


Making Graph Computations Fast, Simple, and Portable

Yunming Zhang and Collaborators




Graphs Are Everywhere



Economy

Affordable rides, all to yourself




\$5.70
3:16pm



\$9.71
3:06pm

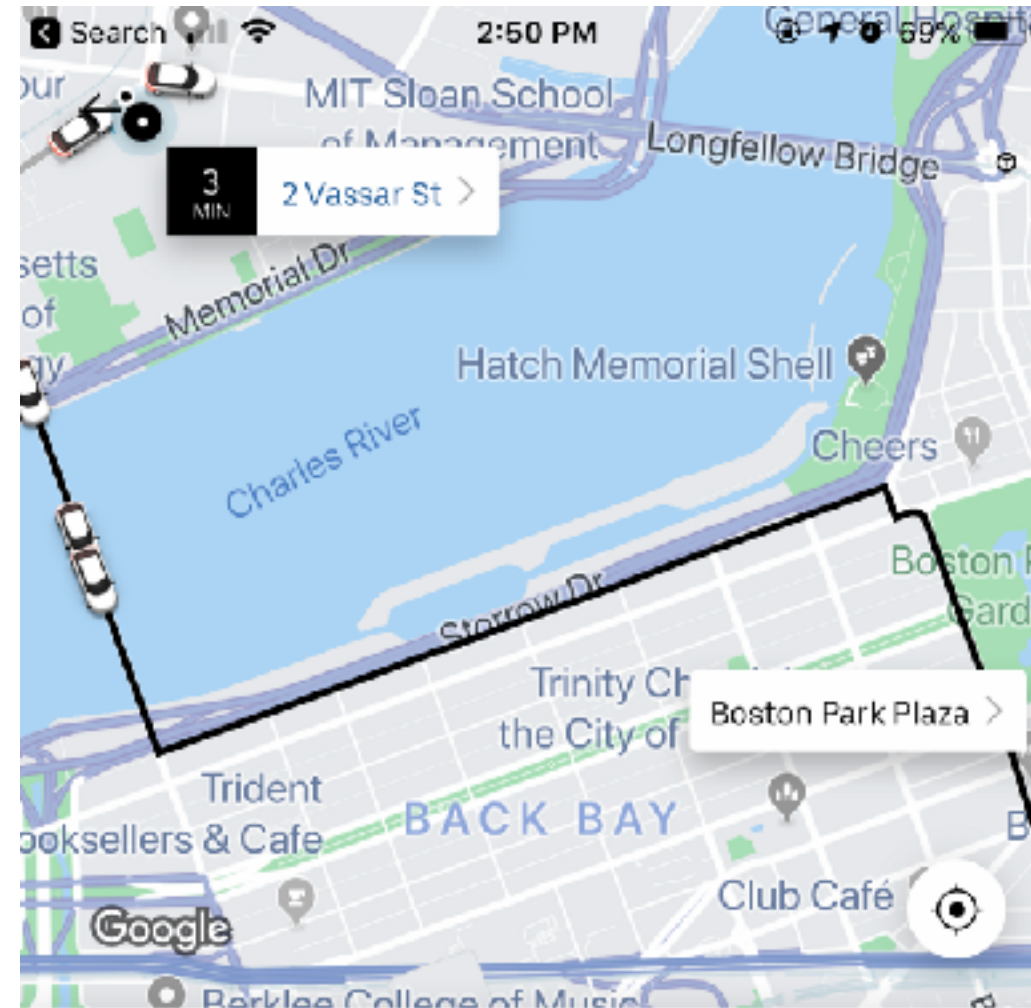
Recommendations for You, Yunming



Google Search



I'm Feeling Lucky

Performance is Important


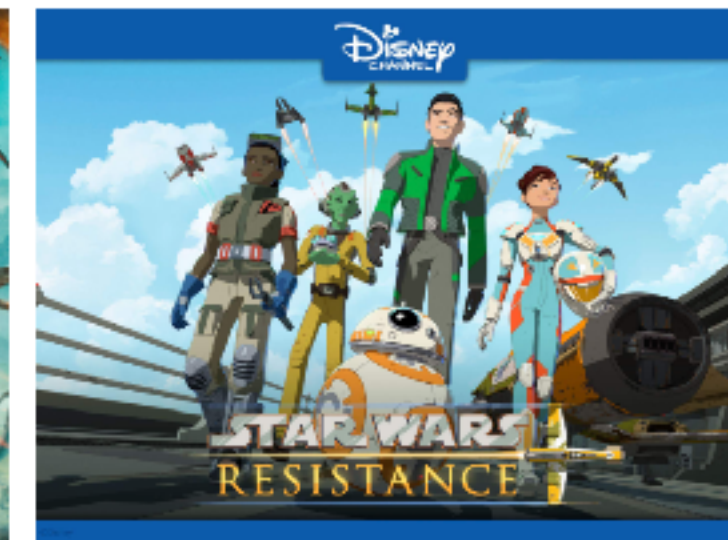


Economy Premiu

Affordable rides, all to yourself

 Pool Options	 UberX
\$5.70 3:16pm	\$9.71 3:06pm

Recommendations for You, Yunming



World is Built for Dense

- Hardware Utilization
- Programming System

World is Built for Dense

- Hardware Utilization
 - Peak Performance (GEMM)
 - 70-80% of CPU
 - 80-90% of GPU
 - Optimizations
 - Prefetching, Branch Predictions, TLB, cache, ..
- Programming System

World is Built for Dense

- Hardware Utilization
 - Peak Performance (GEMM)
 - 70-80% of CPU
 - 80-90% of GPU
 - Optimizations
 - Prefetching, Branch Predictions, TLB, cache, ..
- Programming System
 - Abstractions that Work across Different Algorithms (Dense Linear Algebra, Image Processing, Deep Learning, ..)
 - BLAS, Halide, TensorFlow, ...
 - Optimizing Compilers
 - Tiling, Vectorization, Unrolling, ..

Not Ready for Graphs

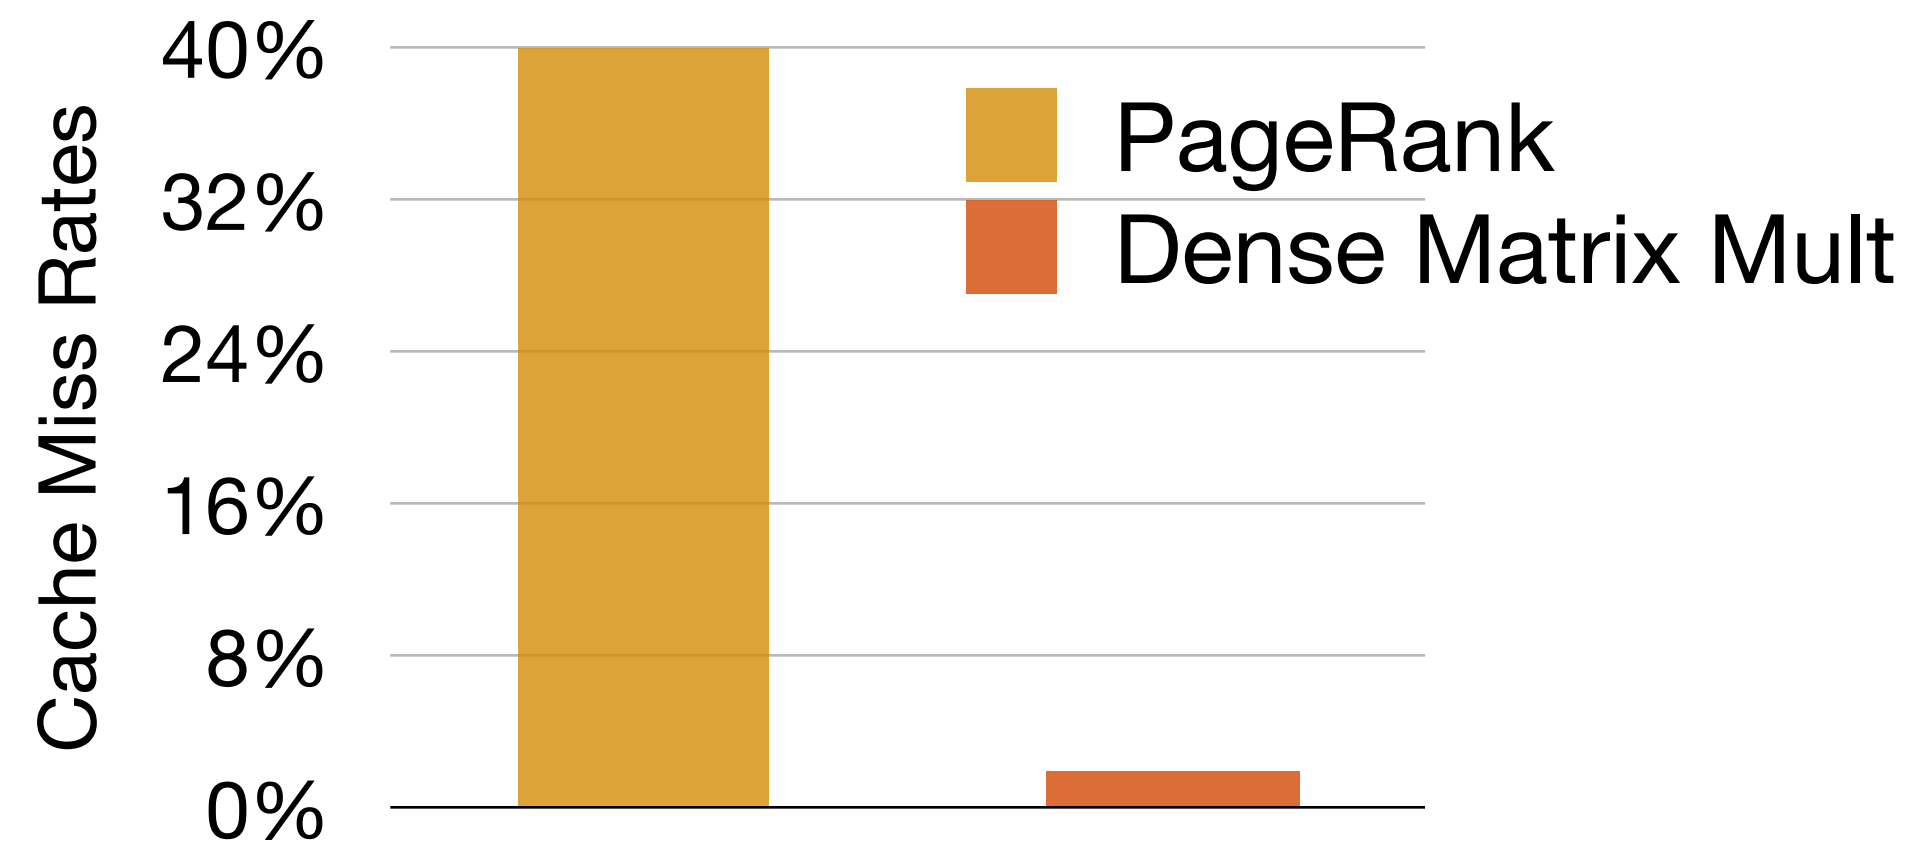
- Hardware Utilization
- Programming System

Not Ready for Graphs

- Hardware Utilization
- Peak Performance (PageRank, SpMv)
- < 10% Peak of CPU and GPU
- Programming System


Not Ready for Graphs

- Hardware Utilization
- Programming System
- Peak Performance (PageRank, SpMv)
- < 10% Peak of CPU and GPU



Not Ready for Graphs

- Hardware Utilization
- Peak Performance (PageRank, SpMv)
- < 10% Peak of CPU and GPU



- Programming System


```
template<typename APPLY_FUNC>
void edgeset_apply_pull_parallel(Graph &g, APPLY_FUNC apply_func) {
    int64_t numVertices = g.num_nodes(), numEdges = g.num_edges();
    parallel_for(int n = 0; n < numVertices; n++) {
        for (int socketId = 0; socketId < omp_get_num_places(); socketId++) {
            local_new_rank[socketId][n] = new_rank[n]; } }
    int numPlaces = omp_get_num_places();
    int numSegments = g.getNumSegments("s1");
    int segmentsPerSocket = (numSegments + numPlaces - 1) / numPlaces;
    #pragma omp parallel num_threads(numPlaces) proc_bind(spread){
    int socketId = omp_get_place_num();
    for (int i = 0; i < segmentsPerSocket; i++) {
        int segmentId = socketId + i * numPlaces;
        if (segmentId >= numSegments) break;
        auto sg = g.getSegmentedGraph(std::string("s1"), segmentId);
        #pragma omp parallel num_threads(omp_get_place_num_procs(socketId)) proc_bind(close){
        #pragma omp for schedule(dynamic, 1024)
        for (NodeID localId = 0; localId < sg->numVertices; localId++) {
            NodeID d = sg->graphId[localId];
            for (int64_t ngh = sg->vertexArray[localId]; ngh < sg->vertexArray[localId +
1]; ngh++) {
                NodeID s = sg->edgeArray[ngh];
                local_new_rank[socketId][d] += contrib[s]; }}}}
    parallel_for(int n = 0; n < numVertices; n++) {
        for (int socketId = 0; socketId < omp_get_num_places(); socketId++) {
            new_rank[n] += local_new_rank[socketId][n]; }}}}
    struct updateVertex {
    void operator() (NodeID v) {
        double old_score = old_rank[v];
        new_rank[v] = (base_score + (damp * new_rank[v]));
        error[v] = fabs((new_rank[v] - old_rank[v]));
        old_rank[v] = new_rank[v];
        new_rank[v] = ((float) 0); } };
    void pagerank(Graph &g, double *new_rank, double *old_rank, int *out_degree, int max_iter) {
        for (int i = (0); i < (max_iter); i++) {
            parallel_for(int v_iter = 0; v_iter < builtin_getVertices(edges); v_iter++) {
                contrib[v] = (old_rank[v] / out_degree[v]);};
            edgeset_apply_pull_parallel(edges, updateEdge());
            parallel_for(int v_iter = 0; v_iter < builtin_getVertices(edges); v_iter++) {
                updateVertex()(v_iter); } };
```

Optimized PageRank for Multi-Core CPU

Graph Computations have Variety


Data

Social Networks, Web
Graphs, Road Networks,
Engineering Meshes,
Transaction Graphs, Network
Traffic Graphs, Email Networks,
Similarity Graphs, ...



Algorithms

Breadth-first search, betweenness
centrality, Bellman-Ford, Delta-stepping,
collaborative filtering, Page Rank, Page Rank
Delta, connected components, k-core
decomposition, triangle counting, local
clustering, structural clustering minimum
spanning forest, eccentricity estimation, graph
coloring, k-truss decomposition, nuclei
decomposition, biconnectivity, set cover,
maximum flow, butterfly counting, strongly
connected components, graph partitioning, RDF
queries, random walks, point-to-point shortest
paths, A* search, low-diameter decomposition,
densest subgraph, multi-source BFS, maximal
independent set, maximal matching, etc...



Hardware

CPU, GPU, KNL, Distributed
Environment, FPGA,
HammerBlade, Symphony,...

Outline

Hardware Utilization

Making Caches Work
for Graph Analytics
(BigData17)
Zhang, et al.

Programming System to Handle Variety in Data and Algorithms

GraphIt: a High-
Performance Graph DSL
(OOPSLA18)
Zhang, et al.

Optimizing Ordered
Graph Algorithms with
GraphIt
(CGO2020)
Zhang, et al.

Variety in Hardware

Universal Graph
Framework
(Under Submission)
Brahmakshatriya, Zhang, et al.

Outline

Hardware Utilization

Making Caches Work
for Graph Analytics
(BigData17)
Zhang, et al.

- **Frequency-based Reordering**
- **Cache-aware Partitioning**

Programming System to Handle Variety in Data and Algorithms

GraphIt: a High-Performance Graph DSL
(OOPSLA18)
Zhang, et al.

Optimizing Ordered Graph Algorithms with GraphIt
(CGO2020)
Zhang, et al.

Variety in Hardware

Universal Graph Framework
(Under Submission)
Brahmakshatriya, Zhang, et al.

Outline

Hardware Utilization

Making Caches Work for Graph Analytics (BigData17)
Zhang, et al.

- **Frequency-based Reordering**
- **Cache-aware Partitioning**

Programming System to Handle Variety in Data and Algorithms

GraphIt: a High-Performance Graph DSL (OOPSLA18)
Zhang, et al.

- **GraphIt Compiler and DSL that Decouples**
 - **Algorithm**
 - **Optimization**
 - **Hardware**
- **for Graph Applications**

Optimizing Ordered Graph Algorithms with GraphIt (CGO2020)
Zhang, et al.

Variety in Hardware

Universal Graph Framework (Under Submission)
Brahmakshatriya, Zhang, et al.

Outline

Hardware Utilization

Making Caches Work for Graph Analytics (BigData17)
Zhang, et al.

- **Frequency-based Reordering**
- **Cache-aware Partitioning**

Programming System to Handle Variety in Data and Algorithms

GraphIt: a High-Performance Graph DSL (OOPSLA18)
Zhang, et al.

- GraphIt Compiler and DSL that Decouples**
- **Algorithm**
 - **Optimization**
 - **Hardware**
- for Graph Applications**

Optimizing Ordered Graph Algorithms with GraphIt (CGO2020)
Zhang, et al.

Variety in Hardware

Universal Graph Framework (Under Submission)
Brahmakshatriya, Zhang, et al.

PageRank

```
while ...  
  for node : graph.vertices  
    for ngh : graph.getInNeighbors(node)  
      newRanks[node] += ranks[ngh]/outDegree[ngh];  
  for node : graph.vertices  
    newRanks[node] = baseScore + damping*newRanks[node];  
swap ranks and newRanks
```


PageRank

```
while ...  
    for node : graph.vertices  
        for ngh : graph.getInNeighbors(node)  
            newRanks[node] += ranks[ngh]/outDegree[ngh];  
    for node : graph.vertices  
        newRanks[node] = baseScore + damping*newRanks[node];  
    swap ranks and newRanks
```

PageRank

```
while ...  
  for node : graph.vertices  
    for ngh : graph.getInNeighbors(node)  
      newRanks[node] += ranks[ngh]/outDegree[ngh];  
  for node : graph.vertices  
    newRanks[node] = baseScore + damping*newRanks[node];  
  swap ranks and newRanks
```

PageRank

```
while ...  
  for node : graph.vertices  
    for ngh : graph.getInNeighbors(node)  
      newRanks[node] += ranks[ngh]/outDegree[ngh];  
  for node : graph.vertices  
    newRanks[node] = baseScore + damping*newRanks[node];  
swap ranks and newRanks
```

PageRank

while ...

for node : graph.vertices

for ngh : graph.getInNeighbors(node)

newRanks[node] += ranks[ngh]/outDegree[ngh];

for node : graph.vertices

newRanks[node] = baseScore + damping*newRanks[node];

swap ranks and newRanks

PageRank

```
while ...  
    for node : graph.vertices  
        for ngh : graph.getInNeighbors(node)  
            newRanks[node] += ranks[ngh]/outDegree[ngh];  
    for node : graph.vertices  
        newRanks[node] = baseScore + damping*newRanks[node];  
    swap ranks and newRanks
```

PageRank

while ...

for node : graph.vertices


for ngh : graph.getInNeighbors(node)

newRanks[node] += ranks[ngh]/outDegree[ngh];

for node : graph.vertices

newRanks[node] = baseScore + damping*newRanks[node];

swap ranks and newRanks



PageRank

while ...

for node : graph.vertices

for ngh : graph.getInNeighbors(node)

newRanks[node] += ranks[ngh]/outDegree[ngh];

for node : graph.vertices

newRanks[node] = baseScore + damping*newRanks[node];

swap ranks and newRanks


Compressed Sparse Row (CSR)

Vertex
Array

0	1	3	5	7
---	---	---	---	---

Edge
Array

2	0	2	0	3	0	2
---	---	---	---	---	---	---



PageRank

while ...

for node : graph.vertices

for ngh : graph.getInNeighbors(node)


newRanks[node] += ranks[ngh]/outDegree[ngh];

for node : graph.vertices


newRanks[node] = baseScore + damping*newRanks[node];

swap ranks and newRanks

Compressed Sparse Row (CSR)



Vertex Array stores indices into the Edge Array. Edge Array stores neighbors' ID in the CSR



PageRank

while ...

for node : graph.vertices

for ngh : graph.getInNeighbors(node)

`newRanks[node]` += ranks[ngh]/outDegree[ngh];


for node : graph.vertices

newRanks[node] = baseScore + damping*newRanks[node];



swap ranks and newRanks

Sequential access on node when
scanning through vertex array

Vertex
Array



Edge
Array



PageRank

while ...

for node : graph.vertices

for ngh : graph.getInNeighbors(node)


newRanks[node] += ranks[ngh]/outDegree[ngh];

for node : graph.vertices


newRanks[node] = baseScore + damping*newRanks[node];

swap ranks and newRanks


Vertex
Array



Edge
Array



Irregular access on ngh's rank and
outDegree data when scanning
through the edge array



PageRank

while ...

```
for node : graph.vertices
```

```
  for ngh : graph.getInNeighbors(node)
```

```
    newRanks[node] += ranks[ngh]/outDegree[ngh];
```

```
for node : graph.vertices
```

```
  newRanks[node] = baseScore + damping*newRanks[node];
```


```
swap ranks and newRanks
```

Cache



#hits: 0

#misses: 0



PageRank

while ...

for node : graph.vertices

for ngh : graph.getInNeighbors(node)

newRanks[node] += ranks[ngh]/outDegree[ngh];

for node : graph.vertices

newRanks[node] = baseScore + damping*newRanks[node];

swap ranks and newRanks


Focus on the random
memory accesses on
ranks array

Cache



#hits: 0

#misses: 0



PageRank

while ...

for node : graph.vertices

for ngh : graph.getInNeighbors(node)


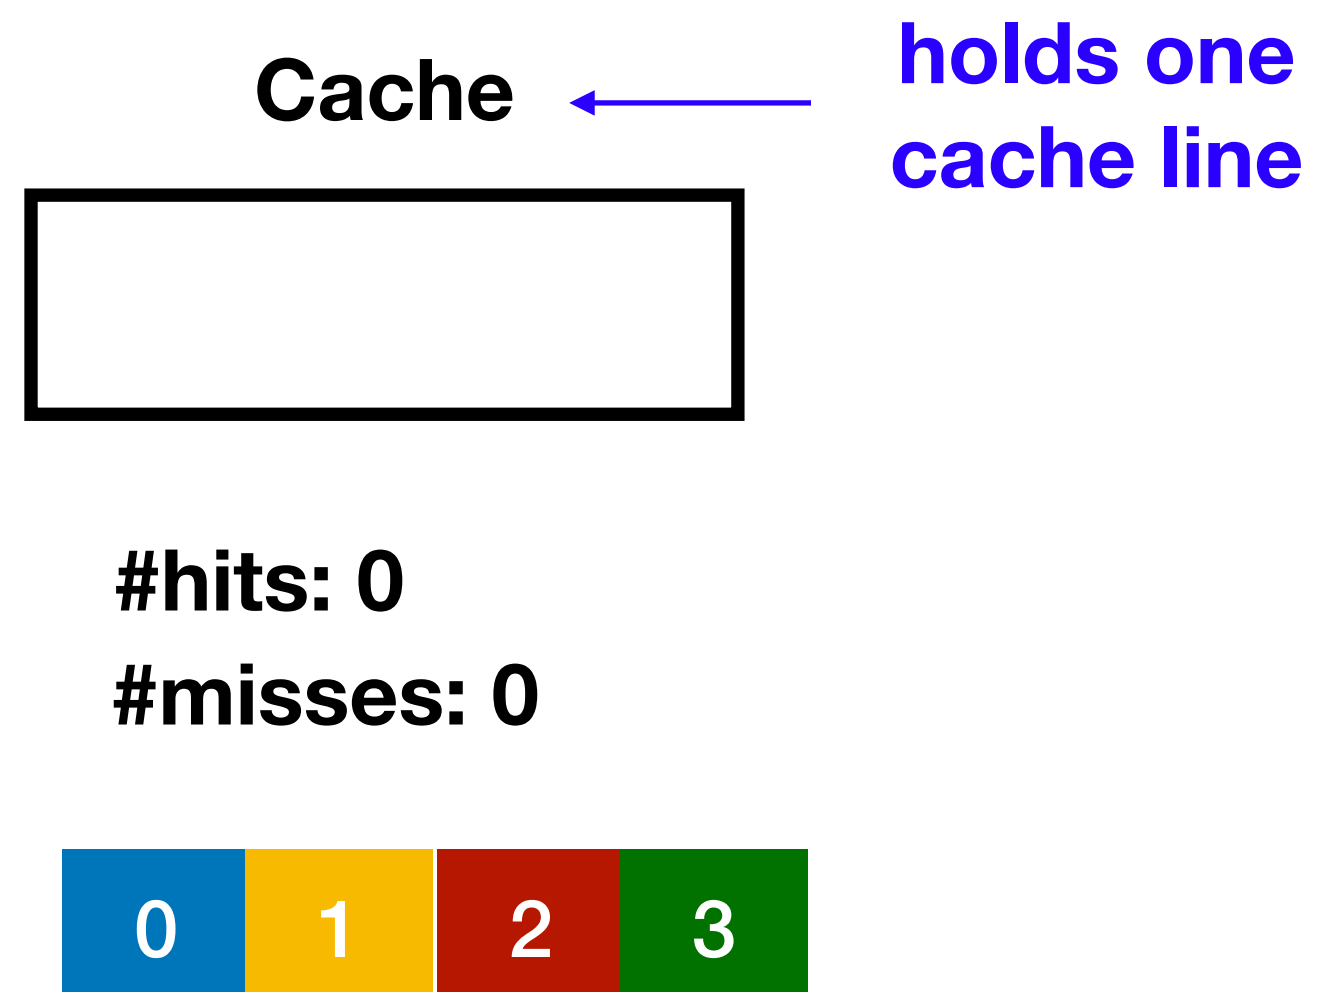
newRanks[node] += ranks[ngh]/outDegree[ngh];

for node : graph.vertices

newRanks[node] = baseScore + damping*newRanks[node];

swap ranks and newRanks

Focus on the random
memory accesses on
ranks array



PageRank

while ...

for node : graph.vertices

for ngh : graph.getInNeighbors(node)

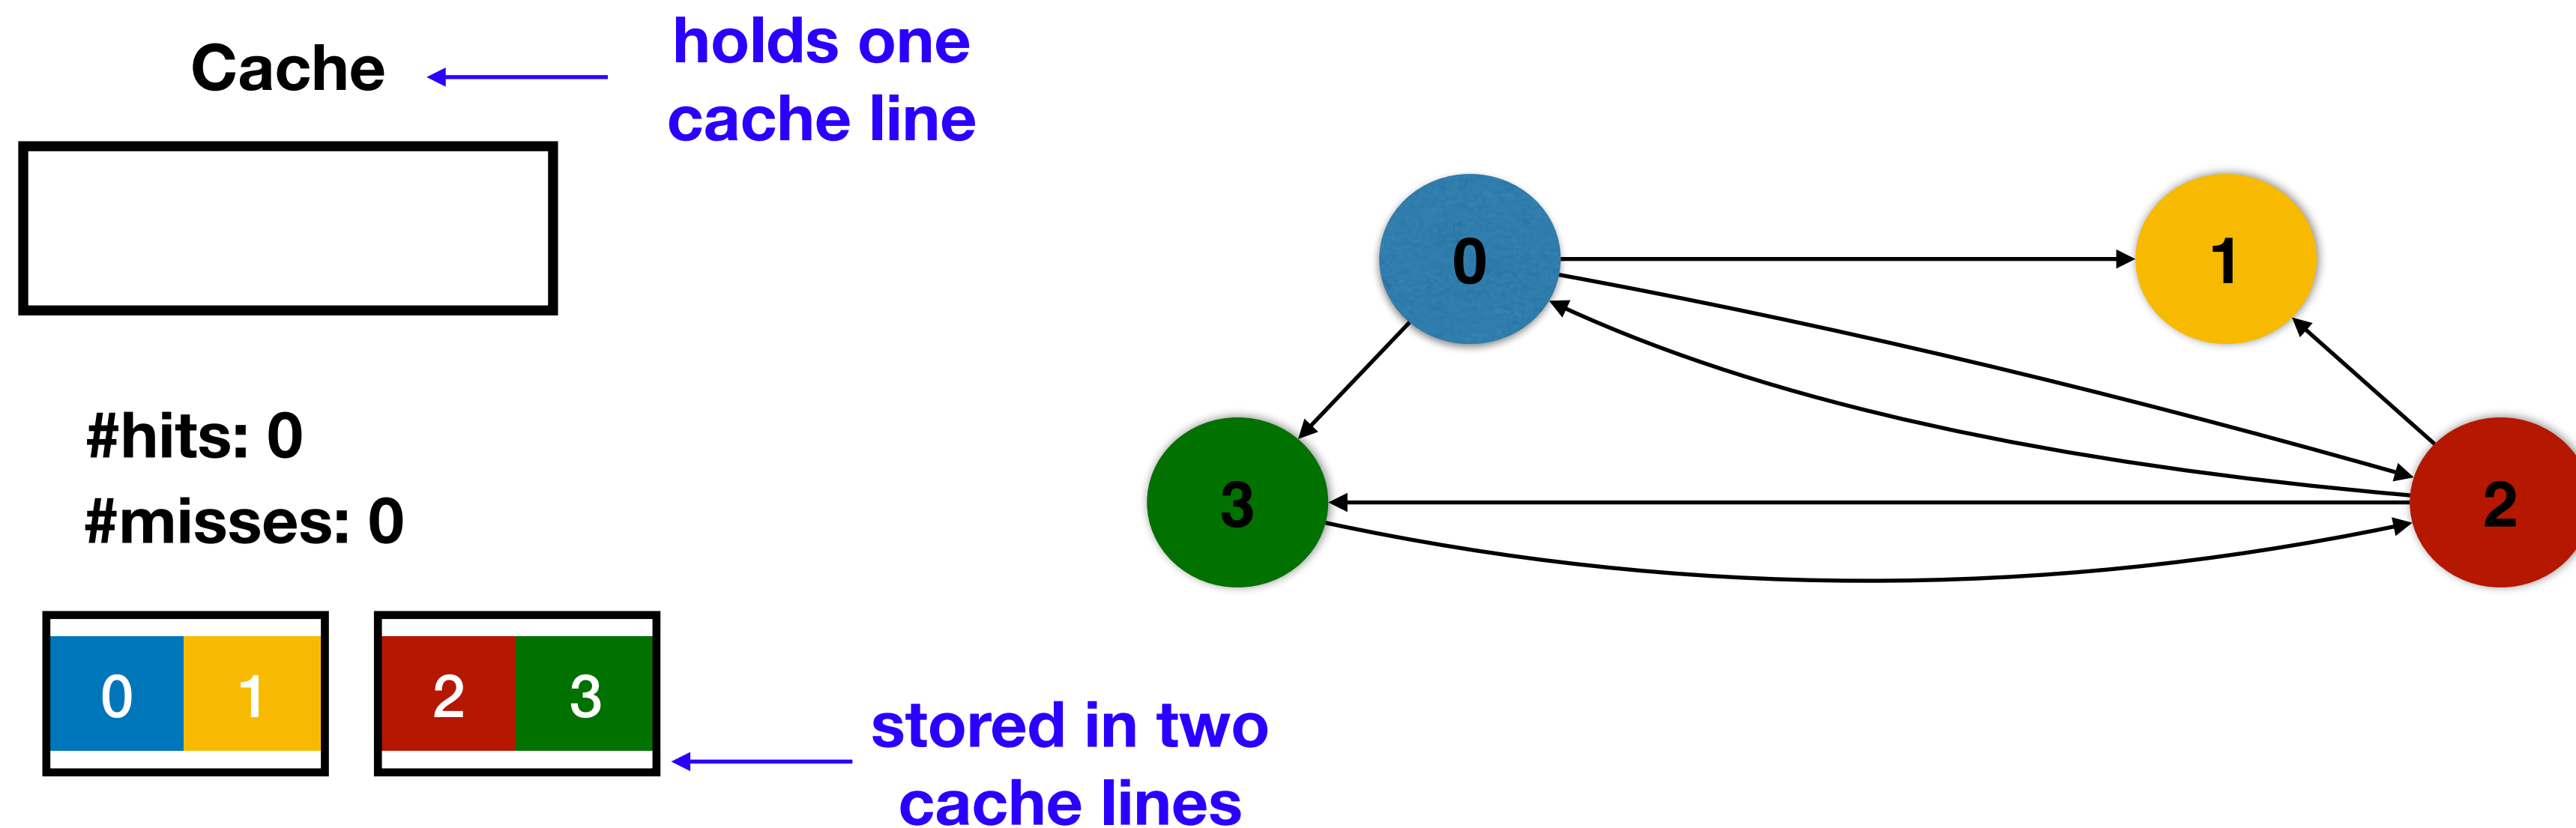
newRanks[node] += ranks[ngh]/outDegree[ngh];

for node : graph.vertices

newRanks[node] = baseScore + damping*newRanks[node];

swap ranks and newRanks

Focus on the random
memory accesses on
ranks array



PageRank

while ...

```
for node : graph.vertices
```

```
for ngh : graph.getInNeighbors(node)
```

```
newRanks[node] += ranks[ngh]/outDegree[ngh];
```

```
for node : graph.vertices
```

```
newRanks[node] = baseScore + damping*newRanks[node];
```


```
swap ranks and newRanks
```

Cache



#hits: 0

#misses: 0



PageRank

while ...

```
for node : graph.vertices
```

```
for ngh : graph.getInNeighbors(node)
```

```
newRanks[node] += ranks[ngh]/outDegree[ngh];
```

```
for node : graph.vertices
```

```
newRanks[node] = baseScore + damping*newRanks[node];
```


```
swap ranks and newRanks
```

Cache



#hits: 0

#misses: 0



PageRank

while ...

```
for node : graph.vertices
```

```
for ngh : graph.getInNeighbors(node)
```


```
newRanks[node] += ranks[ngh]/outDegree[ngh];
```

```
for node : graph.vertices
```

```
newRanks[node] = baseScore + damping*newRanks[node];
```


```
swap ranks and newRanks
```

Cache



#hits: 0

#misses: 1



PageRank

while ...

```
for node : graph.vertices
```

```
for ngh : graph.getInNeighbors(node)
```


```
newRanks[node] += ranks[ngh]/outDegree[ngh];
```

```
for node : graph.vertices
```

```
newRanks[node] = baseScore + damping*newRanks[node];
```

```
swap ranks and newRanks
```

Cache



#hits: 0

#misses: 1



PageRank

while ...

```
for node : graph.vertices
```

```
for ngh : graph.getInNeighbors(node)
```


```
newRanks[node] += ranks[ngh]/outDegree[ngh];
```

```
for node : graph.vertices
```

```
newRanks[node] = baseScore + damping*newRanks[node];
```

```
swap ranks and newRanks
```

Cache



#hits: 0

#misses: 1



PageRank

while ...

```
for node : graph.vertices
```

```
for ngh : graph.getInNeighbors(node)
```

```
newRanks[node] += ranks[ngh]/outDegree[ngh];
```

```
for node : graph.vertices
```

```
newRanks[node] = baseScore + damping*newRanks[node];
```


```
swap ranks and newRanks
```

Cache



#hits: 0

#misses: 2



PageRank

while ...

```
for node : graph.vertices
```

```
for ngh : graph.getInNeighbors(node)
```

```
newRanks[node] += ranks[ngh]/outDegree[ngh];
```

```
for node : graph.vertices
```

```
newRanks[node] = baseScore + damping*newRanks[node];
```


```
swap ranks and newRanks
```

Cache



#hits: 0

#misses: 2



PageRank

while ...

```
for node : graph.vertices
```

```
for ngh : graph.getInNeighbors(node)
```

```
newRanks[node] += ranks[ngh]/outDegree[ngh];
```

```
for node : graph.vertices
```

```
newRanks[node] = baseScore + damping*newRanks[node];
```


```
swap ranks and newRanks
```

Cache



#hits: 0

#misses: 2



PageRank

while ...

```
for node : graph.vertices
```

```
for ngh : graph.getInNeighbors(node)
```

```
newRanks[node] += ranks[ngh]/outDegree[ngh];
```


```
for node : graph.vertices
```

```
newRanks[node] = baseScore + damping*newRanks[node];
```

```
swap ranks and newRanks
```




Cache



#hits: 0

#misses: 3



PageRank

while ...

```
for node : graph.vertices
```

```
for ngh : graph.getInNeighbors(node)
```


```
newRanks[node] += ranks[ngh]/outDegree[ngh];
```

```
for node : graph.vertices
```

```
newRanks[node] = baseScore + damping*newRanks[node];
```


```
swap ranks and newRanks
```

Cache



#hits: 0

#misses: 3



PageRank

while ...

```
for node : graph.vertices
```

```
for ngh : graph.getInNeighbors(node)
```

```
newRanks[node] += ranks[ngh]/outDegree[ngh];
```

```
for node : graph.vertices
```

```
newRanks[node] = baseScore + damping*newRanks[node];
```


```
swap ranks and newRanks
```

Cache



#hits: 0

#misses: 4



PageRank

while ...

```
for node : graph.vertices
```

```
for ngh : graph.getInNeighbors(node)
```

```
newRanks[node] += ranks[ngh]/outDegree[ngh];
```

```
for node : graph.vertices
```

```
newRanks[node] = baseScore + damping*newRanks[node];
```

```
swap ranks and newRanks
```

Cache



#hits: 0

#misses: 4



PageRank

while ...

```
for node : graph.vertices
```

```
for ngh : graph.getInNeighbors(node)
```


```
newRanks[node] += ranks[ngh]/outDegree[ngh];
```

```
for node : graph.vertices
```

```
newRanks[node] = baseScore + damping*newRanks[node];
```


```
swap ranks and newRanks
```

Cache



#hits: 0

#misses: 5



PageRank

while ...

```
for node : graph.vertices
```

```
for ngh : graph.getInNeighbors(node)
```


```
newRanks[node] += ranks[ngh]/outDegree[ngh];
```

```
for node : graph.vertices
```

```
newRanks[node] = baseScore + damping*newRanks[node];
```

```
swap ranks and newRanks
```

Cache



#hits: 0

#misses: 5



PageRank

while ...

```
for node : graph.vertices
```

```
for ngh : graph.getInNeighbors(node)
```


```
newRanks[node] += ranks[ngh]/outDegree[ngh];
```

```
for node : graph.vertices
```

```
newRanks[node] = baseScore + damping*newRanks[node];
```

```
swap ranks and newRanks
```

Cache



#hits: 1

#misses: 5



PageRank

while ...

```
for node : graph.vertices
```

```
for ngh : graph.getInNeighbors(node)
```


```
newRanks[node] += ranks[ngh]/outDegree[ngh];
```

```
for node : graph.vertices
```

```
newRanks[node] = baseScore + damping*newRanks[node];
```

```
swap ranks and newRanks
```

Cache



#hits: 1

#misses: 5



PageRank

while ...

```
for node : graph.vertices
```

```
for ngh : graph.getInNeighbors(node)
```

```
newRanks[node] += ranks[ngh]/outDegree[ngh];
```

```
for node : graph.vertices
```

```
newRanks[node] = baseScore + damping*newRanks[node];
```


```
swap ranks and newRanks
```

Cache



#hits: 1

#misses: 6



PageRank

while ...

```
for node : graph.vertices
```

```
for ngh : graph.getInNeighbors(node)
```

```
newRanks[node] += ranks[ngh]/outDegree[ngh];
```

```
for node : graph.vertices
```

```
newRanks[node] = baseScore + damping*newRanks[node];
```


```
swap ranks and newRanks
```

Cache



#hits: 1
#misses: 6

A very high
miss rate



PageRank

while ...

```
for node : graph.vertices
```

```
for ngh : graph.getInNeighbors(node)
```

```
newRanks[node] += ranks[ngh]/outDegree[ngh];
```

```
for node : graph.vertices
```

```
newRanks[node] = baseScore + damping*newRanks[node];
```



```
swap ranks and newRanks
```

Cache



#hits: 1
#misses: 6

Working set
larger than
cache



PageRank

while ...

```
for node : graph.vertices
```

```
for ngh : graph.getInNeighbors(node)
```

```
newRanks[node] += ranks[ngh]/outDegree[ngh];
```


```
for node : graph.vertices
```

```
newRanks[node] = baseScore + damping*newRanks[node];
```

```
swap ranks and newRanks
```


Often only use
part of the
cache line

Cache



#hits: 1

#misses: 6



Performance Bottleneck

- Working set much larger than cache size
- Access pattern is irregular
 - Often uses part of the cache line
 - Hard to benefit from hardware prefetching
 - TLB miss, DRAM row miss (hundreds of cycles)

Performance Bottleneck

**Real-world graphs often have working set
10-200x larger than cache size**

- Working set much larger than cache size
- Access pattern is irregular
 - Often uses part of the cache line
 - Hard to benefit from hardware prefetching
 - TLB miss, DRAM row miss (hundreds of cycles)

Performance Bottleneck


- Working set much larger than cache size
- Access pattern is irregular
 - Often uses part of the cache line
 - Hard to benefit from hardware prefetching
 - TLB miss, DRAM row miss (hundreds of cycles)

Performance Bottleneck

- Working set much larger than cache size
- Access pattern is irregular **Often only use 1/16 - 1/8 of a cache line in modern hardware**
- Often uses part of the cache line
- Hard to benefit from hardware prefetching
- TLB miss, DRAM row miss (hundreds of cycles)


PageRank

```
while ...  
  for node : graph.vertices  
    for ngh : graph.getInNeighbors(node)  
      newRanks[node] += ranks[ngh]/outDegree[ngh];  
  for node : graph.vertices  
    newRanks[node] = baseScore + damping*newRanks[node];  
  swap ranks and newRanks
```



