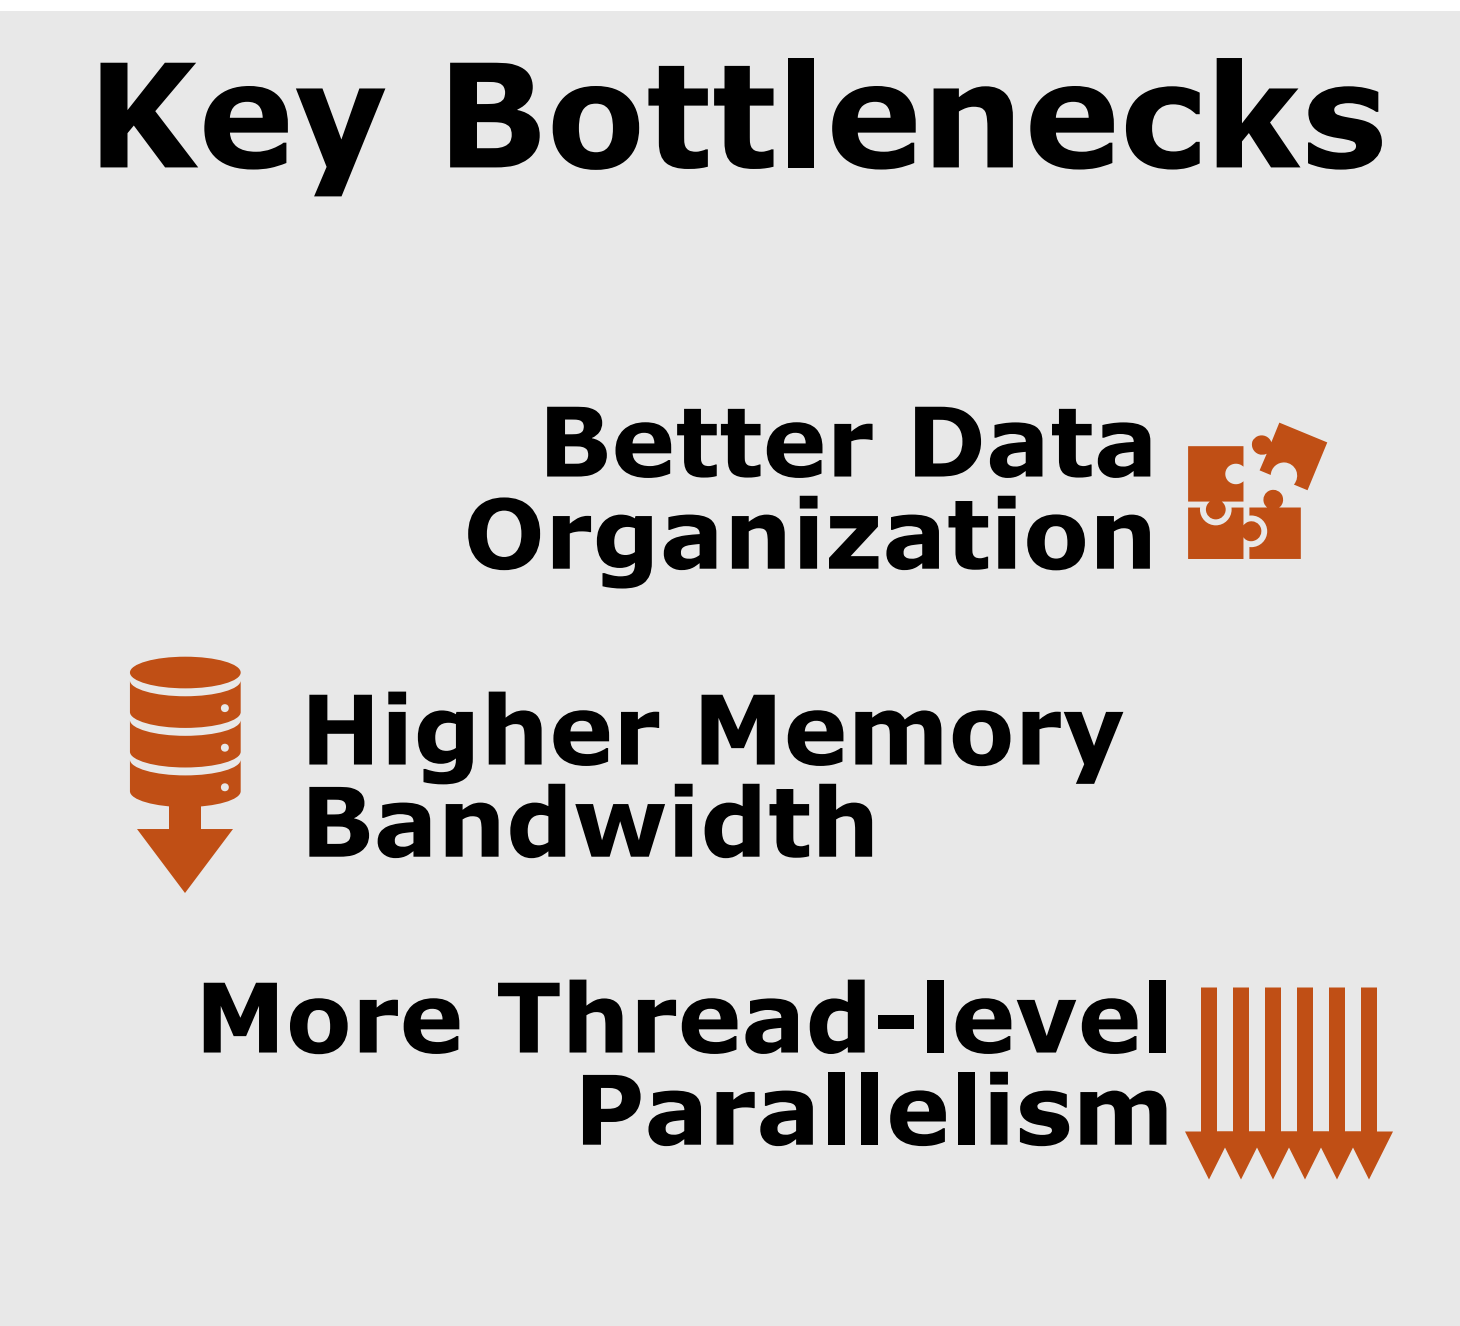
