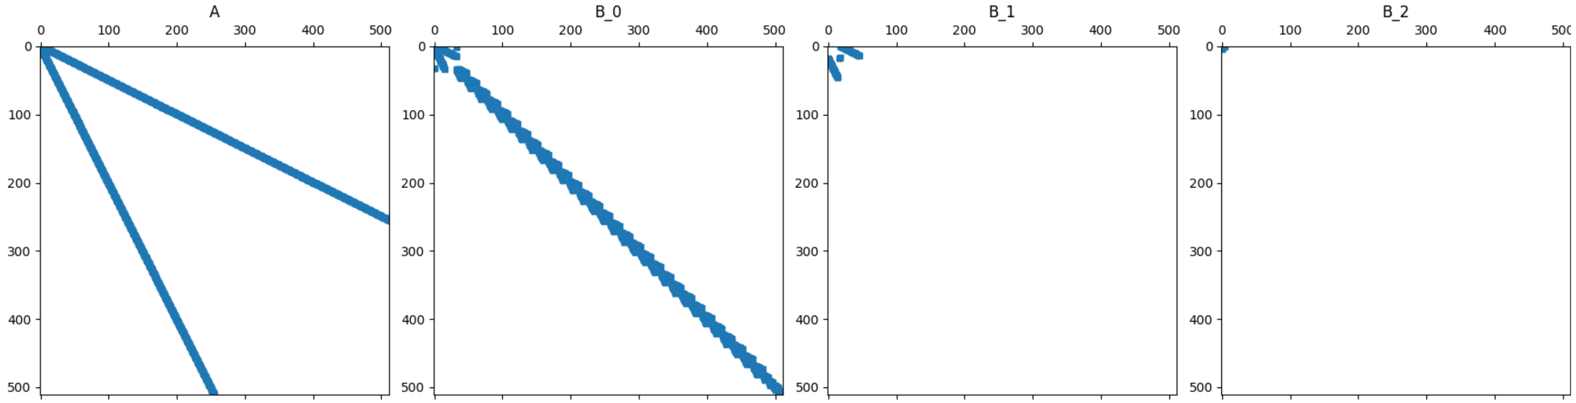

$$A = \sum_i P_i B_i P_i^\intercal$$

Challenge 2: *High Degree*

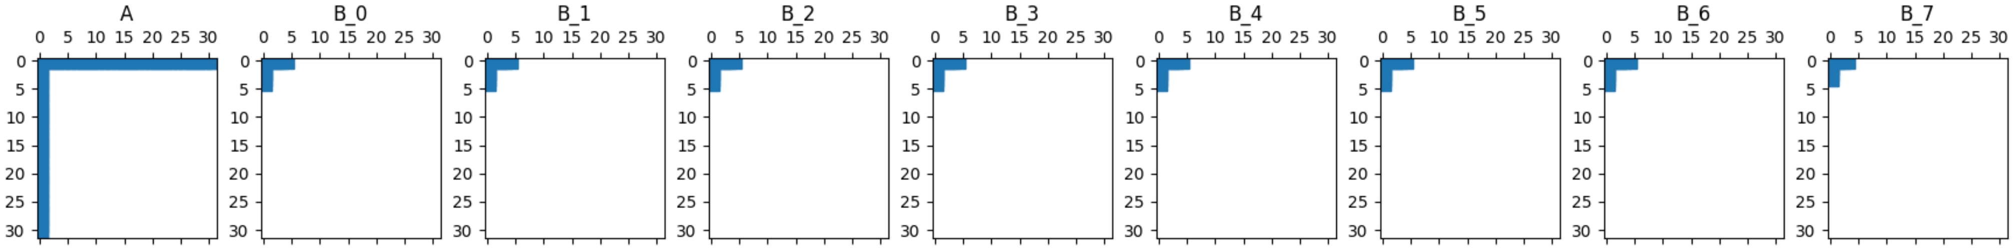

$$A = \sum_i P_i B_i P_i^\intercal$$

Challenge 2: *High Degree*

Idea: Arrow Shape

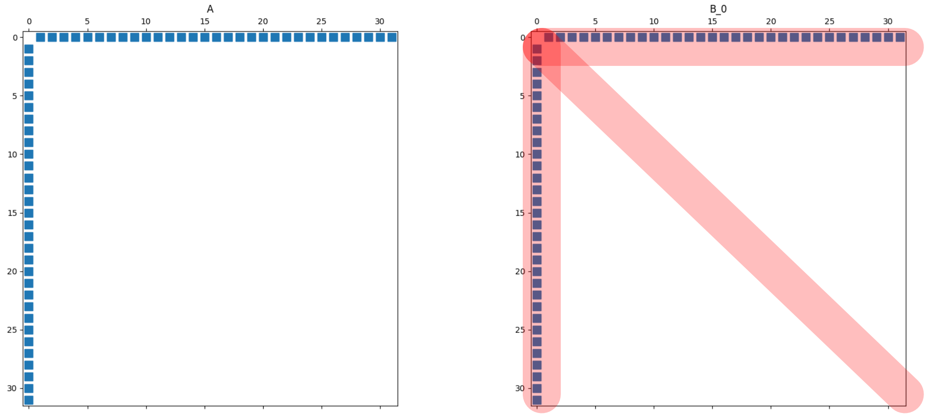

Provably addresses the **degree** limitation

$$A = \sum_i P_i B_i P_i^\intercal$$

Challenge 1: *Low Diameter*

Idea: Multiple Matrices

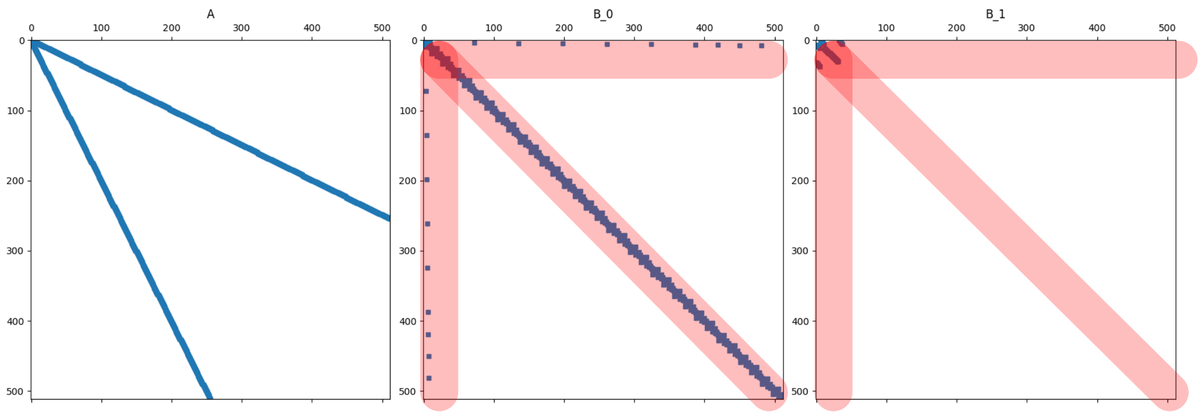

Challenge 2: *High Degree*

Idea: Arrow Shape

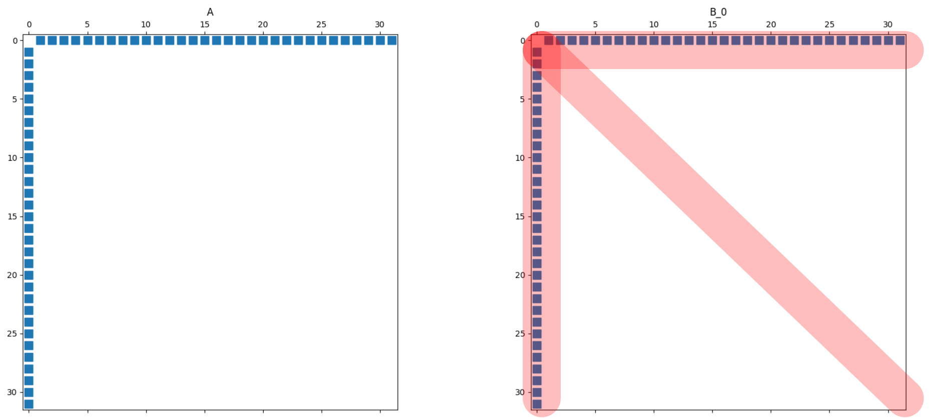

$$A = \sum_i P_i B_i P_i^\intercal$$

Provably addresses the **degree** limitation

