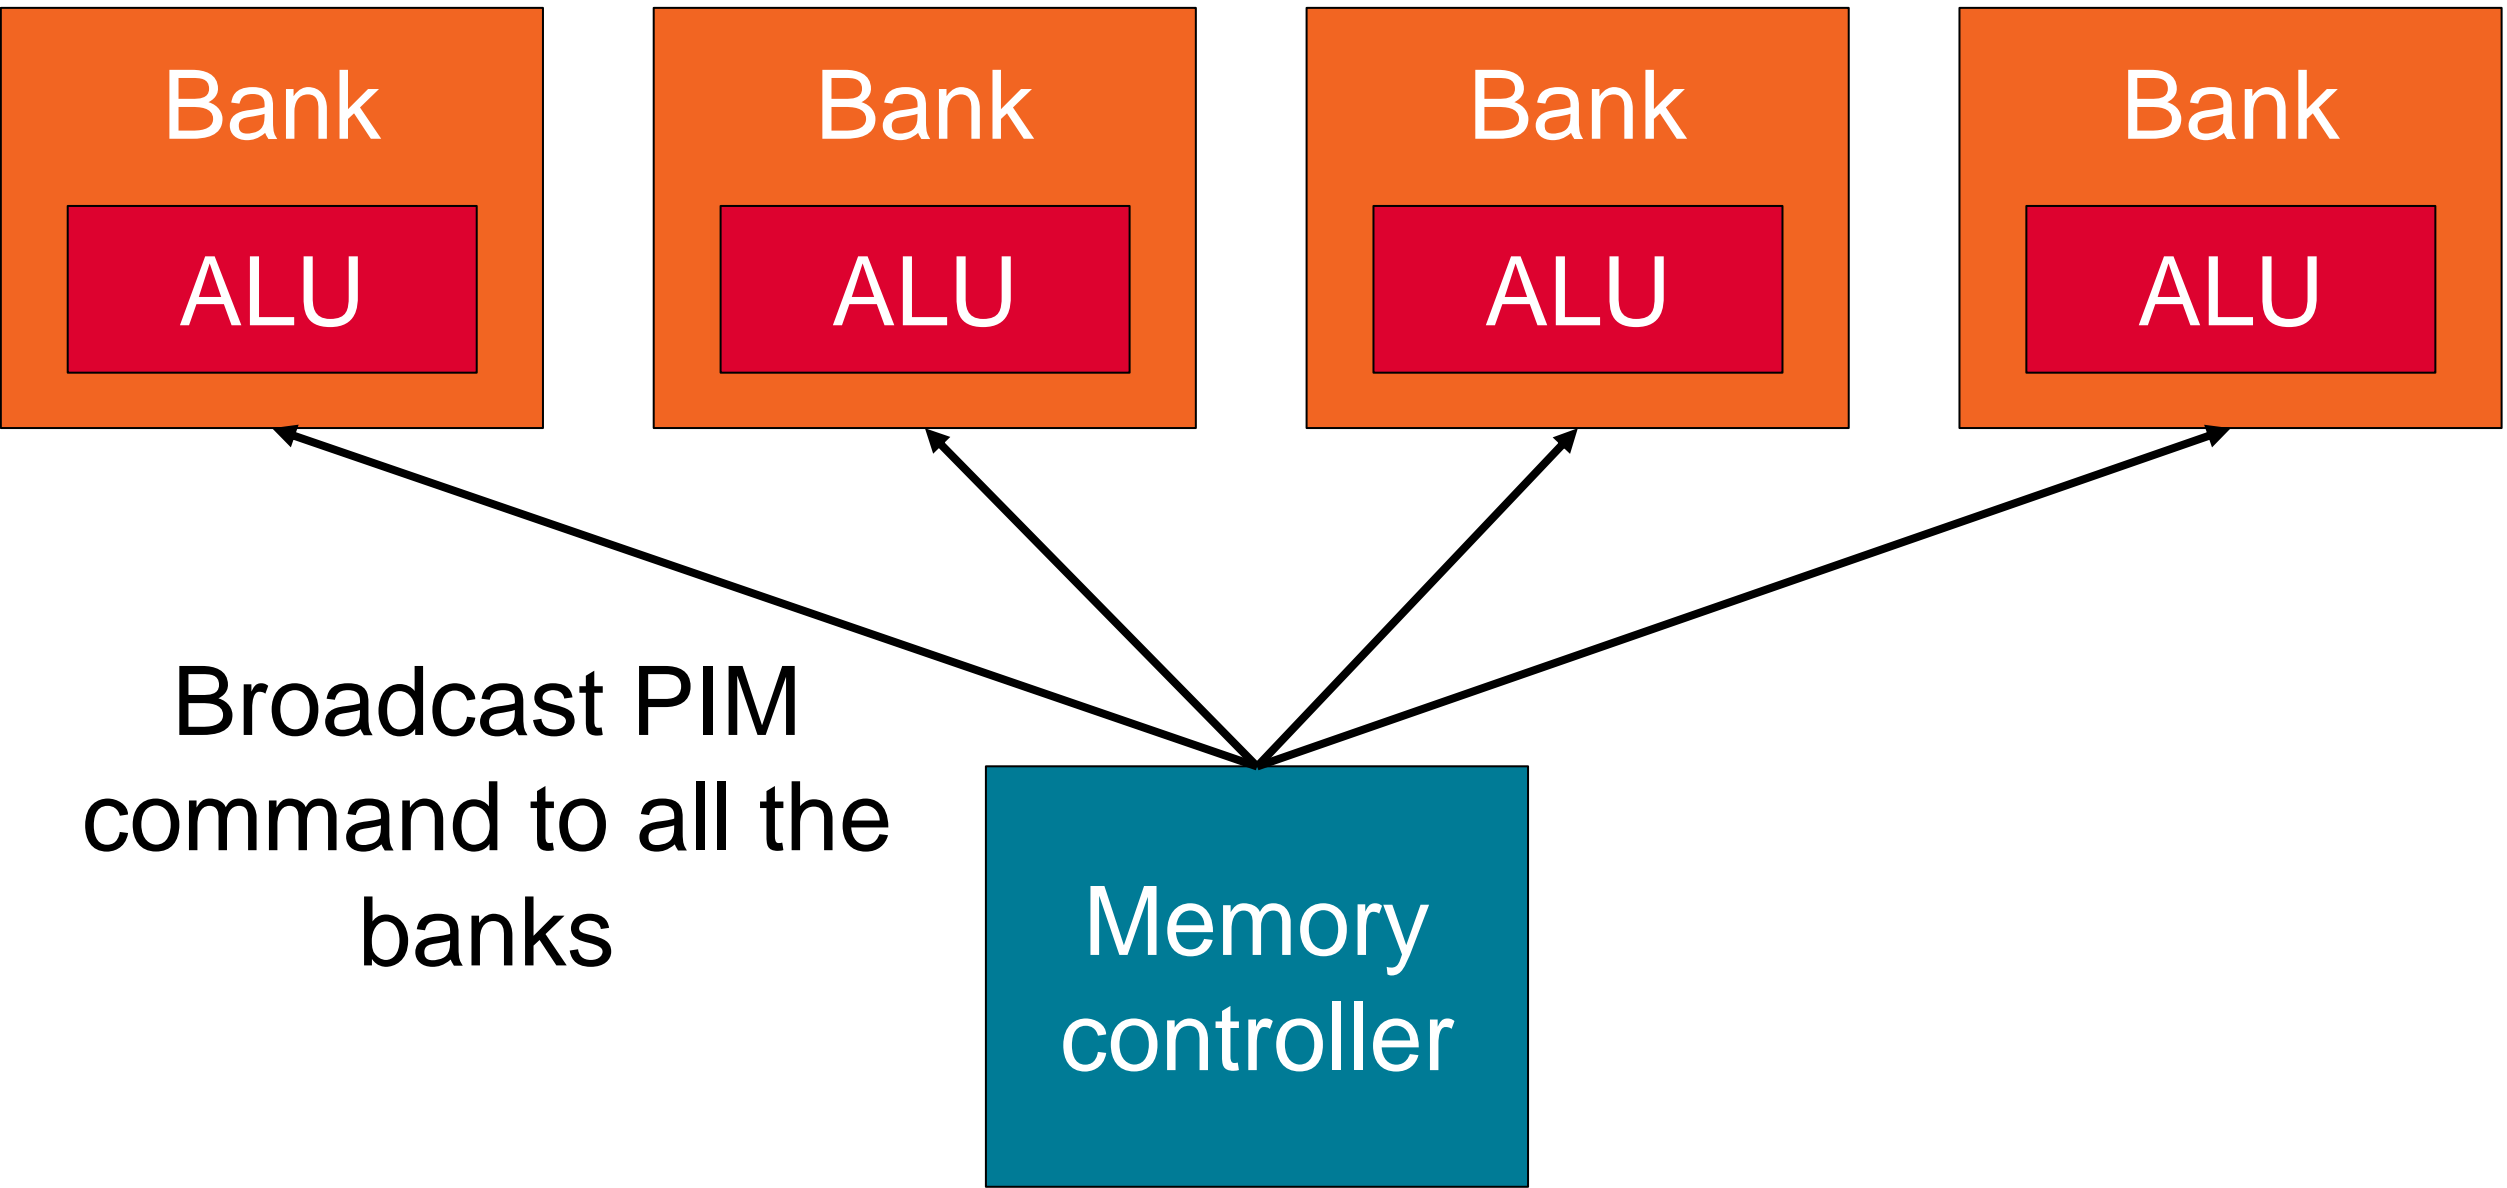


Accelerating Ray Tracing via Processing in Memory

Onur Kayiran, Mohamed A. Ibrahim, Shaizeen Aga
Advanced Micro Devices, Inc.