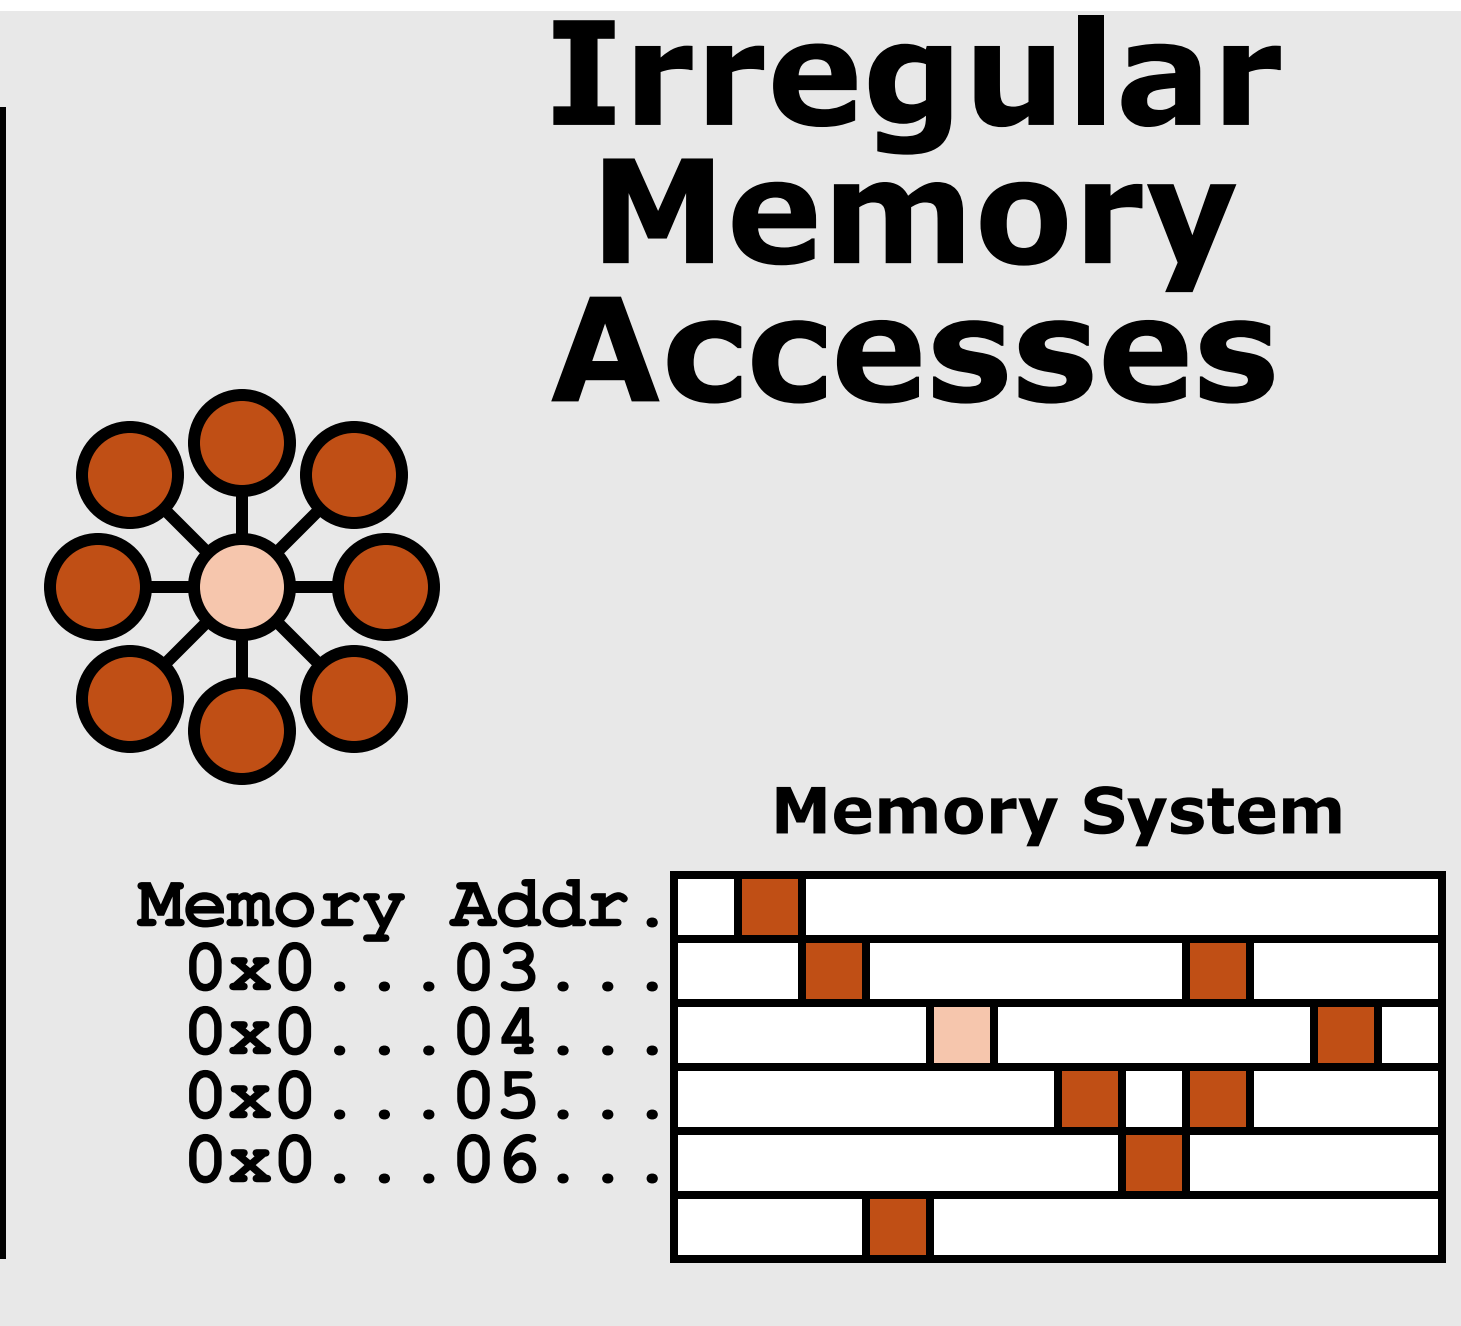