PageRank

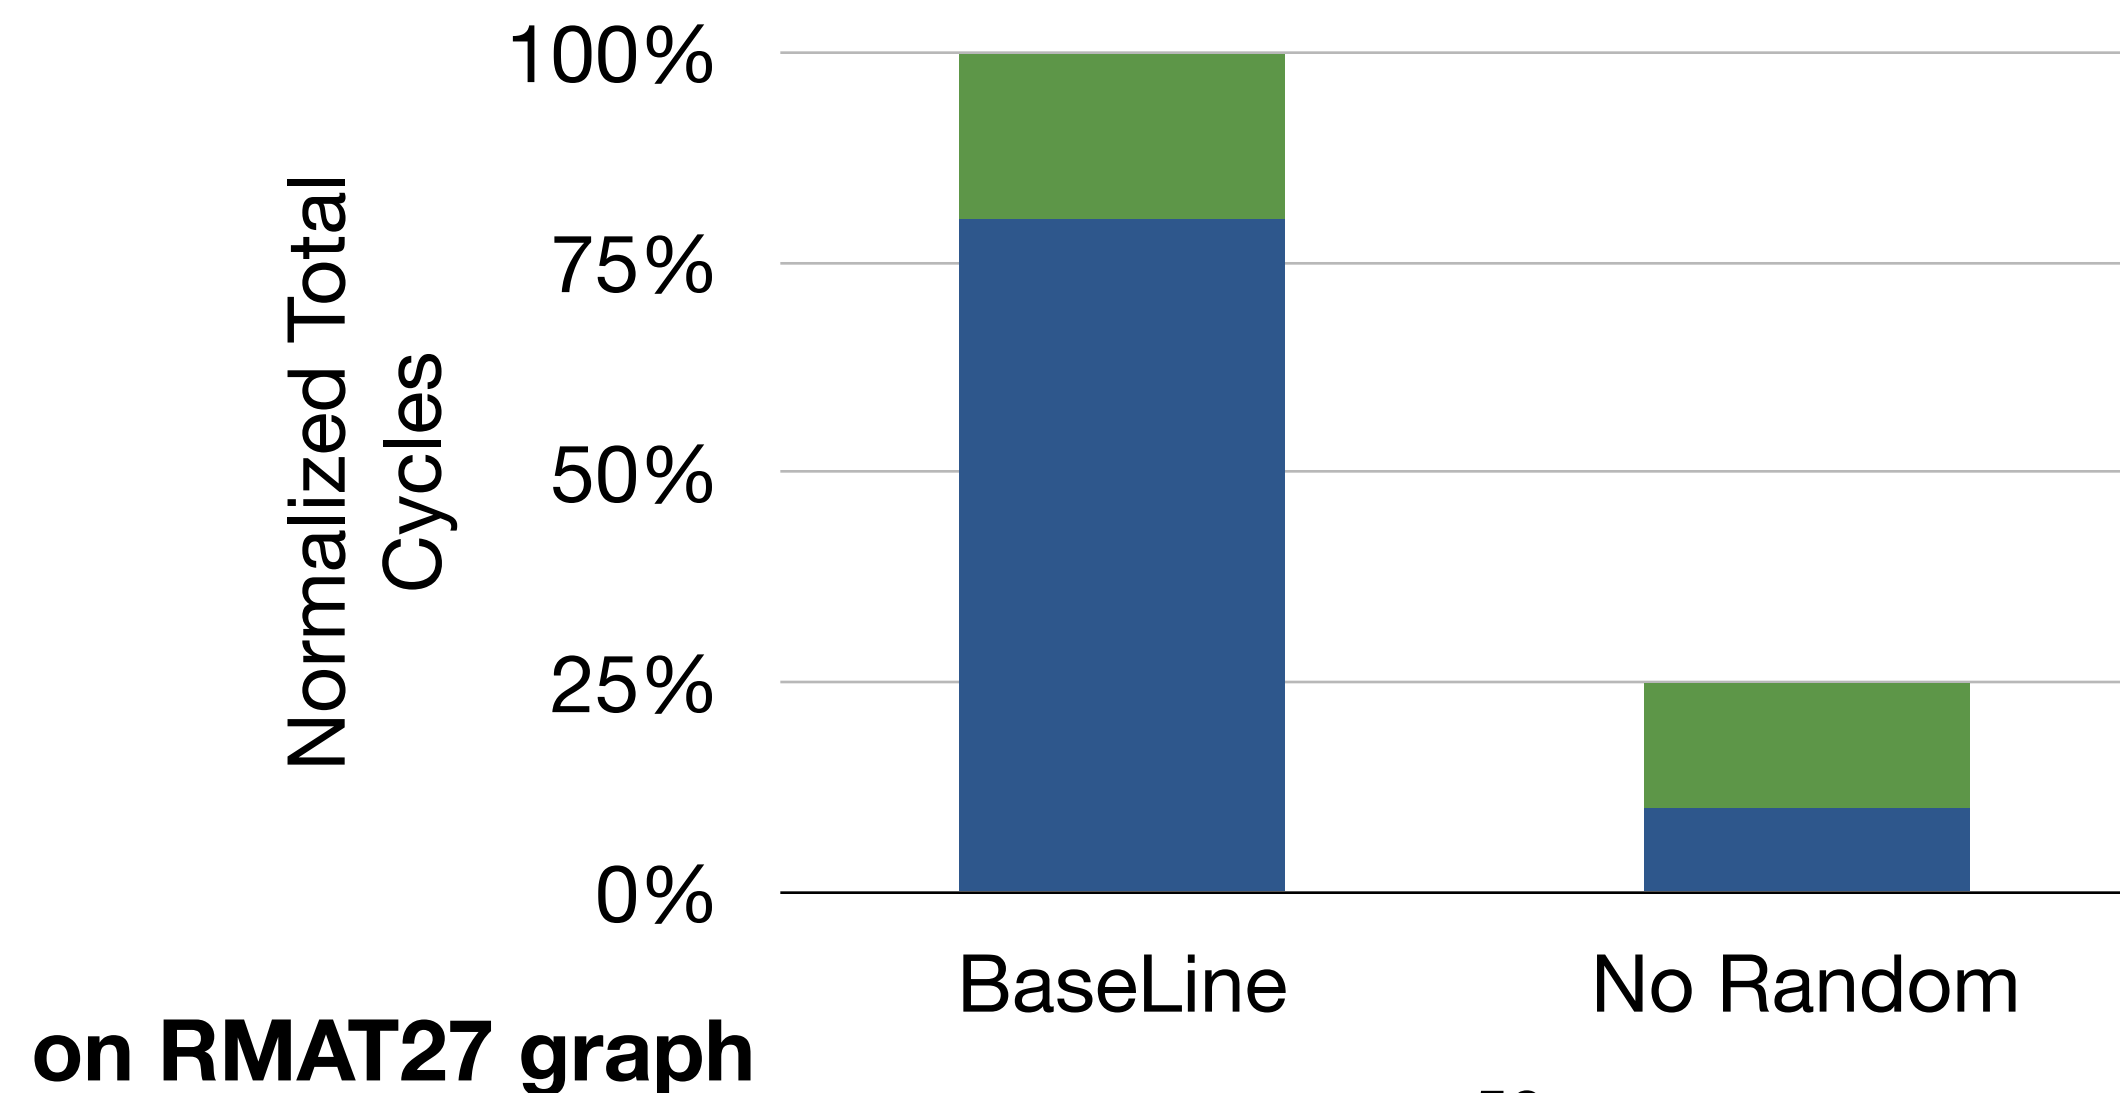
```
while ...  
  for node : graph.vertices  
    for ngh : graph.getInNeighbors(node)  
      newRanks[node] += ranks[ngh]/outDegree[ngh];  
  for node : graph.vertices  
    newRanks[node] = baseScore + damping*newRanks[node];  
  swap ranks and newRanks
```



PageRank

```
while ...  
  for node : graph.vertices  
    for ngh : graph.getInNeighbors(node)  
      newRanks[node] += ranks[0]/outDegree[0];  
  for node : graph.vertices  
    newRanks[node] = baseScore + damping*newRanks[node];  
  swap ranks and newRanks
```


Removing
Random Accesses
(Incorrect)



PageRank


```
while ...  
  for node : graph.vertices  
    for ngh : graph.getInNeighbors(node)  
      newRanks[node] += ranks[0]/outDegree[0];  
  for node : graph.vertices  
    newRanks[node] = baseScore + damping*newRanks[node];  
  swap ranks and newRanks
```

Removing
Random Accesses
(Incorrect)



PageRank

```
while ...  
  for node : graph.vertices  
    for ngh : graph.getInNeighbors(node)  
      newRanks[node] += ranks[ngh]/outDegree[ngh];  
  for node : graph.vertices  
    newRanks[node] = baseScore + damping*newRanks[node];  
  swap ranks and newRanks
```




Frequency-based Vertex Reordering

- Key Observations
 - Cache lines are underutilized
 - Certain vertices are much more likely to be accessed than other vertices


Frequency-based Vertex Reordering

- Key Observations
 - Cache lines are underutilized
 - Certain vertices are much more likely to be accessed than other vertices
- Design
 - Group together the frequently accessed nodes
 - Keep the ordering of average degree nodes