PIM Overview and Key Properties



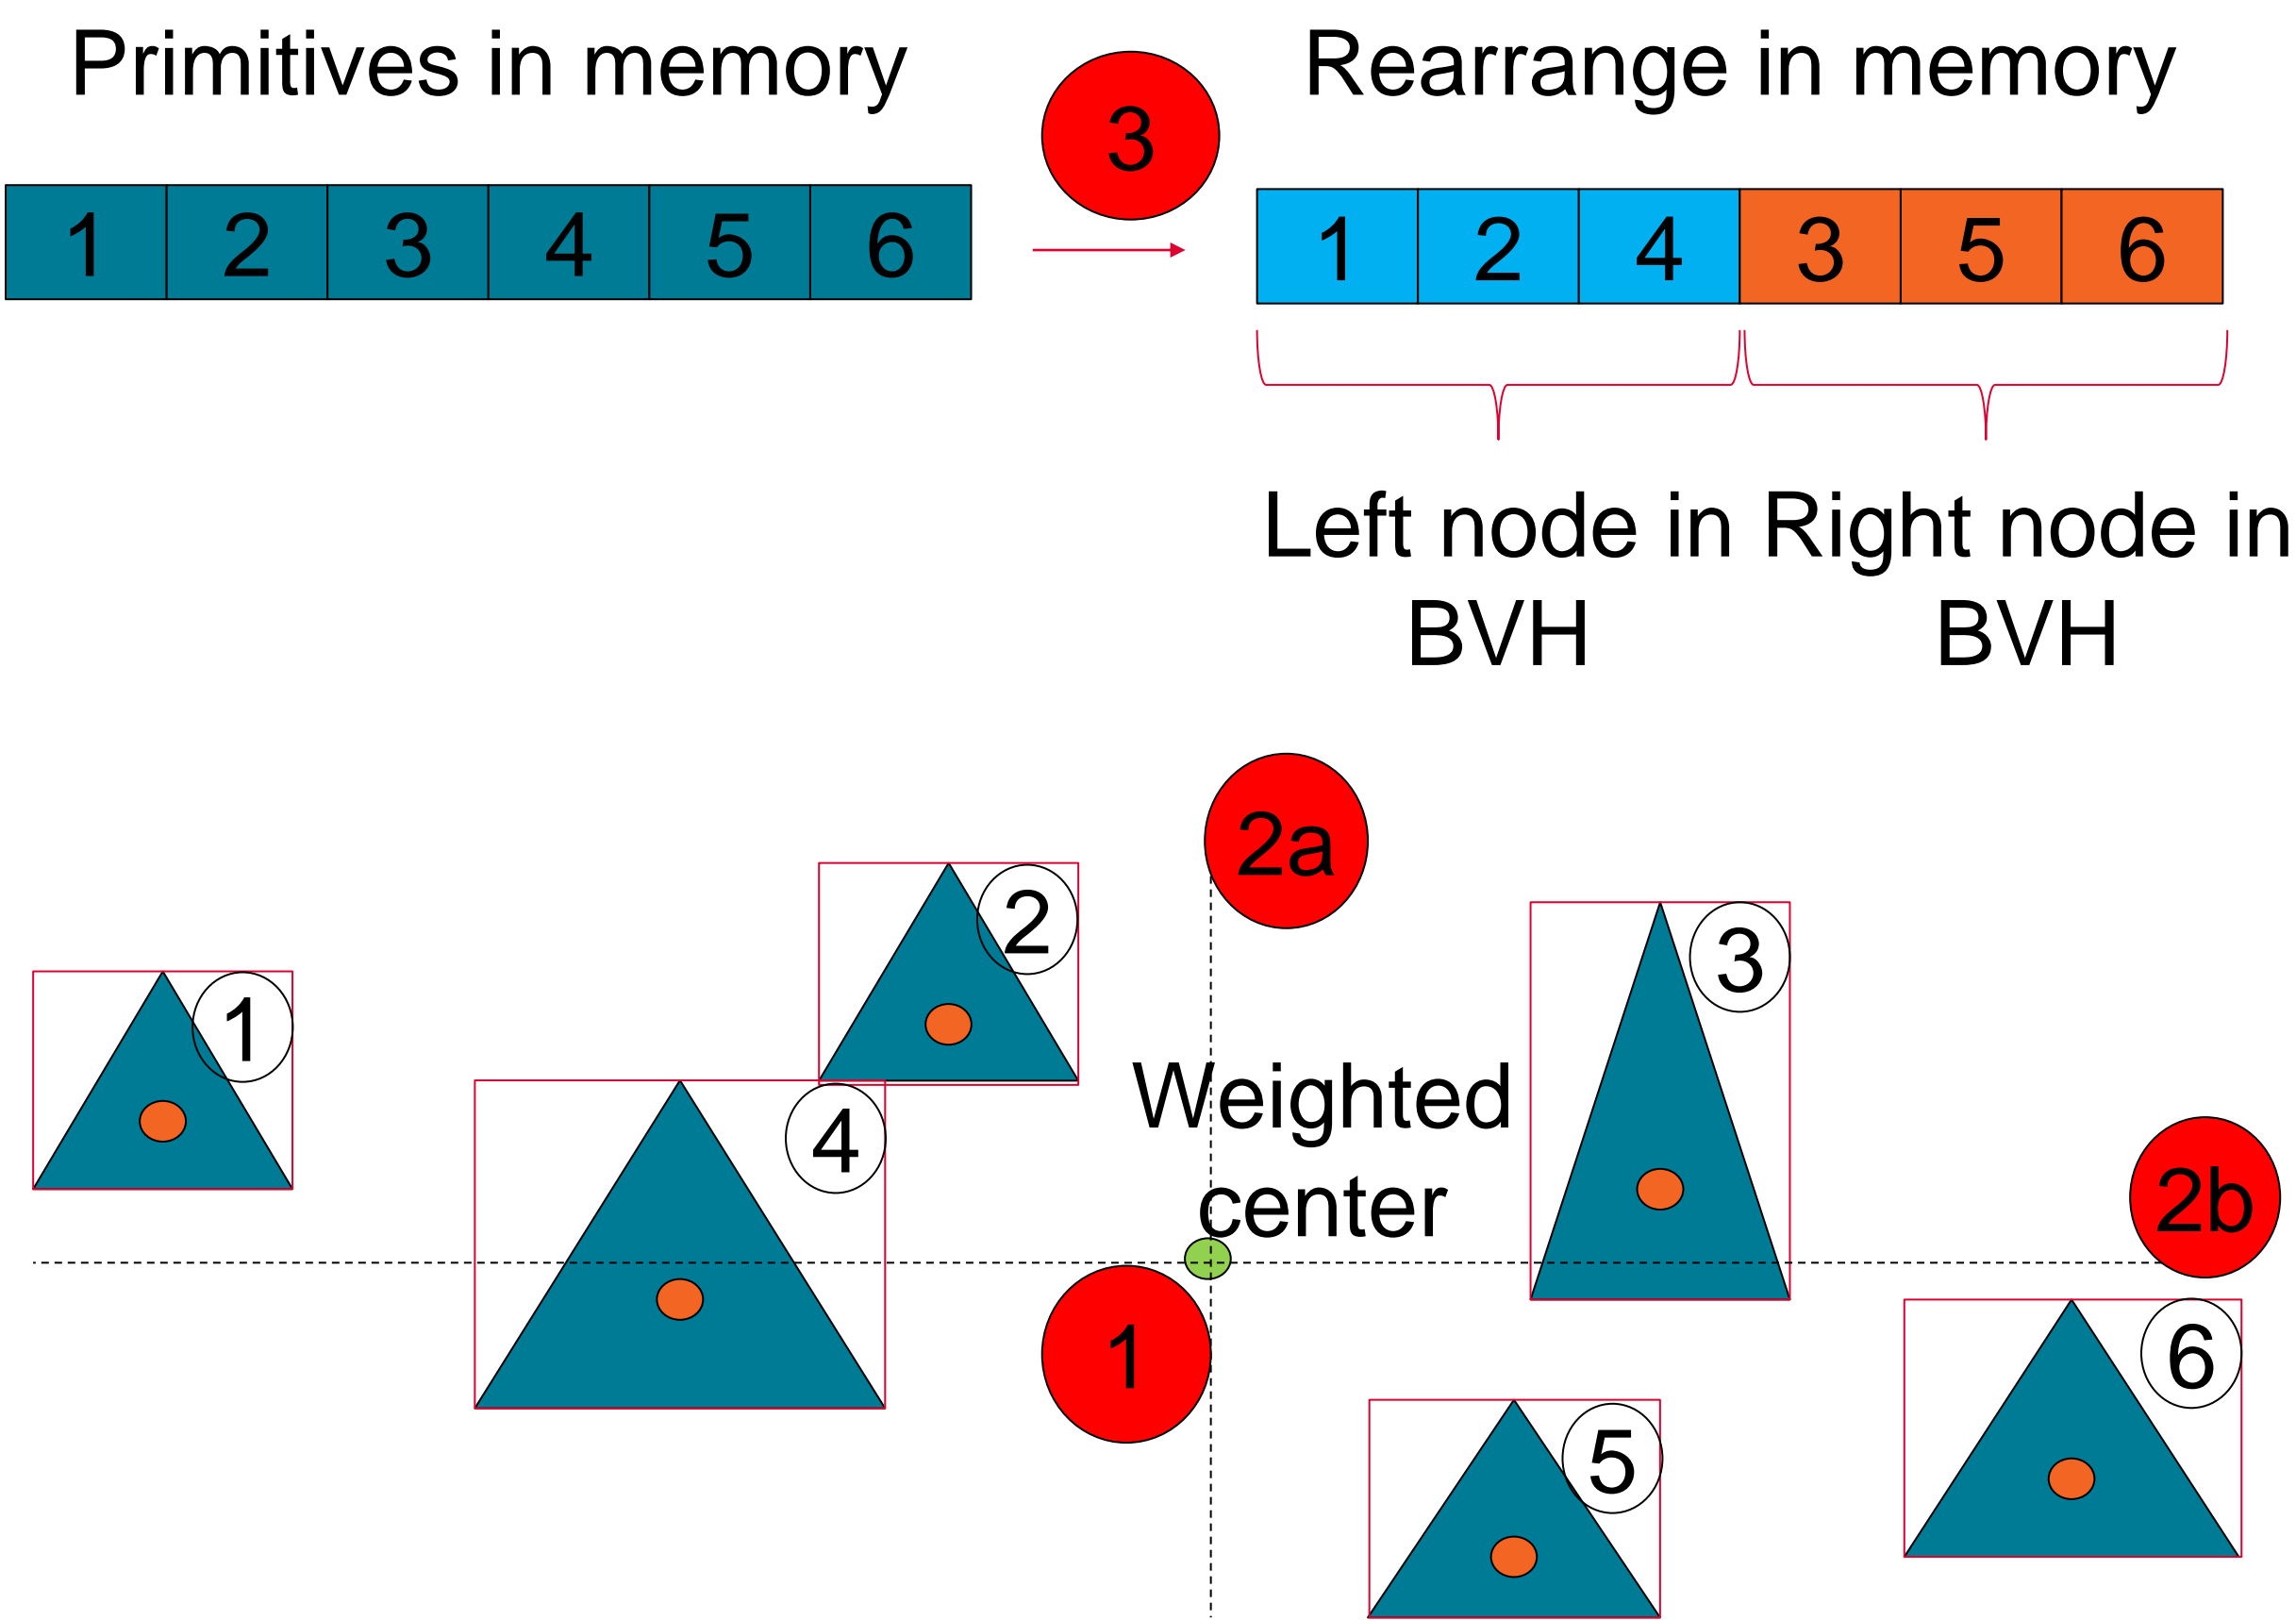
This PIM design can accelerate workloads with:

1. High memory bandwidth consumption
2. Low data reuse
3. Regular memory access patterns

Key Computations in Ray Tracing


1. Ray – triangle/bounding box intersection
 - Acceleration by fixed-function hardware
2. Bounding volume hierarchy (BVH) traversal
 - Irregular memory accesses
3. Image denoising
 - High data reuse
4. BVH build
 - Good candidate for PIM acceleration

BVH Build Overview




Memory Mapping Schemes

Naïve memory mapping




Row-aware memory mapping



- Banks contain primitives from a single tree node
 - If tree node is smaller than one row buffer → host processing
 - The spillover data → host processing & at most one row per bank
- Pros:
 - Improved command and memory bandwidth
- Cons:
 - Increased memory allocation (bounded)
 - Data movement reduction per row buffer → limited benefits

Row+bank-aware memory mapping



- Banks contain primitives from a single tree node & same tree node is mapped to the same bank
 - If tree node is smaller than one row buffer → host processing
 - The spillover data → host processing & at most one row per bank
- Pros:
 - Improved command and memory bandwidth
 - Data movement reduction per tree node → significant benefits
- Cons:
 - Increased memory allocation (bounded)

Benefits