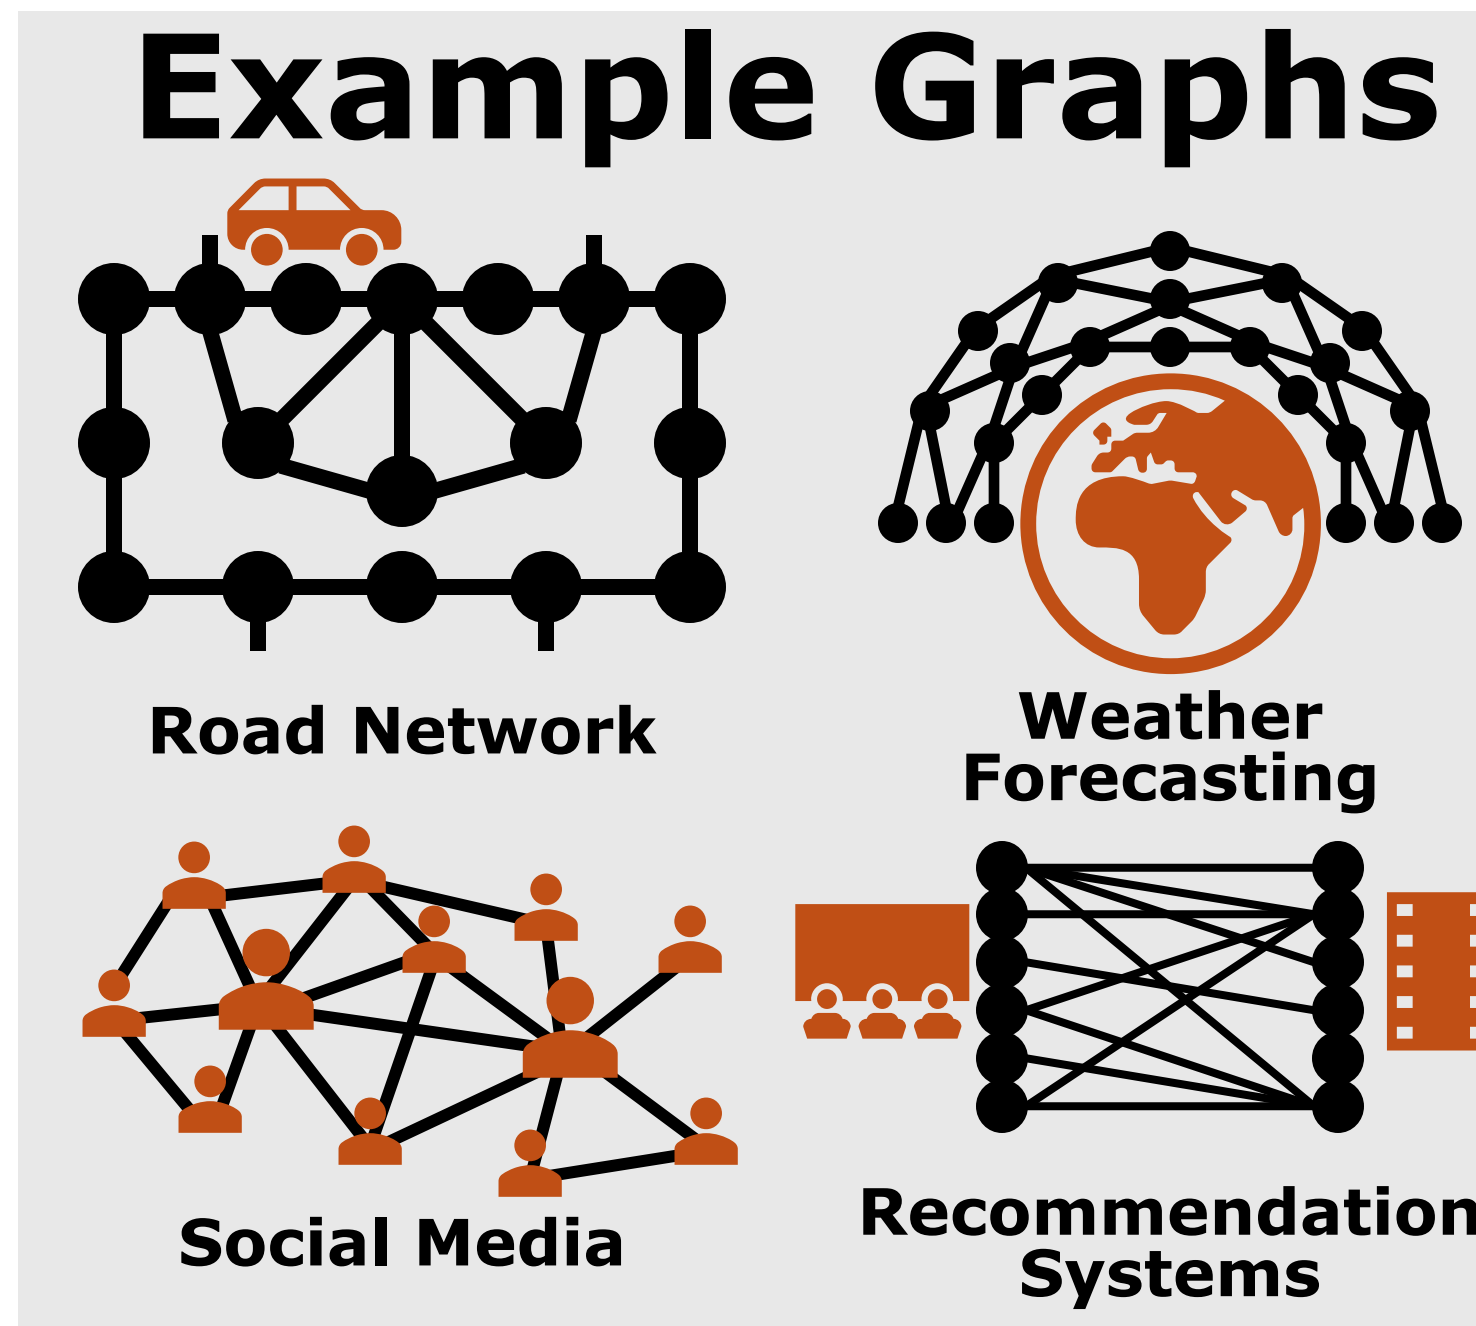
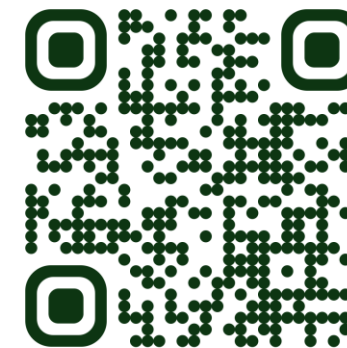