Challenge 1: *Low Diameter*

Idea: Multiple Matrices

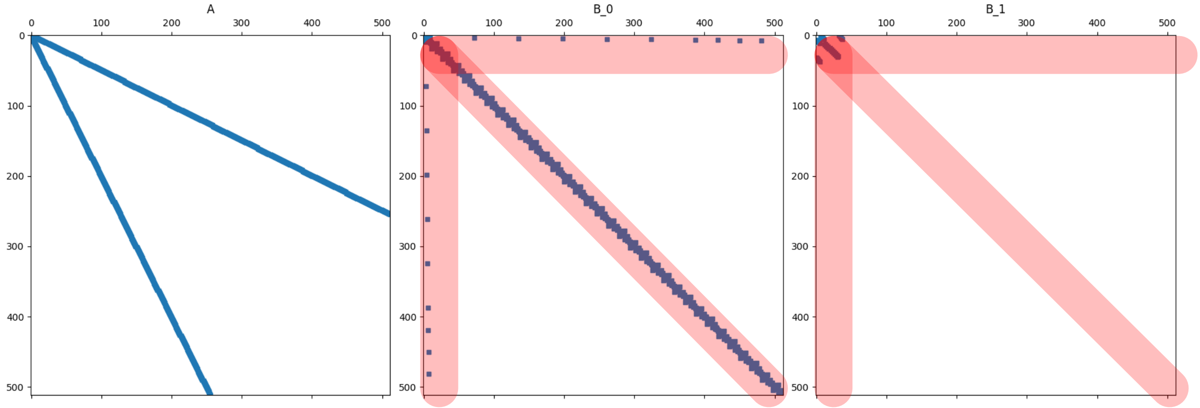

Challenge 2: *High Degree*

Idea: Arrow Shape

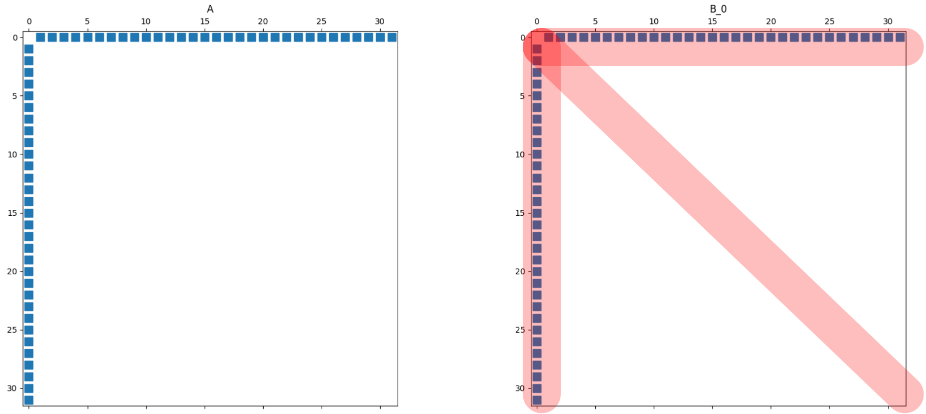

1. Decompose

$$AX = \sum_i P_i B_i (P_i^\intercal X)$$

2. Arrow Multiply

3. Aggregate

Provably addresses the **degree** limitation

$$AX = \sum_i P_i B_i (P_i^\intercal X)$$

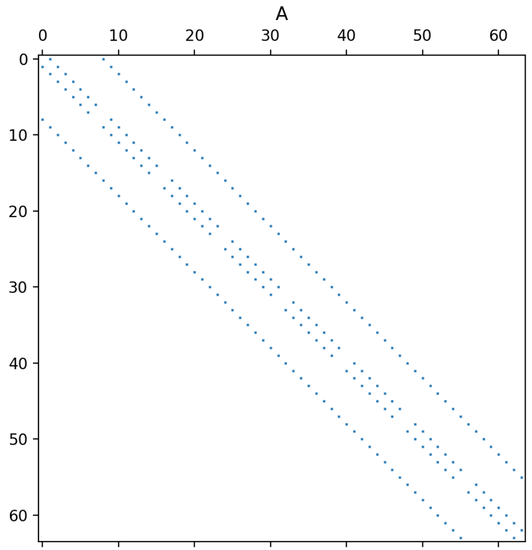

Matrix

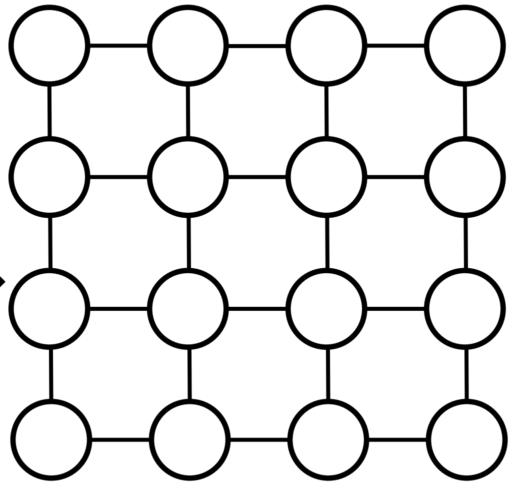

Graph

The Minimum Linear Arrangement

Linear arrangement of the vertices that ***minimizes*** *the total* ***distance*** *between any neighbor*