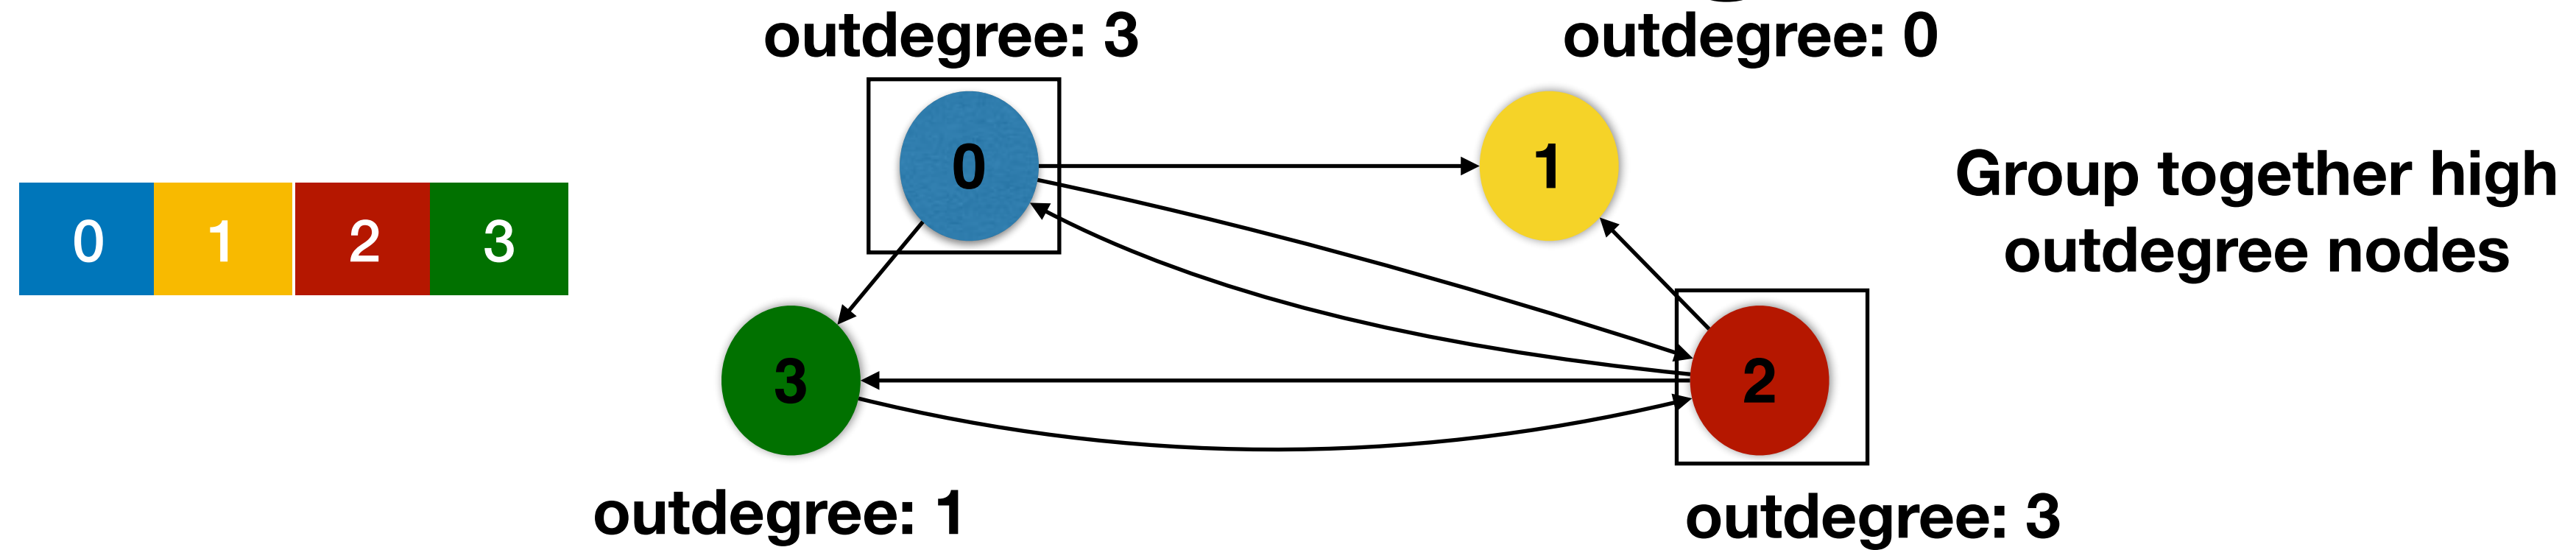
Frequency-based Vertex Reordering



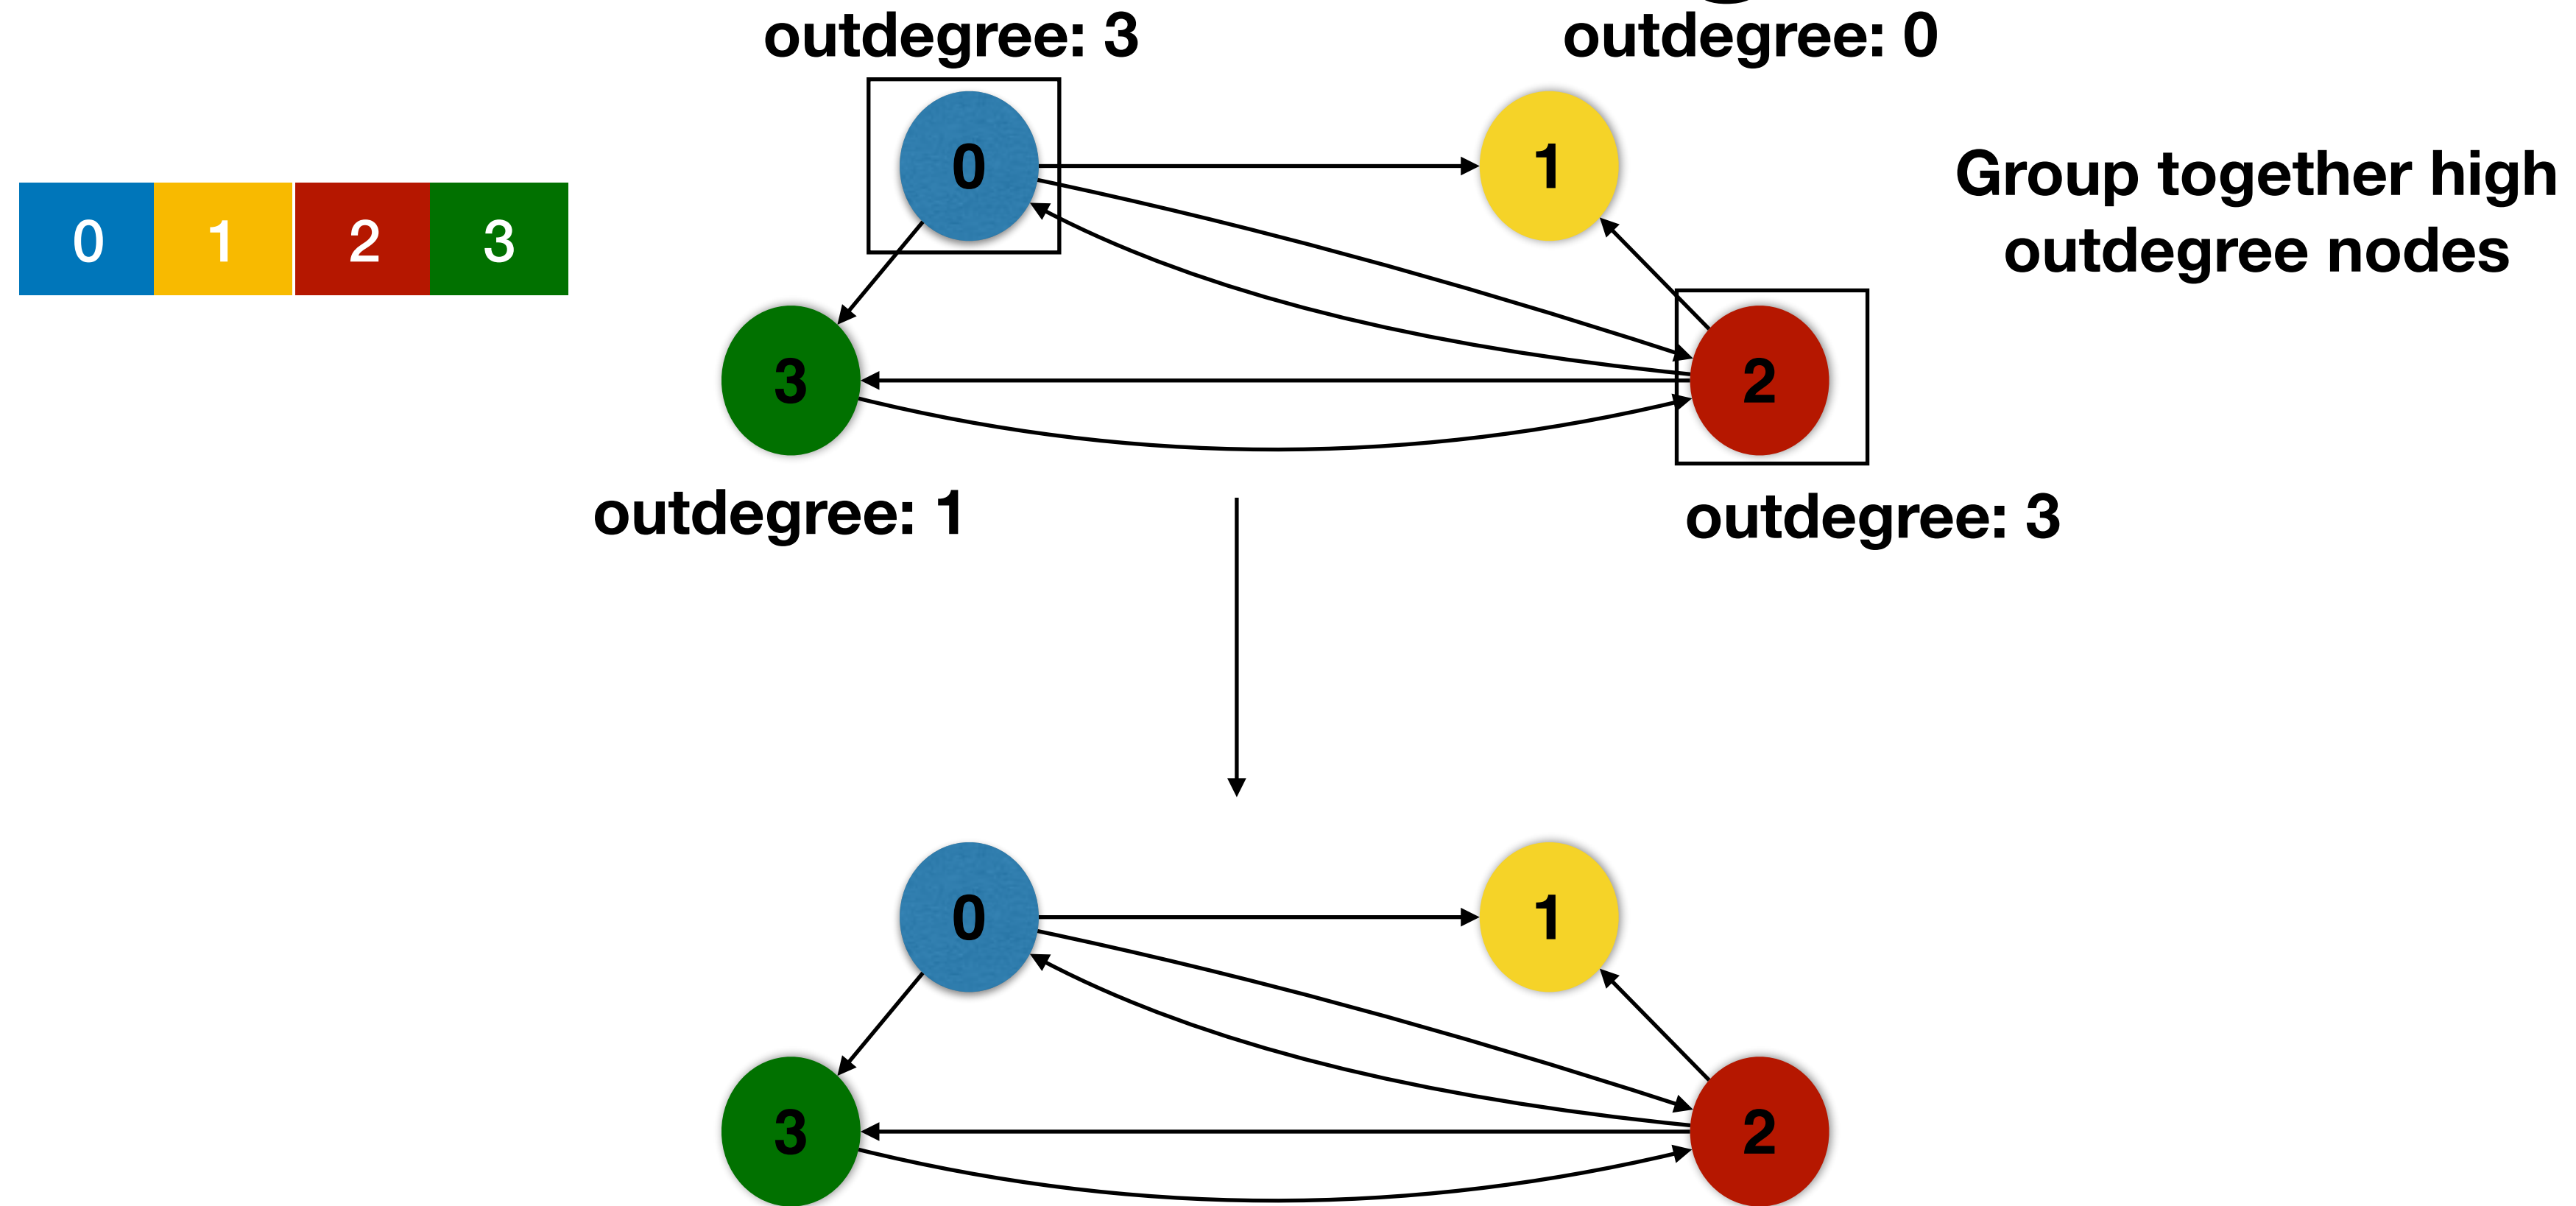
Frequency-based Vertex Reordering



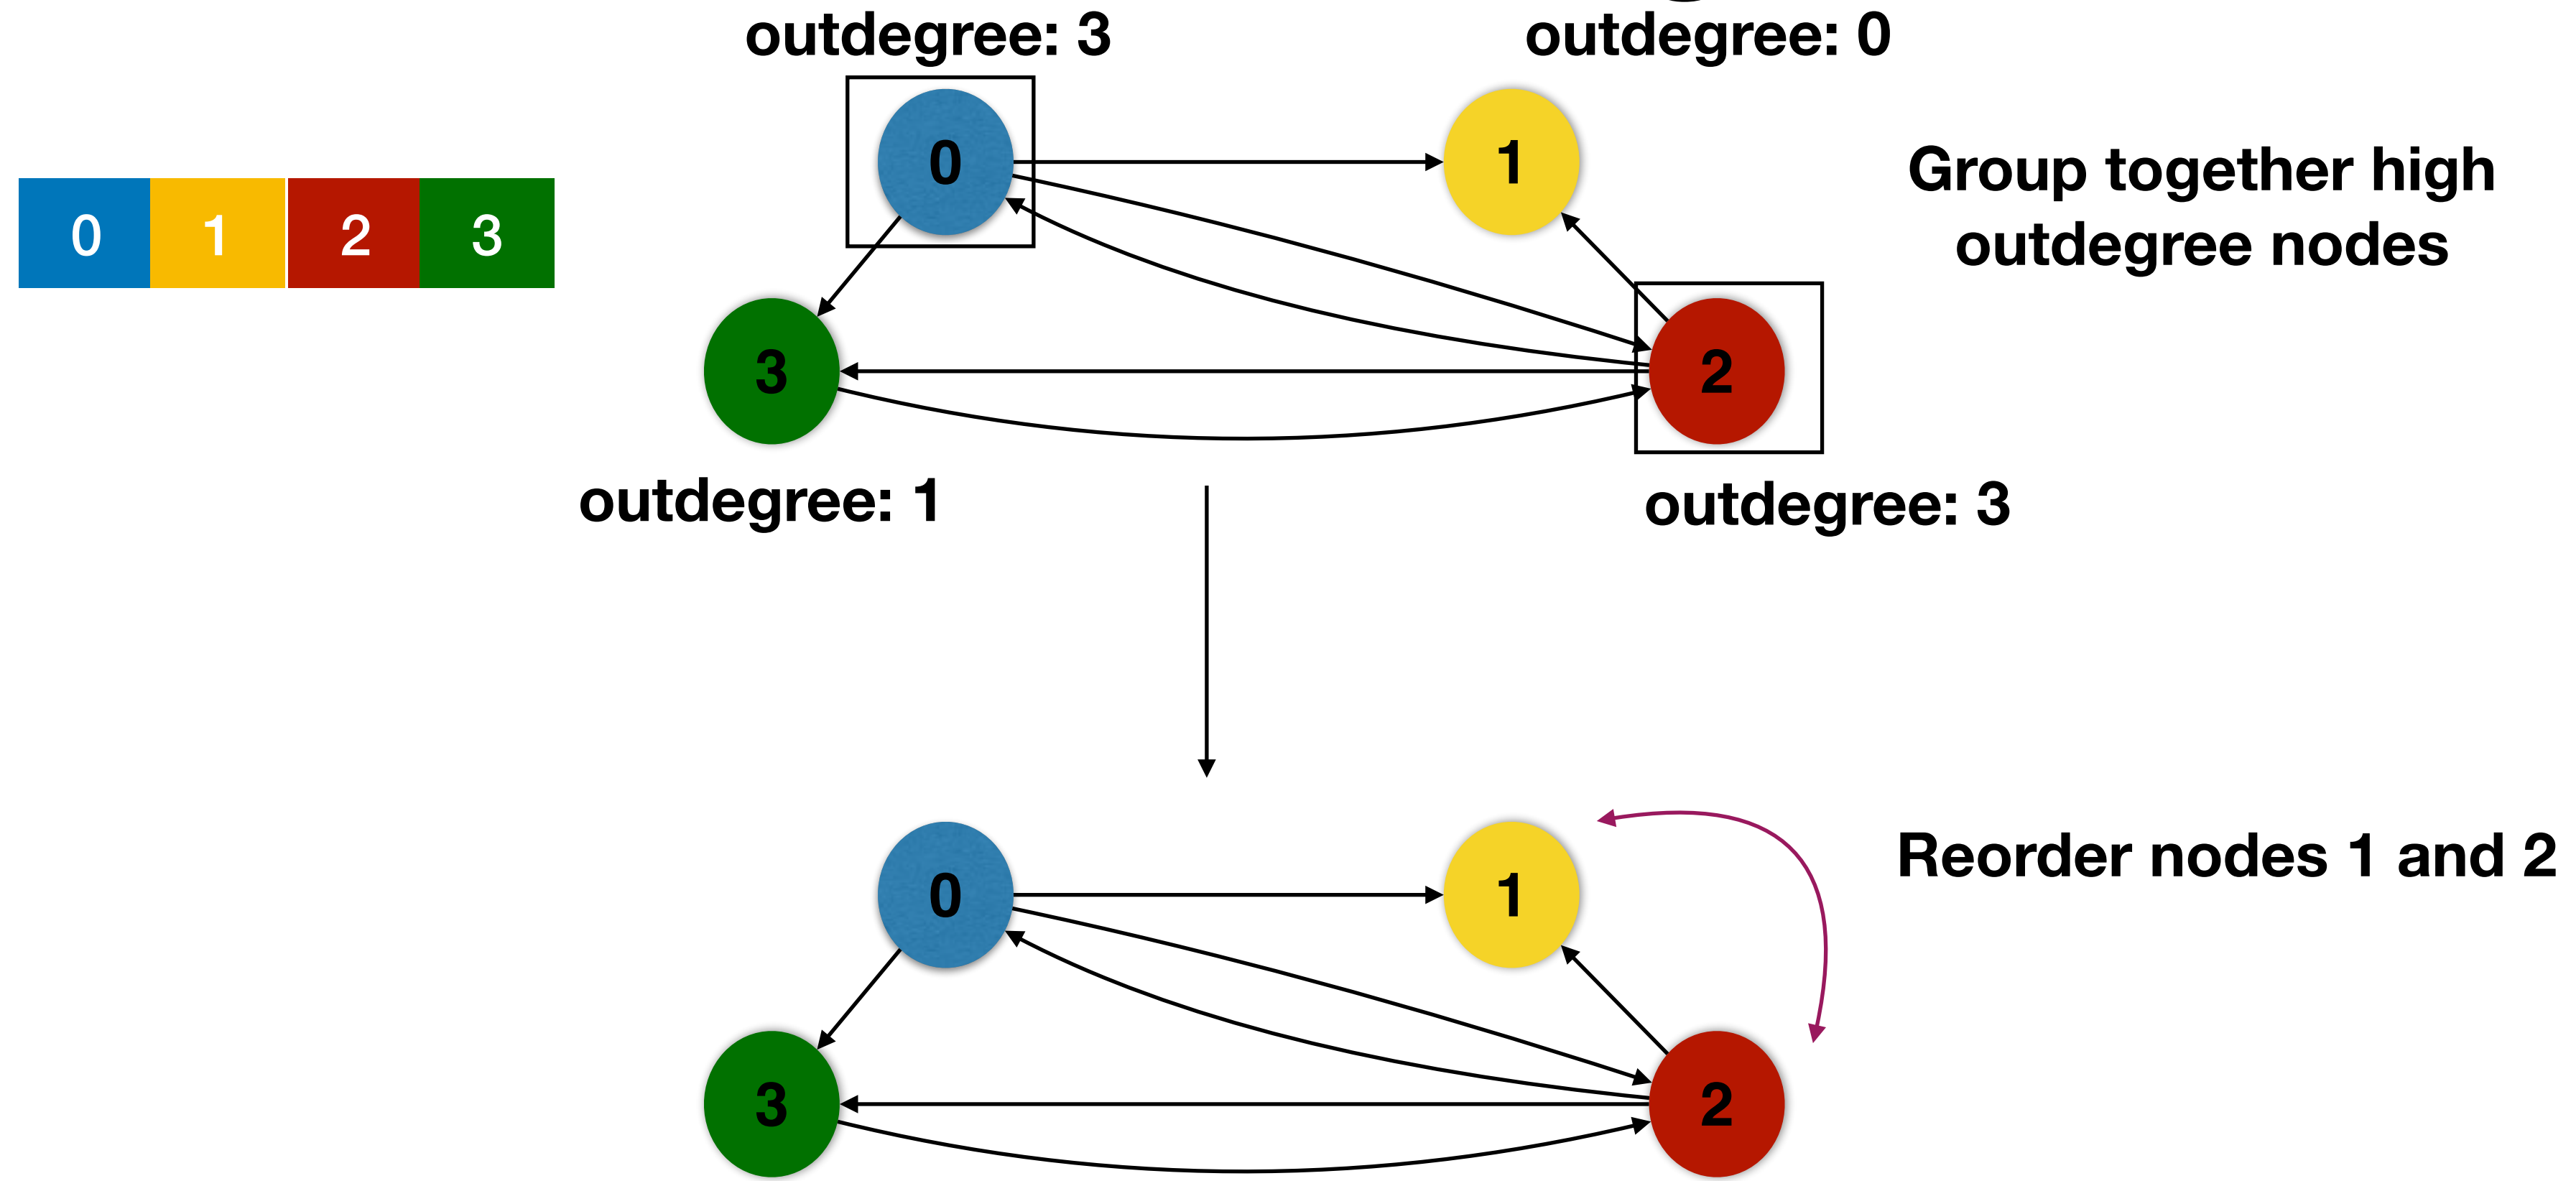
Frequency-based Vertex Reordering



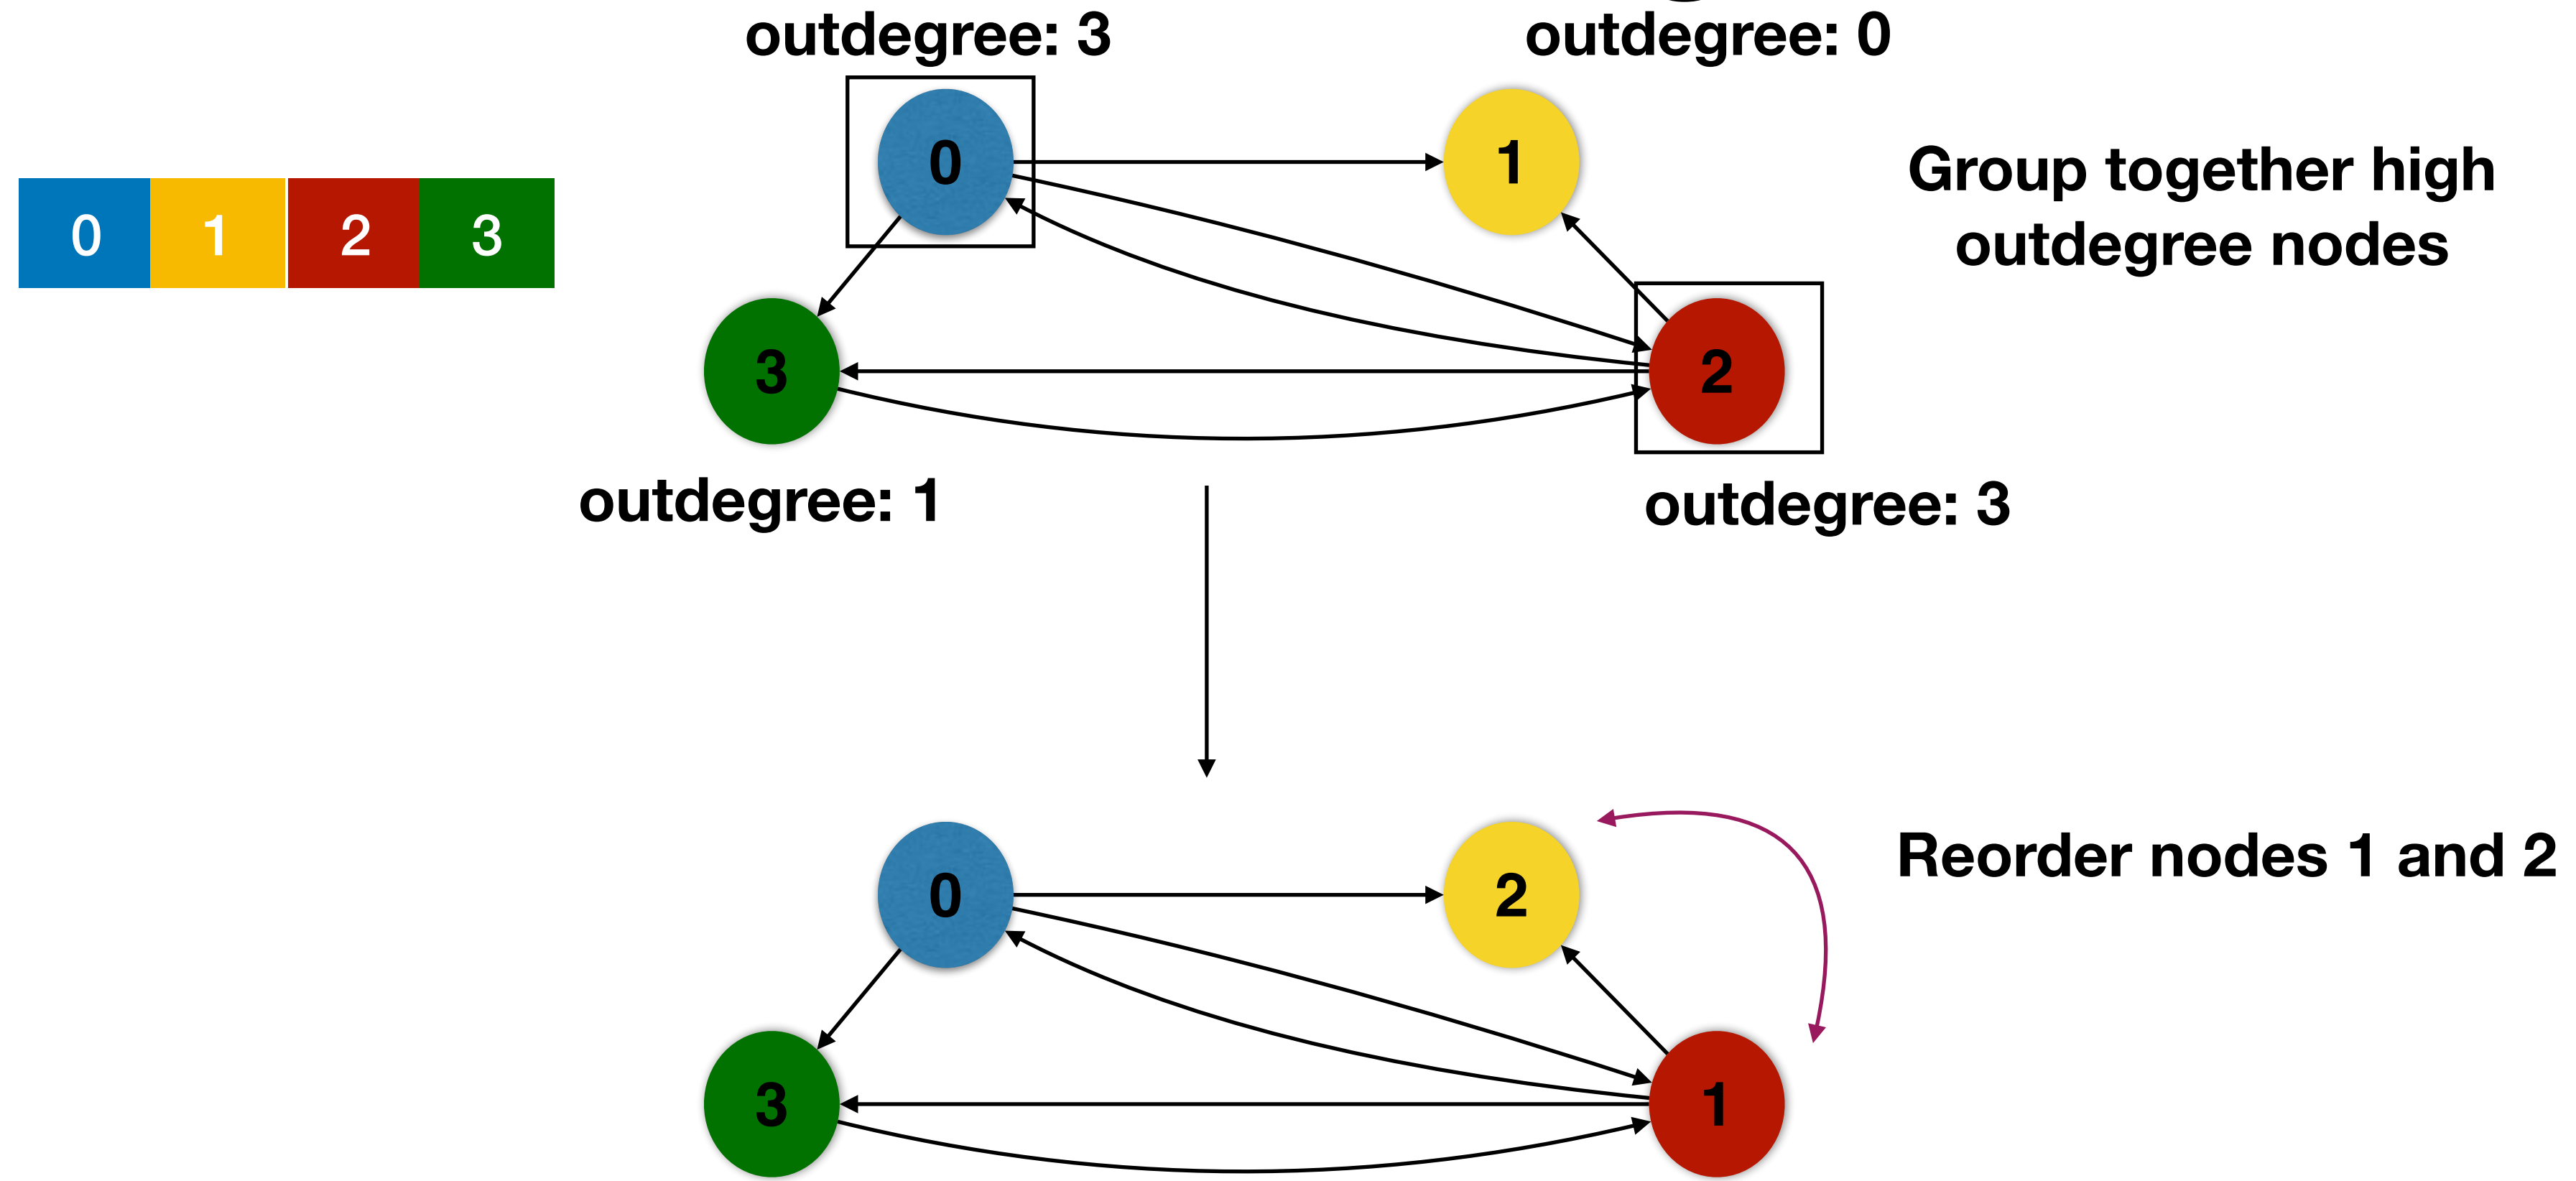
Frequency-based Vertex Reordering



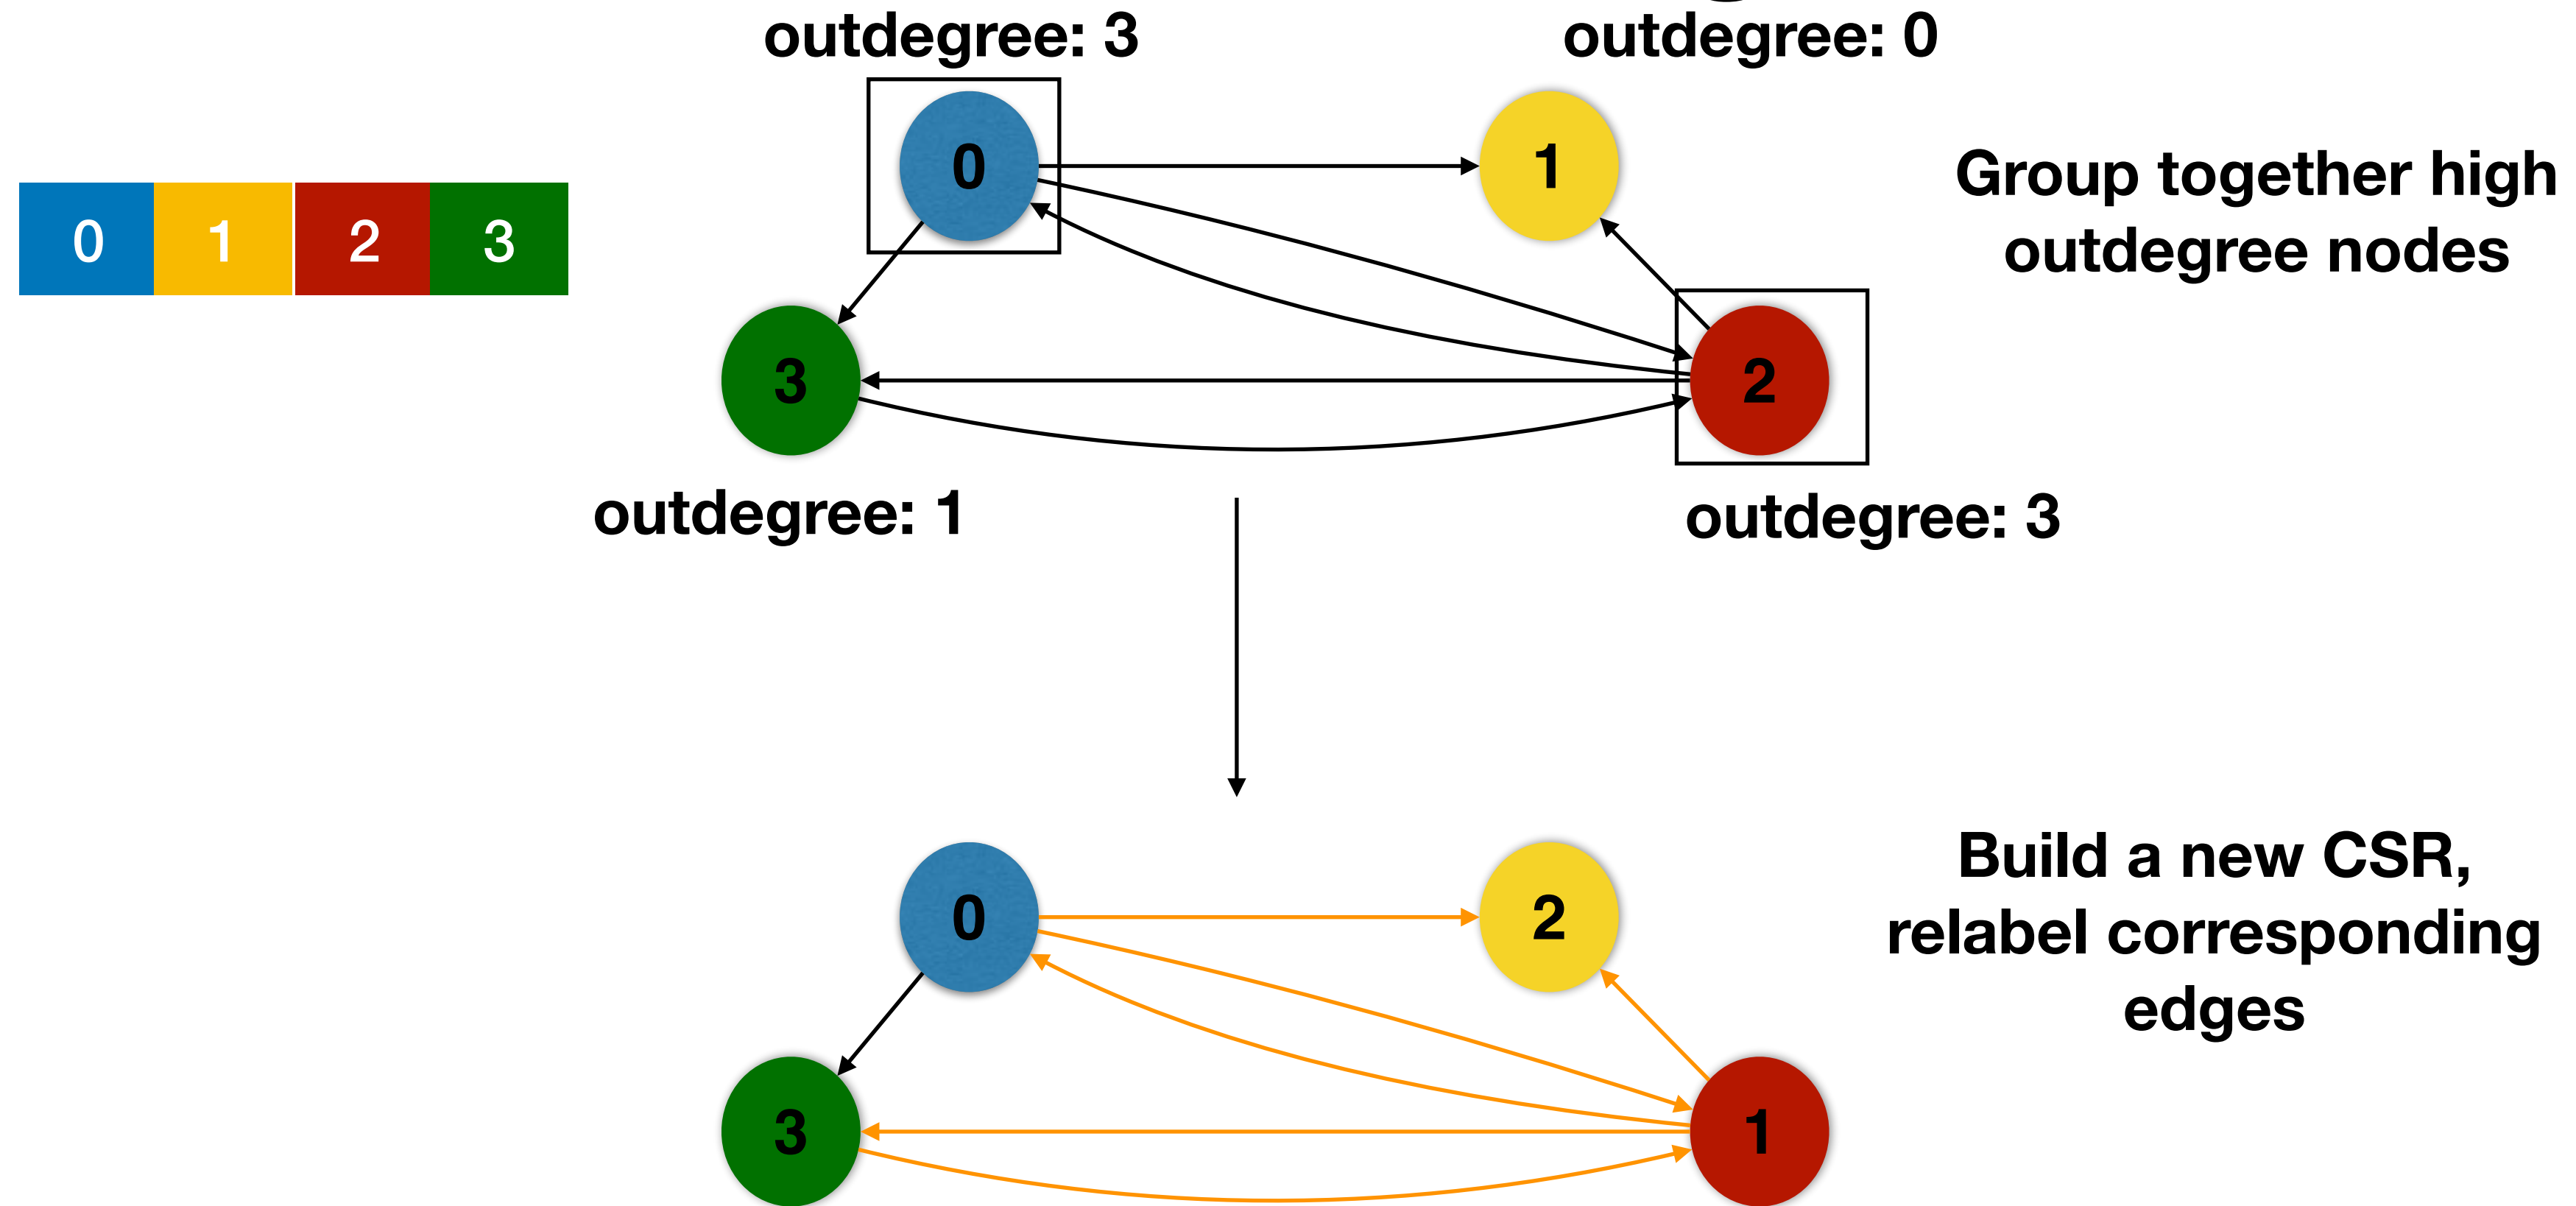
Frequency-based Vertex Reordering



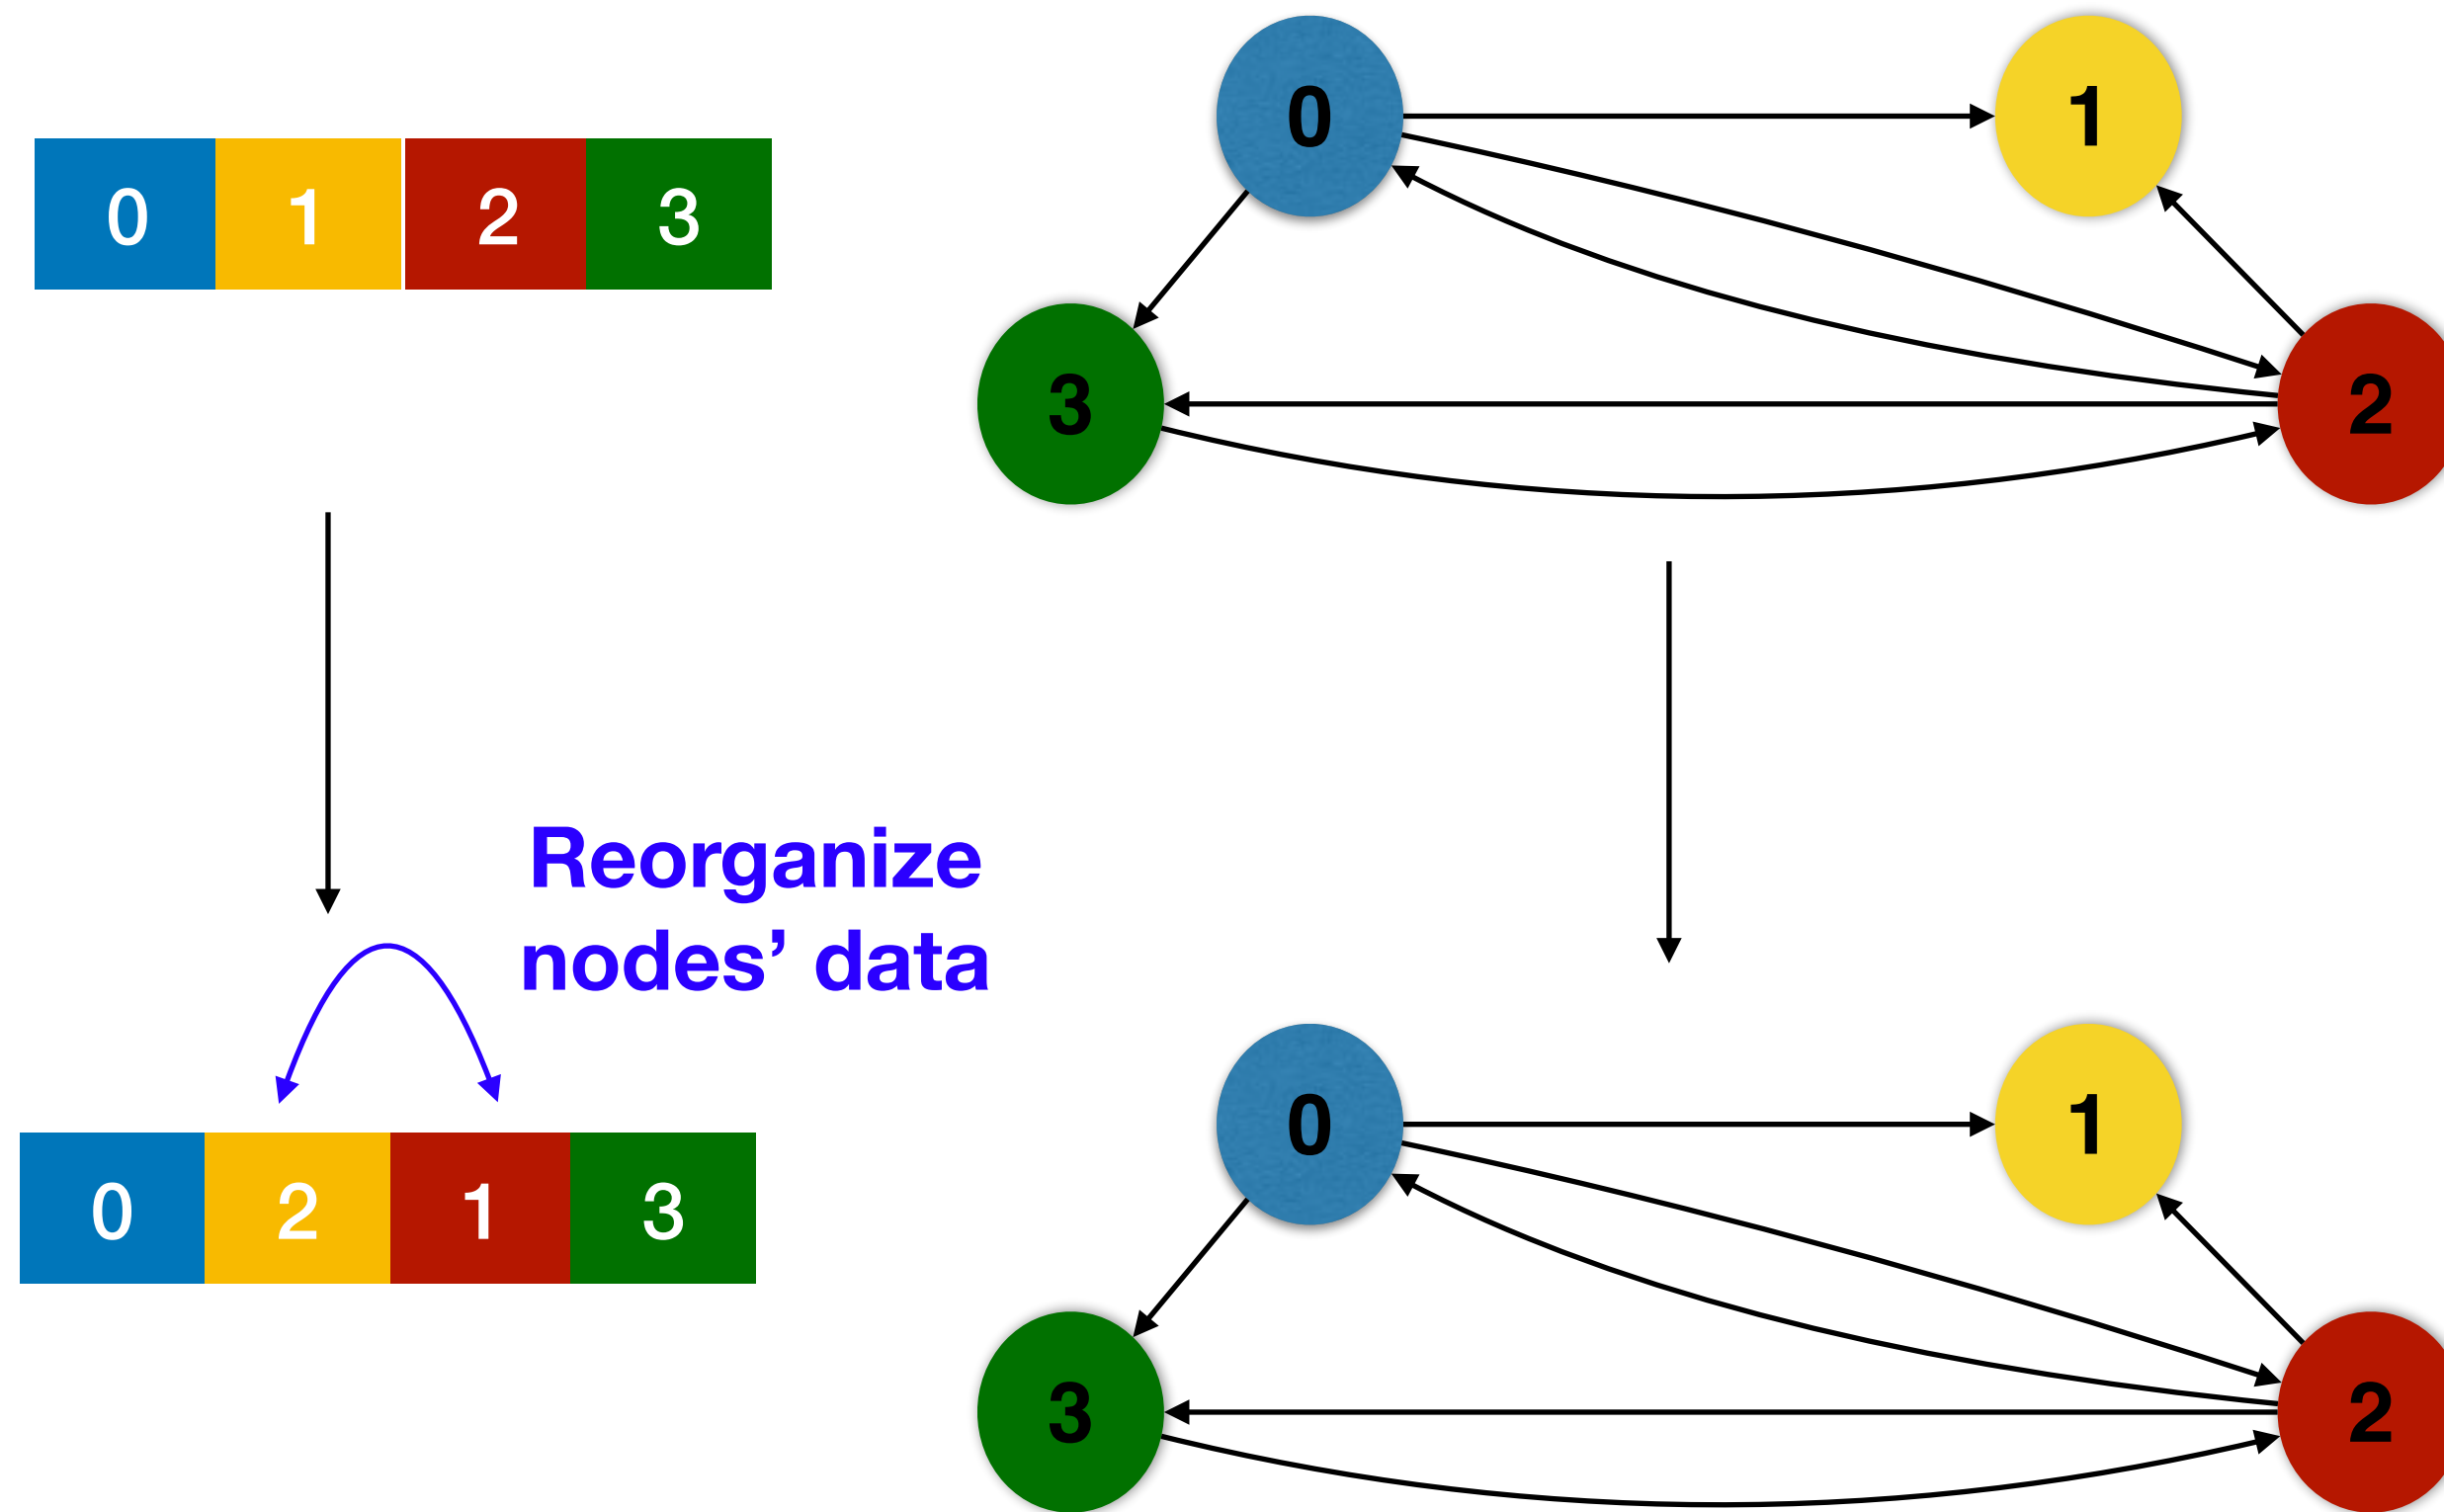
Frequency-based Vertex Reordering



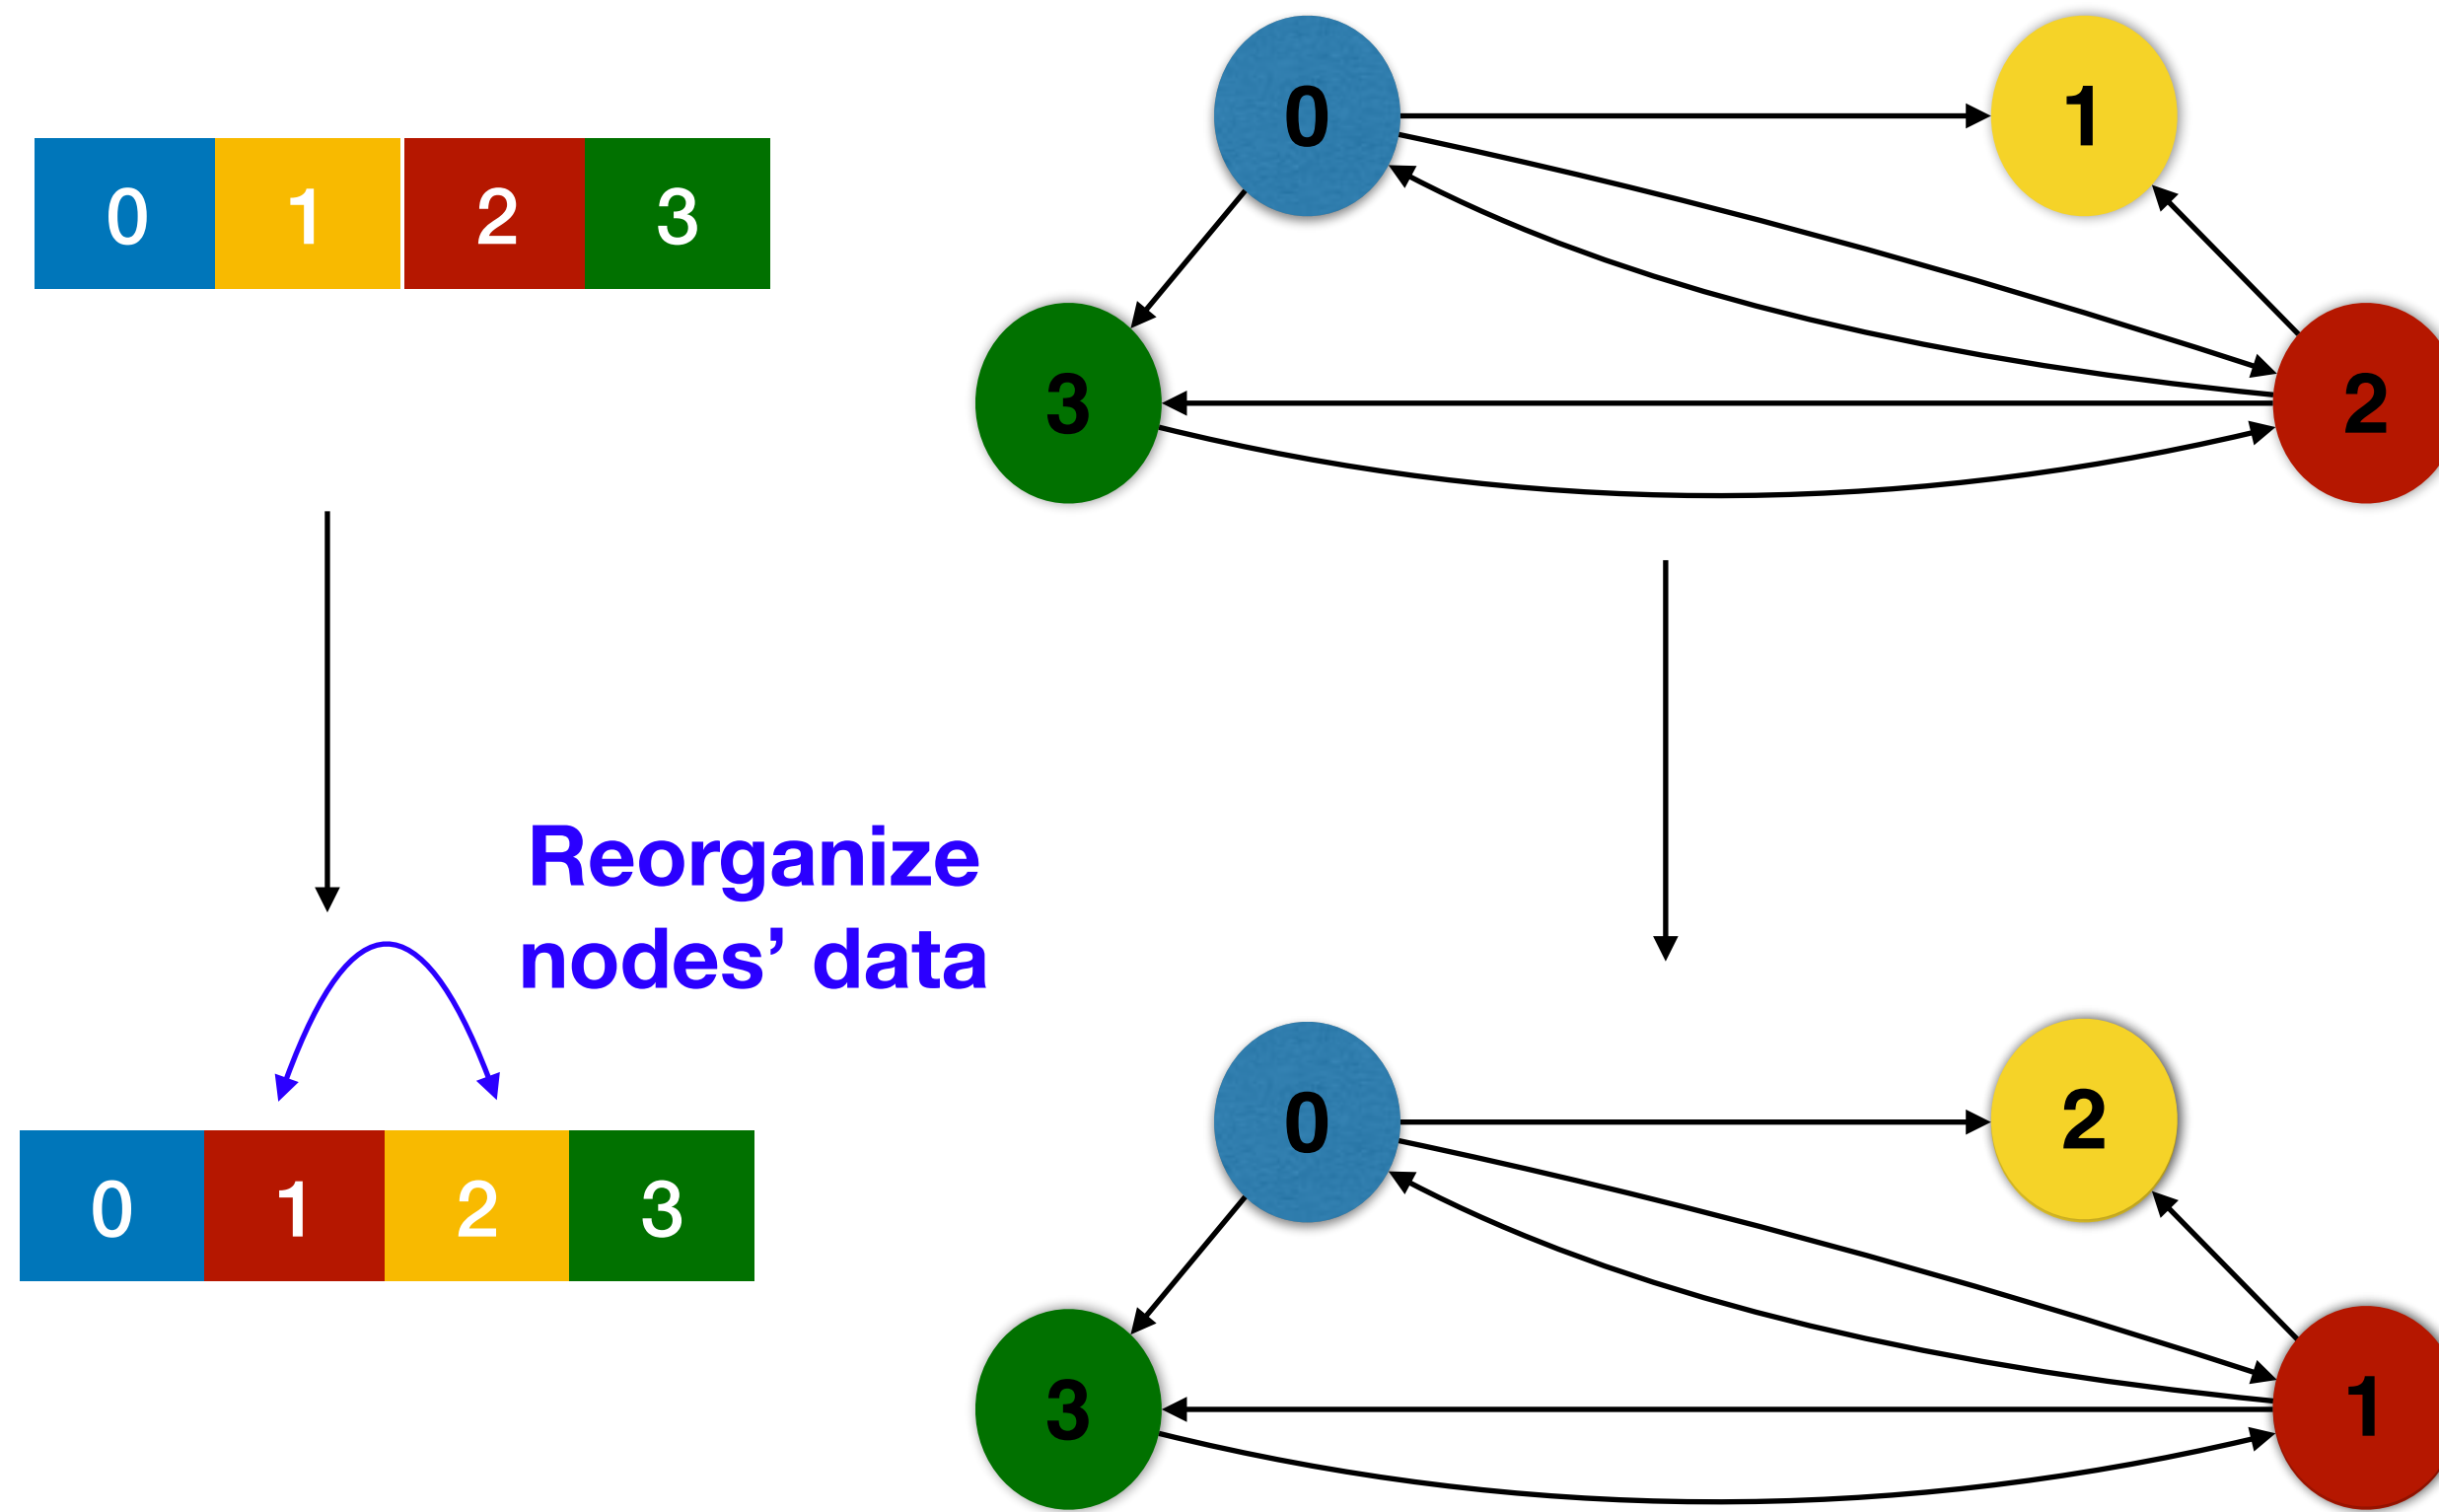
Frequency-based Vertex Reordering



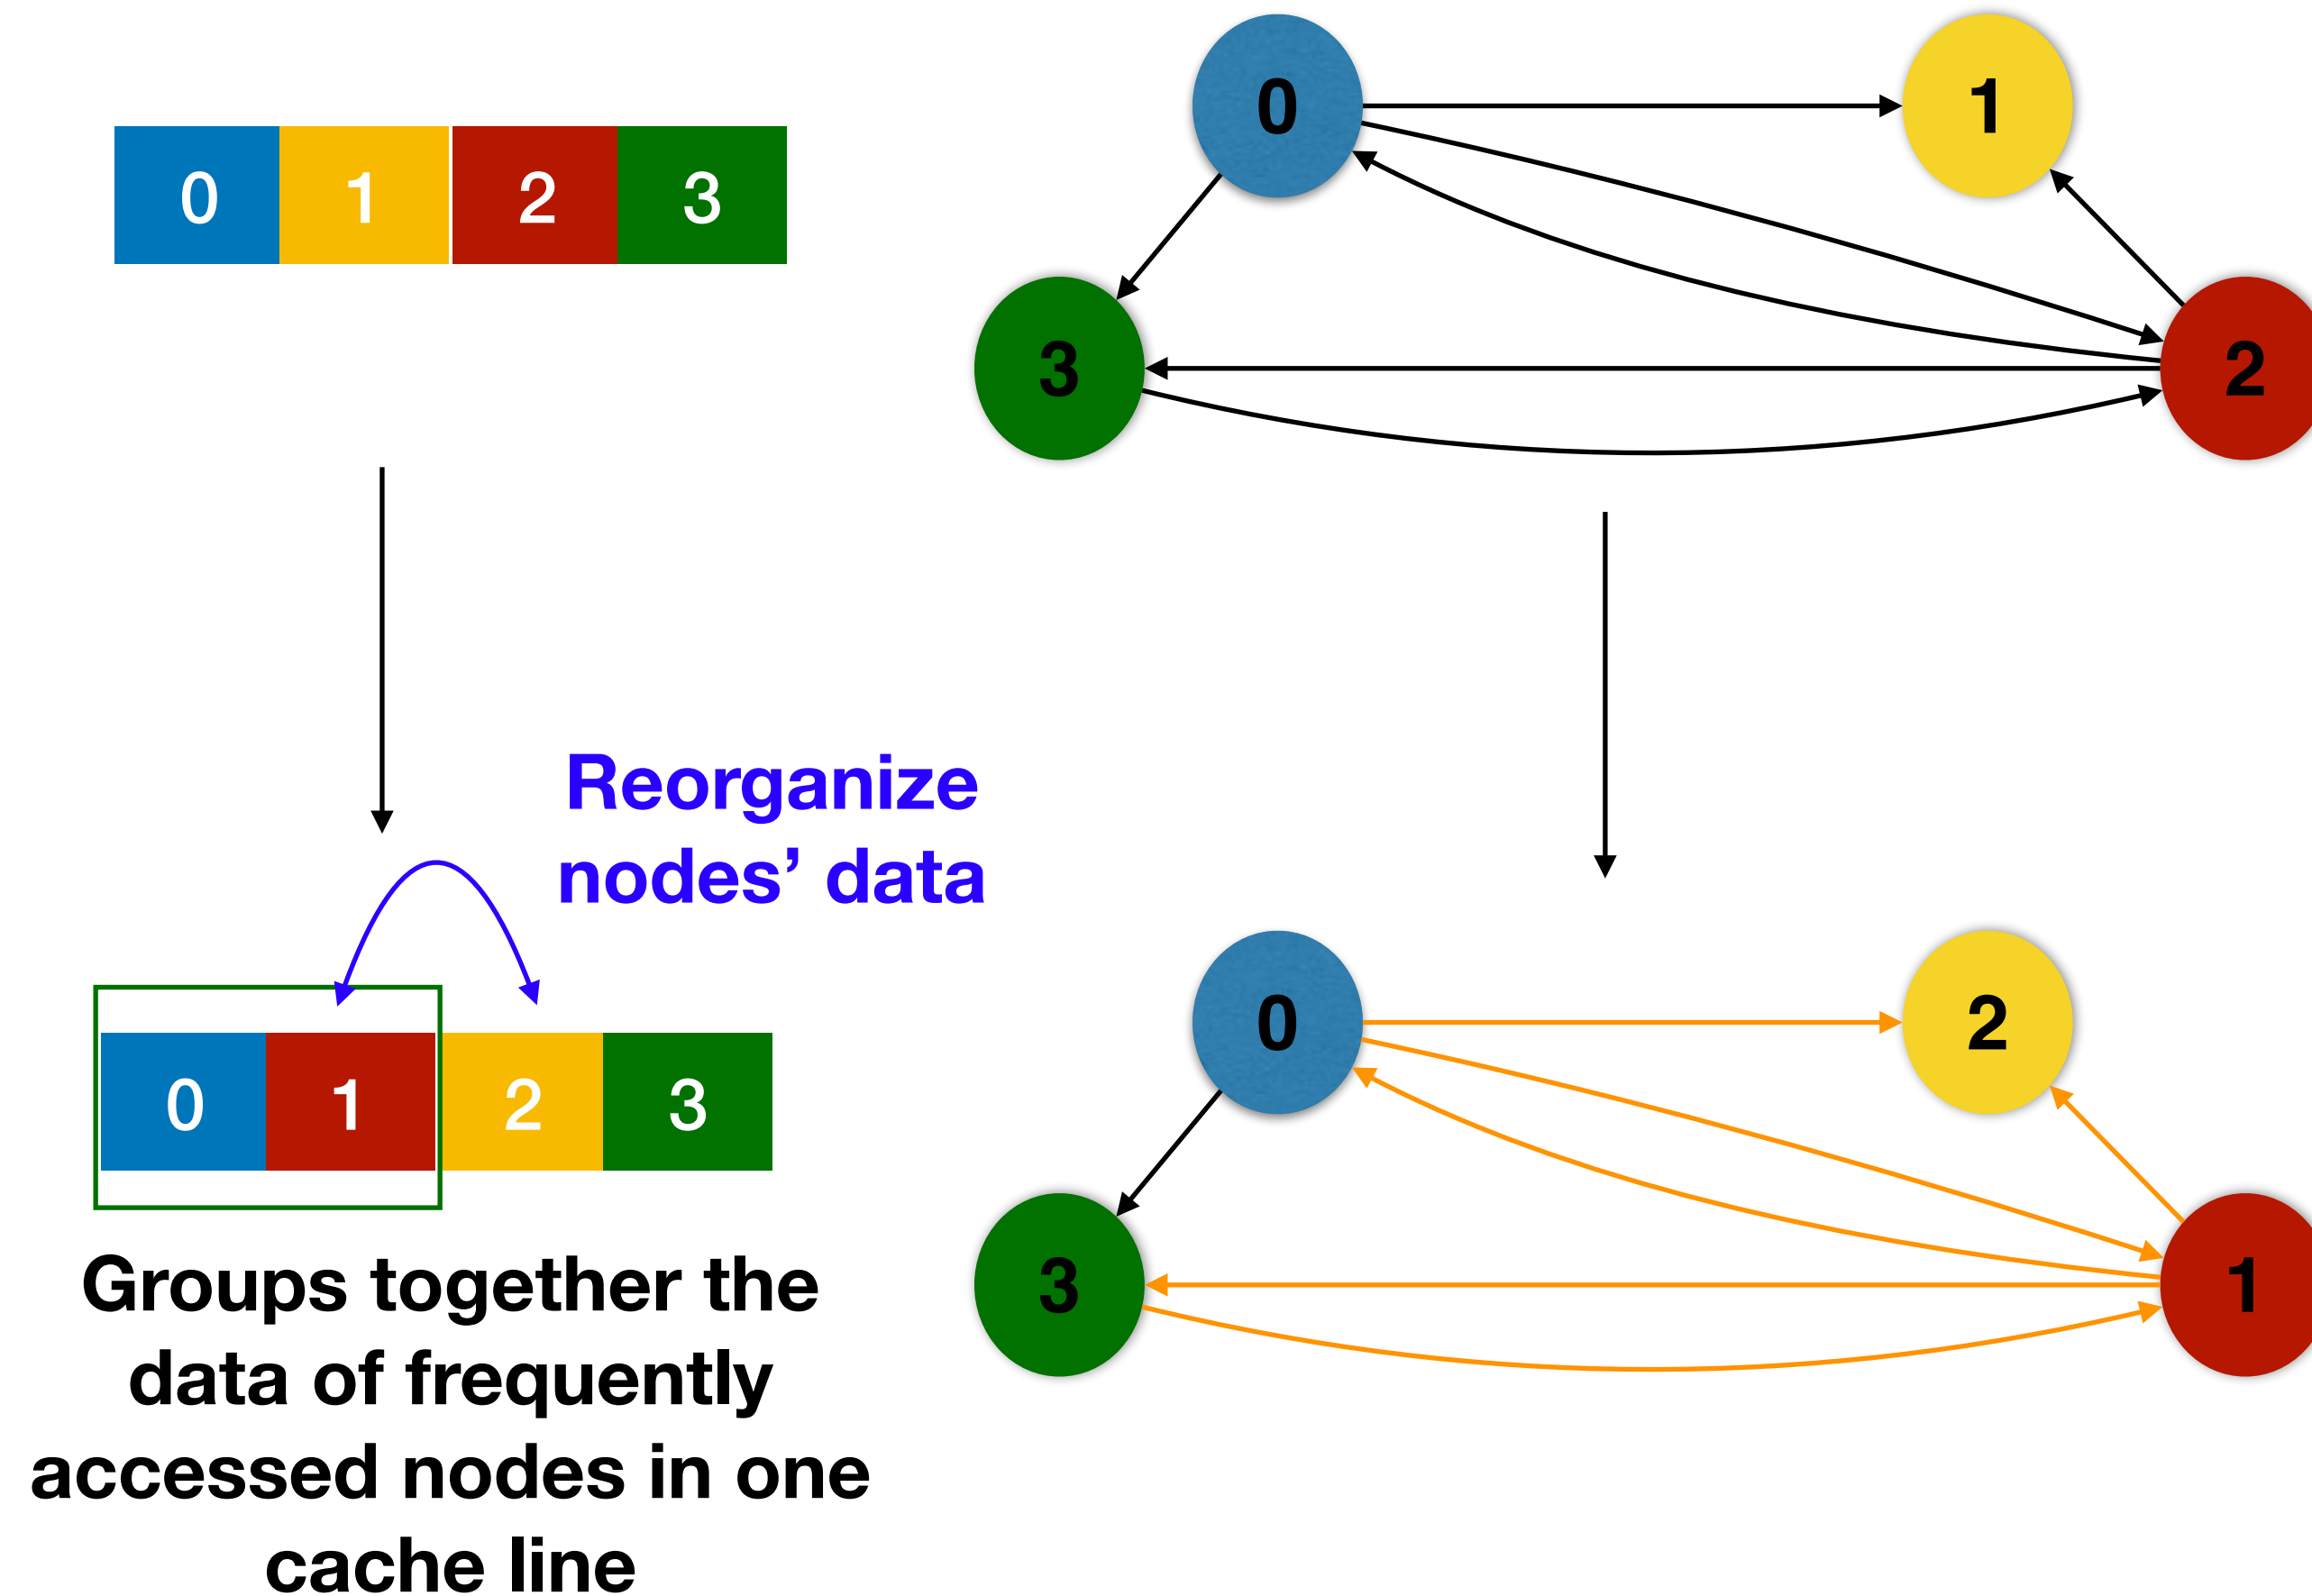
Frequency-based Vertex Reordering



Frequency-based Vertex Reordering



Frequency-based Vertex Reordering



PageRank

while ...