Processing-in-Memory hardware accelerators can improve the Performance & Energy Efficiency for Graph-based Processing applications.

 Andrew McCrabb
Advisor: Valeria Bertacco
 DATE 2025 - PhD Forum



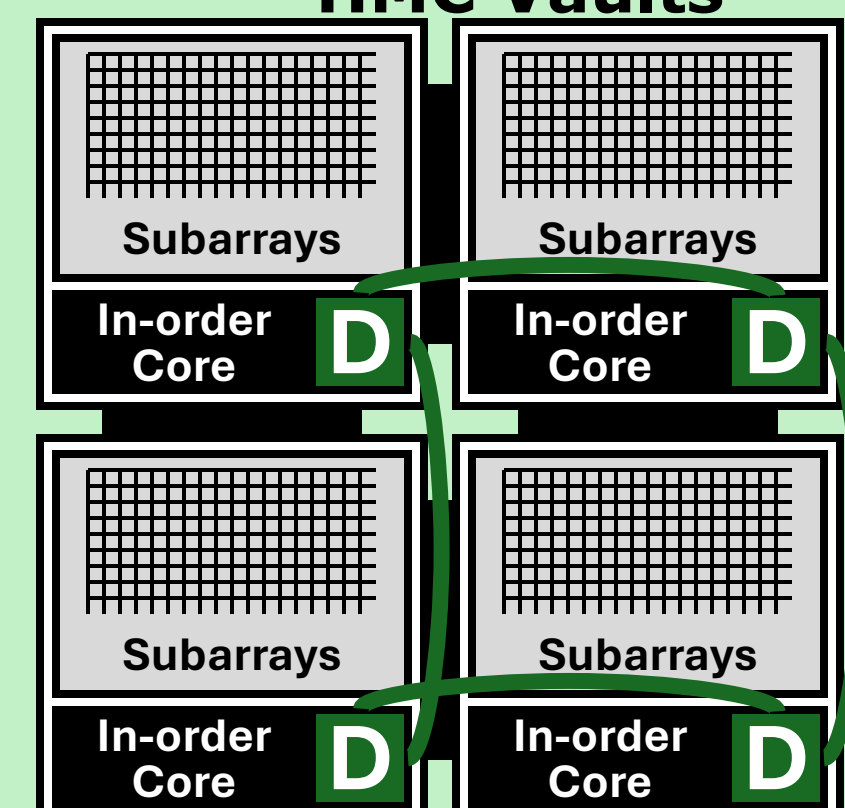
Graph as Data Framework



DREDGE

Dynamic Graph Processing can be accelerated by executing data reorganization functions in PIM.
 (e.g., Dynamic BFS, Dynamic PageRank, etc.)

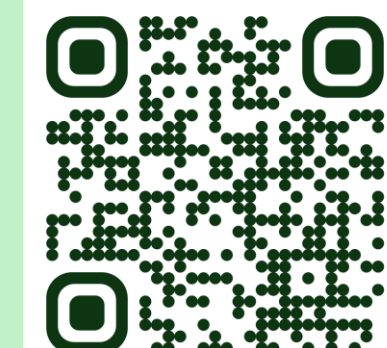
HMC Vaults



- Better Data Organization
- Higher Memory Bandwidth
- No More Thread-level Parallelism