NP-Hard

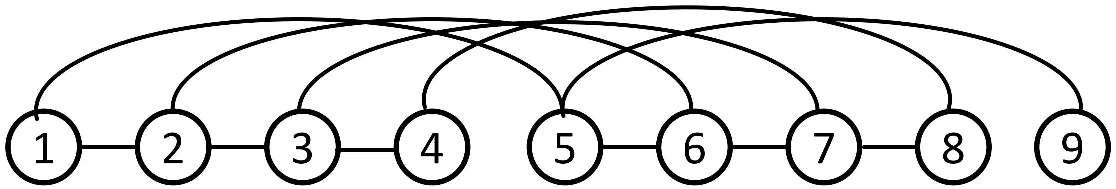

Linear Arrangement

We prove *connection between linear arrangement* and size of decomposition

We propose *new linear-time MLA* heuristic

$$AX = \sum_i P_i B_i (P_i^\intercal X)$$

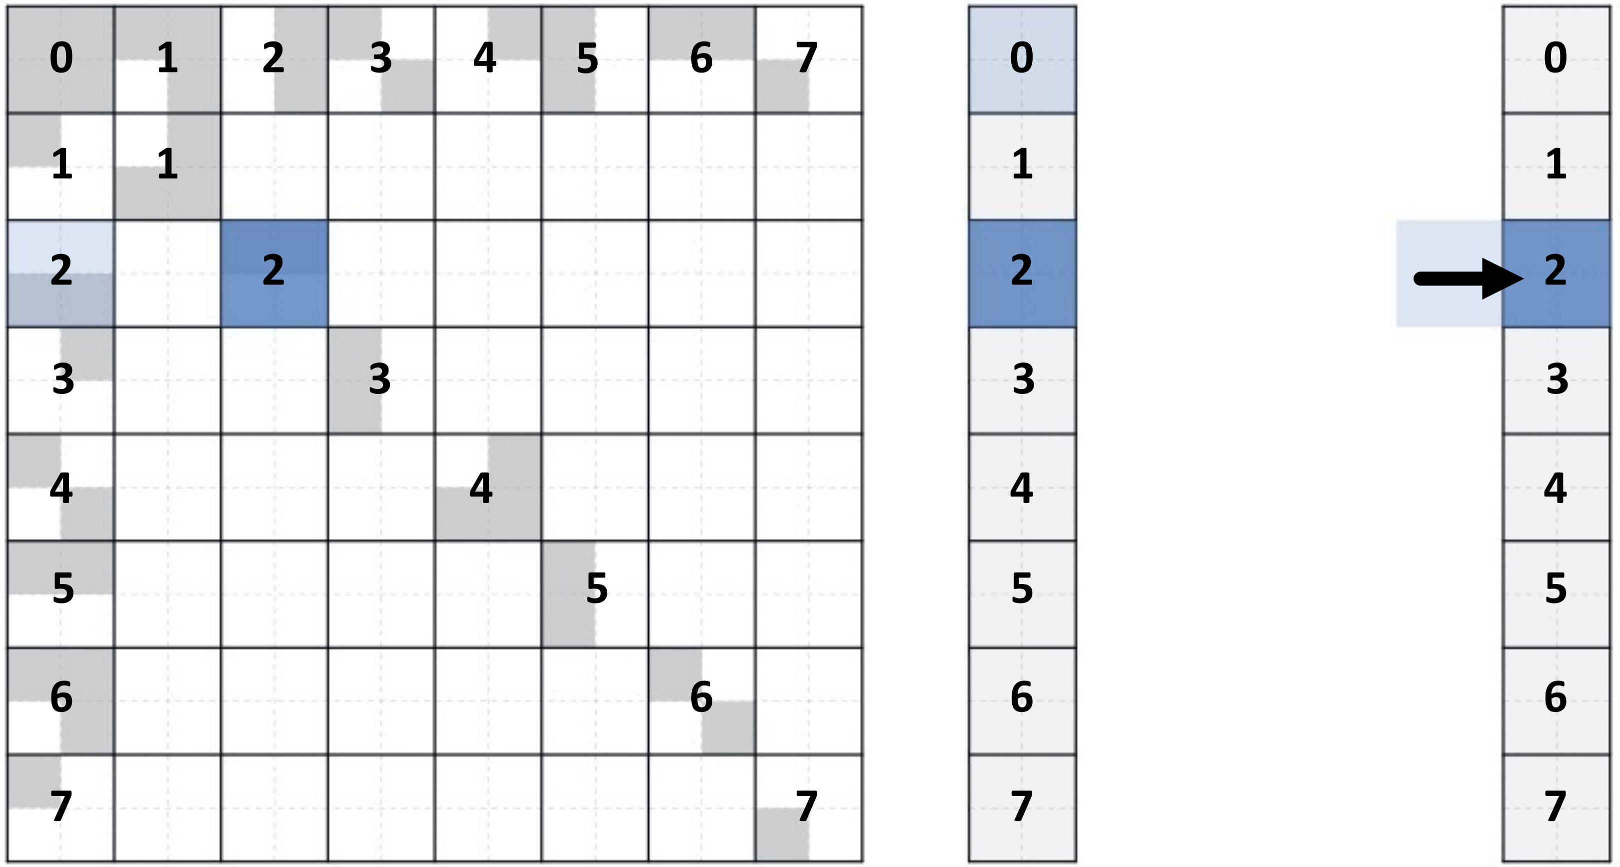

B_i X' = Y'

METHOD – 2. ARROW MULTIPLY

$$AX = \sum_i P_i B_i (P_i^\intercal X)$$

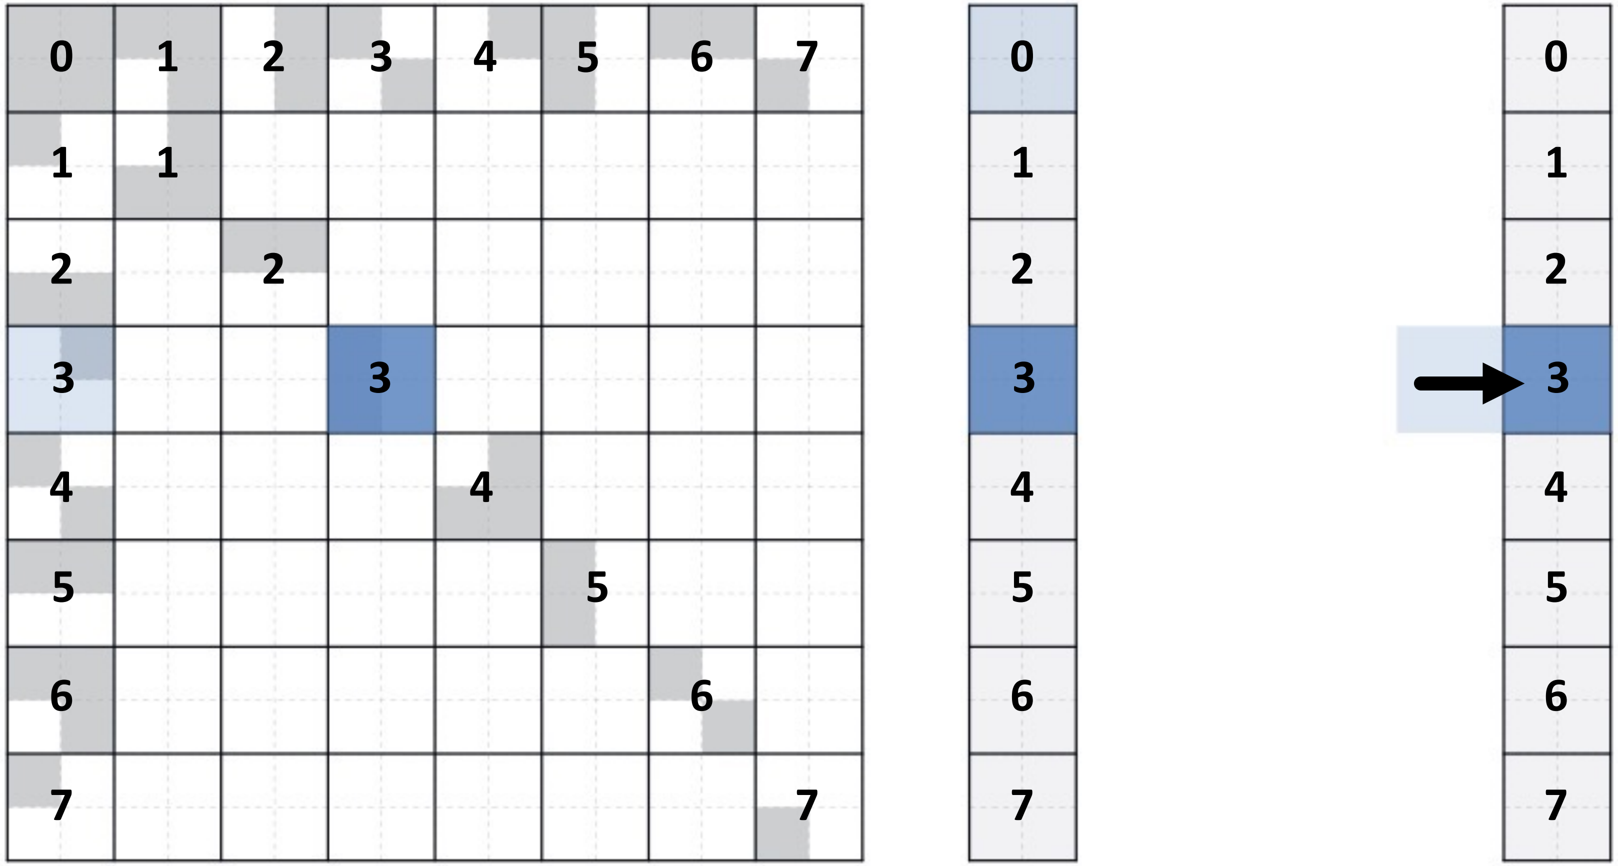

B_i X' = Y'