for node : graph.vertices

for ngh : graph.getInNeighbors(node)

newRanks[node] += ranks[ngh]/outDegree[ngh];

for node : graph.vertices

newRanks[node] = baseScore + damping*newRanks[node];


swap ranks and newRanks

Cache



#hits: 0

#misses: 0



PageRank

while ...

```
for node : graph.vertices
```

```
for ngh : graph.getInNeighbors(node)
```

```
newRanks[node] += ranks[ngh]/outDegree[ngh];
```

```
for node : graph.vertices
```

```
newRanks[node] = baseScore + damping*newRanks[node];
```

```
swap ranks and newRanks
```


Focus on the random
memory accesses on
ranks array

Cache



#hits: 0

#misses: 0



PageRank

while ...

```
for node : graph.vertices
```

```
for ngh : graph.getInNeighbors(node)
```

```
newRanks[node] += ranks[ngh]/outDegree[ngh];
```

```
for node : graph.vertices
```

```
newRanks[node] = baseScore + damping*newRanks[node];
```


```
swap ranks and newRanks
```

Cache



#hits: 0

#misses: 0



PageRank

while ...

```
for node : graph.vertices
```

```
for ngh : graph.getInNeighbors(node)
```


```
newRanks[node] += ranks[ngh]/outDegree[ngh];
```

```
for node : graph.vertices
```

```
newRanks[node] = baseScore + damping*newRanks[node];
```


```
swap ranks and newRanks
```

Cache



#hits: 0

#misses: 1



PageRank

while ...

```
for node : graph.vertices
```

```
for ngh : graph.getInNeighbors(node)
```


```
newRanks[node] += ranks[ngh]/outDegree[ngh];
```

```
for node : graph.vertices
```

```
newRanks[node] = baseScore + damping*newRanks[node];
```


```
swap ranks and newRanks
```

Cache



#hits: 1

#misses: 1



PageRank

while ...

```
for node : graph.vertices
```

```
for ngh : graph.getInNeighbors(node)
```

```
newRanks[node] += ranks[ngh]/outDegree[ngh];
```

```
for node : graph.vertices
```

```
newRanks[node] = baseScore + damping*newRanks[node];
```


```
swap ranks and newRanks
```

Cache



#hits: 1

#misses: 2



PageRank

while ...

```
for node : graph.vertices
```

```
for ngh : graph.getInNeighbors(node)
```

```
newRanks[node] += ranks[ngh]/outDegree[ngh];
```

```
for node : graph.vertices
```

```
newRanks[node] = baseScore + damping*newRanks[node];
```


```
swap ranks and newRanks
```

Cache



#hits: 1

#misses: 2



PageRank

while ...

```
for node : graph.vertices
```

```
for ngh : graph.getInNeighbors(node)
```


```
newRanks[node] += ranks[ngh]/outDegree[ngh];
```

```
for node : graph.vertices
```

```
newRanks[node] = baseScore + damping*newRanks[node];
```


```
swap ranks and newRanks
```

Cache



#hits: 1

#misses: 3



PageRank

while ...

```
for node : graph.vertices
```

```
for ngh : graph.getInNeighbors(node)
```


```
newRanks[node] += ranks[ngh]/outDegree[ngh];
```

```
for node : graph.vertices
```

```
newRanks[node] = baseScore + damping*newRanks[node];
```


```
swap ranks and newRanks
```

Cache



#hits: 2

#misses: 3



PageRank

while ...

```
for node : graph.vertices
```

```
for ngh : graph.getInNeighbors(node)
```


```
newRanks[node] += ranks[ngh]/outDegree[ngh];
```

```
for node : graph.vertices
```

```
newRanks[node] = baseScore + damping*newRanks[node];
```


```
swap ranks and newRanks
```

Cache



#hits: 3

#misses: 3



PageRank

while ...

```
for node : graph.vertices
```

```
for ngh : graph.getInNeighbors(node)
```


```
newRanks[node] += ranks[ngh]/outDegree[ngh];
```

```
for node : graph.vertices
```

```
newRanks[node] = baseScore + damping*newRanks[node];
```


```
swap ranks and newRanks
```

Cache



#hits: 4

#misses: 3



PageRank

while ...

```
for node : graph.vertices
```

```
for ngh : graph.getInNeighbors(node)
```


```
newRanks[node] += ranks[ngh]/outDegree[ngh];
```

```
for node : graph.vertices
```

```
newRanks[node] = baseScore + damping*newRanks[node];
```

```
swap ranks and newRanks
```


Cache



#hits: 4
#misses: 3

Much better
than

#hits: 1
#misses: 6



PageRank

while ...

```
for node : graph.vertices
```

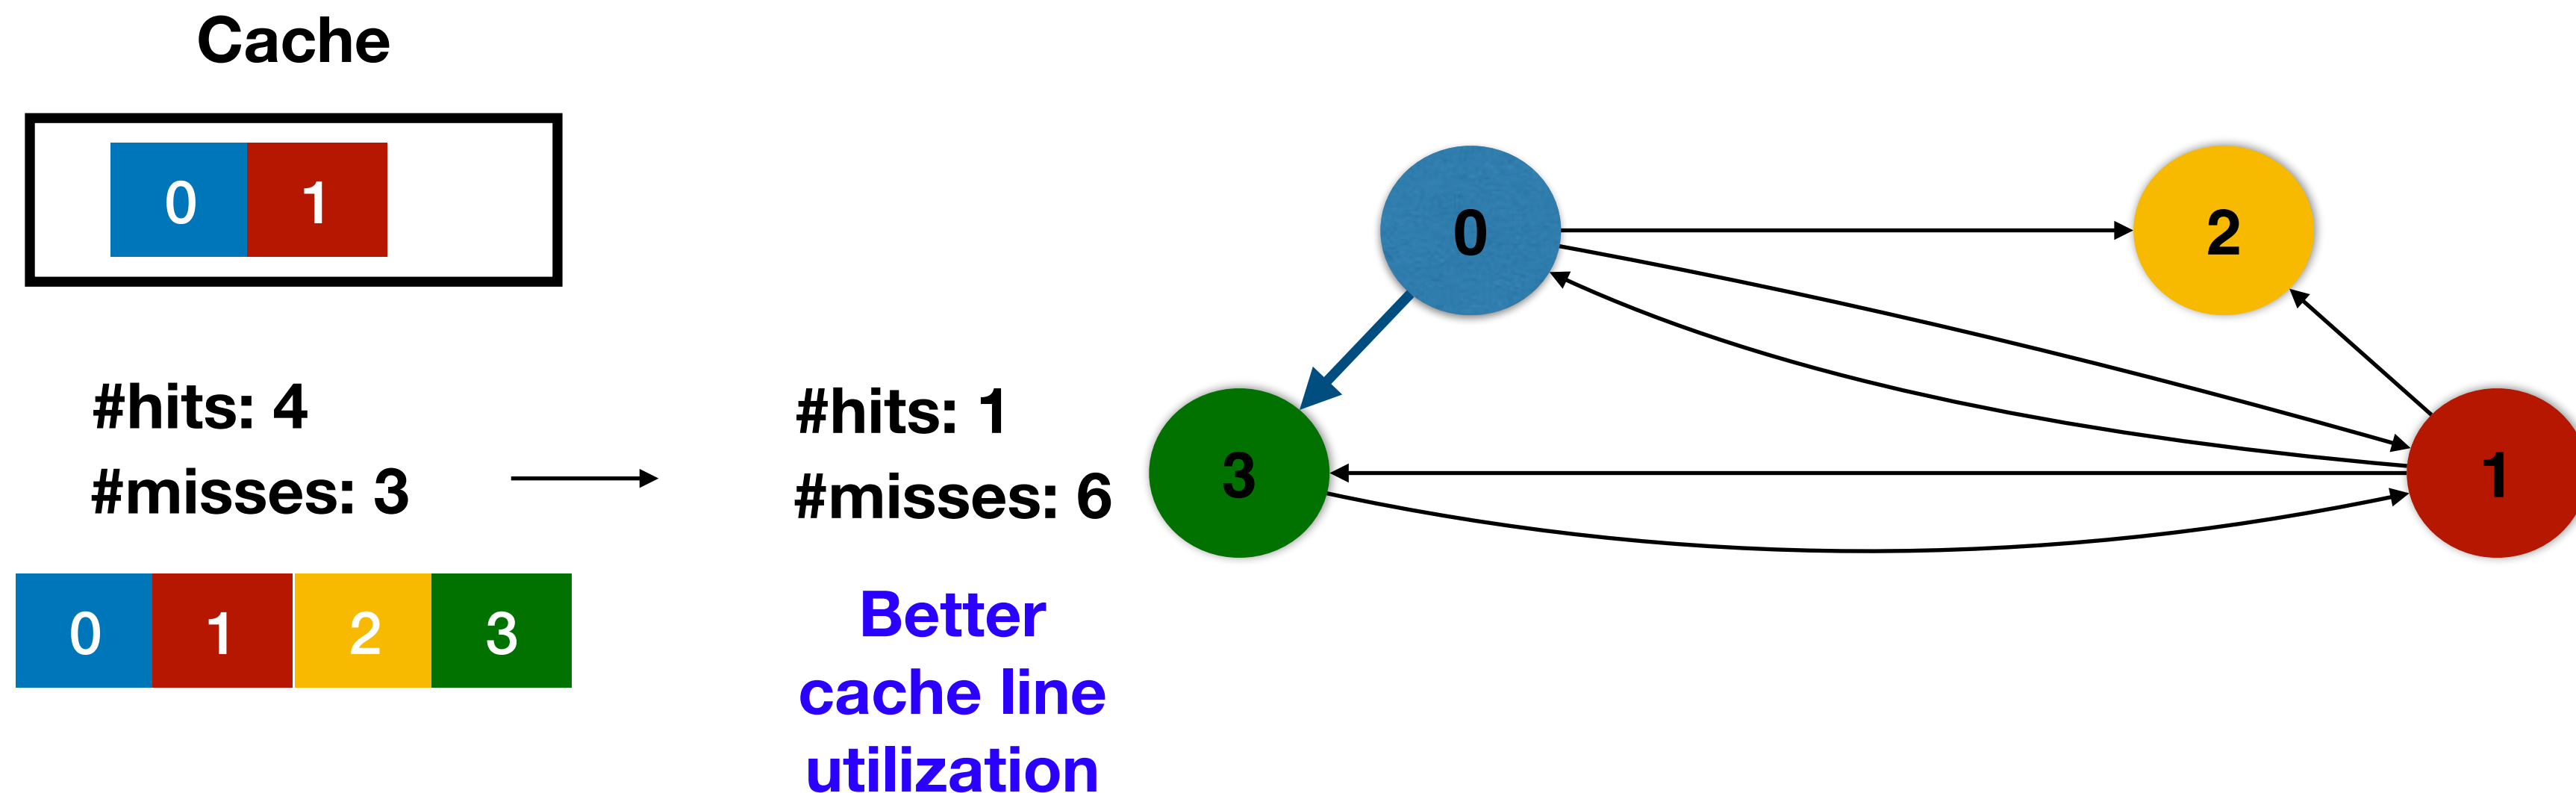
```
for ngh : graph.getInNeighbors(node)
```

```
newRanks[node] += ranks[ngh]/outDegree[ngh];
```

```
for node : graph.vertices
```

```
newRanks[node] = baseScore + damping*newRanks[node];
```


```
swap ranks and newRanks
```




Cache-aware Partitioning

- Design
 - Partition the graph into subgraphs where the random access are limited to LLC
 - Process each partition sequentially and merge rank contributions for each partition

Graph Partitioning




Graph Partitioning



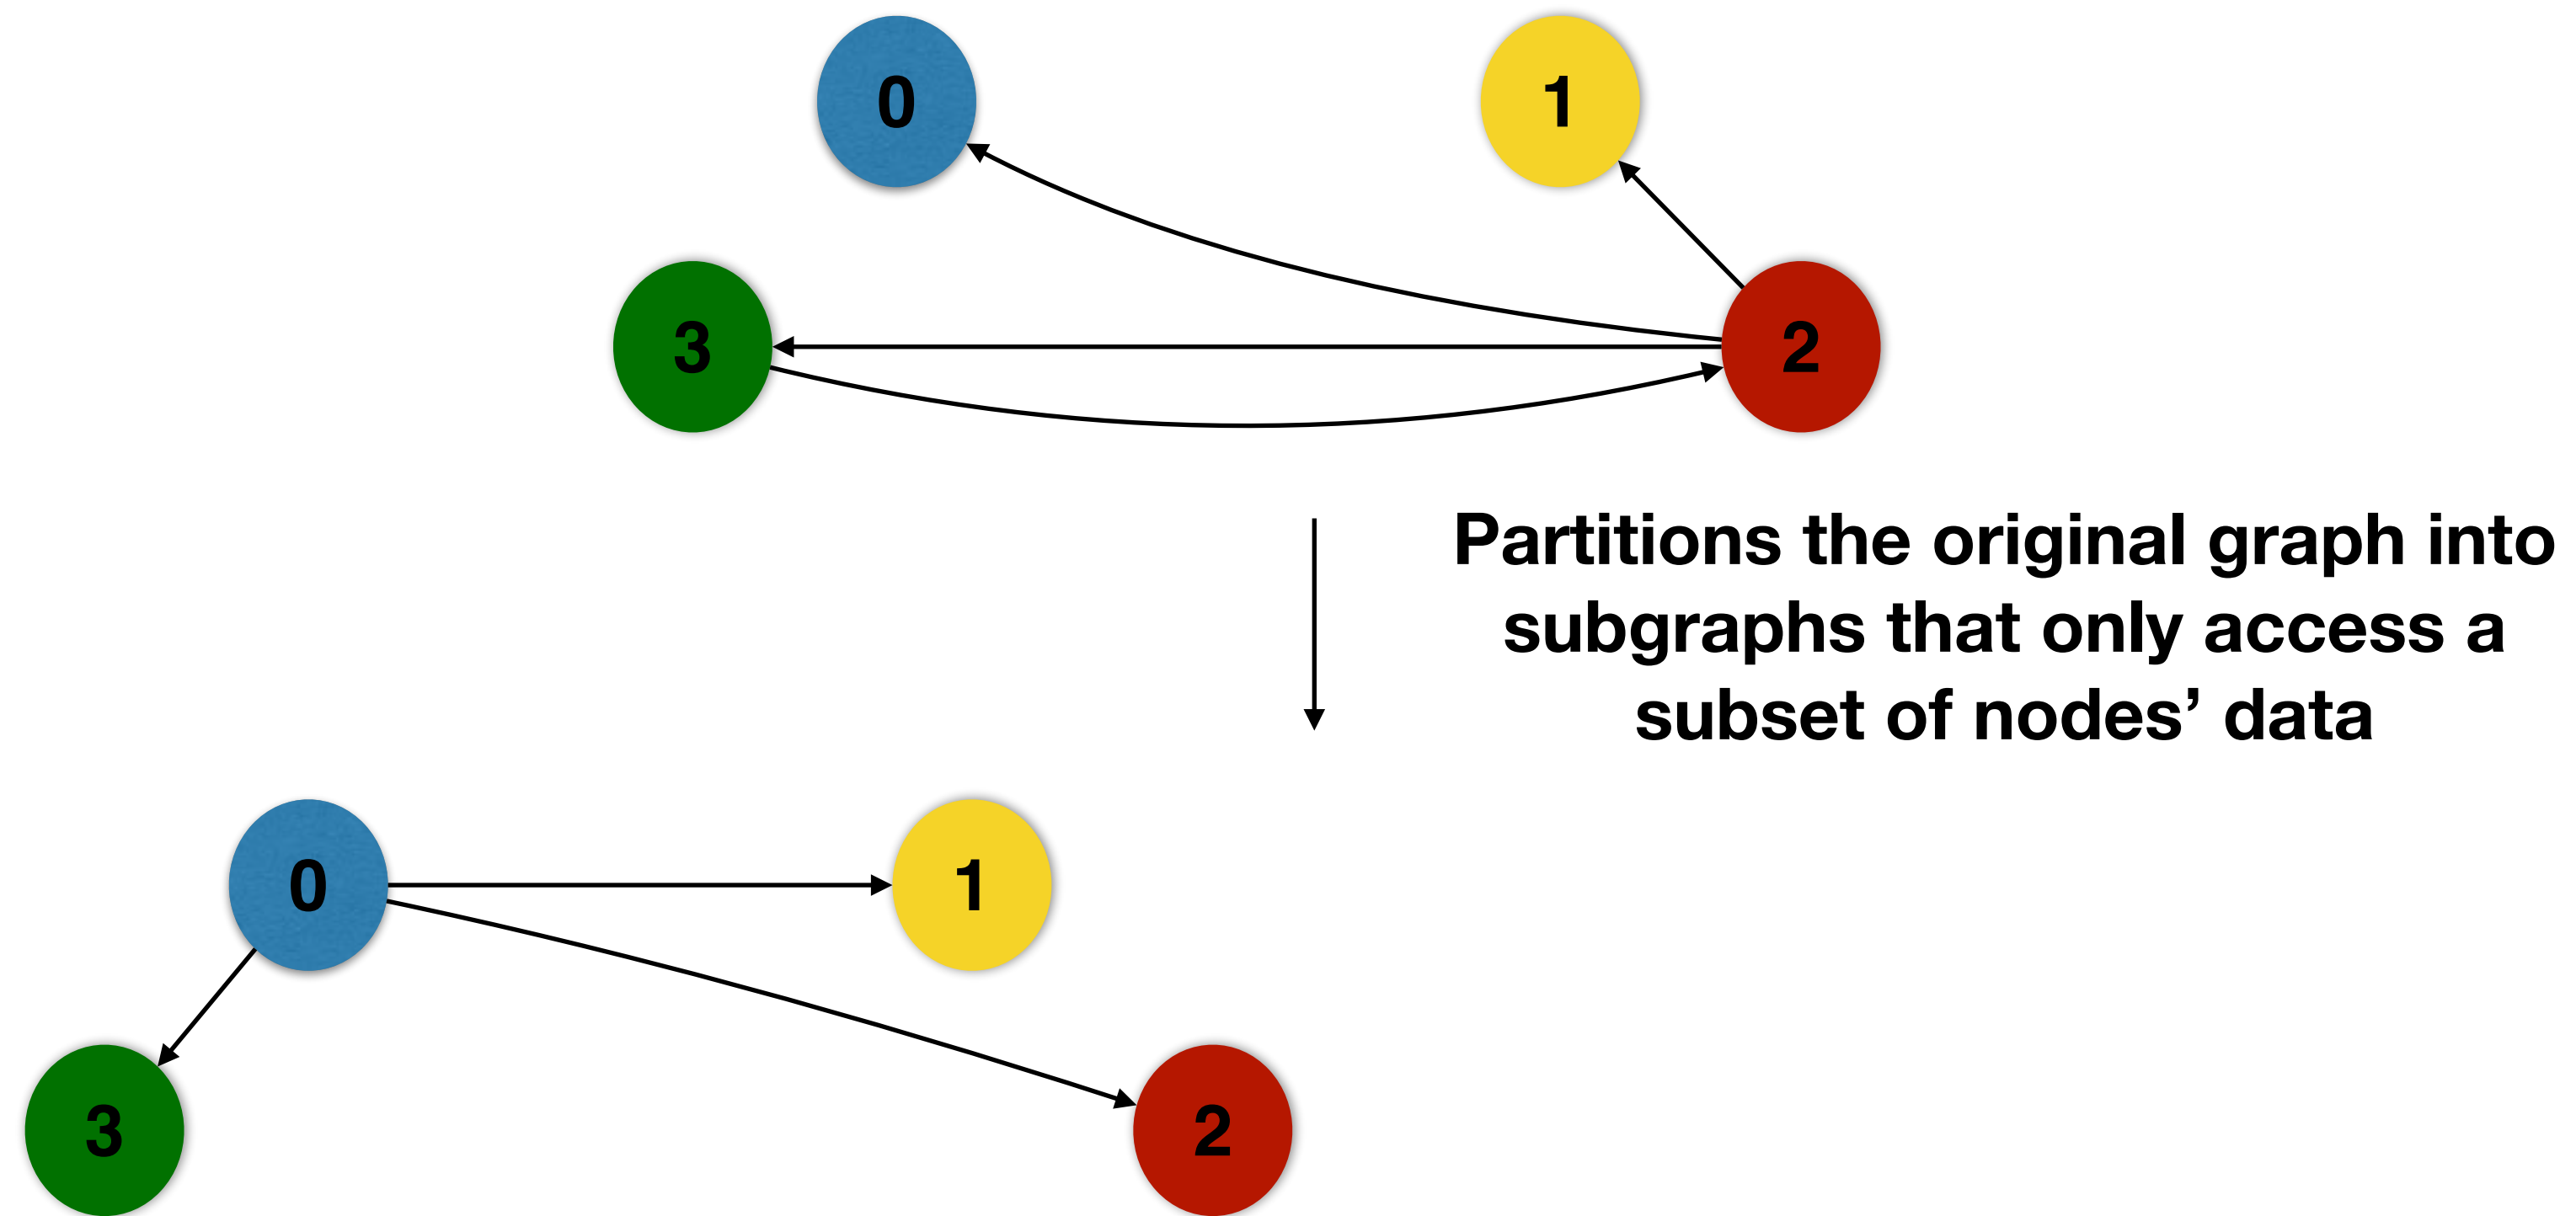
↓
Partitions the original graph into subgraphs that only access a subset of nodes' data

Graph Partitioning

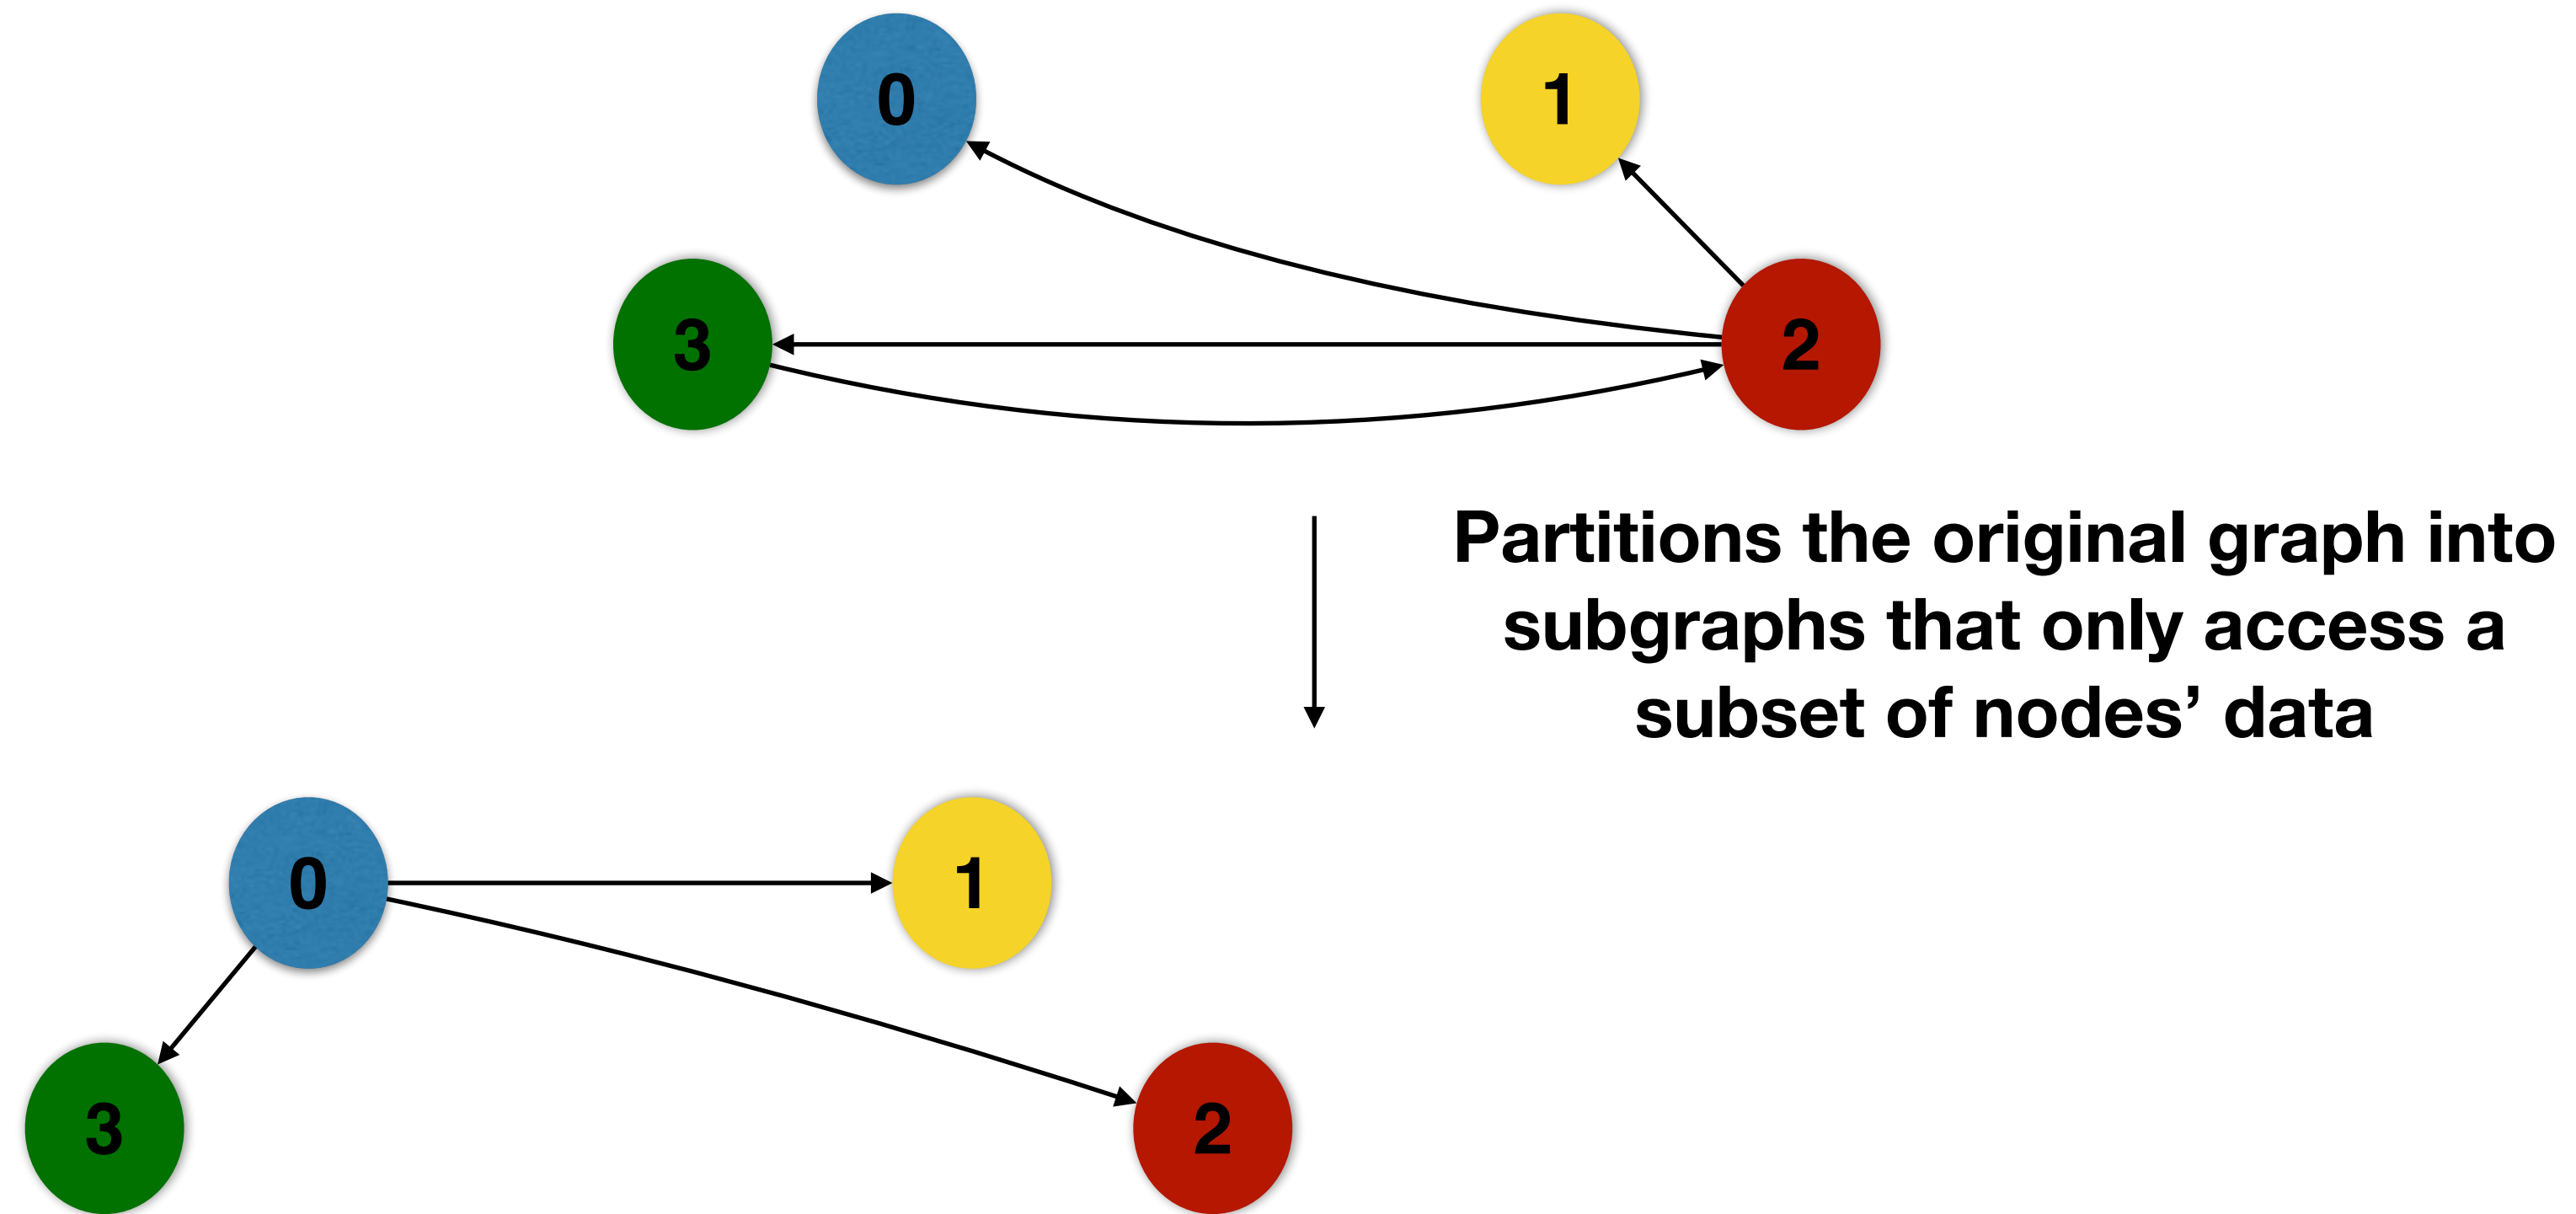


↓
Partitions the original graph into subgraphs that only access a subset of nodes' data