Dynamic Graph Application
Performance Gain: 3-15x
DREDGE Power: 50mW

Benchmark Suite



DyGraph

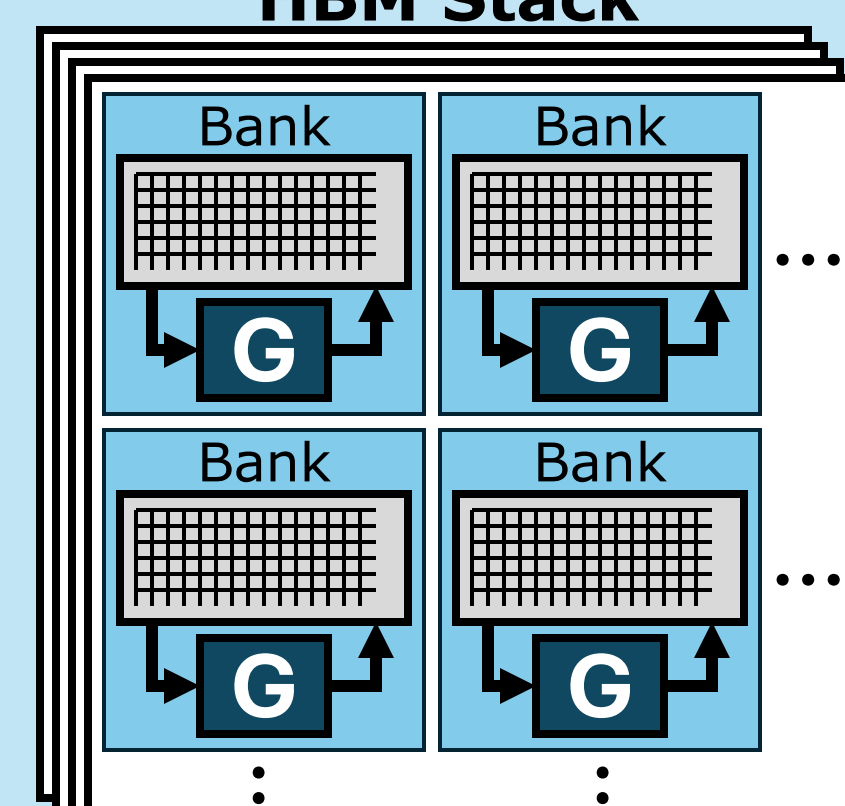
Graph as Both



GLEAM

Graph Neural Network Processing can be accelerated by executing neuron aggregation functions in PIM.
 (e.g., GraphSAGE, Graph Convolution Networks, Graph Attention Networks, etc.)

HBM Stack




- Better Data Organization
- Higher Memory Bandwidth
- More Thread-level Parallelism

Graph Neural Network Application
Performance Gain: 10x
Power Savings: 77%

Benchmark Suite

[Abundant Prior Work]

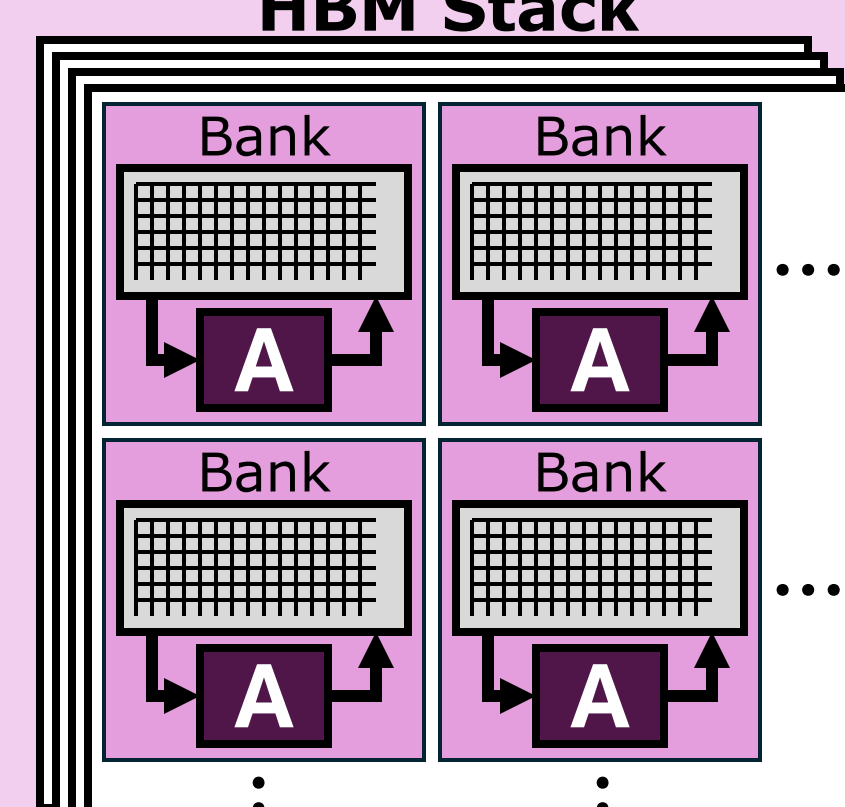
Graph as Algorithmic Framework



ACRE

Explainable AI Applications can be accelerated by executing Comparison functions in PIM.
 (e.g., Random Forest, XGBoost, etc.)


HBM Stack



- No Better Data Organization
- Higher Memory Bandwidth
- More Thread-level Parallelism

Explainable Random Forest Applications
Performance Gain: 10x
Power Savings: 93%

Benchmark Suite



BeXAI