$$AX = \sum_i P_i B_i (P_i^\intercal X)$$

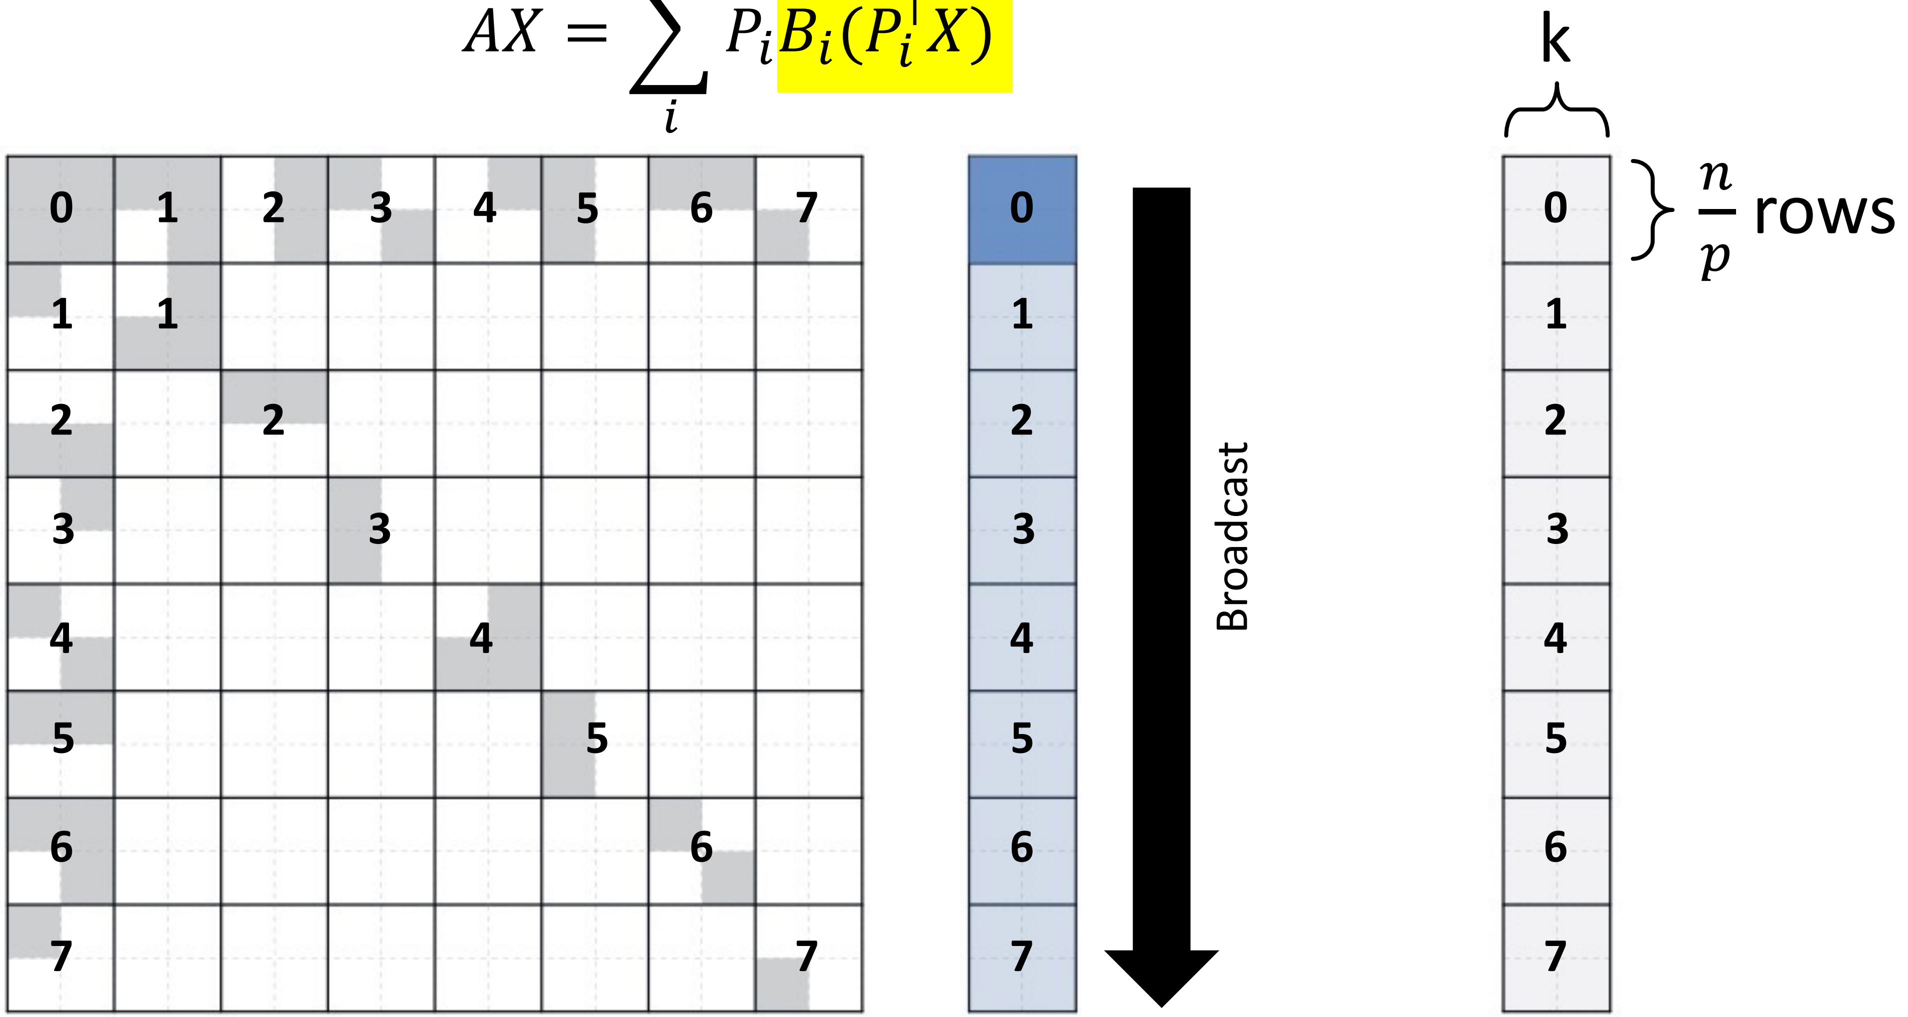

$O\left(\frac{nk}{p}\right)$ communication

B_i X' = Y'