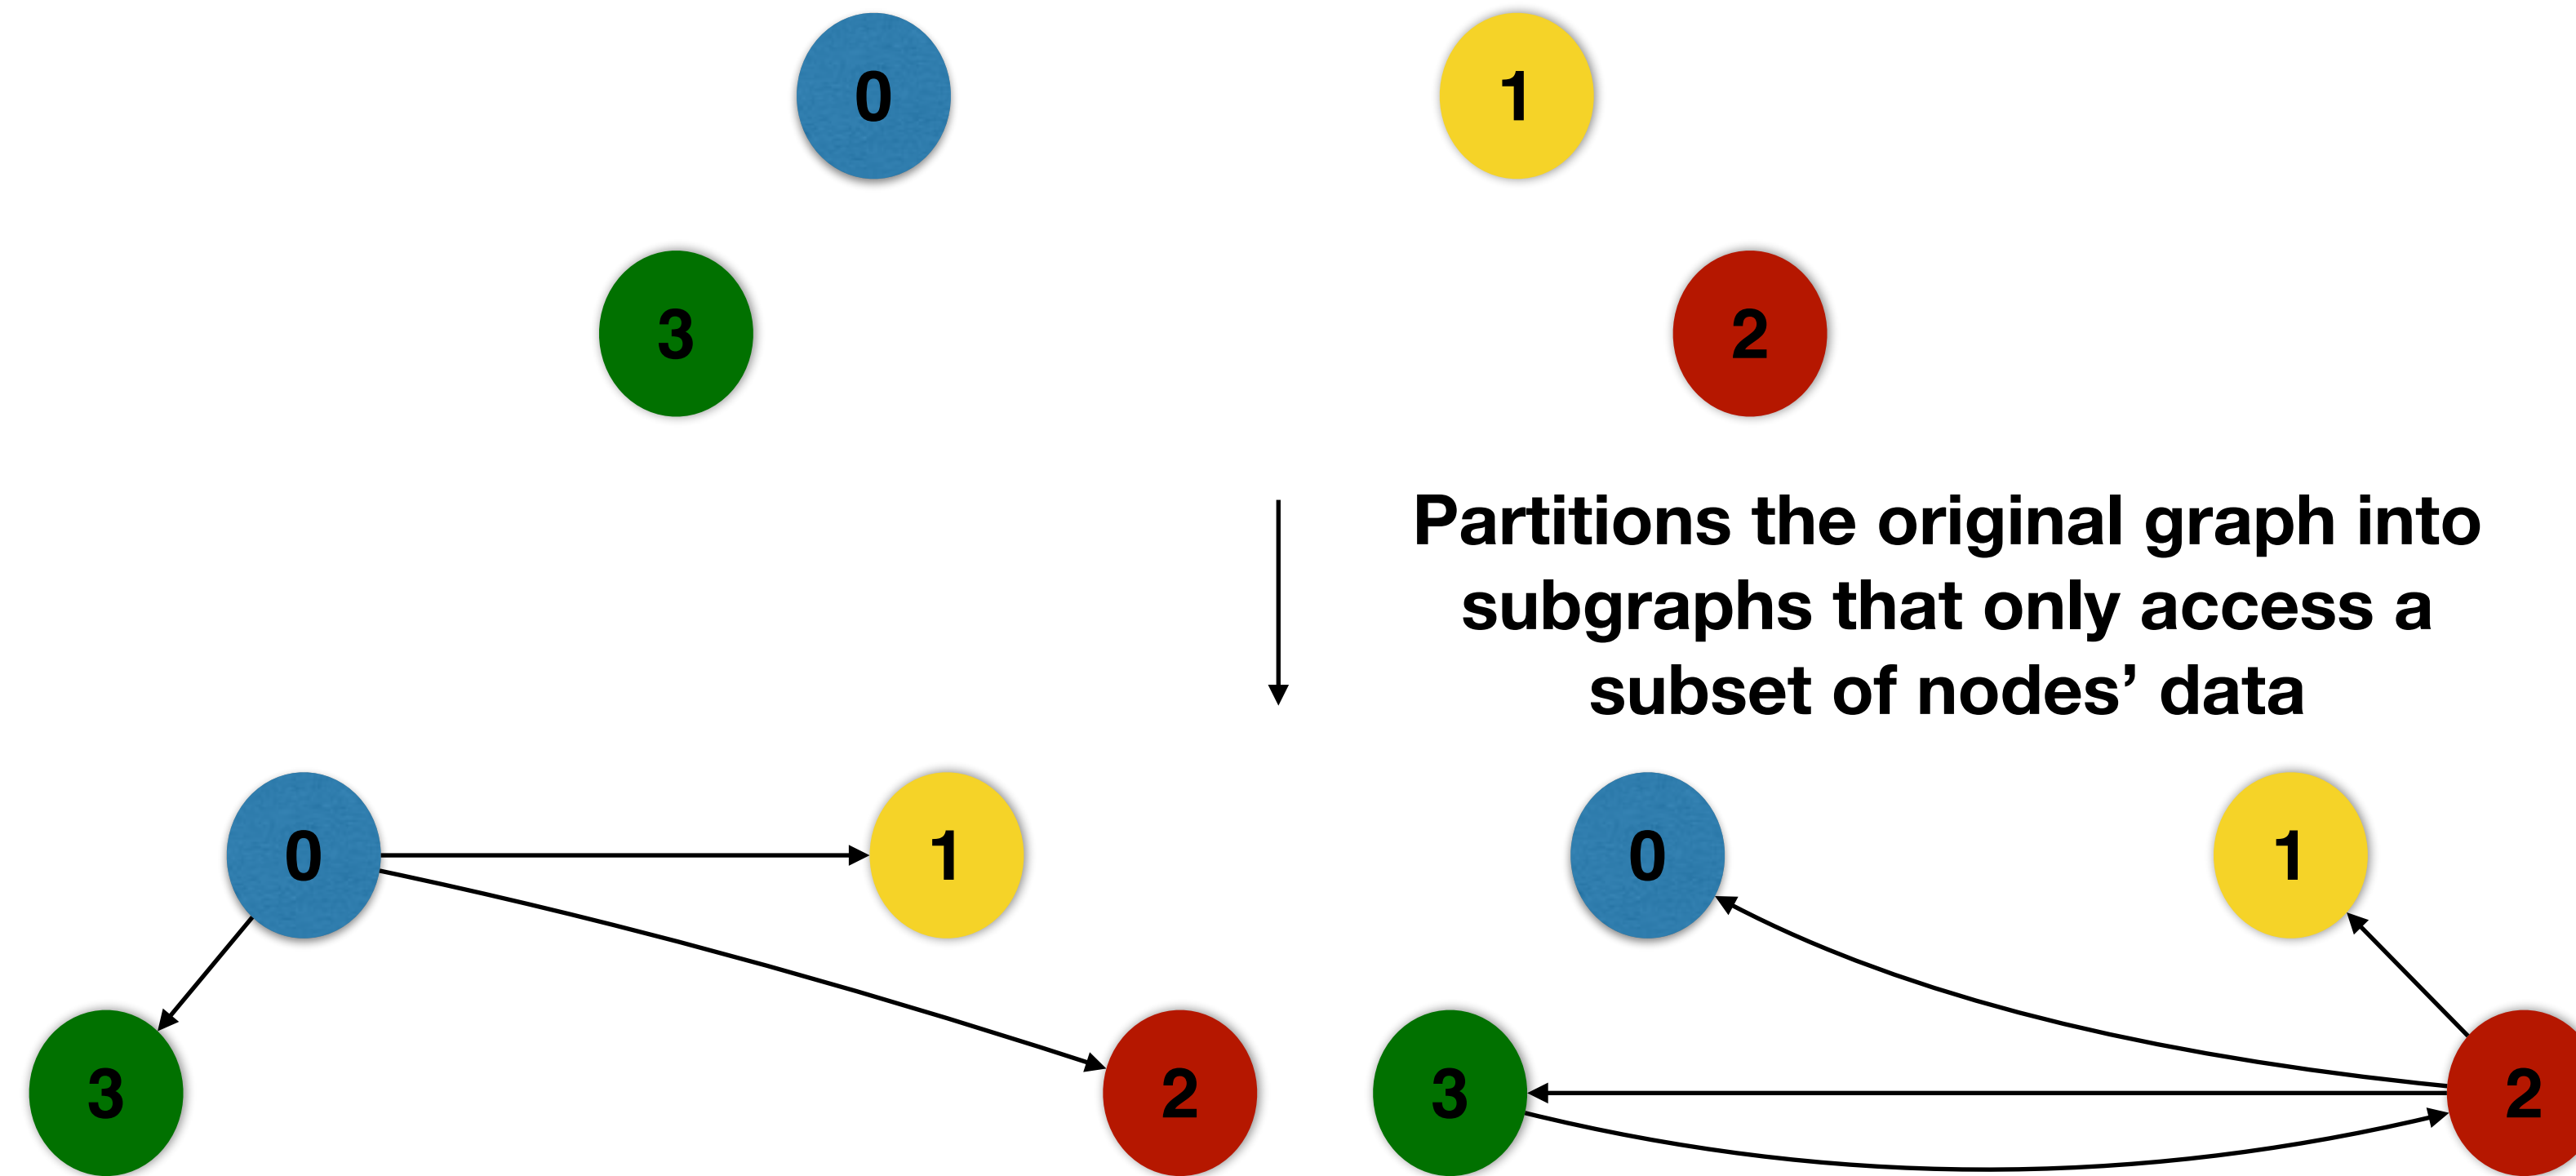
Graph Partitioning



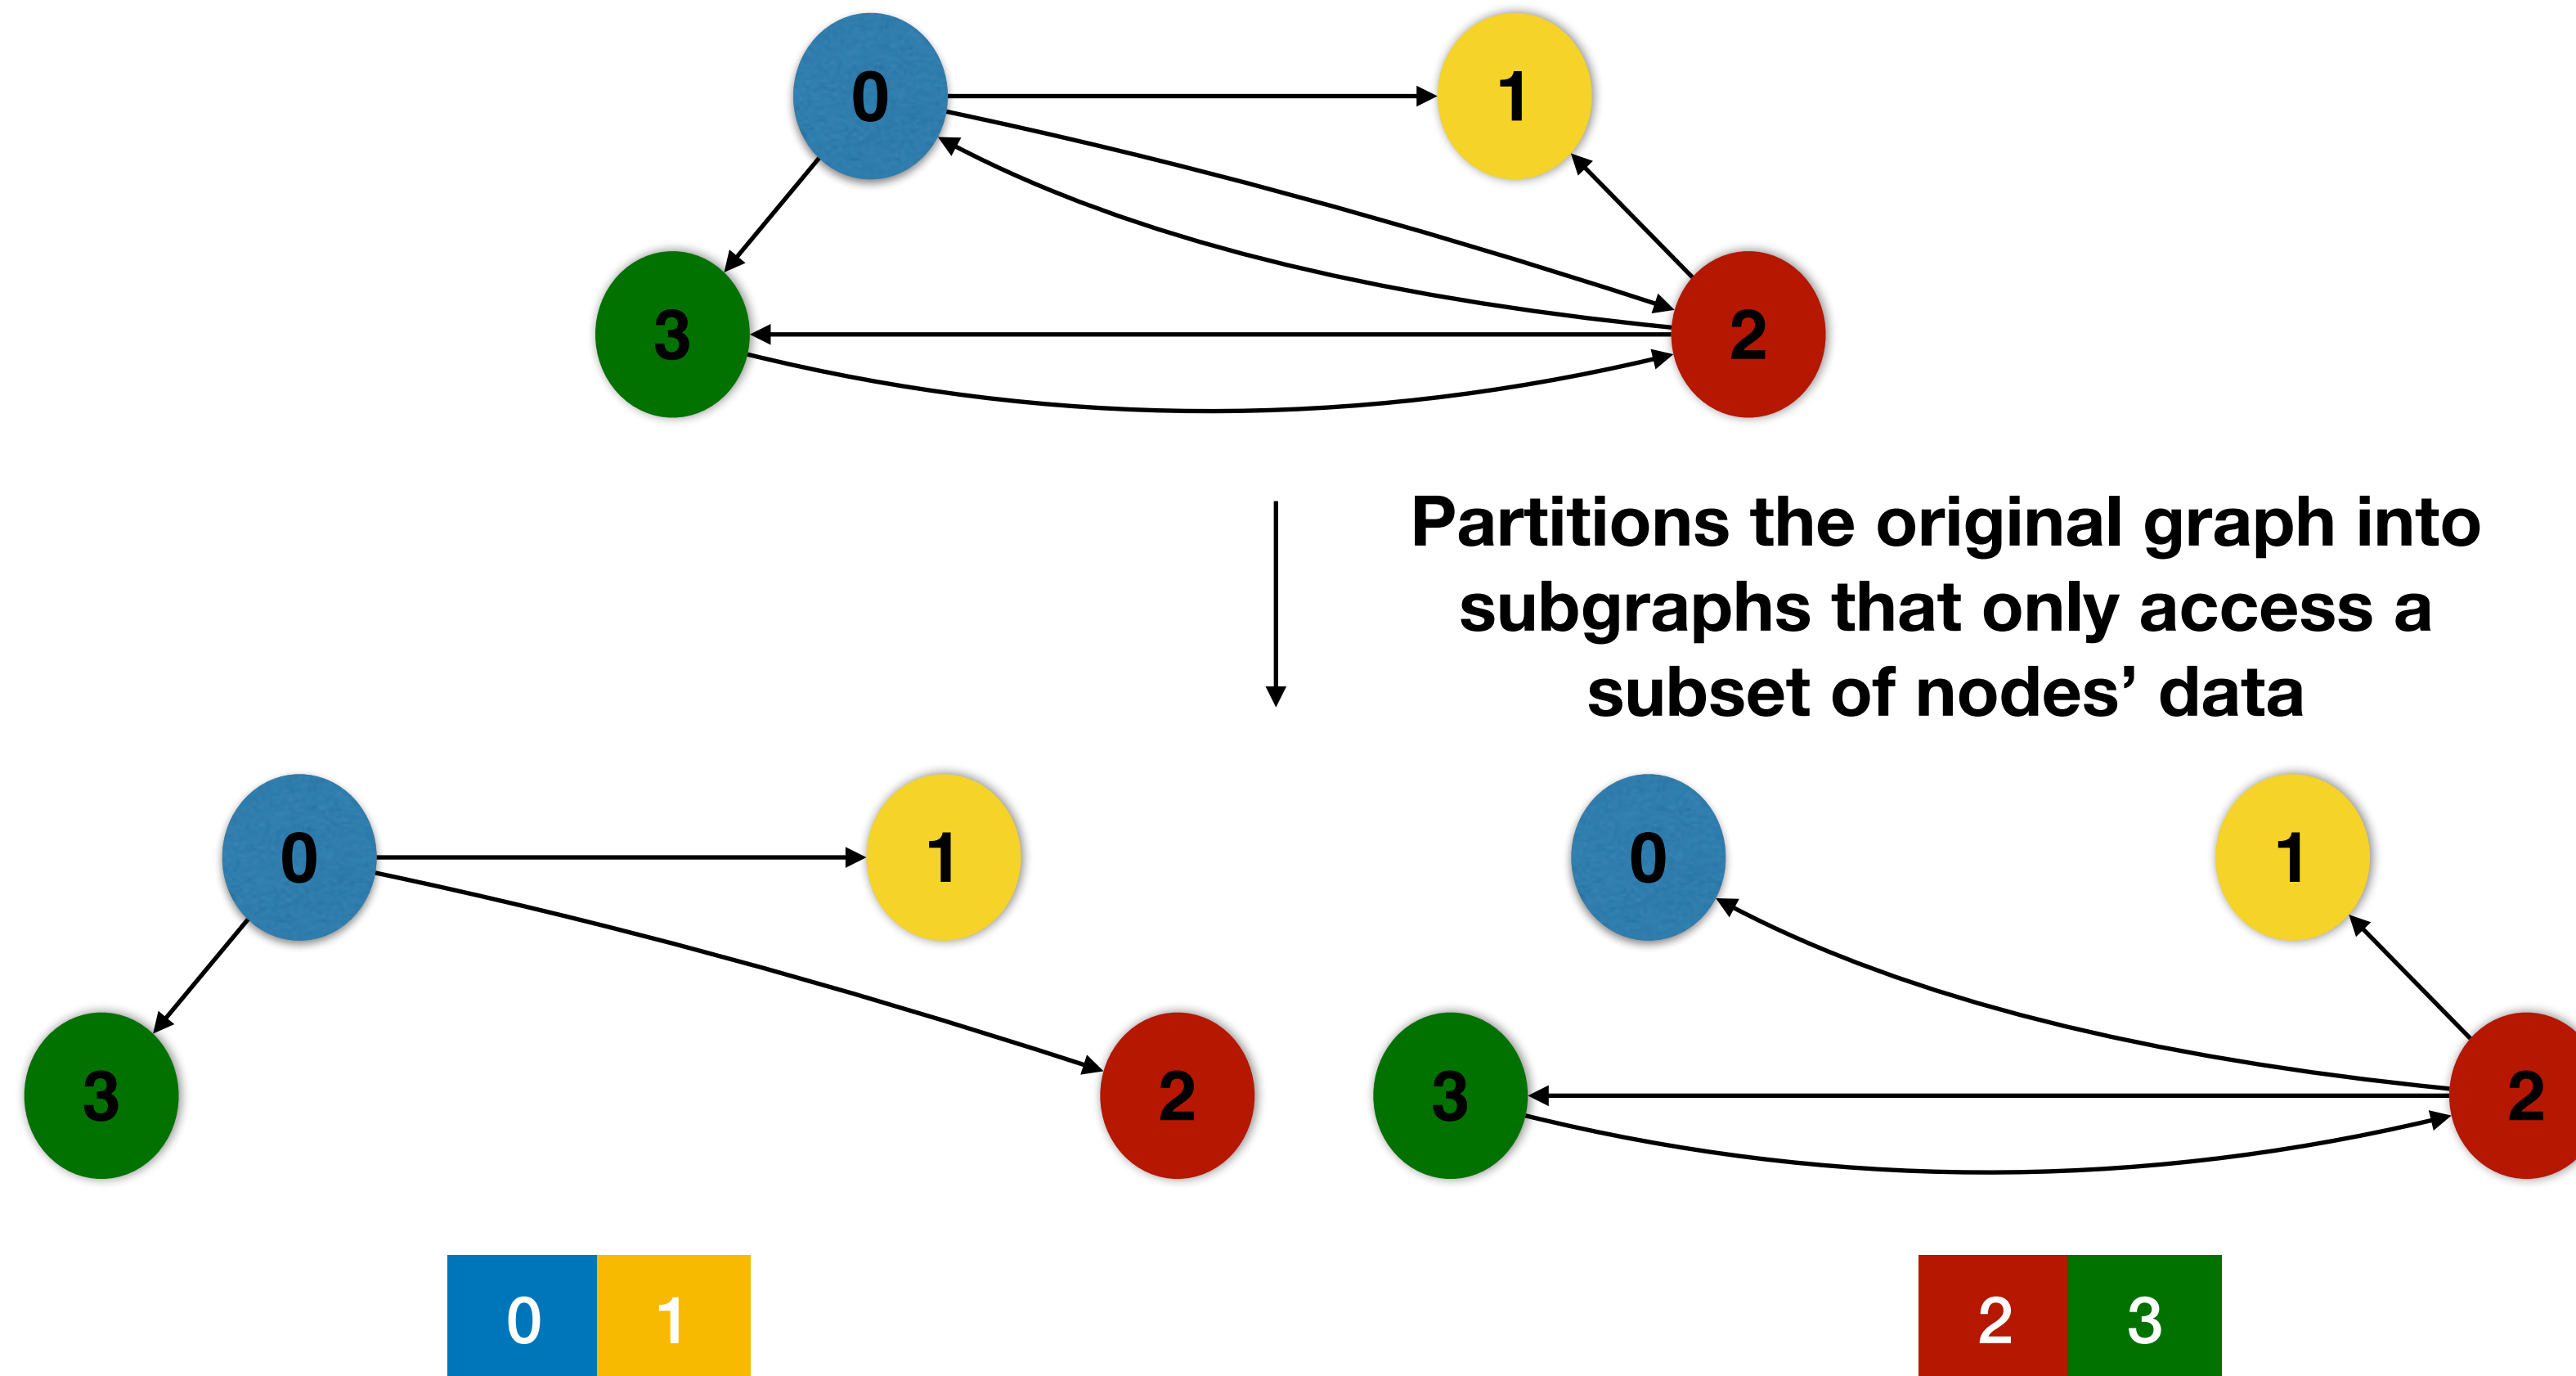
Graph Partitioning



Graph Partitioning



Graph Partitioning





Graph Processing

Cache



#hits: 0

#misses: 0





Graph Processing

Cache



#hits: 0

#misses: 0





Graph Processing

Cache



#hits: 0

#misses: 1





Graph Processing

Cache



#hits: 0

#misses: 1





Graph Processing

Cache




#hits: 1

#misses: 1





Graph Processing

Cache




#hits: 2

#misses: 1





Graph Processing

Cache




#hits: 2

#misses: 1





Graph Processing

Cache



#hits: 2

#misses: 1





Graph Processing

Cache




#hits: 2

#misses: 1





Graph Processing

Cache



#hits: 2


#misses: 1



Graph Processing





Cache




#hits: 2

#misses: 2





Graph Processing

Cache




#hits: 3

#misses: 2





Graph Processing

Cache




#hits: 4

#misses: 2





Graph Processing

Cache




#hits: 5

#misses: 2





Graph Processing

Cache



Only have 2 misses

#hits: 5
#misses: 2




Graph Processing



**Better than
Frequency-based
Reordering**

#hits: 4
#misses: 3

Cache

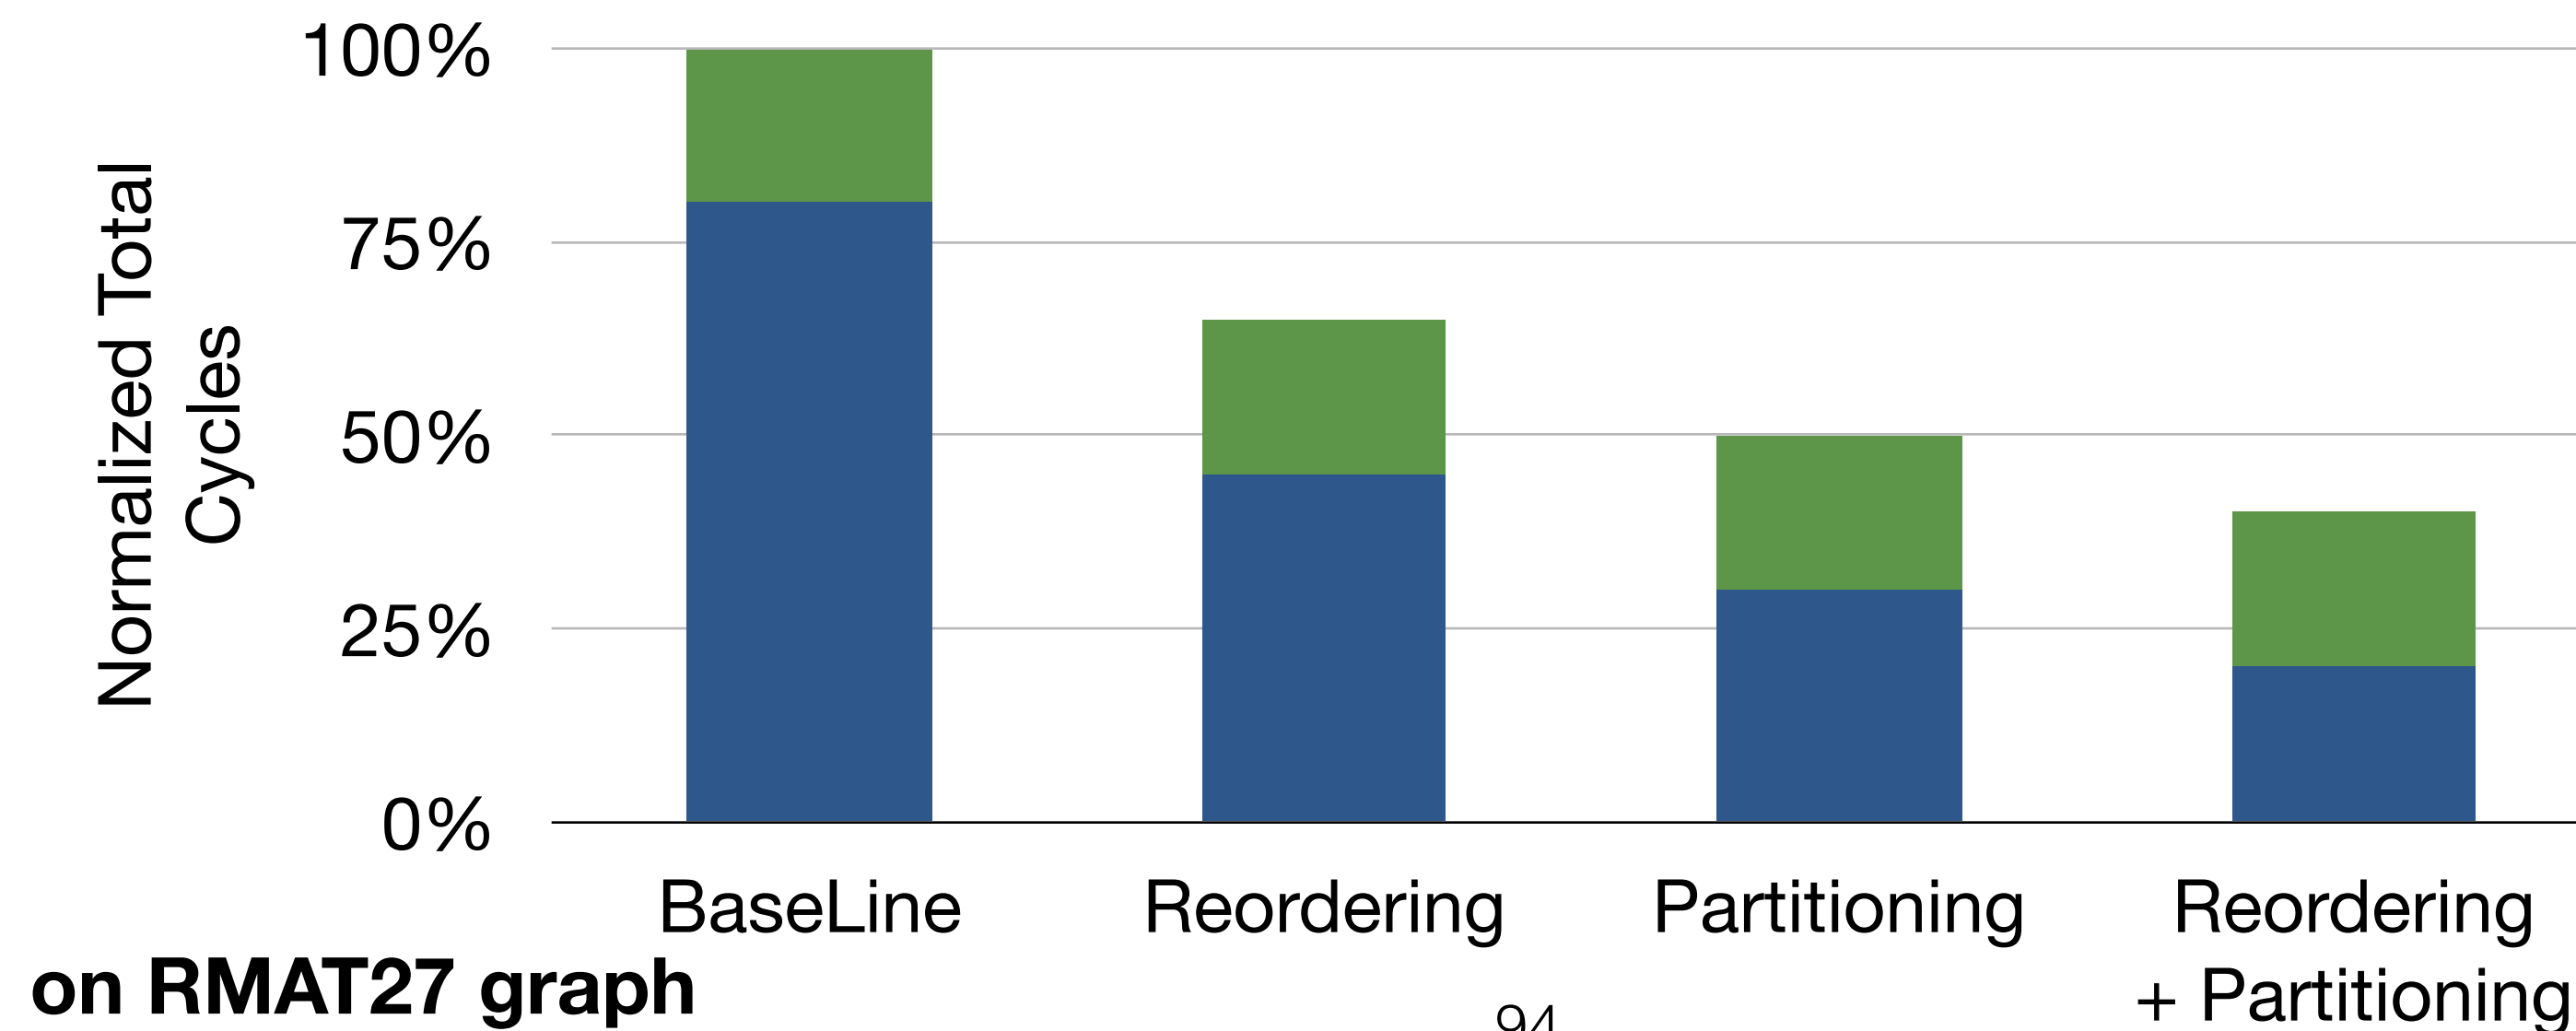


#hits: 5
#misses: 2




PageRank

```
while ...  
  for node : graph.vertices  
    for ngh : graph.getInNeighbors(node)  
      newRanks[node] += ranks[ngh]/outDegree[ngh];  
  for node : graph.vertices  
    newRanks[node] = baseScore + damping*newRanks[node];  
  swap ranks and newRanks
```




PageRank

```
while ...  
  for node : graph.vertices  
    for ngh : graph.getInNeighbors(node)  
      newRanks[node] += ranks[ngh]/outDegree[ngh];  
  for node : graph.vertices  
    newRanks[node] = baseScore + damping*newRanks[node];  
  swap ranks and newRanks
```




PageRank

```
while ...  
  for node : graph.vertices  
    for ngh : graph.getInNeighbors(node)  
      newRanks[node] += ranks[ngh]/outDegree[ngh];  
  for node : graph.vertices  
    newRanks[node] = baseScore + damping*newRanks[node];  
  swap ranks and newRanks
```



PageRank


```
while ...  
  for node : graph.vertices  
    for ngh : graph.getInNeighbors(node)  
      newRanks[node] += ranks[ngh]/outDegree[ngh];  
  for node : graph.vertices  
    newRanks[node] = baseScore + damping*newRanks[node];  
  swap ranks and newRanks
```



Impact of Partitioning

PageRank


Cycle Reduction



Impact of Partitioning


PageRank

Cycle Reduction



Breadth First Search

Cycle Overhead



on RMAT27 graph




Impact of Partitioning

PageRank

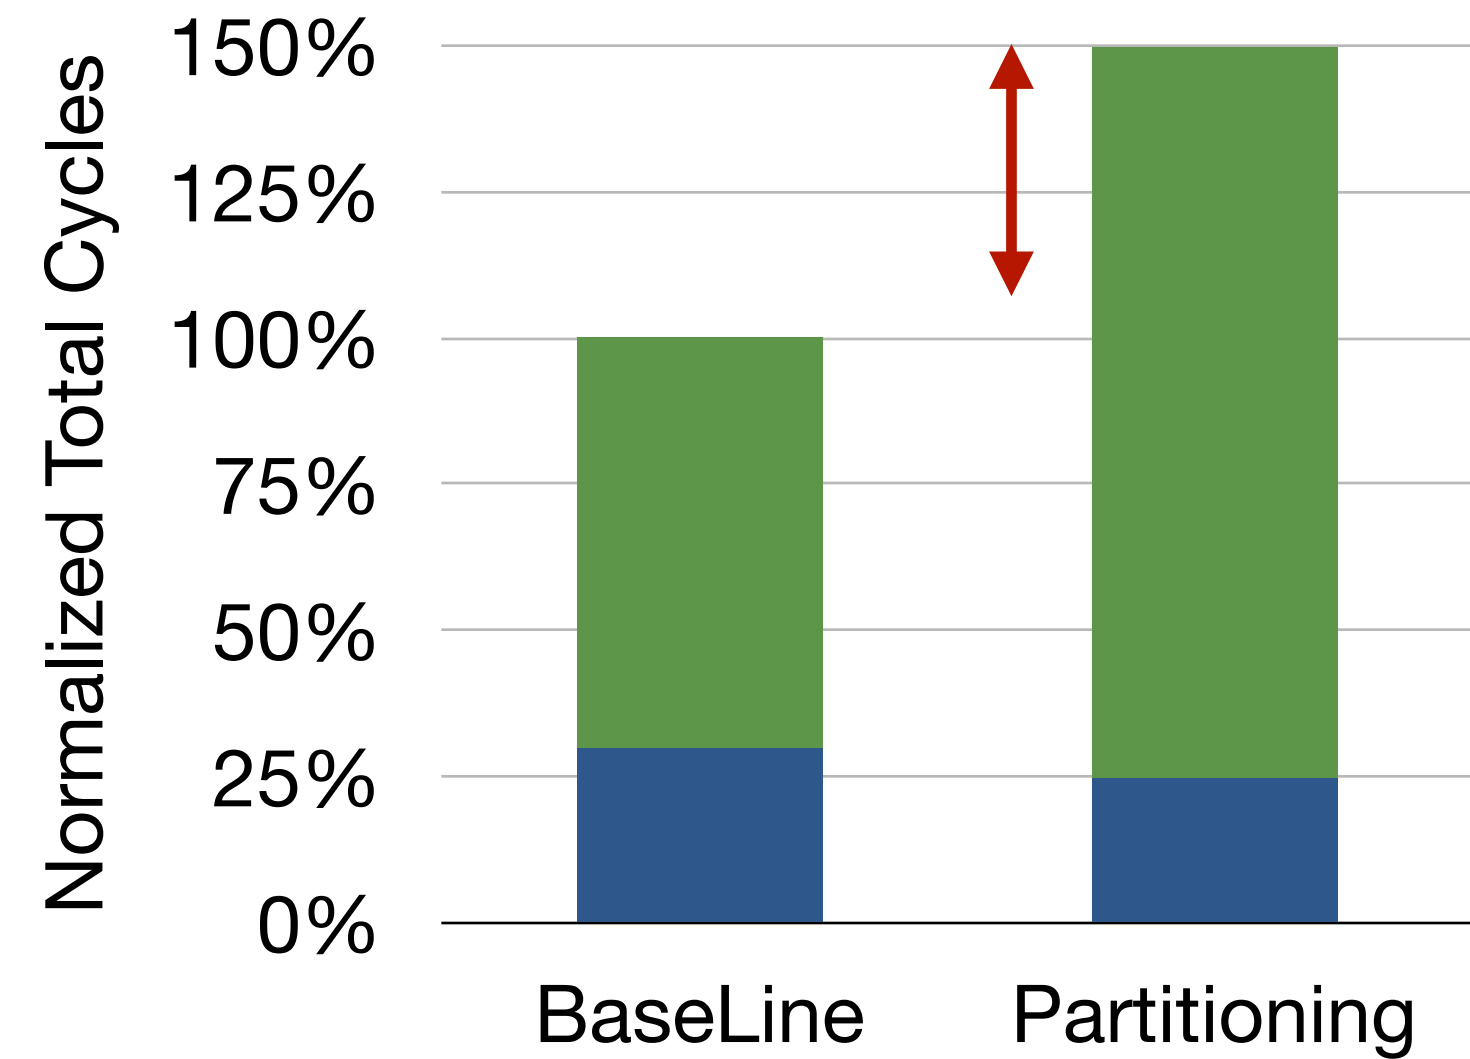
Optimizations Don't Always Work across Different Algorithms and Data

Breadth First Search

Cycle Reduction



Cycle Overhead




on RMAT27 graph

Everything Else
Random Memory Accesses

Graph Computations have Variety


Data

Social Networks, Web
Graphs, Road Networks,
Engineering Meshes,
Transaction Graphs, Network
Traffic Graphs, Email Networks,
Similarity Graphs, ...



Algorithms

Breadth-first search, betweenness
centrality, Bellman-Ford, Delta-stepping,
collaborative filtering, Page Rank, Page Rank
Delta, connected components, k-core
decomposition, triangle counting, local
clustering, structural clustering minimum
spanning forest, eccentricity estimation, graph
coloring, k-truss decomposition, nuclei
decomposition, biconnectivity, set cover,
maximum flow, butterfly counting, strongly
connected components, graph partitioning, RDF
queries, random walks, point-to-point shortest
paths, A* search, low-diameter decomposition,
densest subgraph, multi-source BFS, maximal
independent set, maximal matching, etc...



Hardware

CPU, GPU, KNL, Distributed
Environment, FPGA,
HammerBlade, Symphony,...

Outline

Hardware Utilization

Making Caches Work for Graph Analytics (BigData17)
Zhang, et al.

- Frequency-based Reordering
- Cache-aware Partitioning

Programming System to Handle Variety in Data and Algorithms

GraphIt: a High-Performance Graph DSL (OOPSLA18)
Zhang, et al.

GraphIt Compiler and DSL that Decouples


- **Algorithm**
 - **Optimization**
 - **Hardware**
- ### **for Graph Applications**

Optimizing Ordered Graph Algorithms with GraphIt (CGO2020)
Zhang, et al.


Variety in Hardware

Universal Graph Framework (Under Submission)
Brahmakshatriya, Zhang, et al.

Power-Law Graphs




World Wide Web




**Power-Law Degree Distribution,
Small Diameter, Poor Locality**



Social Networks



Maps




Engineering Meshes

1. <http://googlesystem.blogspot.com/2007/05/world-wide-web-as-seen-by-google.html> 2. <http://www.facebookfever.com/introducing-facebook-new-graph-api-explorer-features/> 3. <http://maps.google.com> 4. https://en.wikipedia.org/wiki/Polygon_mesh#/media/File:Dolphin_triangle_mesh.png

Bounded-Degree Graphs




World Wide Web




Power-Law Degree Distribution,
Small Diameter, Poor Locality




Social Networks



Maps




Bounded Degree Distribution
Large Diameter, Excellent Locality




Engineering Meshes

1. <http://googlesystem.blogspot.com/2007/05/world-wide-web-as-seen-by-google.html> 2. <http://www.facebook.com/introducing-facebook-new-graph-api-explorer-features/> 3. <http://maps.google.com> 4. https://en.wikipedia.org/wiki/Polygon_mesh#/media/File:Dolphin_triangle_mesh.png

Graph Algorithms





Topology-Driven Algorithms



Work on All Edges and Vertices


Topology-Driven Algorithms




Recommendations for You, Yunming




Data-Driven Algorithms



Data-Driven Algorithms




Data-Driven Algorithms



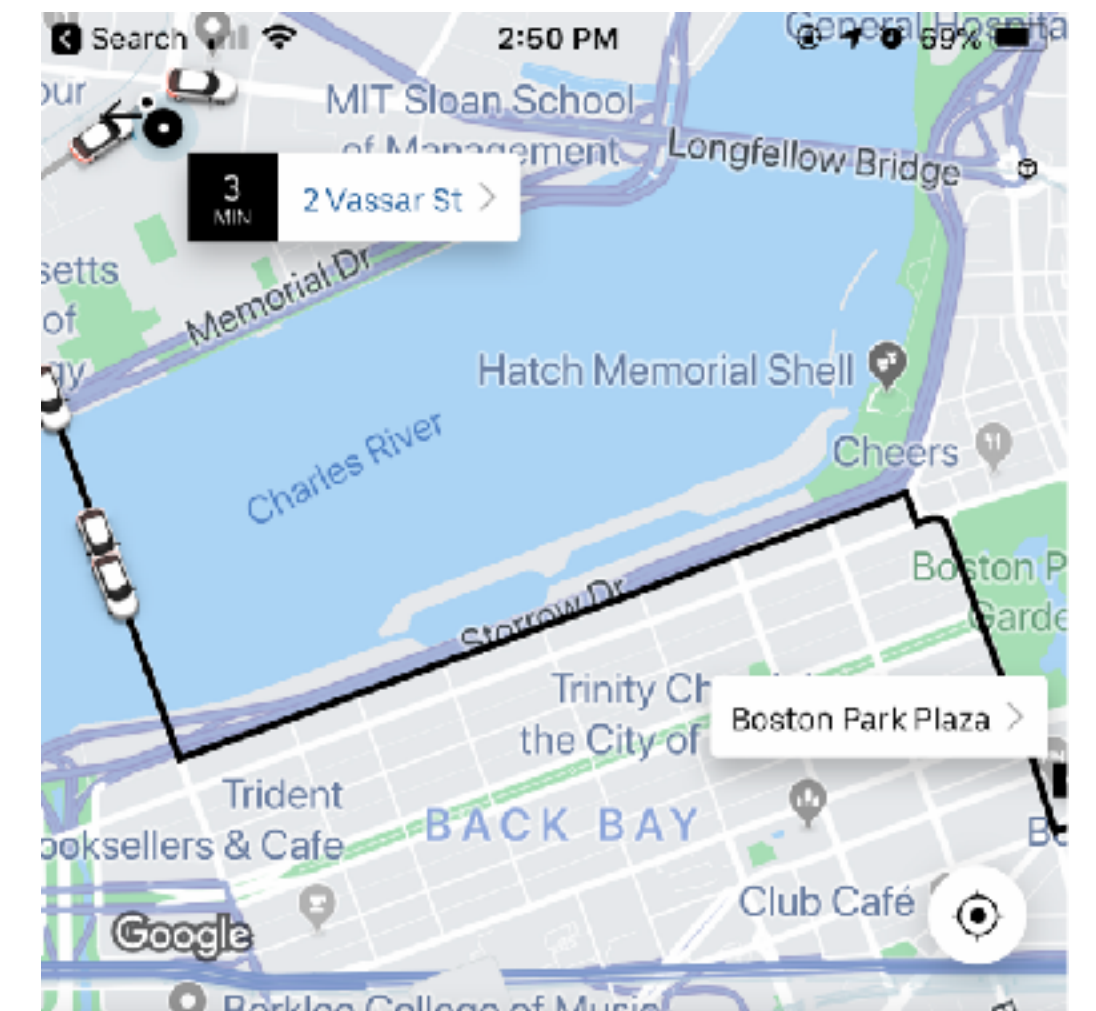

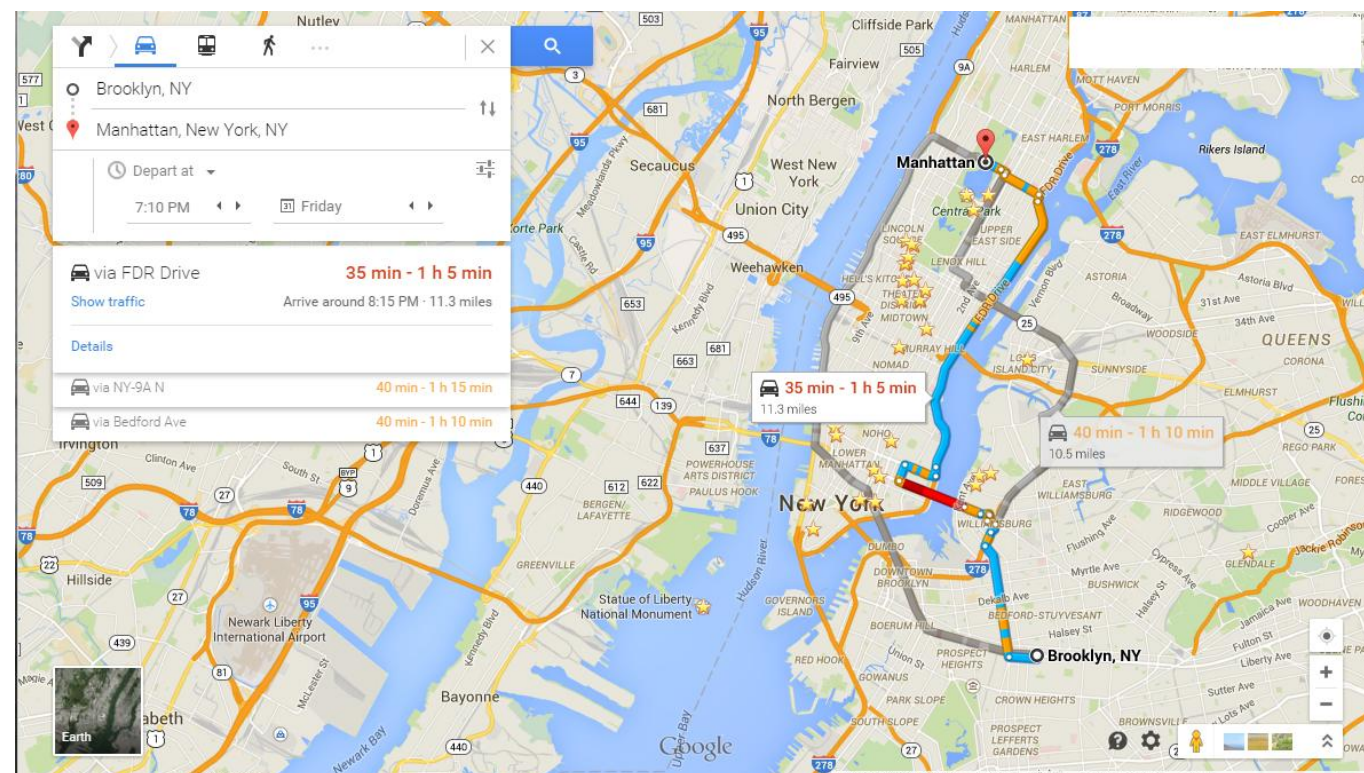
**Work on a subset of vertices and edges
(Data-Driven)**