Method – 2. Arrow Multiply

$$AX = \sum_i P_i B_i (P_i^\intercal X)$$

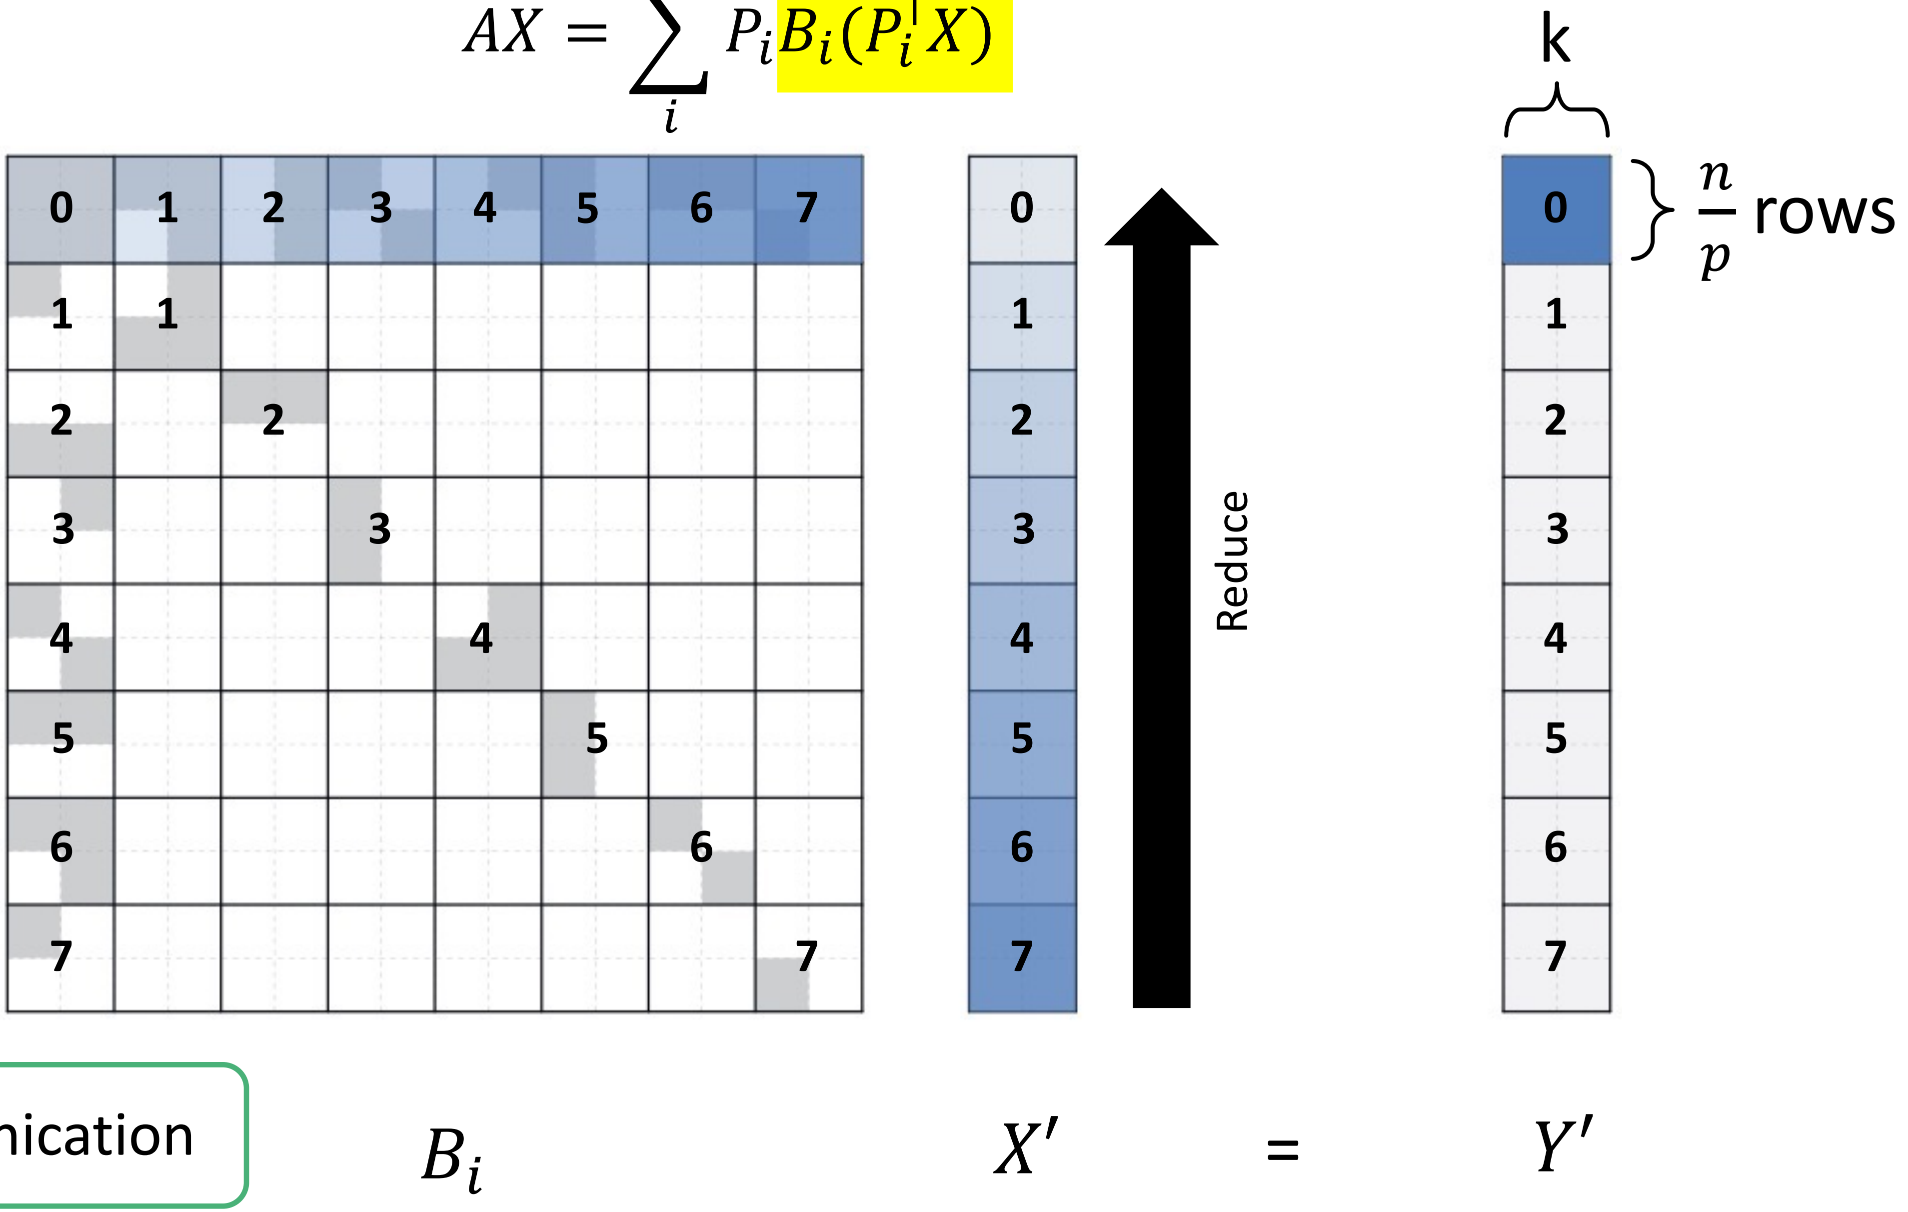

$O\left(\frac{nk}{p}\right)$ communication

$$AX = \sum_i P_i B_i (P_i^\intercal X)$$

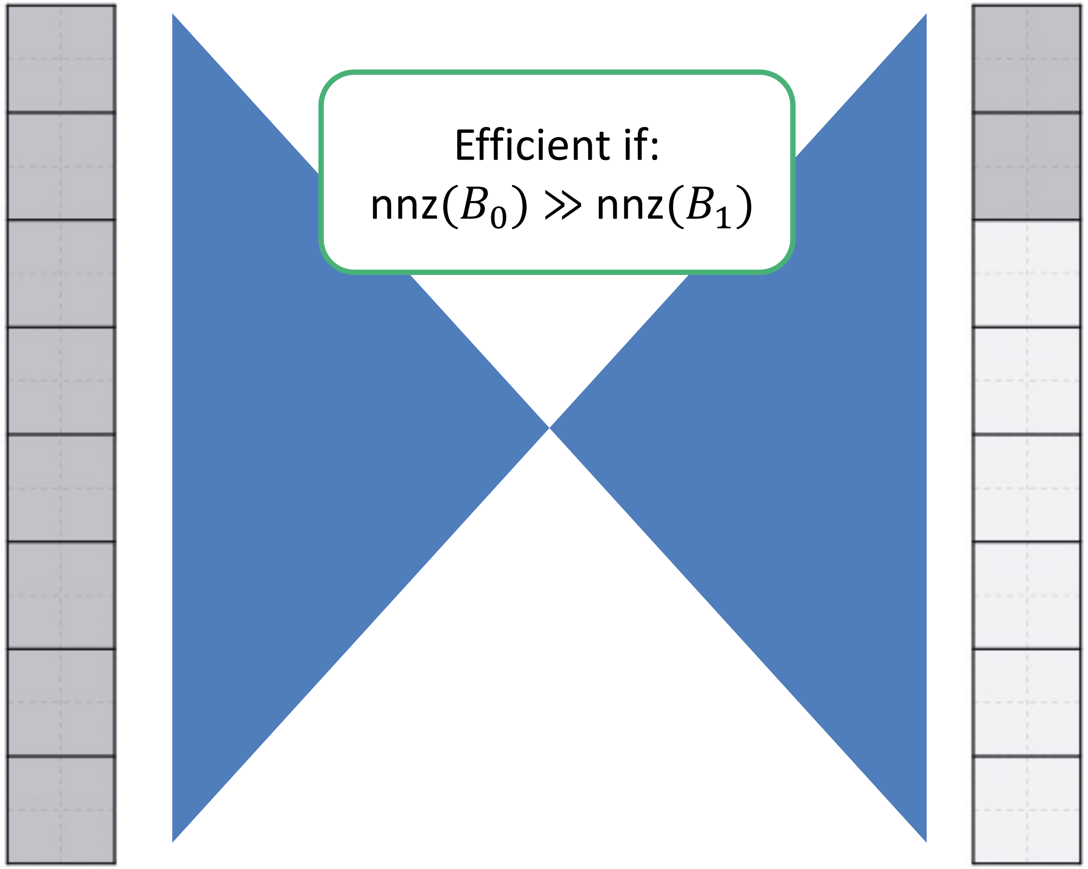

Efficient if:
$\text{nnz}(B_0) \gg \text{nnz}(B_1)$

$$B_0(P_0^\intercal X) + (P_0^\intercal P_1) B_1 (P_1^\intercal X) \qquad B_1(P_1^\intercal X)$$

Baselines:
1.5D = 1.5D algorithm [1]
HP = Hypergraph Partitioning [2] + PETSc SpMV-style (nonblocking)

All algorithms use cuSPARSE CSR kernel, MPICH
Piz Daint, up to 512 GPUs (Nvidia P100)

SuiteSparse Matrix Collection:

Dataset	Vertices n	$\frac{\mathrm{nnz}(A)}{n}$	Max. Degree Δ
MAWI 226M	226,196,185	2.12	210,795,477
MAWI 69M	68,863,315	2.08	63,040,326
GenBank 214M	214,005,017	2.17	8
GenBank 68M	67,716,231	2.05	35
WebBase 118M	118,142,155	8.63	816,127
OSM Europe	50,912,018	2.12	13
GAP-twitter 62M	61,578,415	23.85	770,155
sk-2005 51M	50,636,154	38.50	8,563,808

[1] Tripati et al., SC 2020

[2] Mayer et al., IEEE BigData 2018

Street Network, 51M rows
Max. degree ≤ 3

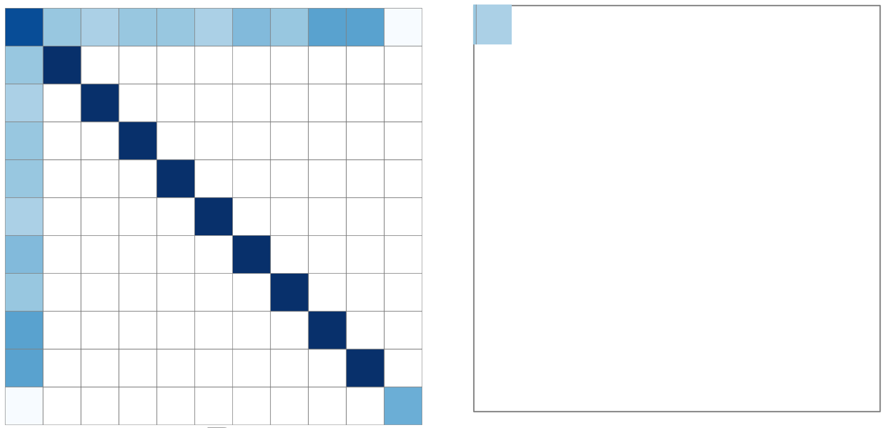

B_0 B_1

Protein Network, 68M rows
Max. degree ≤ 3

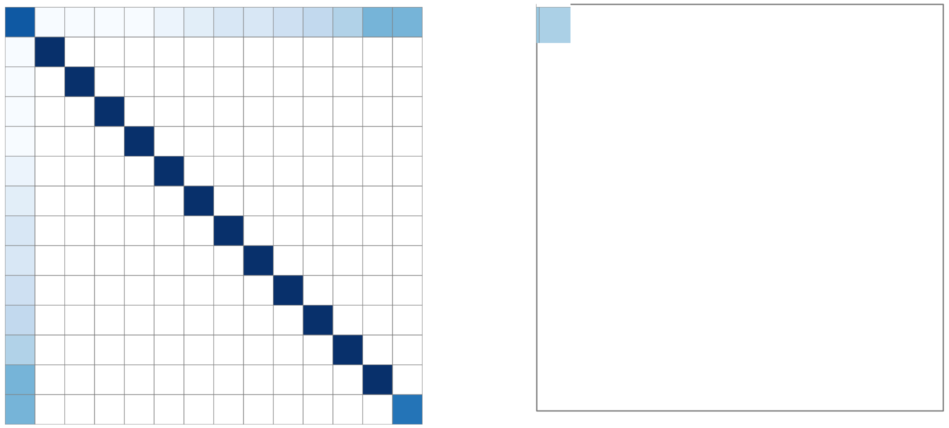

B_0 B_1

Network Anomalies, 69M rows
Max. degree ~ 63M

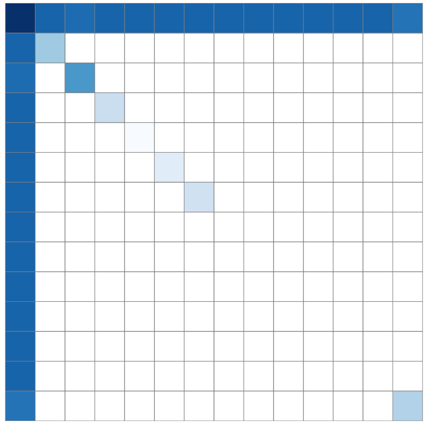

B_0

Web Graph, 118M rows