Data-Driven Algorithms




**Work on a subset of vertices and edges
(Data-Driven)**

Data-Driven Algorithms



Economy Premium

Affordable rides, all to yourself

Pool Options  UberX

\$5.70
3:16pm

\$9.71
3:06pm

Graph Execution Hardware



CPU



GPU




Xeon Phi




Distributed Cluster

Push Traversal




Push Traversal




Incurs overhead with atomics

Traverses no extra edges

Pull Traversal




Pull Traversal





Incurs no overhead from atomics

Traverses extra edges



Partitioning



Partitioning




Partitioning




Improves locality

Needs extra instructions to traverse two graphs


Optimization Tradeoff Space



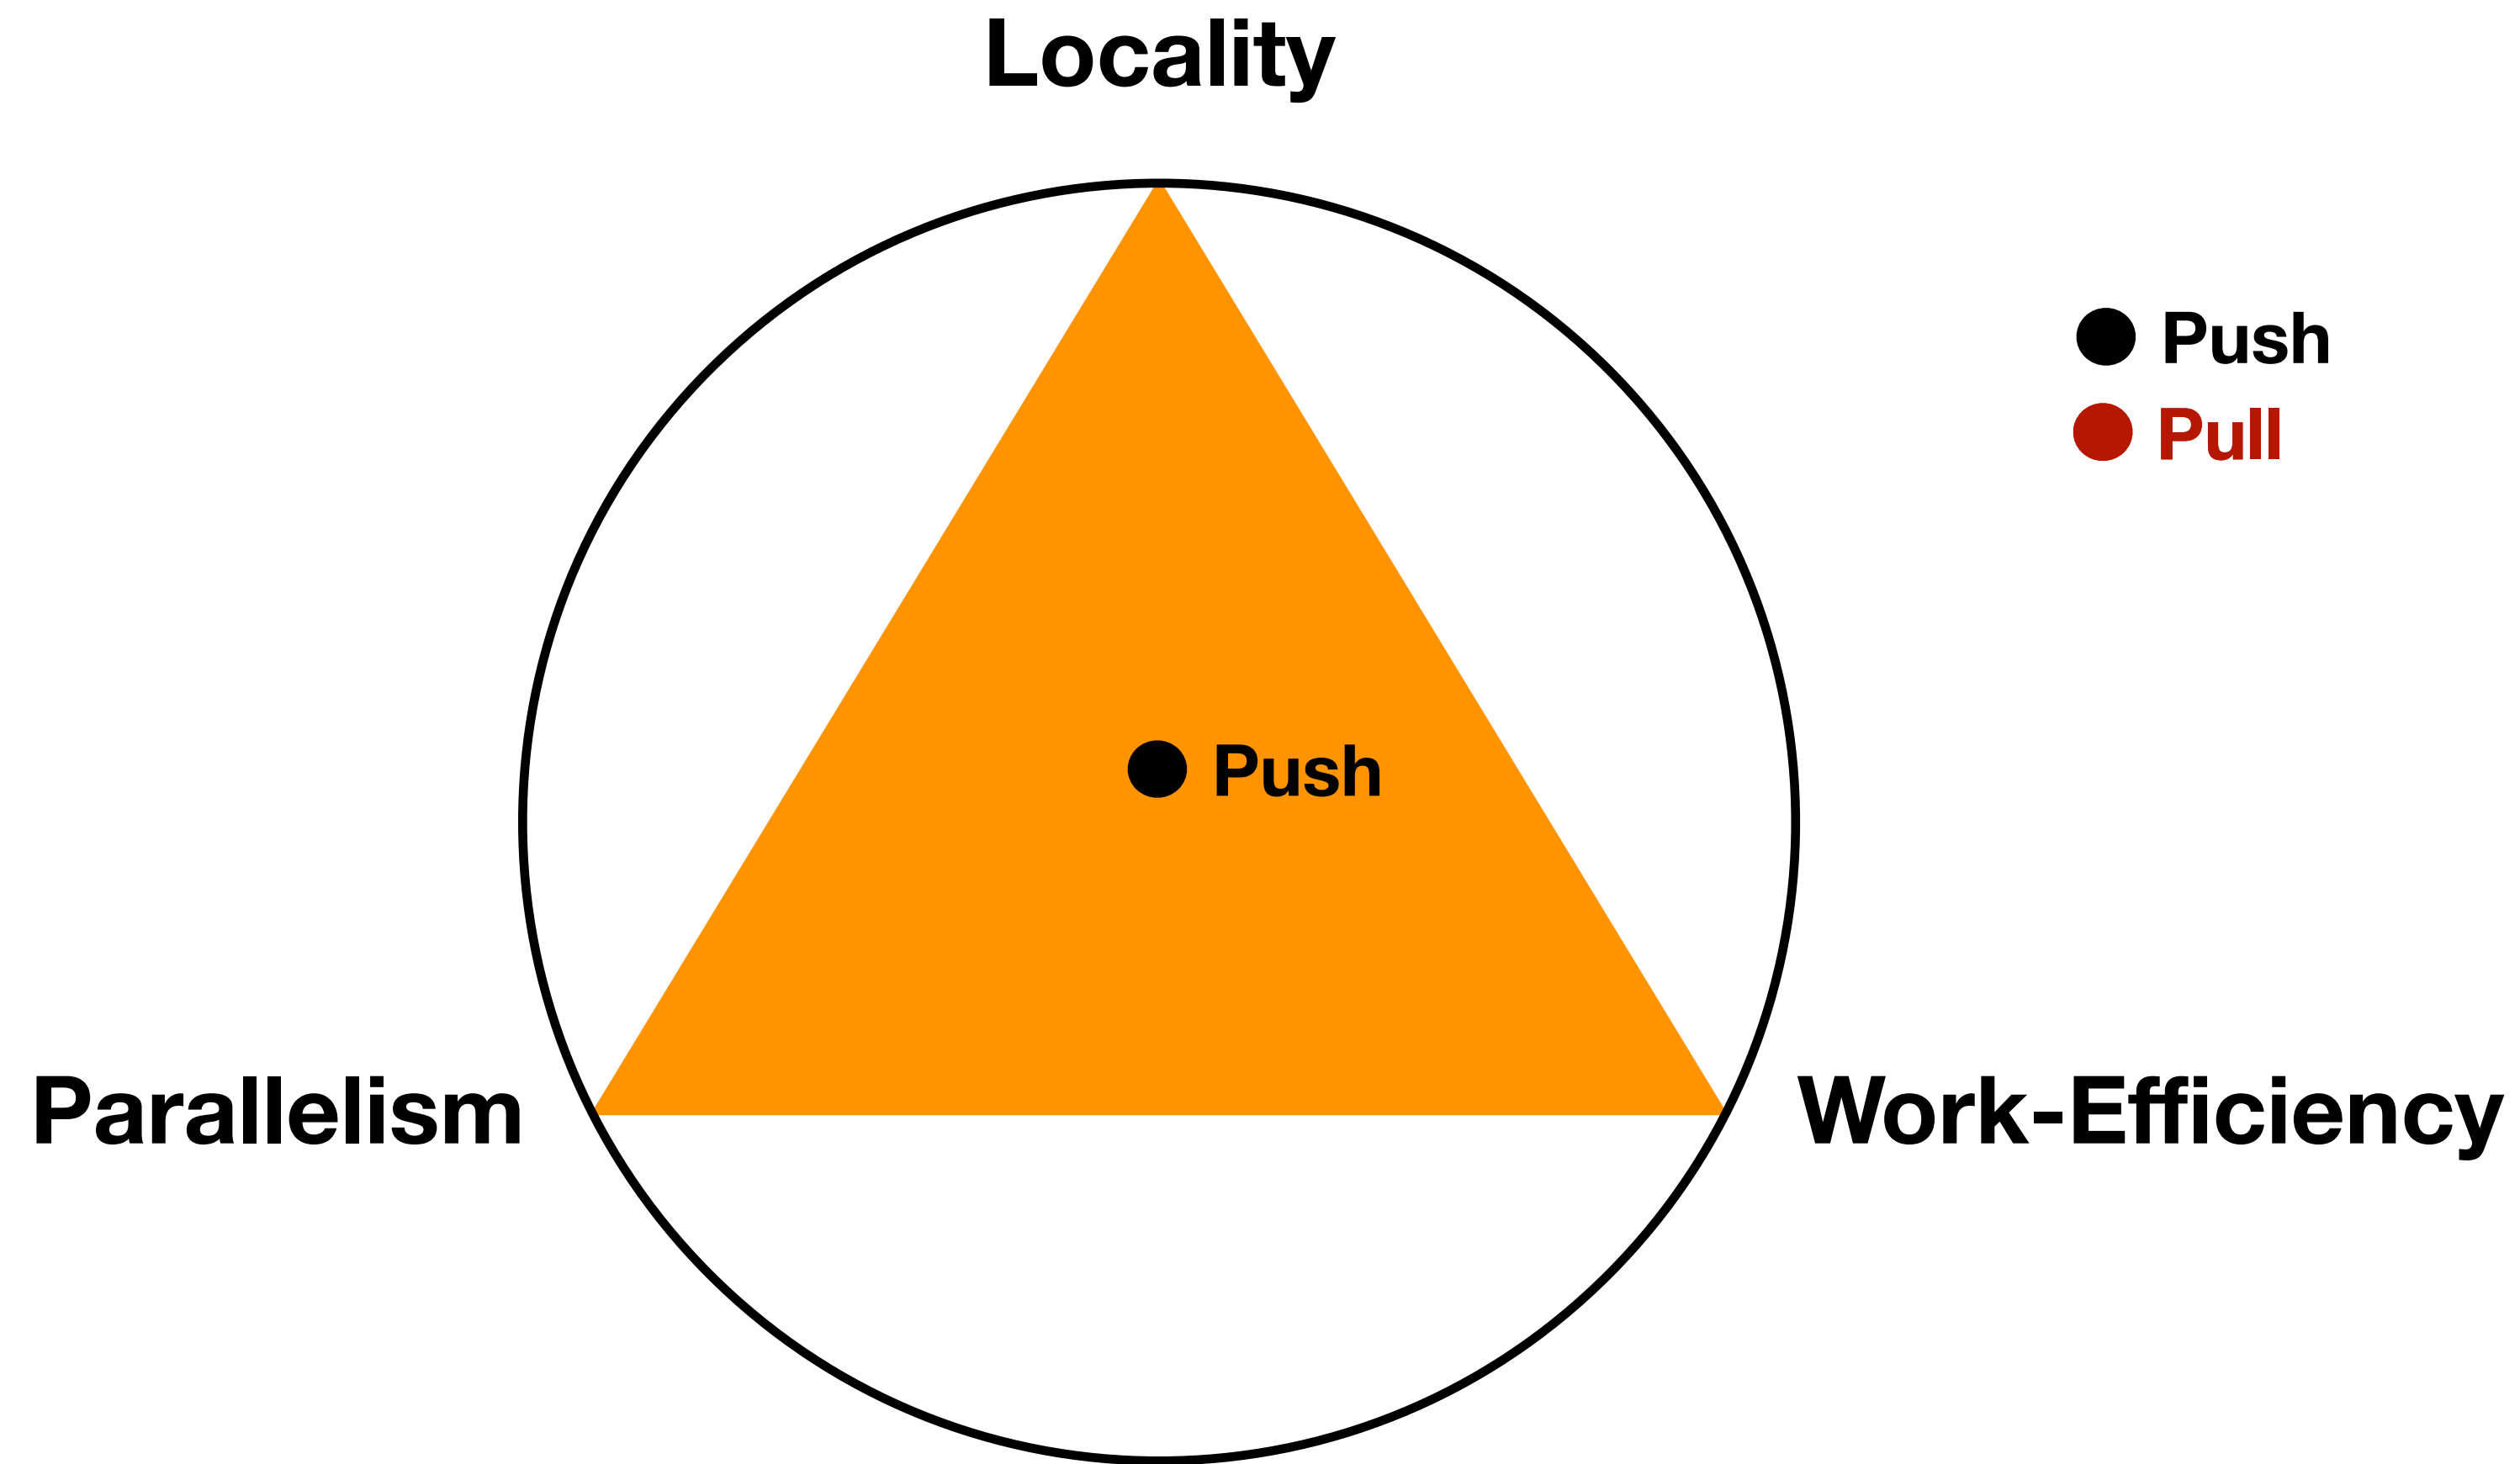
Optimization Tradeoff Space



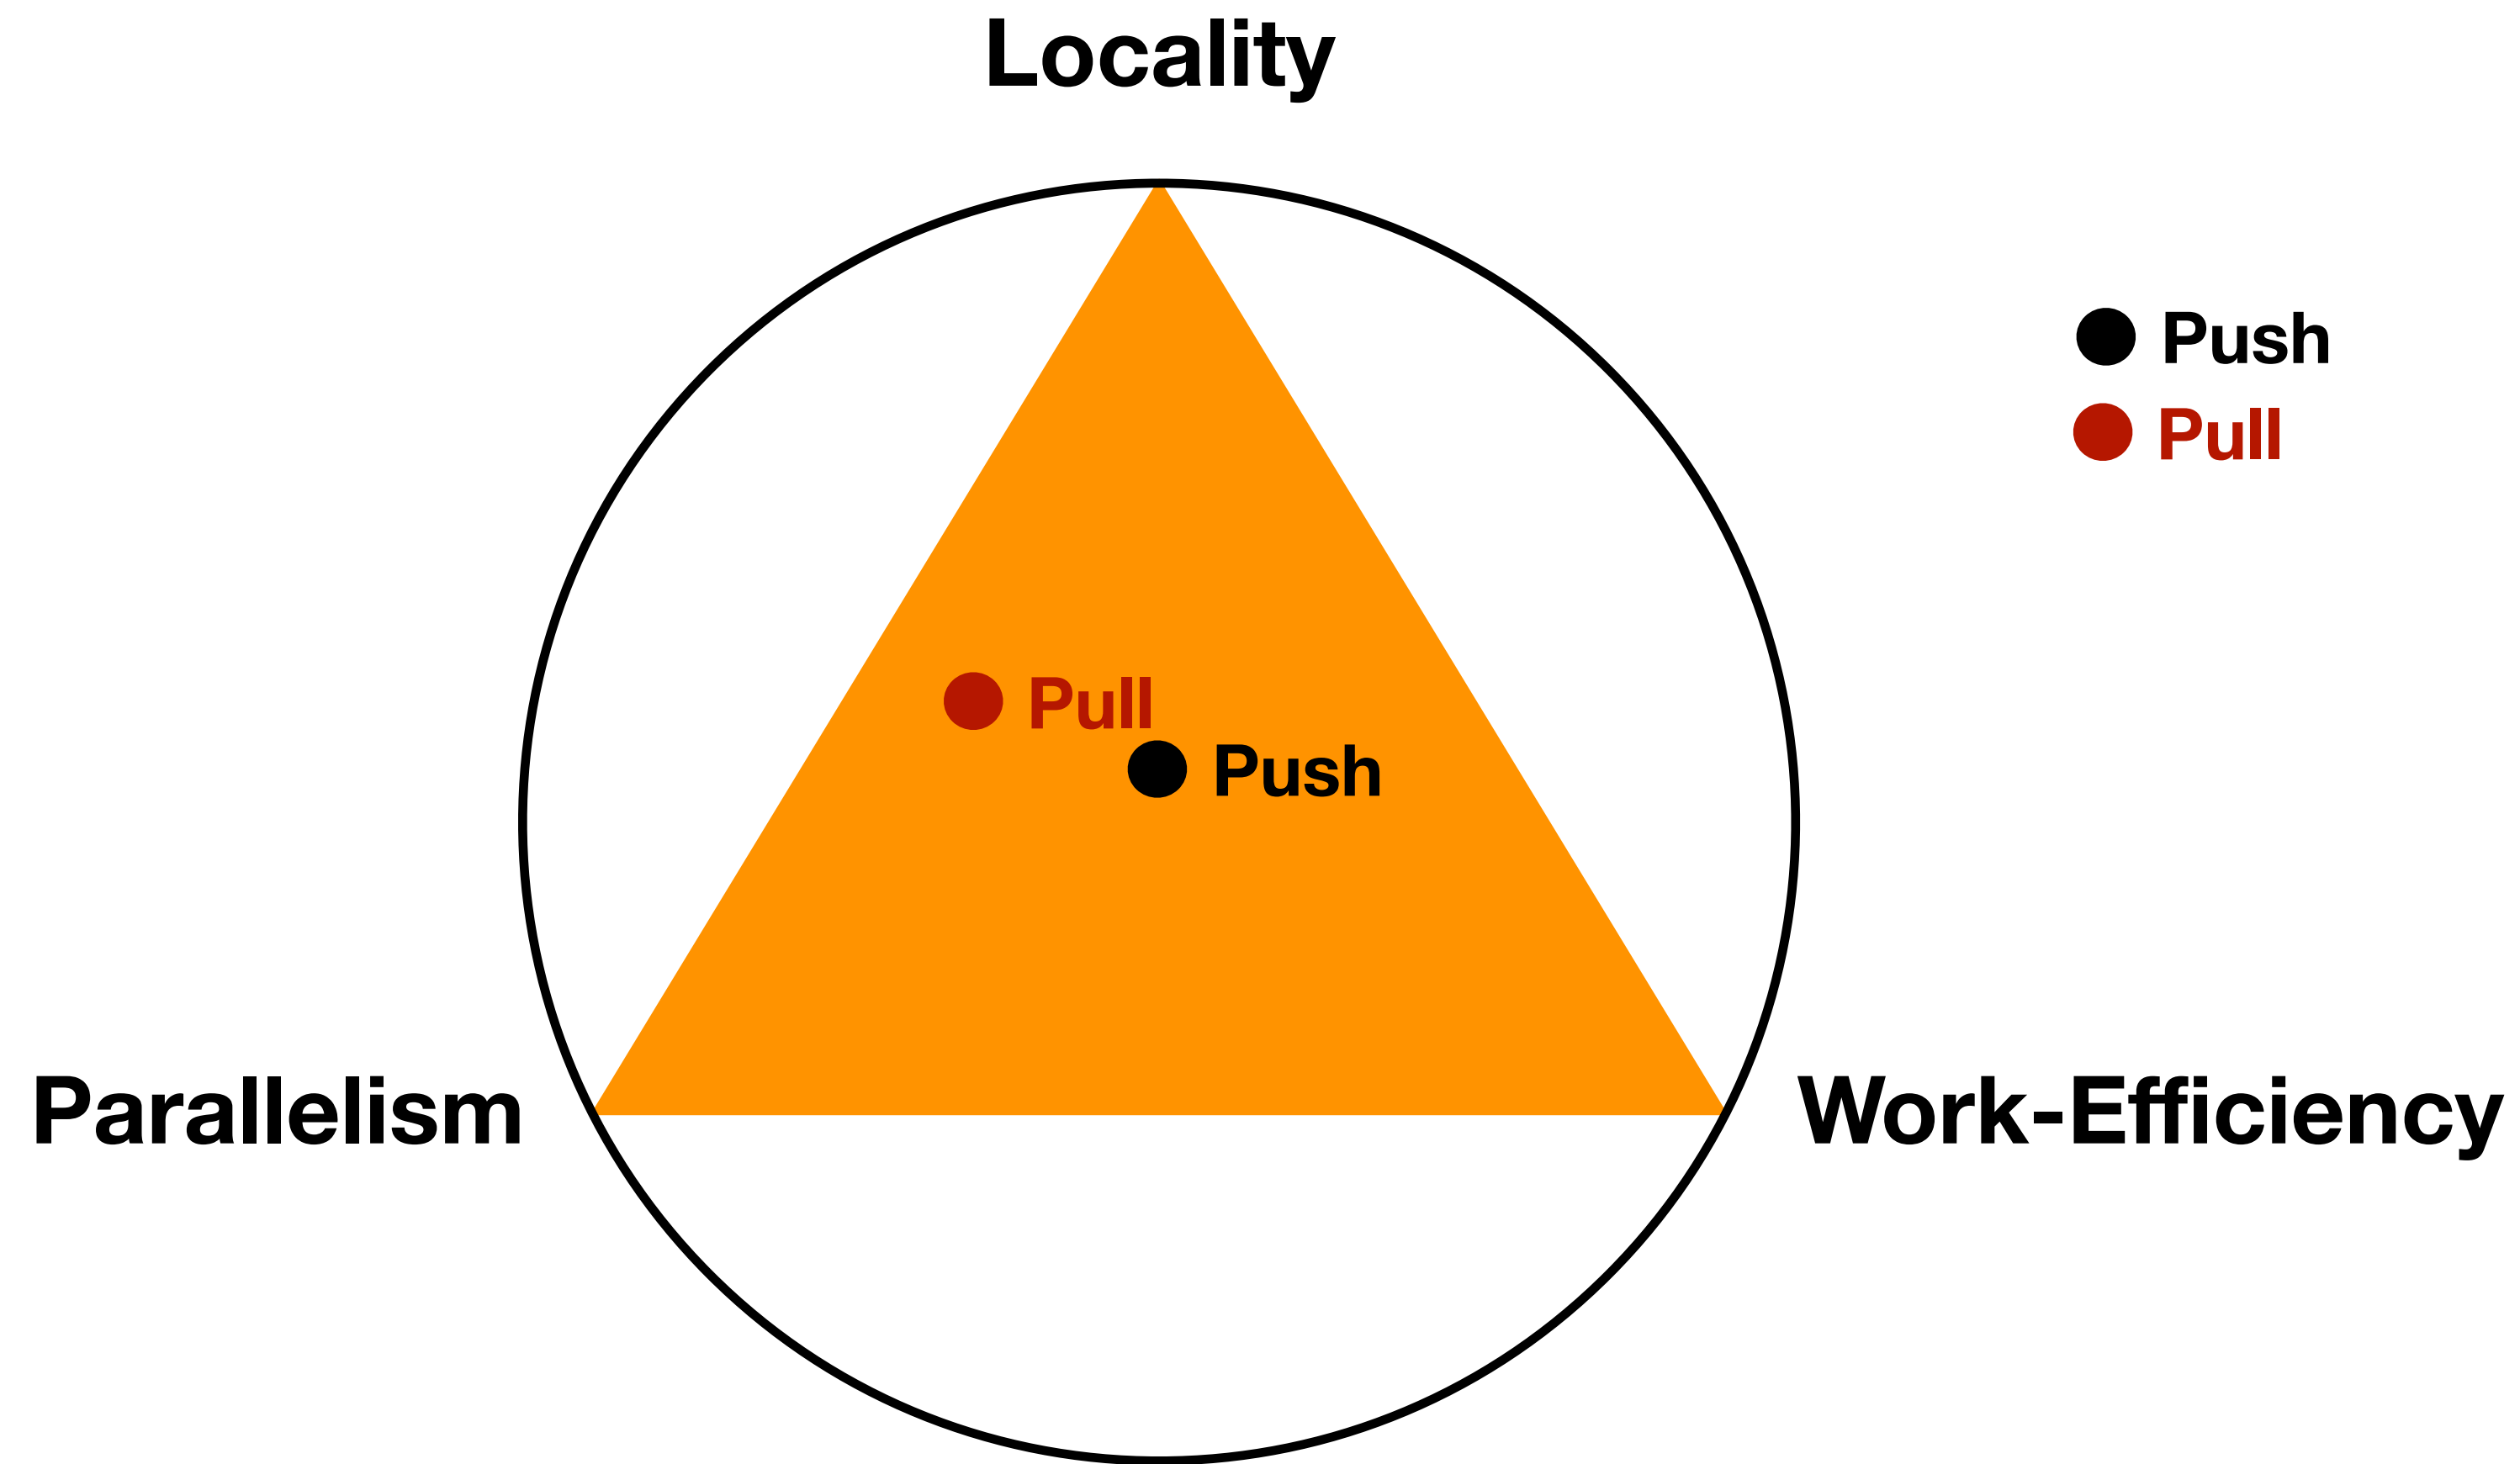
Optimization Tradeoff Space




Optimization Tradeoff Space



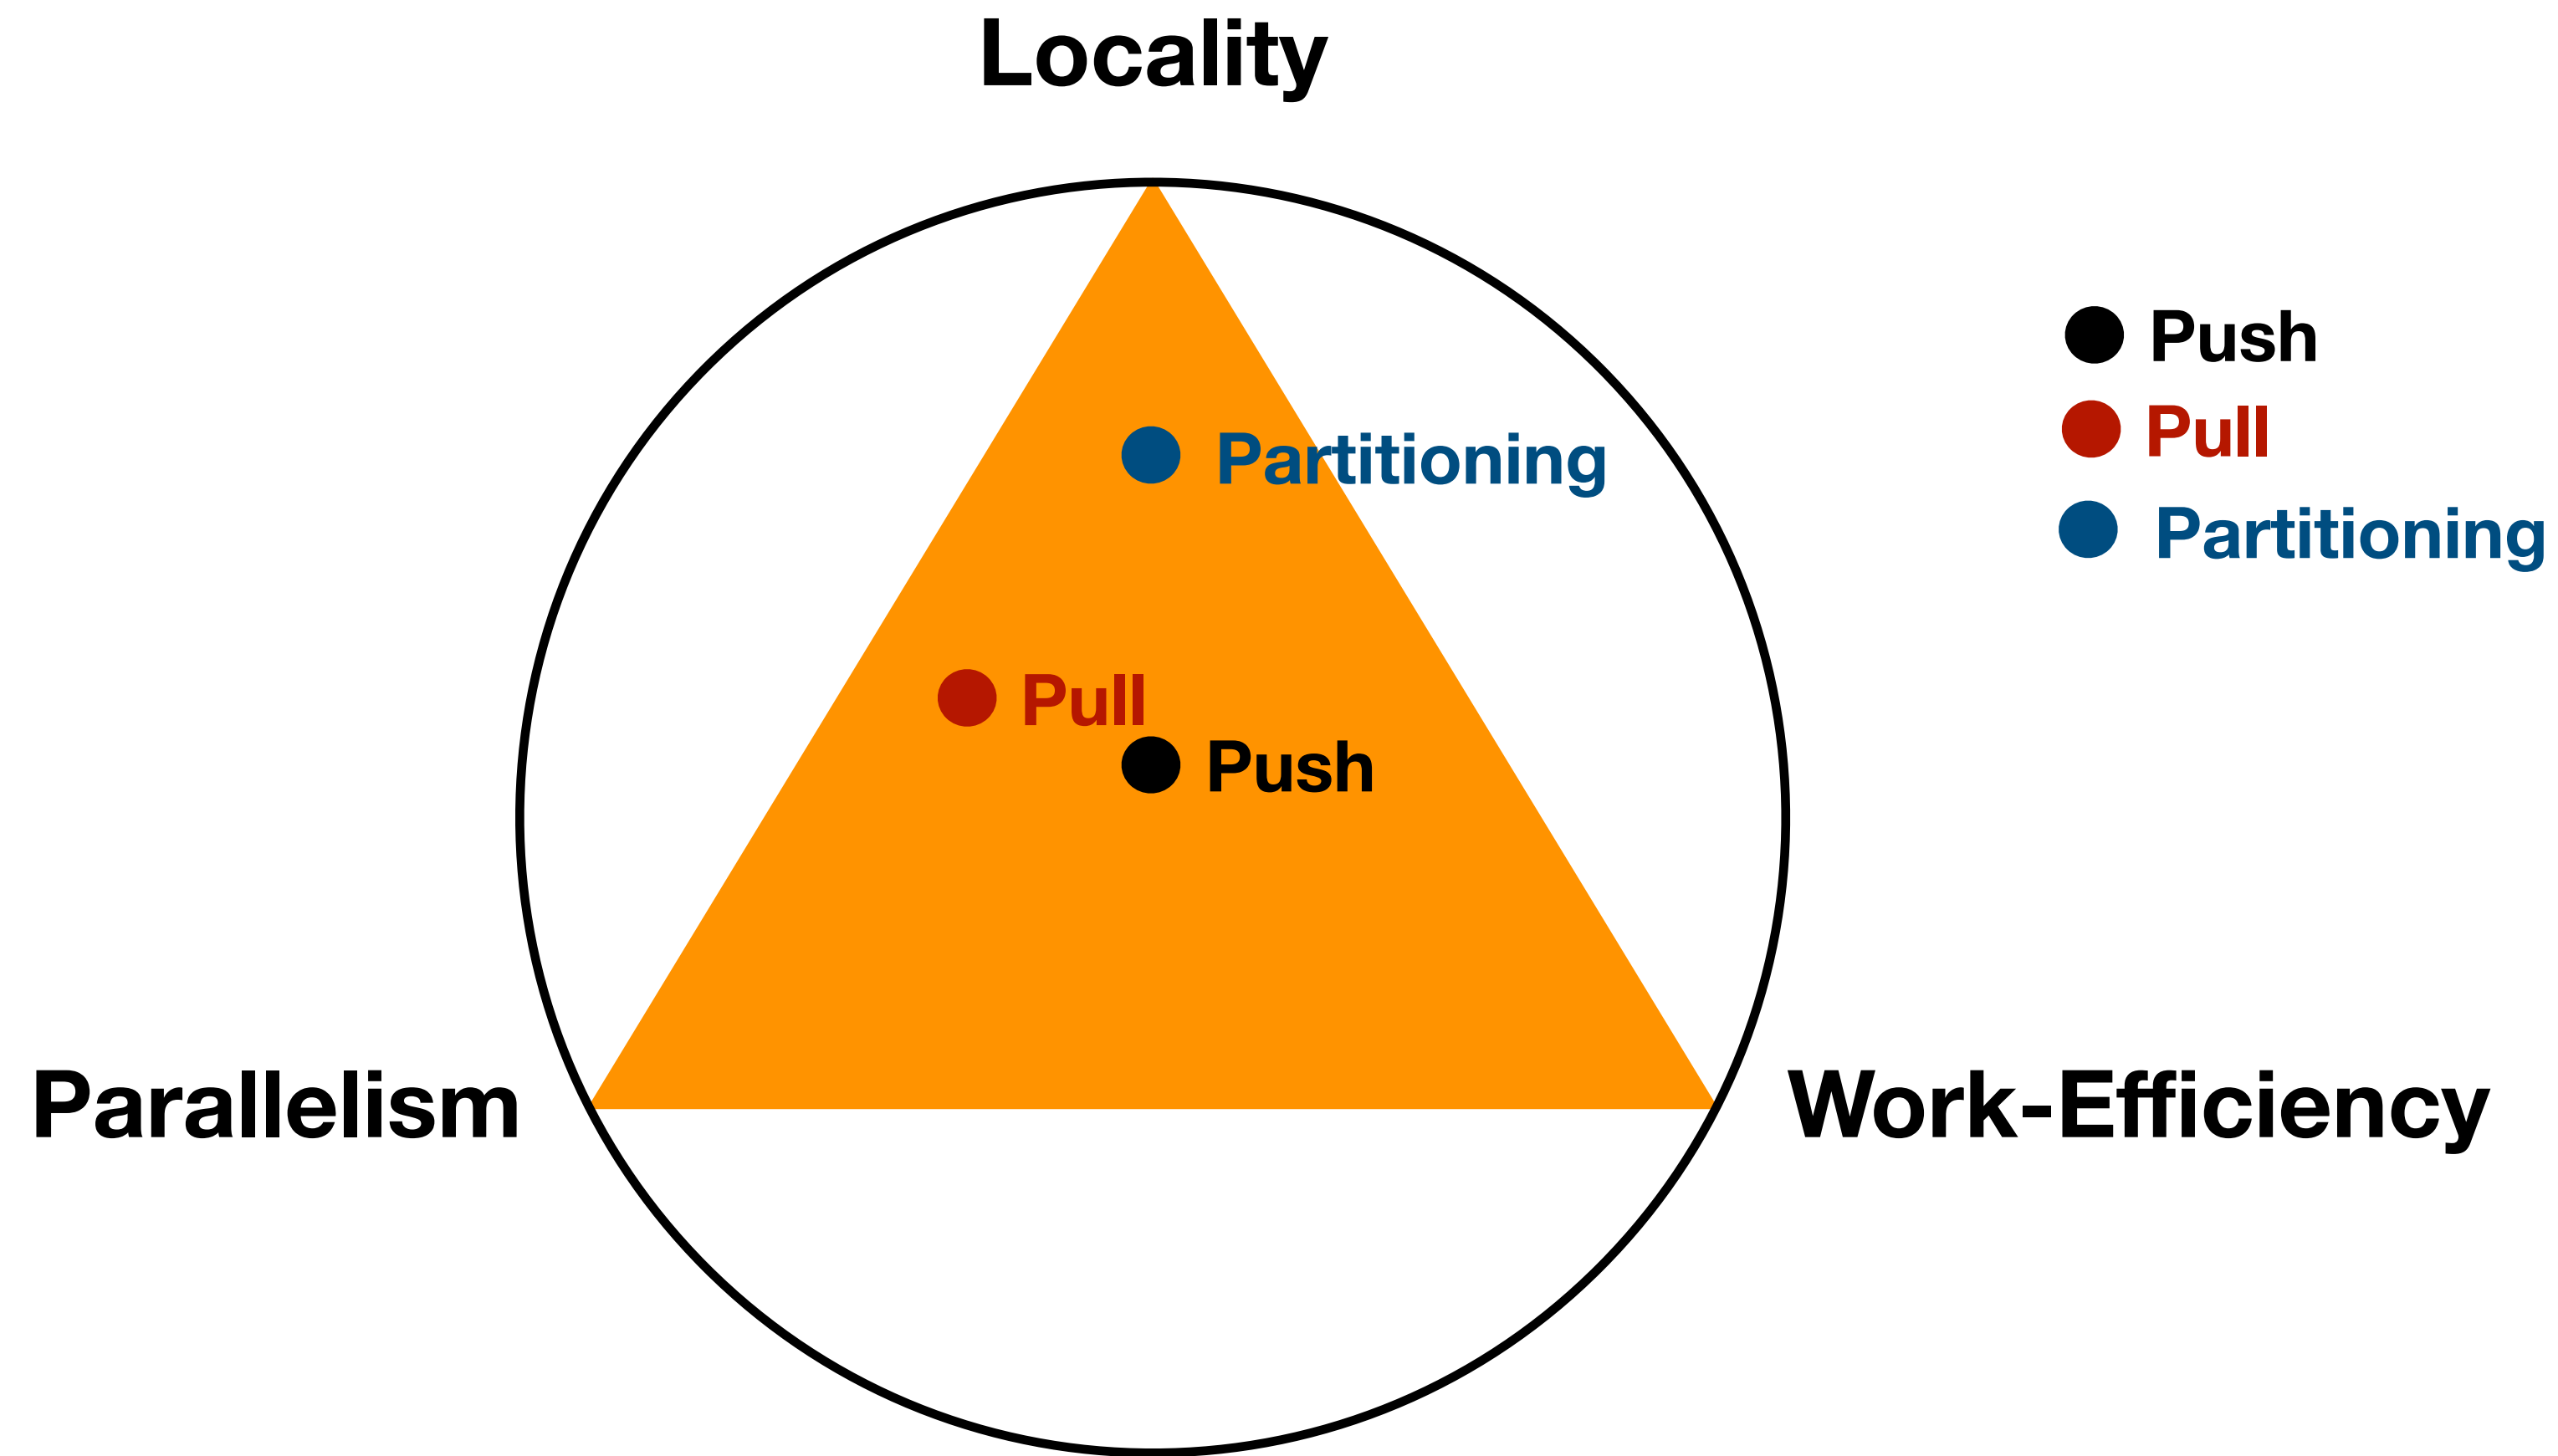
Optimization Tradeoff Space




Optimization Tradeoff Space




Optimization Tradeoff Space




Optimization Tradeoff Space







Optimizations




Graphs




Optimizations




Graphs




Algorithms




Optimizations







Graphs



Algorithms




Optimizations




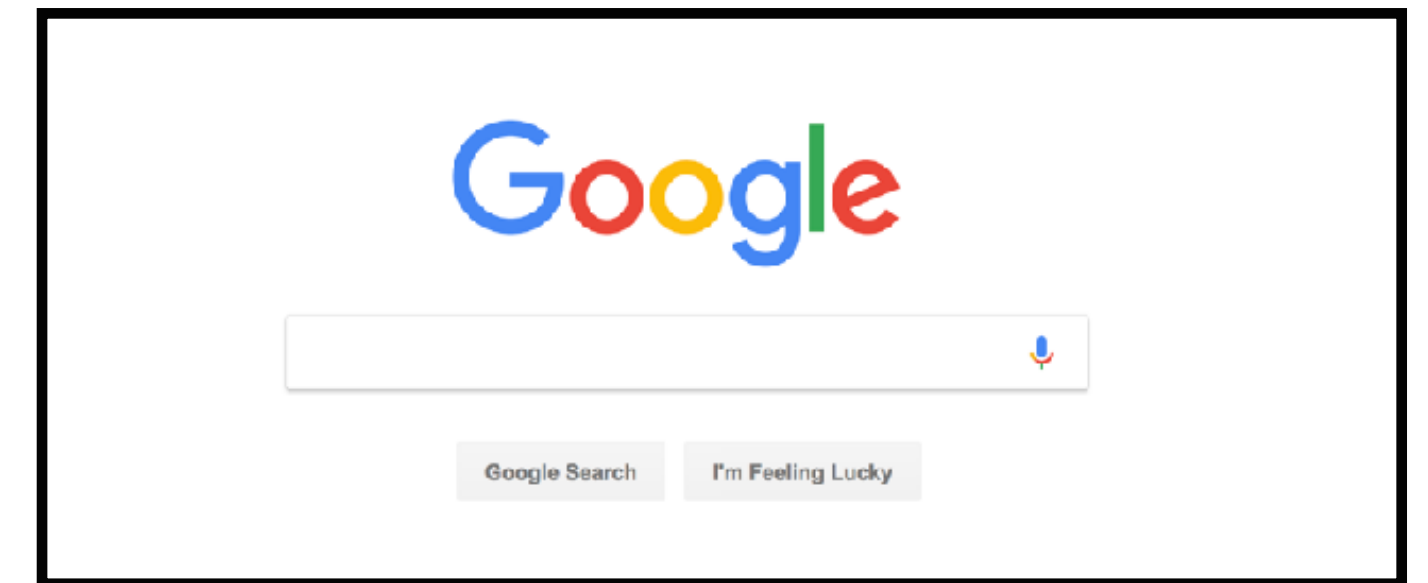
Hardware






Graphs

PageRank






Algorithms




Optimizations


- Pull**
- Partitioning**
- Vertex-Parallel**

Hardware

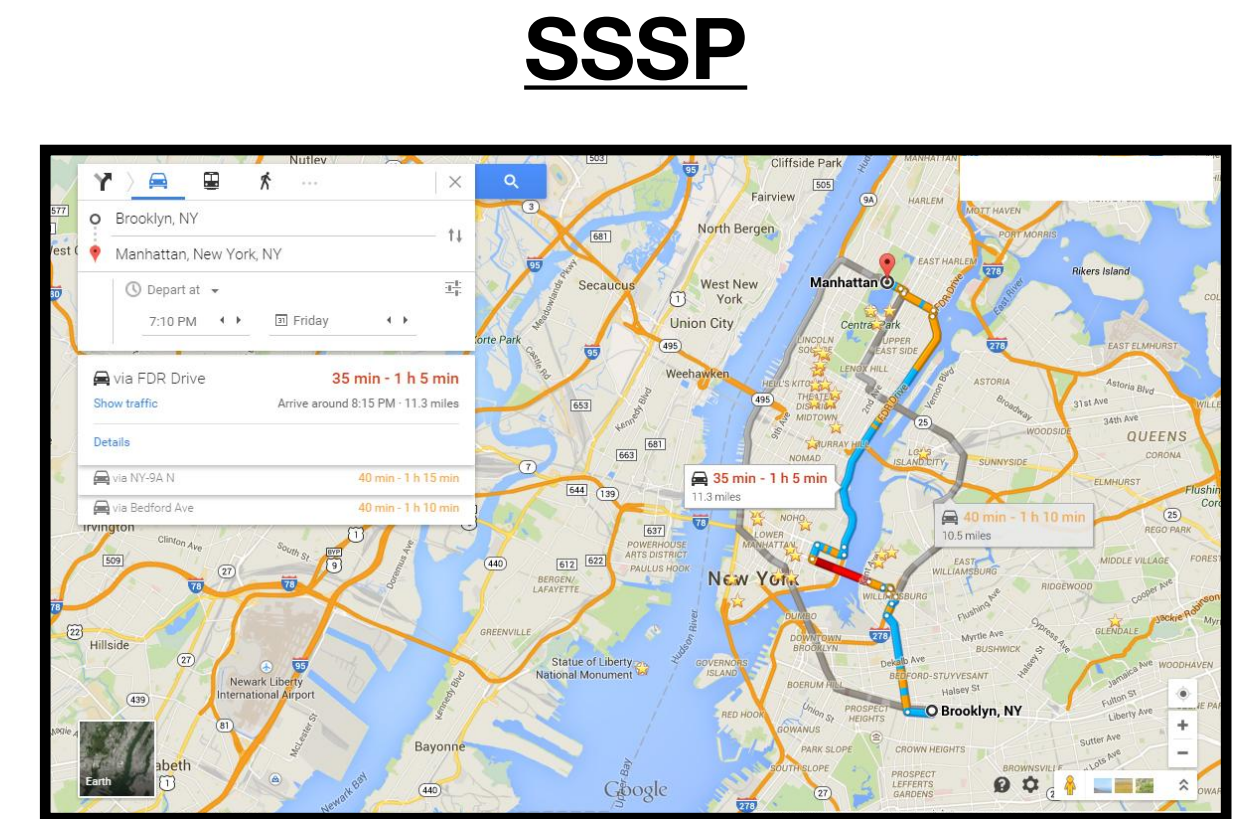







Graphs



Algorithms




Hardware




Optimizations

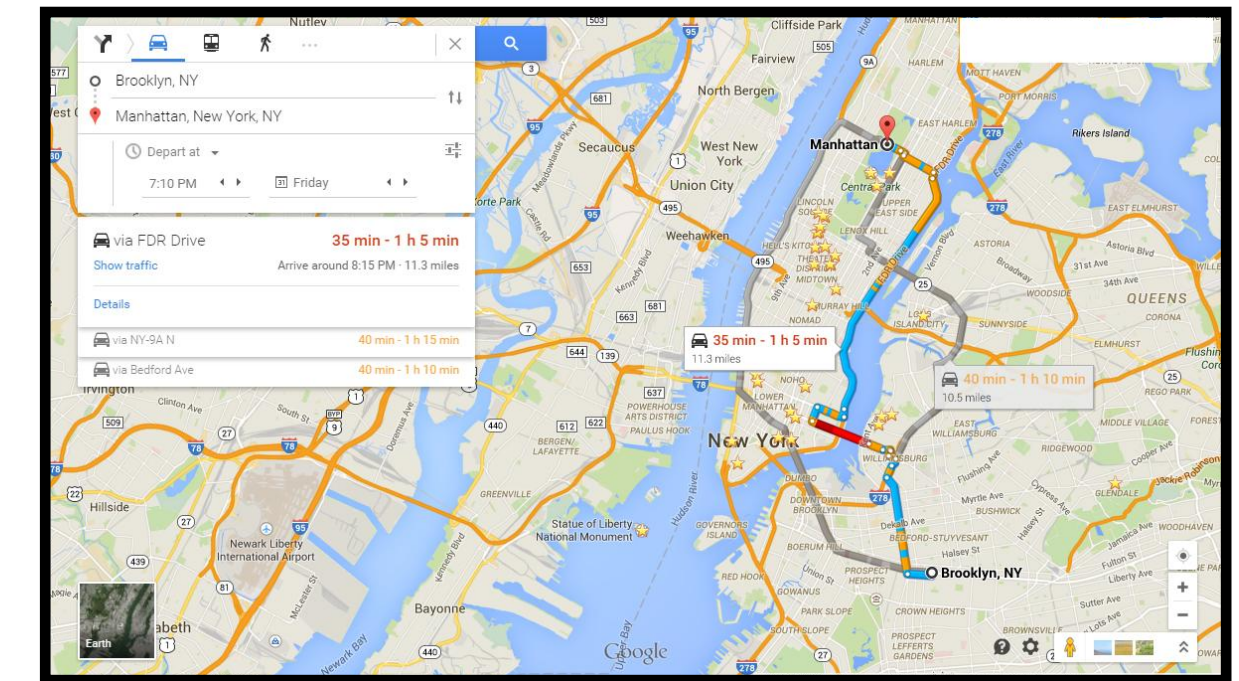

Push
Vertex-Parallel



Graphs





Algorithms




SSSP

Hardware




Optimizations


Push
Edge-Parallel




Graphs



Algorithms




**Bad
optimizations
(schedules) can
be > 100x slower**





Optimizations

**Pull
Partitioning
Vertex-Parallel**

SSSP



Hardware



GraphIt

A Domain-Specific Language for Graph Applications


- **Decouple algorithm from optimization for graph applications**
 - **Algorithm:** What to Compute
 - **Optimization (schedule):** How to Compute

GraphIt


A Domain-Specific Language for Graph Applications

- **Decouple algorithm from optimization for graph applications**
 - **Algorithm:** What to Compute
 - **Optimization (schedule):** How to Compute
- **DSL and Compiler**
 - **Ease-of-Use:** Improve Productivity for Data Scientists and Library Developers
 - **Performance:** Beat hand-optimized libraries by up to 4.8x on CPU and GPU
 - **Portability:** Working with NVIDIA, UW, Cornell to develop new backends for GPU and Domain-Specific Accelerators

GraphIt DSL



**Algorithm Representation
(Algorithm Language)**




Optimization Representation

- **Scheduling Language**
- **Schedule Representation
(e.g. Graph Iteration Space)**




Autotuner

GraphIt DSL




**Algorithm Representation
(Algorithm Language)**






Optimization Representation

- Scheduling Language
- Schedule Representation
(e.g. Graph Iteration Space)






Autotuner


Algorithm Language




Algorithm Language




Algorithm Language




`edges.apply(func)`




```
edges.from(vertexset)
  .to(vertexset)
  .srcFilter(func)
  .dstFilter(func)
  .apply(func)
```




Algorithm Language



`edges.apply(func)`




`edges.from(vertexset)`
`.to(vertexset)`
`.srcFilter(func)`
`.dstFilter(func)`
`.apply(func)`




`vertices.apply(func)`

PageRank Example




```
func updateEdge (src: Vertex, dst: Vertex)  
    new_rank[dst] += old_rank[src] / out_degree[src]  
end
```


PageRank Example





```
func updateEdge (src: Vertex, dst: Vertex)  
    new_rank[dst] += old_rank[src] / out_degree[src]  
end
```

```
func updateVertex (v: Vertex)  
    new_rank[v] = base_score + 0.85*new_rank[v];  
    old_rank[v] = new_rank[v];  
    new_rank[v] = 0;  
end
```



PageRank Example





```
func updateEdge (src: Vertex, dst: Vertex)
    new_rank[dst] += old_rank[src] / out_degree[src]
end
```

```
func updateVertex (v: Vertex)
    new_rank[v] = base_score + 0.85*new_rank[v];
    old_rank[v] = new_rank[v];
    new_rank[v] = 0;
end
```

```
func main()
    for i in 1:max_iter
        #s1# edges.apply(updateEdge);
        vertices.apply(updateVertex);
    end
end
```

PageRank Example





```
func updateEdge (src: Vertex, dst: Vertex)
  new_rank[dst] += old_rank[src] / out_degree[src]
end
```

```
func updateVertex (v: Vertex)
  new_rank[v] = base_score + 0.85*new_rank[v];
  old_rank[v] = new_rank[v];
  new_rank[v] = 0;
end
```

```
func main()
  for i in 1:max iter
    #s1# edges.apply(updateEdge);
    vertices.apply(updateVertex);
  end
end
```

PageRank Example




```
func updateEdge (src: Vertex, dst: Vertex)
    new_rank[dst] += old_rank[src] / out_degree[src]
end
```


```
func updateVertex (v: Vertex)
    new_rank[v] = base_score + 0.85*new_rank[v];
    old_rank[v] = new_rank[v];
    new_rank[v] = 0;
end
```

```
func main()
    for i in 1:max_iter
        #s1# edges.apply(updateEdge);
        vertices.apply(updateVertex);
    end
end
```

GraphIt DSL




Algorithm Representation
(Algorithm Language)




Optimization Representation

- Scheduling Language
- Schedule Representation
(e.g. Graph Iteration Space)




Autotuner

GraphIt DSL




Algorithm Representation
(Algorithm Language)



Optimization Representation

- **Scheduling Language**
- **Schedule Representation**
(e.g. Graph Iteration Space)



Autotuner

Scheduling Language

```
func updateEdge (src: Vertex, dst: Vertex)  
    new_rank[dst] += old_rank[src] / out_degree[src]  
end
```

```
func updateVertex (v: Vertex)  
    new_rank[v] = base_score + 0.85*new_rank[v];  
    old_rank[v] = new_rank[v];  
    new_rank[v] = 0;  
end
```

```
func main()  
    for i in 1:max_iter  
        #s1# edges.apply(updateEdge);  
        vertices.apply(updateVertex);  
    end  
end
```

Scheduling Language

```
func updateEdge (src: Vertex, dst: Vertex)
  new_rank[dst] += old_rank[src] / out_degree[src]
end
```

```
func updateVertex (v: Vertex)
  new_rank[v] = base_score + 0.85*new_rank[v];
  old_rank[v] = new_rank[v];
  new_rank[v] = 0;
end
```

```
func main()
  for i in 1:max_iter
    #s1# edges.apply(updateEdge);
    vertices.apply(updateVertex);
  end
end
```


Scheduling Language

Algorithm Specification

```
func updateEdge (src: Vertex, dst: Vertex)
  new_rank[dst] += old_rank[src] / out_degree[src]
end

func updateVertex (v: Vertex)
  new_rank[v] = base_score + 0.85*new_rank[v];
  old_rank[v] = new_rank[v];
  new_rank[v] = 0;
end

func main()
  for i in 1:max_iter
    #s1# edges.apply(updateEdge);
    vertices.apply(updateVertex);
  end
end
```

Schedule 1

Algorithm Specification

```
func updateEdge (src: Vertex, dst: Vertex)
  new_rank[dst] += old_rank[src] / out_degree[src]
end

func updateVertex (v: Vertex)
  new_rank[v] = base_score + 0.85*new_rank[v];
  old_rank[v] = new_rank[v];
  new_rank[v] = 0;
end

func main()
  for i in 1:max_iter
    #s1# edges.apply(updateEdge);
    vertices.apply(updateVertex);
  end
end
```

Scheduling Functions

```
schedule:
  program->configApplyDirection("s1", "SparsePush");
```

Schedule 1

Algorithm Specification

```
func updateEdge (src: Vertex, dst: Vertex)
  new_rank[dst] += old_rank[src] / out_degree[src]
end

func updateVertex (v: Vertex)
  new_rank[v] = base_score + 0.85*new_rank[v];
  old_rank[v] = new_rank[v];
  new_rank[v] = 0;
end

func main()
  for i in 1:max_iter
    #s1# edges.apply(updateEdge);
    vertices.apply(updateVertex);
  end
end
```

Scheduling Functions

```
schedule:
  program->configApplyDirection("s1", "SparsePush");
```

Schedule 1

Algorithm Specification

```
func updateEdge (src: Vertex, dst: Vertex)
  new_rank[dst] += old_rank[src] / out_degree[src]
end

func updateVertex (v: Vertex)
  new_rank[v] = base_score + 0.85*new_rank[v];
  old_rank[v] = new_rank[v];
  new_rank[v] = 0;
end

func main()
  for i in 1:max_iter
    #s1# edges.apply(updateEdge);
    vertices.apply(updateVertex);
  end
end
```

Pseudo Generated Code

```
double * new_rank = new double[num_verts];
double * old_rank = new double[num_verts];
int * out_degree = new int[num_verts];

...

for (NodeID src : vertices) {
  for(NodeID dst : G.getOutNgh(src)){
    new_rank[dst] += old_rank[src] / out_degree[src];
  }
}

....
```

Scheduling Functions

```
schedule:
  program->configApplyDirection("s1", "SparsePush");
```

Schedule 2

Algorithm Specification

```
func updateEdge (src: Vertex, dst: Vertex)
  new_rank[dst] += old_rank[src] / out_degree[src]
end

func updateVertex (v: Vertex)
  new_rank[v] = base_score + 0.85*new_rank[v];
  old_rank[v] = new_rank[v];
  new_rank[v] = 0;
end

func main()
  for i in 1:max_iter
    #s1# edges.apply(updateEdge);
    vertices.apply(updateVertex);
  end
end
```

Pseudo Generated Code

```
double * new_rank = new double[num_verts];
double * old_rank = new double[num_verts];
int * out_degree = new int[num_verts];

...

parallel_for (NodeID src : vertices) {
  for(NodeID dst : G.getOutNgh(src)){
    atomic_add (&new_rank[dst],
               old_rank[src] / out_degree[src] );
  }
}

....
```

Scheduling Functions

```
schedule:
  program->configApplyDirection("s1", "SparsePush");
  program->configApplyParallelization("s1", "dynamic-vertex-parallel");
```

Schedule 3

Algorithm Specification

```
func updateEdge (src: Vertex, dst: Vertex)
  new_rank[dst] += old_rank[src] / out_degree[src]
end

func updateVertex (v: Vertex)
  new_rank[v] = base_score + 0.85*new_rank[v];
  old_rank[v] = new_rank[v];
  new_rank[v] = 0;
end

func main()
  for i in 1:max_iter
    #s1# edges.apply(updateEdge);
    vertices.apply(updateVertex);
  end
end
```

Pseudo Generated Code

```
double * new_rank = new double[num_verts];
double * old_rank = new double[num_verts];
int * out_degree = new int[num_verts];

...

parallel_for (NodeID dst : vertices) {
  for(NodeID src : G.getInNgh(dst)){
    new_rank[dst] += old_rank[src] / out_degree[src];
  }
}

....
```

Scheduling Functions

```
schedule:
  program->configApplyDirection("s1", "DensePull");
  program->configApplyParallelization("s1", "dynamic-vertex-parallel");
```

Schedule 4

Algorithm Specification

```
func updateEdge (src: Vertex, dst: Vertex)
  new_rank[dst] += old_rank[src] / out_degree[src]
end

func updateVertex (v: Vertex)
  new_rank[v] = base_score + 0.85*new_rank[v];
  old_rank[v] = new_rank[v];
  new_rank[v] = 0;
end

func main()
  for i in 1:max_iter
    #s1# edges.apply(updateEdge);
    vertices.apply(updateVertex);
  end
end
```

Pseudo Generated Code


```
double * new_rank = new double[num_verts];
double * old_rank = new double[num_verts];
int * out_degree = new int[num_verts];

...
for (Subgraph sg : G.subgraphs) {
  parallel_for (NodeID dst : verticesa) {
    for(NodeID src : G.getInNgh(dst)){
      new_rank[dst] += old_rank[src] / out_degree[src];
    }
  }
}
....
```

Scheduling Functions

```
schedule:
  program->configApplyDirection("s1", "DensePull");
  program->configApplyParallelization("s1", "dynamic-vertex-parallel");
  program->configApplyNumSSG("s1", "fixed-vertex-count", 10);
```

Speedups of Schedules




Performance of PageRank on Twitter Graph on Intel Xeon E5-2695 v3 CPUs with 12 cores each for a total of 24 cores and 48 hyper-threads.