Max. degree ~ 800K

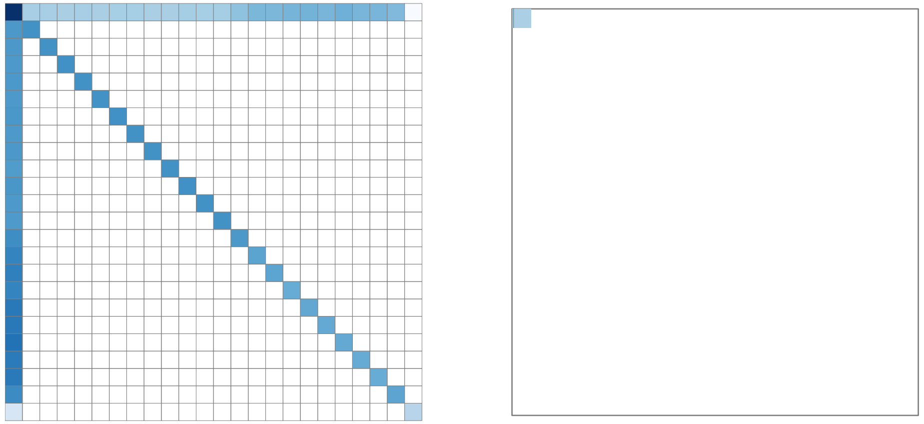

B_0 B_1

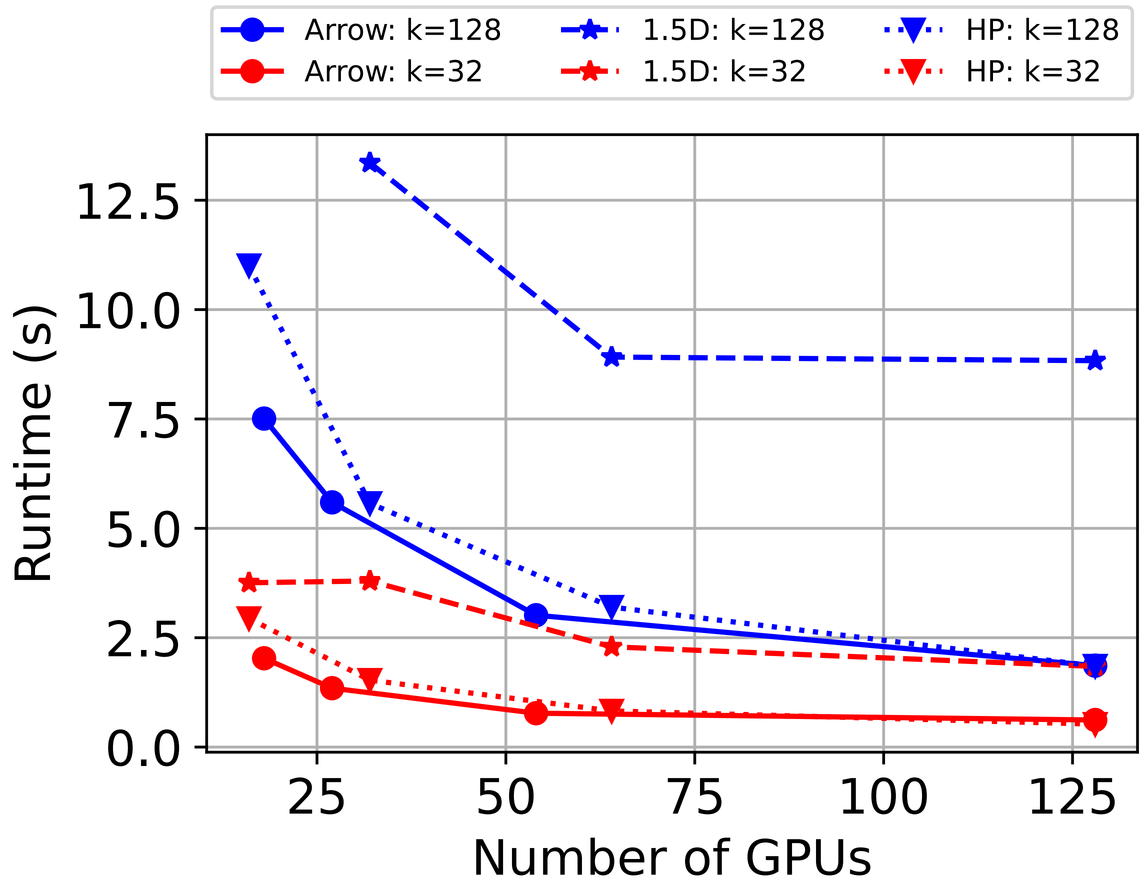

Web Subgraph

sk-2005, 51M rows

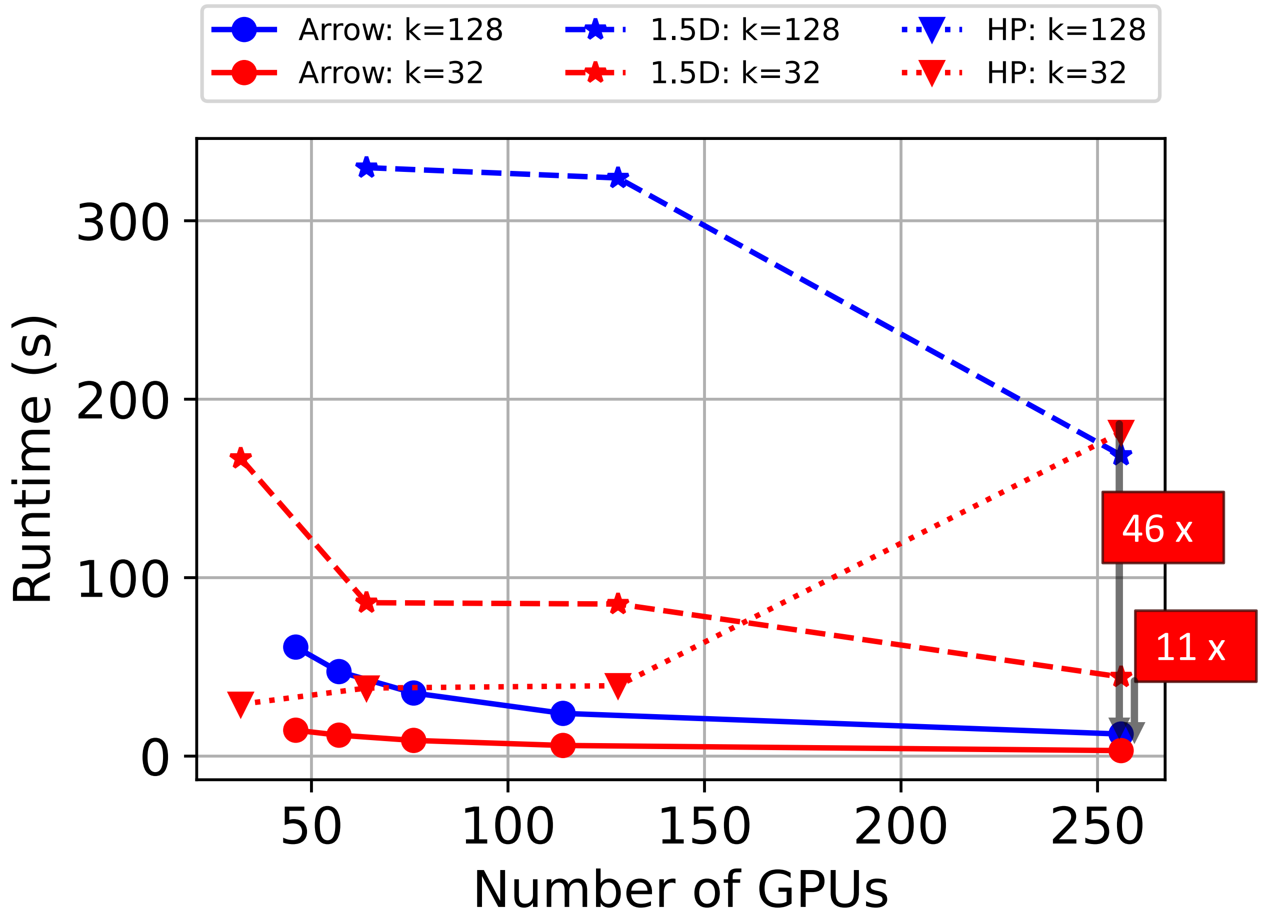

Network Anomalies

MAWI, 226M rows

RESULTS – WEAK SCALING

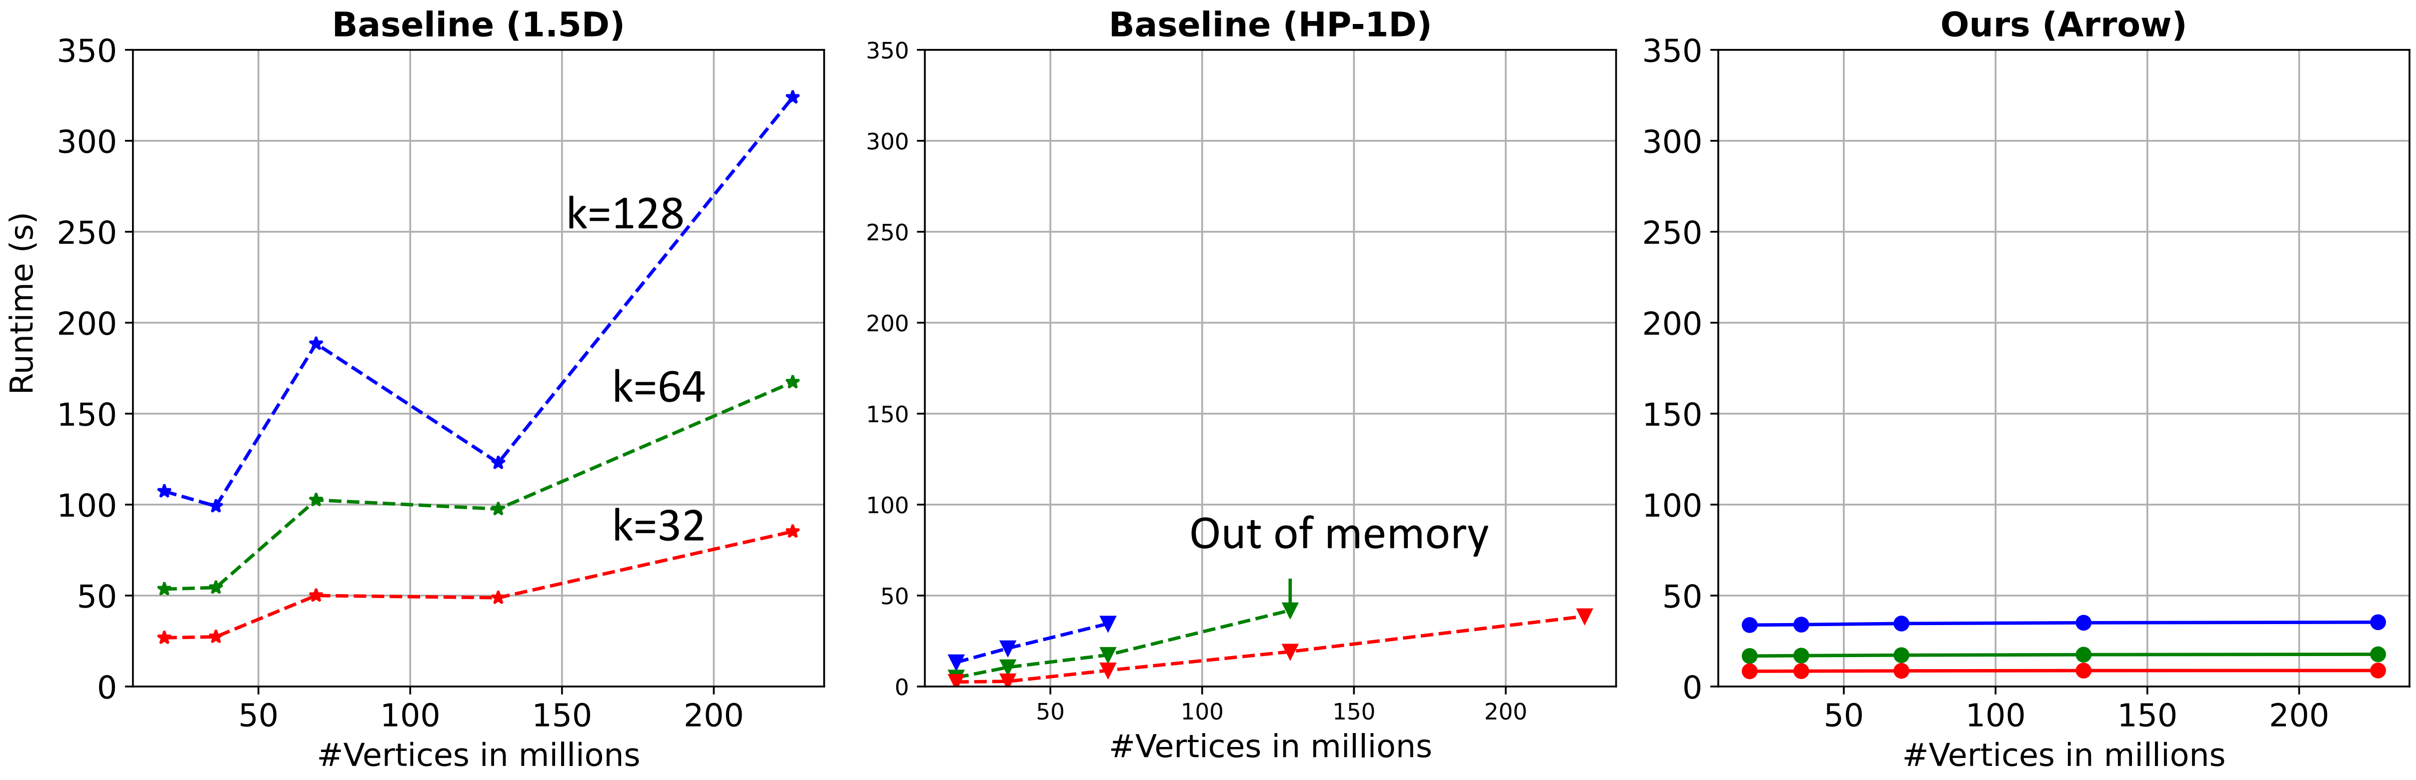

MAWI matrices, |V|/p ~ 3M

CONCLUSIONS

$$A = \sum_i P_i B_i P_i^\intercal$$