Many More Optimizations

- **Direction optimizations (configApplyDirection),**
 - **SparsePush, DensePush, DensePull, DensePull-SparsePush, DensePush-SparsePush**
- **Parallelization strategies (configApplyParallelization)**
 - **serial, dynamic-vertex-parallel, static-vertexparallel, edge-aware-dynamic-vertex-parallel, edge-parallel**
- **Cache (configApplyNumSSG)**
 - **fixed-vertex-count, edge-aware-vertexcount**
- **NUMA (configApplyNUMA)**
 - **serial, static-parallel, dynamic-parallel**
- **AoS, SoA (fuseFields)**
- **Vertexset data layout (configApplyDenseVertexSet)**
 - **bitvector, boolean array**

GraphIt DSL




Algorithm Representation
(Algorithm Language)



Optimization Representation

- Scheduling Language
- Schedule Representation
(e.g. Graph Iteration Space)



Autotuner

Schedule Representation

Algorithm Specification


```
func updateEdge (src: Vertex, dst: Vertex)
  new_rank[dst] += old_rank[src] / out_degree[src]
end

func updateVertex (v: Vertex)
  new_rank[v] = base_score + 0.85*new_rank[v];
  old_rank[v] = new_rank[v];
  new_rank[v] = 0;
end

func main()
  for i in 1:11
    #s1# edges.apply(updateEdge);
    vertices.apply(updateVertex);
  end
end
```

Scheduling Functions

```
schedule:
  program->configApplyDirection("s1", "SparsePush");
  program->configApplyParallelization("s1", "dynamic-vertex-parallel");
```



Schedule Representation

Algorithm Specification

```
func updateEdge (src: Vertex, dst: Vertex)
  new_rank[dst] += old_rank[src] / out_degree[src]
end

func updateVertex (v: Vertex)
  new_rank[v] = base_score + 0.85*new_rank[v];
  old_rank[v] = new_rank[v];
  new_rank[v] = 0;
end

func main()
  for i in 1:11
    #s1# edges.apply(updateEdge);
    vertices.apply(updateVertex);
  end
end
```

Scheduling Functions

```
schedule:
  program->configApplyDirection("s1", "SparsePush");
  program->configApplyParallelization("s1", "dynamic-vertex-parallel");
```

GraphIt Compiler

```
graph TD
    A[Algorithm Specification] --> C[GraphIt Compiler]
    B[Scheduling Functions] --> C
    C --> D[Pseudo Generated Code]
```

Pseudo Generated Code

```
...
parallel_for (NodeID src : vertices) {
  for(NodeID dst : G.getOutNgh(src)){
    new_rank[dst] = atomic_add ( new_rank[dst] ,
                                old_rank[src] / out_degree[src] );
  }
}
....
```

Schedule Representation

Algorithm Specification

```
func updateEdge (src: Vertex, dst: Vertex)
  new_rank[dst] += old_rank[src] / out_degree[src]
end

func updateVertex (v: Vertex)
  new_rank[v] = base_score + 0.85*new_rank[v];
  old_rank[v] = new_rank[v];
  new_rank[v] = 0;
end

func main()
  for i in 1:11
    #s1# edges.apply(updateEdge);
    vertices.apply(updateVertex);
  end
end
```

Scheduling Functions

```
schedule:
  program->configApplyDirection("s1", "SparsePush");
  program->configApplyParallelization("s1", "dynamic-vertex-parallel");
```

Internal Schedule Representation

Pseudo Generated Code

```
...
parallel_for (NodeID src : vertices) {
  for(NodeID dst : G.getOutNgh(src)){
    new_rank[dst] = atomic_add ( new_rank[dst] ,
                                old_rank[src] / out_degree[src] );
  }
}
....
```

Schedule Representation

Algorithm Specification

```
func updateEdge (src: Vertex, dst: Vertex)
  new_rank[dst] += old_rank[src] / out_degree[src]
end

func updateVertex (v: Vertex)
  new_rank[v] = base_score + 0.85*new_rank[v];
  old_rank[v] = new_rank[v];
  new_rank[v] = 0;
end

func main()
  for i in 1:11
    #s1# edges.apply(updateEdge);
    vertices.apply(updateVertex);
  end
end
```

Scheduling Functions

```
schedule:
  program->configApplyDirection("s1", "SparsePush");
  program->configApplyParallelization("s1", "dynamic-vertex-parallel");
```

Internal Schedule Representation


Reason about Optimization Composition, Validity, and Code Generation

Pseudo Generated Code

```
...
parallel_for (NodeID src : vertices) {
  for(NodeID dst : G.getOutNgh(src)){
    new_rank[dst] = atomic_add ( new_rank[dst] ,
                                old_rank[src] / out_degree[src] );
  }
}
....
```

Iteration Spaces


- Very versatile representation of dense loops and arrays
- Used for:
 - Program analyses
 - Composition of complex loop transformations
- Framework for code generation



Iteration Spaces

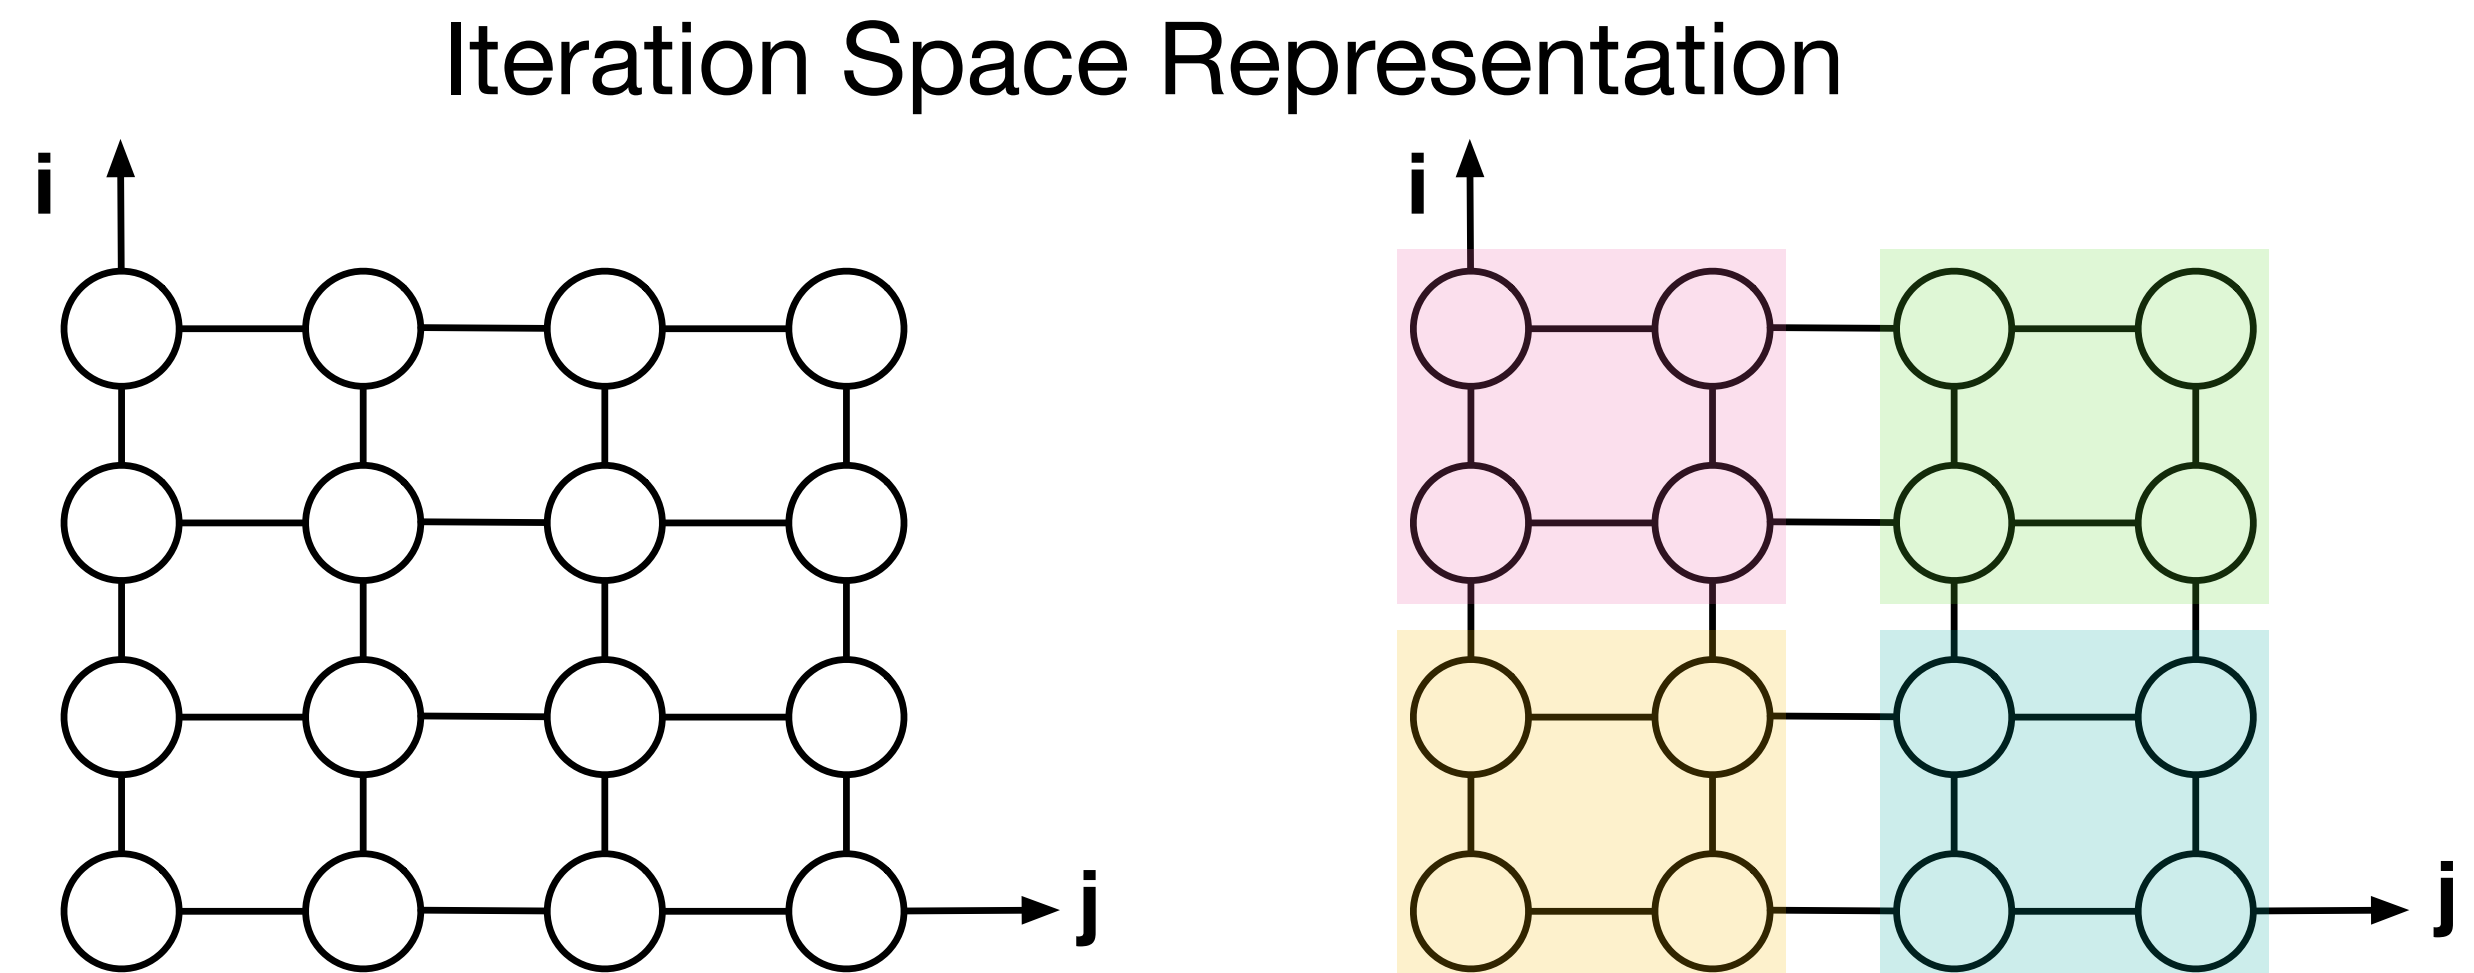
- Very versatile representation of dense loops and arrays
- Used for:
 - Program analyses
 - Composition of complex loop transformations
 - Framework for code generation

We extend it to sparse loops



Graph Iteration Space

- Very versatile representation of dense loops and arrays
- Used for:
 - Program analyses
 - Composition of complex loop transformations
 - Framework for code generation



Iteration Space Vectors

$\langle i, j \rangle$


$\langle i_1, j_1, i_2, j_2 \rangle$

We extend it to sparse loops


$\langle \text{SSG} [\text{tags}], \text{BSG} [\text{tags}], \text{OuterIter} [\text{dst}, \text{tags}], \text{InnerIter} [\text{src}, \text{tags}] \rangle$

Augmented with Parallelization, Partitioning, Data Layout Tags

GraphIt DSL




Algorithm
Representation
(Algorithm Language)



Optimization Representation

- Scheduling Language
- Schedule Representation
(e.g. Graph Iteration Space)



Autotuner

Schedule 4

Algorithm Specification

```
func updateEdge (src: Vertex, dst: Vertex)
  new_rank[dst] += old_rank[src] / out_degree[src]
end

func updateVertex (v: Vertex)
  new_rank[v] = base_score + 0.85*new_rank[v];
  old_rank[v] = new_rank[v];
  new_rank[v] = 0;
end

func main()
  for i in 1:11
    #s1# edges.apply(updateEdge);
    vertices.apply(updateVertex);
  end
end
```

Scheduling Functions

```
schedule:
  program->configApplyDirection("s1", "DensePull");
  program->configApplyParallelization("s1", "dynamic-vertex-parallel");
  program->configApplyNumSSG("s1", "fixed-vertex-count", 10);
```

Schedule 4

Algorithm Specification

```
func updateEdge (src: Vertex, dst: Vertex)
  new_rank[dst] += old_rank[src] / out_degree[src]
end

func updateVertex (v: Vertex)
  new_rank[v] = base_score + 0.85*new_rank[v];
  old_rank[v] = new_rank[v];
  new_rank[v] = 0;
end

func main()
  for i in 1:11
    #s1# edges.apply(updateEdge);
    vertices.apply(updateVertex);
  end
end
```

Finding the best schedule can be hard for non-experts.

Scheduling Functions

```
schedule:
  program->configApplyDirection("s1", "DensePull");
  program->configApplyParallelization("s1", "dynamic-vertex-parallel");
  program->configApplyNumSSG("s1", "fixed-vertex-count", 10);
```

Goal

Algorithm Specification

```
func updateEdge (src: Vertex, dst: Vertex)
  new_rank[dst] += old_rank[src] / out_degree[src]
end

func updateVertex (v: Vertex)
  new_rank[v] = base_score + 0.85*new_rank[v];
  old_rank[v] = new_rank[v];
  new_rank[v] = 0;
end

func main()
  for i in 1:11
    #s1# edges.apply(updateEdge);
    vertices.apply(updateVertex);
  end
end
```

**Ideally, the user only need
to write the algorithm**


Autotuner

Algorithm Specification


```
func updateEdge (src: Vertex, dst: Vertex)
  new_rank[dst] += old_rank[src] / out_degree[src]
end

func updateVertex (v: Vertex)
  new_rank[v] = base_score + 0.85*new_rank[v];
  old_rank[v] = new_rank[v];
  new_rank[v] = 0;
end

func main()
  for i in 1:11
    #s1# edges.apply(updateEdge);
    vertices.apply(updateVertex);
  end
end
```



Input Graphs



Autotuner

Autotuner

Algorithm Specification


```
func updateEdge (src: Vertex, dst: Vertex)
  new_rank[dst] += old_rank[src] / out_degree[src]
end
```

```
func updateVertex (v: Vertex)
  new_rank[v] = base_score + 0.85*new_rank[v];
  old_rank[v] = new_rank[v];
  new_rank[v] = 0;
end
```


```
func main()
  for i in 1:11
    #s1# edges.apply(updateEdge);
    vertices.apply(updateVertex);
  end
end
```

Scheduling Functions

```
schedule:
  program->configApplyDirection("s1", "DensePull");
  program->configApplyParallelization("s1", "dynamic-vertex-parallel");
  program->configApplyNumSSG("s1", "fixed-vertex-count", 10);
```




Input Graphs




Autotuner

Autotuner




Autotuner


Uses an ensemble
of search methods.
Build on top of
OpenTuner
[PACT14]



Autotuner




State of the Art and GraphIt



Intel Xeon E5-2695 v3 CPUs with 12 cores each for a total of 24 cores and 48 hyper-threads.

State of the Art and GraphIt



Intel Xeon E5-2695 v3 CPUs with 12 cores each for a total of 24 cores and 48 hyper-threads.

State of the Art and GraphIt

LJ	3.48	1	1	1
TW	5.63	1.13	3.12	1.14
WB	4.15	1.42	2.96	1.13
RD	2.69	4.81	2.16	4.57
FT	6.17	1.38	4.94	2.77
	PR	BFS	CC	SSSP

Ligra
(PPoPP13)

LJ	1.64	3.7	5.98	1.86
TW	2.34	9.4	11	1.62
WB	2.14	7.44	9.13	2.98
RD	1.61	9.06	7.04	151
FT				
	PR	BFS	CC	SSSP

GraphMat
(VLDB15)

LJ	1.51	1.83	3.06	1.82
TW	2.42	6.03	5.78	1.41
WB	2.59	2.84	5.96	2.54
RD	1.26	2.45	8.99	328
FT				
	PR	BFS	CC	SSSP

GreenMarl
(ASPLOS12)

LJ	1	1.3	1.11	1.07
TW	1	1	1	1
WB	1	1	1	1
RD	1.23	1	1.43	1
FT	1	1	1	1
	PR	BFS	CC	SSSP

GraphIt
(OOPSLA18)

LJ	8.15	1.41	2.05	1.78
TW	3.53	4.49	5.68	1.43
WB	2.82	1.83	8.07	1.36
RD	13	1.02	1.05	3.25
FT	3.61	7.02	7.05	1.08
	PR	BFS	CC	SSSP

Galois
(SOSP13)

LJ	1.26	2.22	2.46	1.57
TW	1.26	1.64	4.33	1
WB	1	1.52	4.93	1.67
RD	1.49	48.8	7.08	26.1
FT	1.37	1.49	5.24	1.43
	PR	BFS	CC	SSSP

Gemini
(OSDI16)

LJ	1.08	1.93	1.38	
TW	1.8	1.17	1.94	
WB	1.26	1.28	1.64	
RD	1	8.26	1	
FT	1.67	1.04	2.24	
	PR	BFS	CC	SSSP

Grazelle
(PPoPP18)

Intel Xeon E5-2695 v3 CPUs with 12 cores each for a total of 24 cores and 48 hyper-threads.

State of the Art and GraphIt

Most frameworks are good at certain applications and graphs

LJ	3.48	1	1	1
FT	5.65	1.13	3.12	1.14
RD	4.15	1.42	2.96	1.13
WB	2.69	4.81	2.16	4.57
TW	6.17	1.38	4.94	2.77
	PR	BFS	CC	SSSP

Ligra
(PPoPP13)

LJ	1.64	3.7	5.98	1.86
FT	2.34	9.4	11	1.62
RD	2.14	7.44	9.13	2.98
WB	1.61	9.06	7.04	151
TW				
	PR	BFS	CC	SSSP

GraphMat
(VLDB15)

LJ	1.51	1.83	3.06	1.82
FT	2.42	6.03	5.78	1.41
RD	2.59	2.84	5.96	2.54
WB	1.26	2.45	8.99	328
TW				
	PR	BFS	CC	SSSP

GreenMarl
(ASPLOS12)

LJ	8.15	1.41	2.05	1.78
FT	3.53	4.49	5.68	1.43
RD	2.82	1.83	8.07	1.36
WB	13	1.02	1.05	3.25
TW	3.61	7.02	7.03	1.08
	PR	BFS	CC	SSSP

Galois
(SOSP13)

LJ	1.26	2.22	2.46	1.57
FT	1.26	1.64	4.33	1
RD	1	1.52	4.93	1.67
WB	1.49	48.8	7.08	26.1
TW	1.37	1.49	5.24	1.43
	PR	BFS	CC	SSSP

Gemini
(OSDI16)

LJ	1.08	1.93	1.38	
FT	1.8	1.17	1.94	
RD	1.26	1.28	1.64	
WB	1	8.26	1	
TW	1.67	1.04	2.24	
	PR	BFS	CC	SSSP

Grazelle
(PPoPP18)

LJ	1	1.3	1.11	1.07
FT	1	1	1	1
RD	1	1	1	1
WB	1.23	1	1.43	1
TW	1	1	1	1
	PR	BFS	CC	SSSP

GraphIt
(OOPSLA18)

Intel Xeon E5-2695 v3 CPUs with 12 cores each for a total of 24 cores and 48 hyper-threads.

State of the Art and GraphIt

Most frameworks are bad at certain applications and graphs

LJ	3.48	1	1	1
TW	5.63	1.13	3.12	1.14
WB	4.15	1.42	2.96	1.13
RD	2.69	4.81	2.16	4.57
FT	6.17	1.38	4.94	2.77
	PR	BFS	CC	SSSP

Ligra
(PPoPP13)

LJ	1.64	3.7	5.98	1.86
TW	2.34	9.4	11	1.62
WB	2.14	7.44	9.13	2.98
RD	1.61	9.06	7.04	151
FT				
	PR	BFS	CC	SSSP

GraphMat
(VLDB15)

LJ	1.51	1.83	3.06	1.82
TW	2.42	6.03	5.78	1.41
WB	2.59	2.84	5.96	2.54
RD	1.26	2.45	8.99	328
FT				
	PR	BFS	CC	SSSP

GreenMarl
(ASPLOS12)

LJ	1	1.3	1.11	1.07
TW	1	1	1	1
WB	1	1	1	1
RD	1.23	1	1.43	1
FT	1	1	1	1
	PR	BFS	CC	SSSP

GraphIt
(OOPSLA18)

LJ	8.15	1.41	2.05	1.78
TW	3.53	4.49	5.68	1.43
WB	2.82	1.83	8.07	1.36
RD	13	1.02	1.05	3.25
FT	3.61	7.02	7.05	1.08
	PR	BFS	CC	SSSP

Galois
(SOSP13)

LJ	1.26	2.22	2.46	1.57
TW	1.26	1.64	4.33	1
WB	1	1.52	4.93	1.67
RD	1.49	48.8	7.08	26.1
FT	1.37	1.49	5.24	1.43
	PR	BFS	CC	SSSP

Gemini
(OSDI16)

LJ	1.08	1.93	1.38	
TW	1.8	1.17	1.94	
WB	1.26	1.28	1.64	
RD	1	8.26	1	
FT	1.67	1.04	2.24	
	PR	BFS	CC	SSSP

Grazelle
(PPoPP18)

Intel Xeon E5-2695 v3 CPUs with 12 cores each for a total of 24 cores and 48 hyper-threads.

State of the Art and GraphIt

Previous work support a subset of optimizations

LJ	3.48	1	1	1
TW	5.63	1.13	3.12	1.14
WB	4.15	1.42	2.96	1.13
D	2.69	4.81	2.16	4.57
FT	6.17	1.38	4.94	2.77
	PR	BFS	CC	SSSP

Ligra
(PPoPP13)

LJ	1.64	3.7	5.98	1.86
TW	2.34	9.4	11	1.62
WB	2.14	7.44	9.13	2.98
D	1.61	9.06	7.04	151
FT				
	PR	BFS	CC	SSSP

GraphMat
(VLDB15)

LJ	1.51	1.83	3.06	1.82
TW	2.42	6.03	5.78	1.41
WB	2.59	2.84	5.96	2.54
D	1.26	2.45	8.99	328
FT				
	PR	BFS	CC	SSSP

GreenMarl
(ASPLOS12)

LJ	1	1.3	1.11	1.07
TW	1	1	1	1
WB	1	1	1	1
D	1.23	1	1.43	1
FT	1	1	1	1
	PR	BFS	CC	SSSP

GraphIt
(OOPSLA18)

LJ	8.15	1.41	2.05	1.78
TW	3.53	4.49	5.68	1.43
WB	2.82	1.83	8.07	1.36
D	13	1.02	1.05	3.25
FT	3.61	7.02	7.05	1.08
	PR	BFS	CC	SSSP

Galois
(SOSP13)

LJ	1.26	2.22	2.46	1.57
TW	1.26	1.64	4.33	1
WB	1	1.52	4.93	1.67
D	1.49	48.8	7.08	26.1
FT	1.37	1.49	5.24	1.43
	PR	BFS	CC	SSSP

Gemini
(OSDI16)

LJ	1.08	1.93	1.38	
TW	1.8	1.17	1.94	
WB	1.26	1.28	1.64	
D	1	8.26	1	
FT	1.67	1.04	2.24	
	PR	BFS	CC	SSSP

Grazelle
(PPoPP18)

Intel Xeon E5-2695 v3 CPUs with 12 cores each for a total of 24 cores and 48 hyper-threads.

Consistent High-Performance

LJ	3.48	1	1	1
TW	5.63	1.13	3.12	1.14
WB	4.15	1.42	2.96	1.13
RD	2.69	4.81	2.16	4.57
FT	6.17	1.38	4.94	2.77
	PR	BFS	CC	SSSP

Ligra
(PPoPP13)

LJ	1.64	3.7	5.98	1.86
TW	2.34	9.4	11	1.62
WB	2.14	7.44	9.13	2.98
RD	1.61	9.06	7.04	151
FT				
	PR	BFS	CC	SSSP

GraphMat
(VLDB15)

LJ	1.51	1.83	3.06	1.82
TW	2.42	6.03	5.78	1.41
WB	2.59	2.84	5.96	2.54
RD	1.26	2.45	8.99	328
FT				
	PR	BFS	CC	SSSP

GreenMarl
(ASPLOS12)

Good across different applications and graphs

LJ	1	1.3	1.11	1.07
TW	1	1	1	1
WB	1	1	1	1
RD	1.23	1	1.43	1
FT	1	1	1	1
	PR	BFS	CC	SSSP

GraphIt
(OOPSLA18)

LJ	8.15	1.41	2.05	1.78
TW	3.53	4.49	5.68	1.43
WB	2.82	1.83	8.07	1.36
RD	13	1.02	1.05	3.25
FT	3.61	7.02	7.05	1.08
	PR	BFS	CC	SSSP

Galois
(SOSP13)

LJ	1.26	2.22	2.46	1.57
TW	1.26	1.64	4.33	1
WB	1	1.52	4.93	1.67
RD	1.49	48.8	7.08	26.1
FT	1.37	1.49	5.24	1.43
	PR	BFS	CC	SSSP

Gemini
(OSDI16)

LJ	1.08	1.93	1.38	
TW	1.8	1.17	1.94	
WB	1.26	1.28	1.64	
RD	1	8.26	1	
FT	1.67	1.04	2.24	
	PR	BFS	CC	SSSP

Grazelle
(PPoPP18)

Intel Xeon E5-2695 v3 CPUs with 12 cores each for a total of 24 cores and 48 hyper-threads.

Speedup over State of the Art

LJ	3.48	1	1	1
TW	5.63	1.13	3.12	1.14
WB	4.15	1.42	2.96	1.13
RD	2.69	4.81	2.16	4.57
FT	6.17	1.38	4.94	2.77
	PR	BFS	CC	SSSP

Ligra
(PPoPP13)

LJ	1.64	3.7	5.98	1.86
TW	2.34	9.4	11	1.62
WB	2.14	7.44	9.13	2.98
RD	1.61	9.06	7.04	151
FT				
	PR	BFS	CC	SSSP

GraphMat
(VLDB15)

LJ	1.51	1.83	3.06	1.82
TW	2.42	6.03	5.78	1.41
WB	2.59	2.84	5.96	2.54
RD	1.26	2.45	8.99	328
FT				
	PR	BFS	CC	SSSP

GreenMarl
(ASPLOS12)

Finds previously unexplored combinations of optimizations

LJ	1	1.3	1.11	1.07
TW	1	1	1	1
WB	1	1	1	1
RD	1.23	1	1.43	1
FT	1	1	1	1
	PR	BFS	CC	SSSP

GraphIt
(OOPSLA18)

LJ	8.15	1.41	2.05	1.78
TW	3.53	4.49	5.68	1.43
WB	2.82	1.83	8.07	1.36
RD	13	1.02	1.05	3.25
FT	3.61	7.02	7.05	1.08
	PR	BFS	CC	SSSP

Galois
(SOSP13)

LJ	1.26	2.22	2.46	1.57
TW	1.26	1.64	4.33	1
WB	1	1.52	4.93	1.67
RD	1.49	48.8	7.08	26.1
FT	1.37	1.49	5.24	1.43
	PR	BFS	CC	SSSP

Gemini
(OSDI16)

LJ	1.08	1.93	1.38	
TW	1.8	1.17	1.94	
WB	1.26	1.28	1.64	
RD	1	8.26	1	
FT	1.67	1.04	2.24	
	PR	BFS	CC	SSSP

Grazelle
(PPoPP18)

Intel Xeon E5-2695 v3 CPUs with 12 cores each for a total of 24 cores and 48 hyper-threads.

Ease-of-Use

Reduces the lines of code by an order of magnitude compare to the next fastest framework

FT	3.48	1	1	1
RD	5.63	1.13	3.12	1.14
WB	4.15	1.42	2.96	1.13
TW	2.69	4.81	2.16	4.57
LJ	6.17	1.38	4.94	2.77
	PR	BFS	CC	SSSP

Ligra
(PPoPP13)

FT	1.64	3.7	5.98	1.86
RD	2.34	9.4	11	1.62
WB	2.14	7.44	9.13	2.98
TW	1.61	9.06	7.04	151
LJ				
	PR	BFS	CC	SSSP

GraphMat
(VLDB15)

FT	1.51	1.83	3.06	1.82
RD	2.42	6.03	5.78	1.41
WB	2.59	2.84	5.96	2.54
TW	1.26	2.45	8.99	328
LJ				
	PR	BFS	CC	SSSP

GreenMarl
(ASPLOS12)

FT	8.15	1.41	2.05	1.78
RD	3.53	4.49	5.68	1.43
WB	2.82	1.83	8.07	1.36
TW	13	1.02	1.05	3.25
LJ	3.61	7.02	7.05	1.08
	PR	BFS	CC	SSSP

Galois
(SOSP13)

FT	1.26	2.22	2.46	1.57
RD	1.26	1.64	4.33	1
WB	1	1.52	4.93	1.67
TW	1.49	48.8	7.08	26.1
LJ	1.37	1.49	5.24	1.43
	PR	BFS	CC	SSSP

Gemini
(OSDI16)

FT	1.08	1.93	1.38	
RD	1.8	1.17	1.94	
WB	1.26	1.28	1.64	
TW	1	8.26	1	
LJ	1.67	1.04	2.24	
	PR	BFS	CC	SSSP


Grazelle
(PPoPP18)

FT	1	1.3	1.11	1.07
RD	1	1	1	1
WB	1	1	1	1
TW	1.23	1	1.43	1
LJ	1	1	1	1
	PR	BFS	CC	SSSP

GraphIt
(OOPSLA18)


Intel Xeon E5-2695 v3 CPUs with 12 cores each for a total of 24 cores and 48 hyper-threads.

Topology-Driven




Algorithms


Data-Driven




Topology-Driven



Algorithms




Unordered




Active Vertices can be processed in parallel in arbitrary order


Data-Driven




Topology-Driven



Algorithms




Unordered

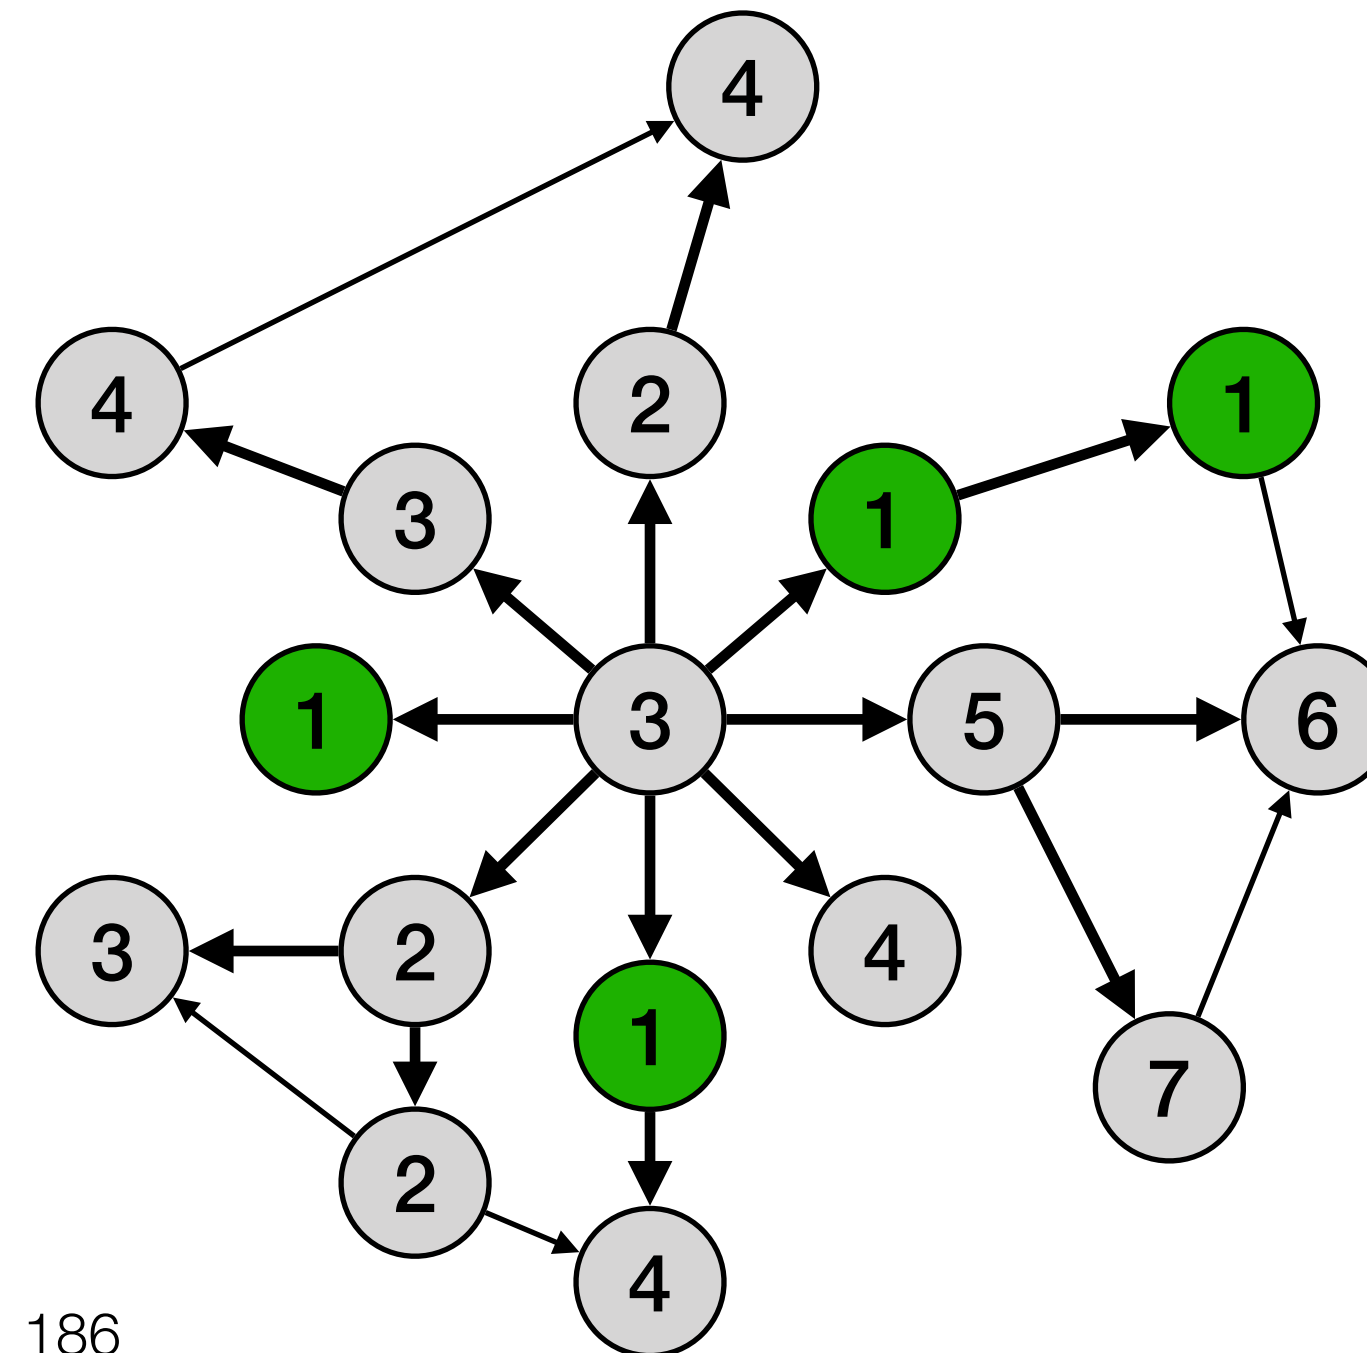


Active Vertices can be processed in parallel in arbitrary order

Data-Driven



Ordered




Vertices are processed according to priorities


Priorities can change dynamically

Vertices of the same priority are processed in parallel


Topology-Driven



Algorithms




Unordered




Active Vertices can be processed in parallel in arbitrary order

Data-Driven



Ordered




Vertices are processed according to priorities

Priorities can change dynamically

Vertices of the same priority are processed in parallel


Ordered vs Unordered

Ordered algorithms can often achieve 2x to 640x speedup over unordered counterparts



Speedup of Ordered vs Unordered on a 24-core CPU

Bellman-Ford (Unordered SSSP)




-  Start Vertex
-  Unreached Vertex
-  Vertex with Shortest Distance
-  Vertex with Suboptimal Distance
-  Active Vertex

Rounds: 0

Updates: 0

A	B	C	D	E	F	G
0	∞	∞	∞	∞	∞	∞

Bellman-Ford (Unordered SSSP)




-  Start Vertex
-  Unreached Vertex
-  Vertex with Shortest Distance
-  Vertex with Suboptimal Distance
-  Active Vertex

Rounds: 1

Updates: 2

A	B	C	D	E	F	G
0	1	∞	50	∞	∞	∞

Bellman-Ford (Unordered SSSP)




-  Start Vertex
-  Unreached Vertex
-  Vertex with Shortest Distance
-  Vertex with Suboptimal Distance
-  Active Vertex

Rounds: 2

Updates: 5

A	B	C	D	E	F	G
0	1	3	50	52	54	∞

Bellman-Ford (Unordered SSSP)




-  Start Vertex
-  Unreached Vertex
-  Vertex with Shortest Distance
-  Vertex with Suboptimal Distance
-  Active Vertex

Rounds: 3

Updates: 7

	A	B	C	D	E	F	G
0		1	3	4	52	54	53

Bellman-Ford (Unordered SSSP)




-  Start Vertex
-  Unreached Vertex
-  Vertex with Shortest Distance
-  Vertex with Suboptimal Distance
-  Active Vertex

Rounds: 4

Updates: 9

A	B	C	D	E	F	G
0	1	3	4	6	8	53

Bellman-Ford (Unordered SSSP)




-  Start Vertex
-  Unreached Vertex
-  Vertex with Shortest Distance
-  Vertex with Suboptimal Distance
-  Active Vertex

Rounds: 5
Updates: 10

	A	B	C	D	E	F	G
0		1	3	4	6	8	7

Delta-Stepping (Ordered SSSP)




-  Start Vertex
-  Unreached Vertex
-  Vertex with Shortest Distance
-  Vertex with Suboptimal Distance
-  Active Vertex

Rounds: 0
 # Updates: 0
 Delta: 10

A	B	C	D	E	F	G
0	∞	∞	∞	∞	∞	∞

Delta-Stepping (Ordered SSSP)




-  Start Vertex
-  Unreached Vertex
-  Vertex with Shortest Distance
-  Vertex with Suboptimal Distance
-  Active Vertex

Rounds: 1
 # Updates: 2
 Delta: 10

A	B	C	D	E	F	G
0	1	∞	50	∞	∞	∞

Delta-Stepping (Ordered SSSP)




-  Start Vertex
-  Unreached Vertex
-  Vertex with Shortest Distance
-  Vertex with Suboptimal Distance
-  Active Vertex

Rounds: 2
 # Updates: 3
 Delta: 10

A	B	C	D	E	F	G
0	1	3	50	∞	∞	∞

Delta-Stepping (Ordered SSSP)




-  Start Vertex
-  Unreached Vertex
-  Vertex with Shortest Distance
-  Vertex with Suboptimal Distance
-  Active Vertex

Rounds: 3
 # Updates: 4
 Delta: 10

A	B	C	D	E	F	G
0	1	3	4	∞	∞	∞

Delta-Stepping (Ordered SSSP)




-  Start Vertex
-  Unreached Vertex
-  Vertex with Shortest Distance
-  Vertex with Suboptimal Distance
-  Active Vertex

Rounds: 4
 # Updates: 6
 Delta: 10

A	B	C	D	E	F	G
0	1	3	4	6	8	∞

Delta-Stepping (Ordered SSSP)




-  Start Vertex
-  Unreached Vertex
-  Vertex with Shortest Distance
-  Vertex with Suboptimal Distance
-  Active Vertex

Rounds: 5
 # Updates: 7
 Delta: 10

A	B	C	D	E	F	G
0	1	3	4	6	8	7

Delta-Stepping (Ordered SSSP)

No Redundant Updates for E, F, and G



-  Start Vertex
-  Unreached Vertex
-  Vertex with Shortest Distance
-  Vertex with Suboptimal Distance
-  Active Vertex


Rounds: 5
Updates: 10

Rounds: 5
Updates: 7
Delta: 10

A	B	C	D	E	F	G
0	1	3	4	6	8	7

Delta-Stepping (Ordered SSSP)

Algorithmic Tradeoff Between Parallelism and Work-Efficiency



-  Start Vertex
-  Unreached Vertex
-  Vertex with Shortest Distance
-  Vertex with Suboptimal Distance
-  Active Vertex


Rounds: 5
Updates: 7
Delta: 10

A	B	C	D	E	F	G
0	1	3	4	6	8	7

Priority-based Extensions to GraphIt

- **Decouple Algorithm from Optimization for Ordered Graph Algorithms**
 - Priority-based Algorithm Language Operators
 - Optimizations for Ordered Parallelism
- **Language and Compiler Extensions Achieve**
 - Ease-of-Use
 - Consistent High-Performance

Priority-Based Extensions



- PriorityQueue
 - dequeueReadySet()
 - getCurrentPriority()
 - finished(), finishedNode()
 - updatePriorityMin, updatePrioritySum, ..