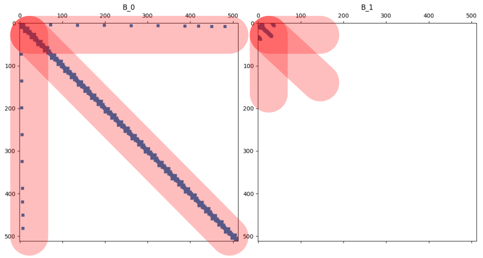

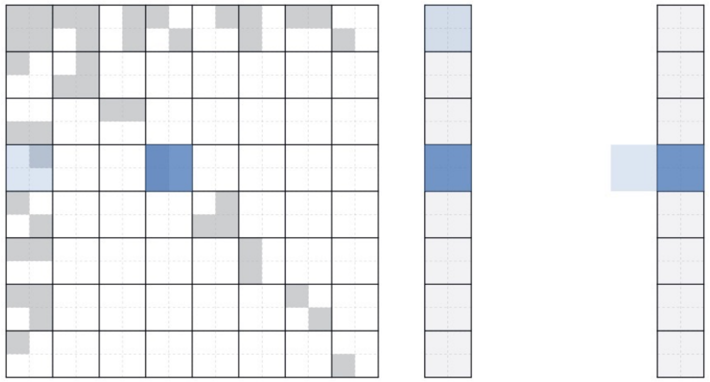

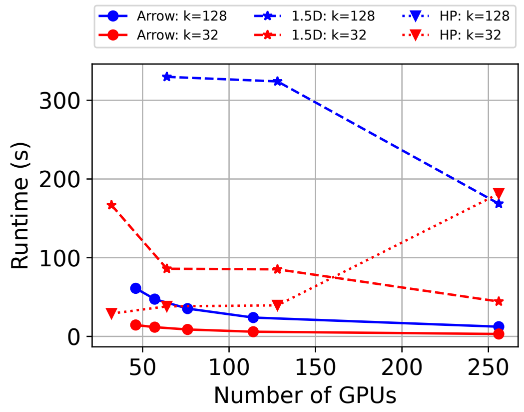

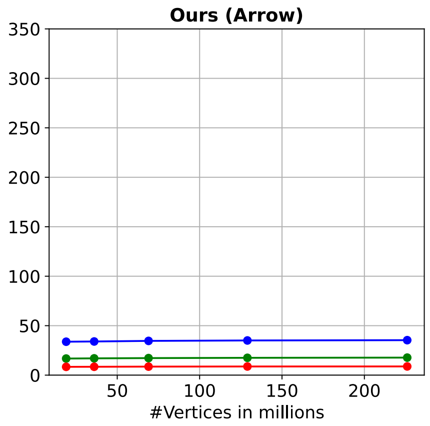

More of SPCL's research:

- youtube.com/@spcl — 180+ Talks
- twitter.com/spcl_eth — 1.4K+ Followers
- github.com/spcl — 3.8K+ Stars

… or spcl.ethz.ch

github.com/spcl/arrow-matrix