Delta-Stepping

```
const pq: priority_queue{Vertex}(int);
```

```
func updateEdge(src : Vertex, dst : Vertex, weight : int)  
    pq.updatePriorityMin(dst, SP[dst], SP[src] + weight);  
end
```

```
func main ()  
    var start_vertex : int = 0;  
    SP[start_vertex] = 0;  
    pq = new priority_queue{Vertex}(int)(true, "lower_first", SP, start_vertex);  
    while (!pq.finished())  
        var frontier: vertexset{Vertex} = pq.dequeueReadySet();  
        #s1# edges.from(frontier).applyUpdatePriority(updateEdge);  
        delete frontier;  
    end  
end
```

Delta-Stepping

```
const pq: priority_queue{Vertex}(int);
```

```
func updateEdge(src : Vertex, dst : Vertex, weight : int)  
    pq.updatePriorityMin(dst, SP[dst], SP[src] + weight);  
end
```

```
func main ()  
    var start_vertex : int = 0;  
    SP[start_vertex] = 0;  
    pq = new priority_queue{Vertex}(int)(true, "lower_first", SP, start_vertex);  
    while (!pq.finished())  
        var frontier: vertexset{Vertex} = pq.dequeueReadySet();  
        #s1# edges.from(frontier).applyUpdatePriority(updateEdge);  
        delete frontier;  
    end  
end
```

Delta-Stepping

```
const pq: priority_queue{Vertex}(int);
```

```
func updateEdge(src : Vertex, dst : Vertex, weight : int)  
    pq.updatePriorityMin(dst, SP[dst], SP[src] + weight);  
end
```

```
func main ()  
    var start_vertex : int = 0;  
    SP[start_vertex] = 0;  
    pq = new priority_queue{Vertex}(int)(true, "lower_first", SP, start_vertex);  
    while (!pq.finished())  
        var frontier: vertexset{Vertex} = pq.dequeueReadySet();  
        #s1# edges.from(frontier).applyUpdatePriority(updateEdge);  
        delete frontier;  
    end  
end
```

Delta-Stepping

```
const pq: priority_queue{Vertex}(int);
```

```
func updateEdge(src : Vertex, dst : Vertex, weight : int)  
    pq.updatePriorityMin(dst, SP[dst], SP[src] + weight);  
end
```

```
func main ()  
    var start_vertex : int = 0;  
    SP[start_vertex] = 0;  
    pq = new priority_queue{Vertex}(int)(true, "lower_first", SP, start_vertex);  
    while (!pq.finished())  
        var frontier: vertexset{Vertex} = pq.dequeueReadySet();  
        #s1# edges.from(frontier).applyUpdatePriority(updateEdge);  
        delete frontier;  
    end  
end
```

Delta-Stepping

Hides Physical
Implementation for
PriorityQueue

```
const pq: priority_queue{Vertex}(int);
```

```
func updateEdge(src : Vertex, dst : Vertex, weight : int)  
    pq.updatePriorityMin(dst, SP[dst], SP[src] + weight);  
end
```

```
func main ()  
    var start_vertex : int = 0;  
    SP[start_vertex] = 0;  
    pq = new priority_queue{Vertex}(int)(true, "lower_first", SP, start_vertex);  
    while (!pq.finished())  
        var frontier: vertexset{Vertex} = pq.dequeueReadySet();  
        #s1# edges.from(frontier).applyUpdatePriority(updateEdge);  
        delete frontier;  
    end  
end
```

Delta-Stepping

```
const pq: priority_queue{Vertex}(int);
```

```
func updateEdge(src : Vertex, dst : Vertex, weight : int)  
    pq.updatePriorityMin(dst, SP[dst], SP[src] + weight);  
end
```

```
func main ()  
    var start_vertex : int = 0;  
    SP[start_vertex] = 0;  
    pq = new priority_queue{Vertex}(int)(true, "lower_first", SP, start_vertex);  
    while (!pq.finished())  
        var frontier: vertexset{Vertex} = pq.dequeueReadySet();  
        #s1# edges.from(frontier).applyUpdatePriority(updateEdge);  
        delete frontier;  
    end  
end
```

Delta-Stepping

```
const pq: priority_queue{Vertex}(int);

func updateEdge(src : Vertex, dst : Vertex, weight : int)
  pq.updatePriorityMin(dst, SP[dst], SP[src] + weight);
end

func main ( )
  var start_vertex : int = 0;
  SP[start_vertex] = 0;
  pq = new priority_queue{Vertex}(int)(true, "lower_first", SP, start_vertex);
  while (!pq.finished())
    var frontier: vertexset{Vertex} = pq.dequeueReadySet();
    #s1# edges.from(frontier).applyUpdatePriority(updateEdge);
    delete frontier;
  end
end
```

Delta-Stepping

```
const pq: priority_queue{Vertex}(int);

func updateEdge(src : Vertex, dst : Vertex, weight : int)
  pq.updatePriorityMin(dst, SP[dst], SP[src] + weight);
end

func main ( )
  var start_vertex : int = 0;
  SP[start_vertex] = 0;
  pq = new priority_queue{Vertex}(int)(true, "lower_first", SP, start_vertex);
  while (!pq.finished())
    var frontier: vertexset{Vertex} = pq.dequeueReadySet();
    #s1# edges.from(frontier).applyUpdatePriority(updateEdge);
    delete frontier;
  end
end
```


Delta-Stepping

```
const pq: priority_queue{Vertex}(int);

func updateEdge(src : Vertex, dst : Vertex, weight : int)
  pq.updatePriorityMin(dst, SP[dst], SP[src] + weight);
end

func main ( )
  var start_vertex : int = 0;
  SP[start_vertex] = 0;
  pq = new priority_queue{Vertex}(int)(true, "lower_first", SP, start_vertex);
  while (!pq.finished())
    var frontier: vertexset{Vertex} = pq.dequeueReadySet();
    #s1# edges.from(frontier).applyUpdatePriority(updateEdge);
    delete frontier;
  end
end
```

Schedule:

program

- >configApplyPriorityUpdate("s1", "lazy")
- >configApplyPriorityUpdateDelta("s1", 4)
- >configApplyDirection("s1", "SparsePush")
- >configApplyParallelization("s1", "dynamic-vertex-parallel")

Delta-Stepping

```
const pq: priority_queue{Vertex}(int);

func updateEdge(src : Vertex, dst : Vertex, weight : int)
  pq.updatePriorityMin(dst, SP[dst], SP[src] + weight);
end

func main ( )
  var start_vertex : int = 0;
  SP[start_vertex] = 0;
  pq = new priority_queue{Vertex}(int)(true, "lower_first", SP, start_vertex);
  while (!pq.finished())
    var frontier: vertexset{Vertex} = pq.dequeueReadySet();
    #s1# edges.from(frontier).applyUpdatePriority(updateEdge);
    delete frontier;
  end
end
```

Schedule:

program

- >configApplyPriorityUpdate("s1", "lazy")
- >configApplyPriorityUpdateDelta("s1", 4)
- >configApplyDirection("s1", "SparsePush")
- >configApplyParallelization("s1", "dynamic-vertex-parallel")

Delta-Stepping

```
const pq: priority_queue{Vertex}(int);

func updateEdge(src : Vertex, dst : Vertex, weight : int)
  pq.updatePriorityMin(dst, SP[dst], SP[src] + weight);
end

func main ( )
  var start_vertex : int = 0;
  SP[start_vertex] = 0;
  pq = new priority_queue{Vertex}(int)(true, "lower_first", SP, start_vertex);
  while (!pq.finished())
    var frontier: vertexset{Vertex} = pq.dequeueReadySet();
    #s1# edges.from(frontier).applyUpdatePriority(updateEdge);
    delete frontier;
  end
end
```

Schedule:

program

- >configApplyPriorityUpdate("s1", "lazy")
- >configApplyPriorityUpdateDelta("s1", 4)
- >configApplyDirection("s1", "SparsePush")
- >configApplyParallelization("s1", "dynamic-vertex-parallel")

```
1 int * dist = new int[num_verts];
2 LazyPriorityQueue* pq;
3 int delta = 4;
4 WGraph* G = loadGraph(argv[1]);
5
6 //simplified snippets of the generated main function
7 ...
8 dist[start_vertex] = 0;
9 pq = new LazyPriorityQueue(true, "lower", dist, delta);
10 While (pq.finished()){
11   VertexSubset * frontier = getNextBucket(pq);
12   uint* outEdges = setupOutputBuffer(g, frontier);
13   uint* offsets = setupOutputBufferOffsets(g, frontier);
14   parallel_for (uint s : frontier.vert_array) {
15     int j = 0;
16     uint offset = offsets[j];
17     for(WNode d : G.getOutNgh(s)){
18       bool tracking_var = false;
19       int new_dist = dist[s.v] + d.weight;
20       tracking_var = atomicWriteMin(&dist[d.v], new_dist);
21       If (tracking_var && CAS(dedup_flags[d.v],0,1)){
22         outEdges[offset + j] = d.v;
23       } else { outEdges[offset + j] = UINT_MAX; }
24       j++;
25     }
26     VertexSubset* nextFrontier = setupFrontier(outEdges);
27     updateBuckets(nextFrontier, pq, delta);
28     ...
29 }
30 ...
```

Delta-Stepping

```
const pq: priority_queue{Vertex}(int);

func updateEdge(src : Vertex, dst : Vertex, weight : int)
  pq.updatePriorityMin(dst, SP[dst], SP[src] + weight);
end

func main ( )
  var start_vertex : int = 0;
  SP[start_vertex] = 0;
  pq = new priority_queue{Vertex}(int)(true, "lower_first", SP, start_vertex);
  while (!pq.finished())
    var frontier: vertexset{Vertex} = pq.dequeueReadySet();
    #s1# edges.from(frontier).applyUpdatePriority(updateEdge);
    delete frontier;
  end
end
```

Schedule:


program

```
->configApplyPriorityUpdate("s1", "eager")
->configApplyPriorityUpdateDelta("s1", 4)
->configApplyDirection("s1", "SparsePush")
->configApplyParallelization("s1", "dynamic-vertex-parallel")
```

```
1 int * dist = new int[num_verts];
2 EagerPriorityQueue* pq;
3 int delta = 4;
4 WGraph* G = loadGraph(argv[1]);
5
6 //simplified snippets of the generated main function
7 ...
8 dist[start_vertex] = 0;
9 frontier[0] = start_vertex;
10 pq = new EagerPriorityQueue(true, "lower", dist, delta);
11 uint* frontier = new uint[G.num_edges()];
12 #pragma omp parallel
13 { vector<vector<uint>> local_bins(0);
14   while (pq.finished()) {
15     #pragma omp for nowait schedule(dynamic, 64)
16     for (size_t i = 0; i < frontier.size(); i++) {
17       uint s = frontier[i];
18       for (WNode d : G.getOutNgh(s)) {
19         int new_dist = dist[s] + d.weight;
20         bool changed = atomicWriteMin(&dist[d.v], new_dist);
21         if (changed == false) {break;}}
22         if (changed) {
23           size_t dest_bin = new_dist/delta;
24           if (dest_bin >= local_bins.size()) {
25             local_bins.resize(dest_bin+1);}
26           local_bins[dest_bin].push_back(d.v);
27         }}// end of for frontier for loop
28     ... //omitted:find next bucket
29   #pragma omp barrier
30   ... //omitted:copy local buckets to global bucket
31   #pragma omp barrier } // end of parallel region
32 ...
```

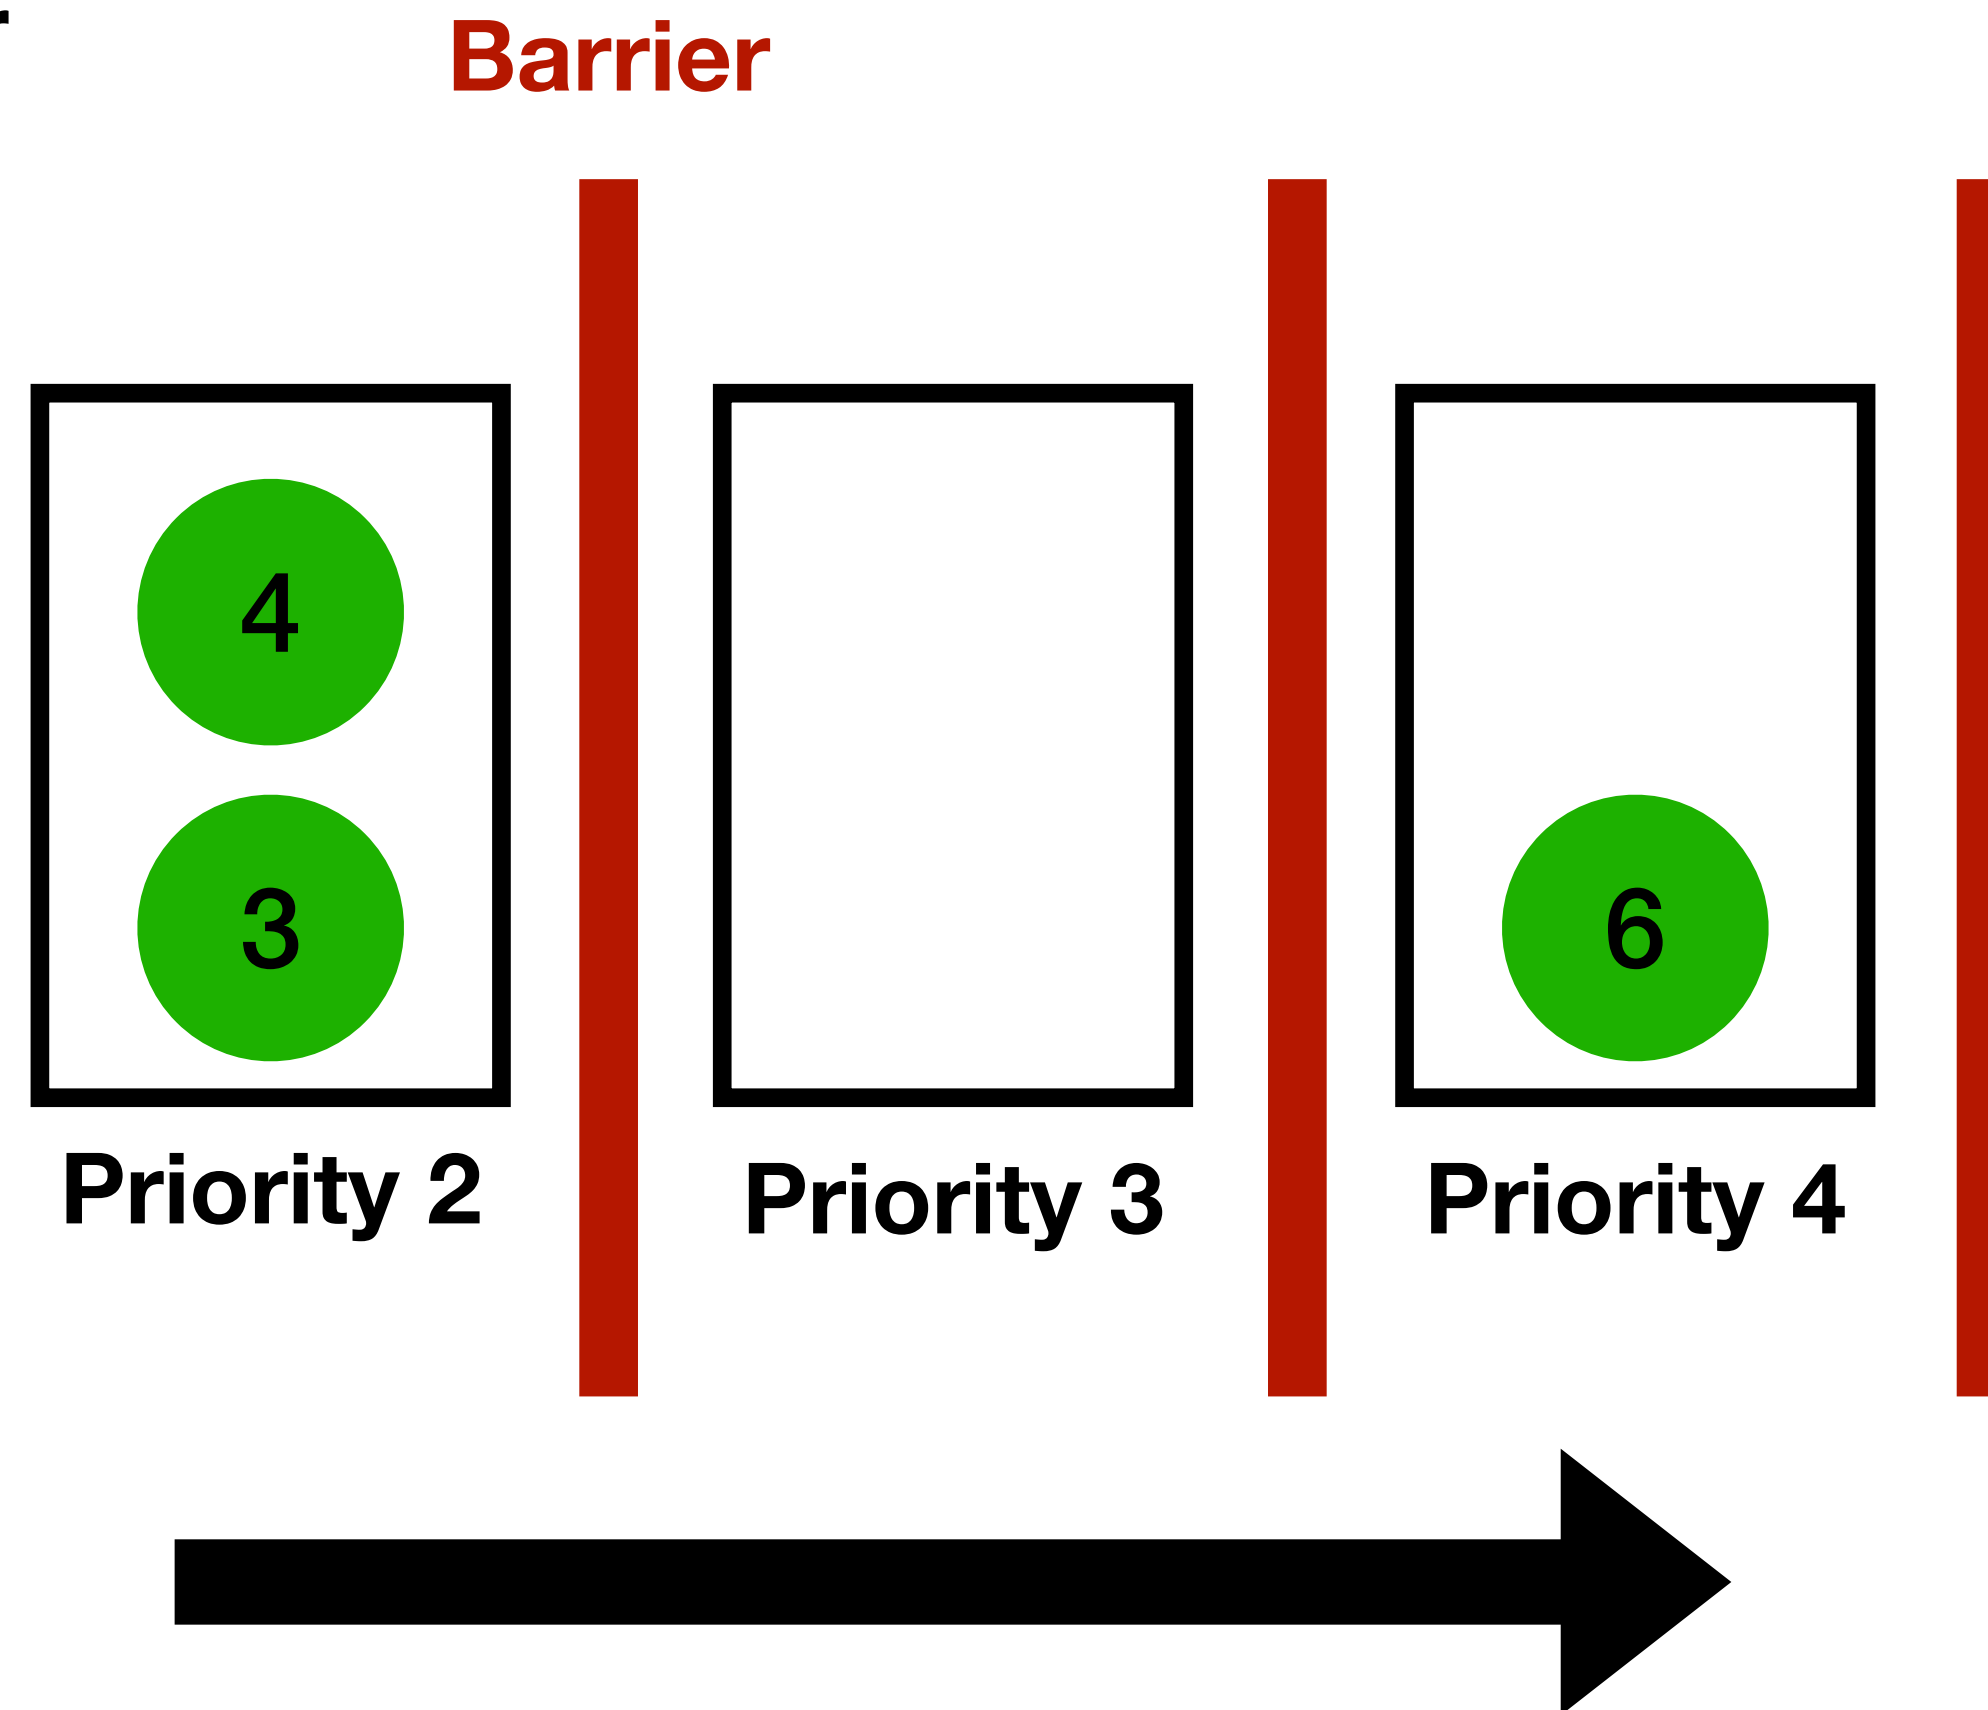
PriorityQueue with Bucketing

Vertices are stored in buckets according to their priority, and are processed in order



PriorityQueue with Bucketing

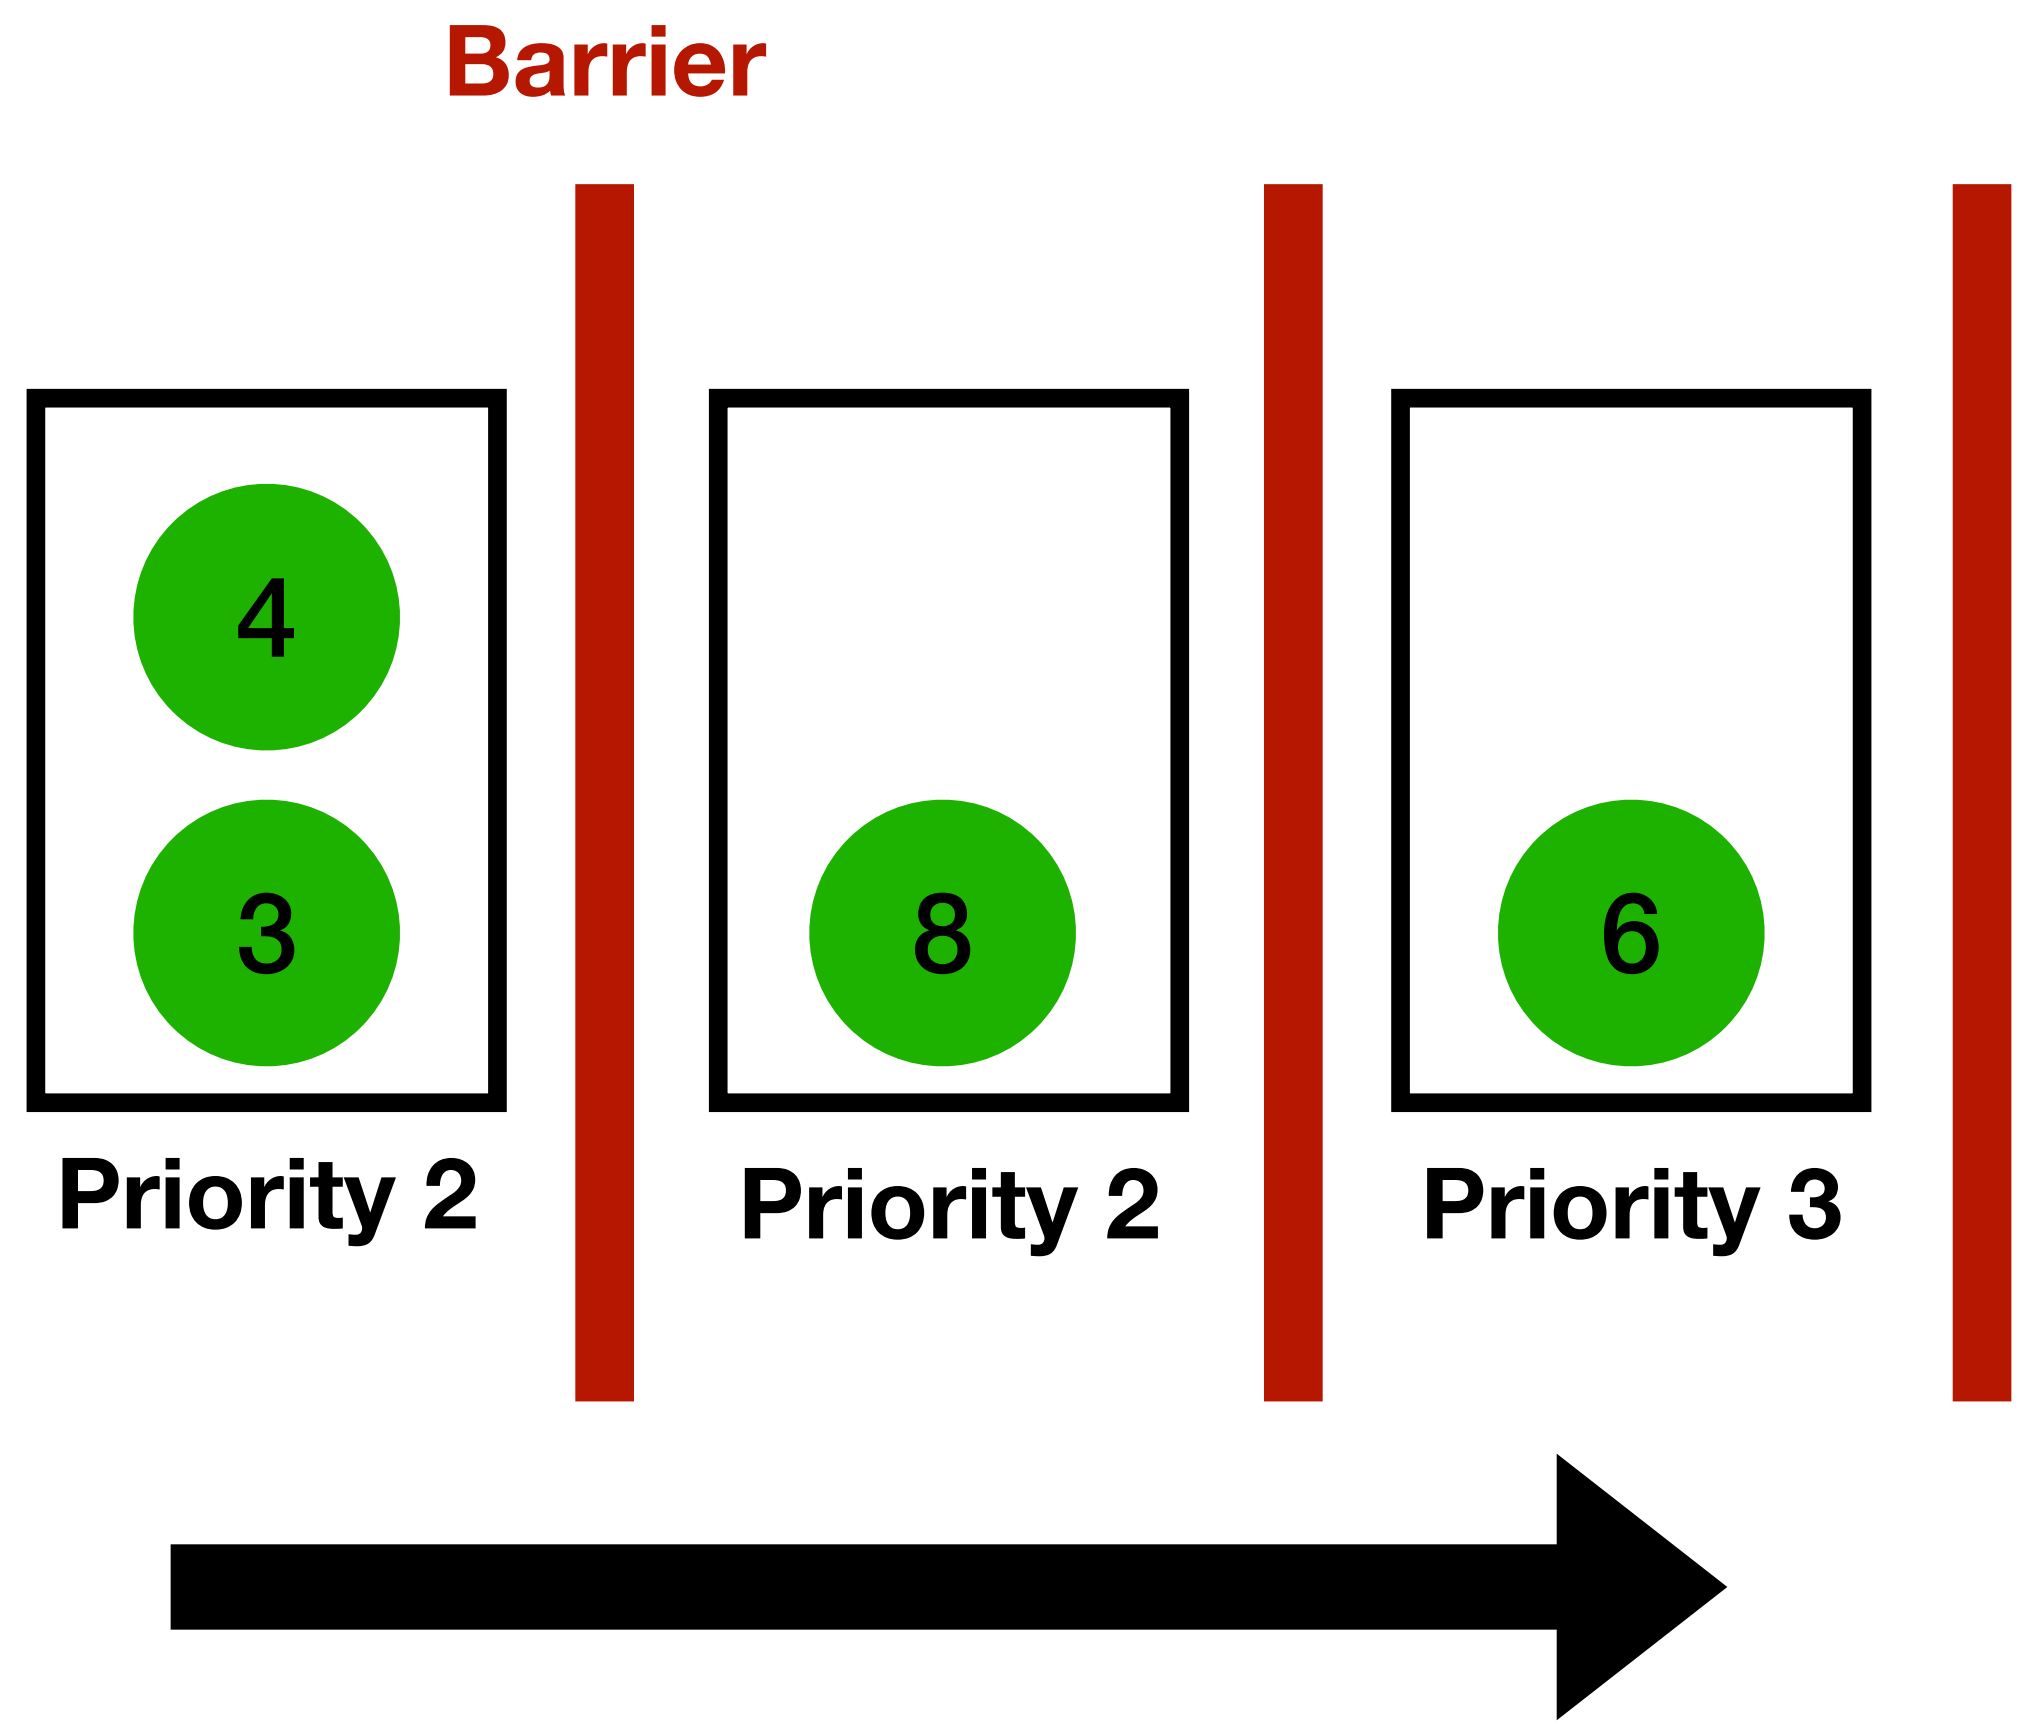
Global synchronization after each bucket (barrier)




Eager vs Lazy

- Eager
 - Update Buckets Immediately after Priority Changes
 - Faster when Synchronization is the Bottleneck
- Lazy
 - Buffer and Reduce Priority Changes before Bucket Updates
 - Faster when Redundant Updates are the Bottleneck

Bucket Fusion



Bucket Fusion



Scheduling Space Extensions

- **configApplyPriorityUpdate**
 - lazy, lazy_const_sum, eager_with_fusion, eager_no_fusion
- **configApplyPriorityUpdateDelta**
 - delta parameter for priority coarsening
- **configBucketFusionThreshold**
- **configNumBuckets** (number of materialized buckets)

Scheduling Space Extensions

- **configApplyPriorityUpdate**
 - lazy, lazy_const_sum, eager_with_fusion, eager_no_fusion
- **configApplyPriorityUpdateDelta**
 - delta parameter for priority coarsening
- **configBucketFusionThreshold**
- **configNumBuckets** (number of materialized buckets)

**Compatible with
existing GraphIt
schedules**

Comparisons with Ordered Frameworks

RD	LJ	1	1	1	1
	TW	1.06	1	1	1
	RD	1	1	1	1
		SSSP	PPSP	k-core	SetCover
Extended GraphIt					

RD	LJ	4	2.41	1.01	1.42
	TW	1.31	1.89	1.03	1.32
	RD	16.9	15.3	1.09	1.2
		SSSP	PPSP	k-core	SetCover
Julienne					

RD	LJ	1.32	1.94		
	TW	1	1.01		
	RD	1.23	1.12		
		SSSP	PPSP	k-core	SetCover
Galois					

**Results for more graphs and algorithms are in the paper
(AStar Search, weighted BFS)**

Comparisons with Ordered Frameworks

LJ	1	1	1	1
TW	1.06	1	1	1
RD	1	1	1	1
	SSSP	PPSP	k-core	SetCover

Extended GraphIt

LJ	4	2.41	1.01	1.42
TW	1.31	1.89	1.03	1.32
RD	16.9	15.3	1.09	1.2
	SSSP	PPSP	k-core	SetCover

Julienne

**No Support for Eager
Bucket Update**

LJ	1.32	1.94		
TW	1	1.01		
RD	1.23	1.12		
	SSSP	PPSP	k-core	SetCover

Galois

Comparisons with Ordered Frameworks

LJ	1	1	1	1
TW	1.06	1	1	1
RD	1	1	1	1
	SSSP	PPSP	k-core	SetCover

Extended GraphIt

LJ	4	2.41	1.01	1.42
TW	1.31	1.89	1.03	1.32
RD	16.9	15.3	1.09	1.2
	SSSP	PPSP	k-core	SetCover

Julienne

LJ	1.32	1.94		
TW	1	1.01		
RD	1.23	1.12		
	SSSP	PPSP	k-core	SetCover

Galois

**No Support for Strict
Priority Ordering Needed
for Correctness**

Comparisons with Ordered Frameworks

LJ	1	1	1	1
TW	1.06	1	1	1
RD	1	1	1	1
	SSSP	PPSP	k-core	SetCover

Extended GraphIt

Achieves consistent high-performance across algorithms and graphs

LJ	4	2.41	1.01	1.42
TW	1.31	1.89	1.03	1.32
RD	16.9	15.3	1.09	1.2
	SSSP	PPSP	k-core	SetCover

Julienne

LJ	1.32	1.94		
TW	1	1.01		
RD	1.23	1.12		
	SSSP	PPSP	k-core	SetCover

Galois

Outline

Hardware Utilization

Making Caches Work for Graph Analytics (BigData17)
Zhang, et al.

- Frequency-based Reordering
- Cache-aware Partitioning

Programming System to Handle Variety in Data and Algorithms

GraphIt: a High-Performance Graph DSL (OOPSLA18)
Zhang, et al.

- **GraphIt Compiler and DSL that Decouples**
 - Algorithm
 - Optimization
 - **Hardware**
- **for Graph Applications**


Optimizing Ordered Graph Algorithms with GraphIt (CGO2020)
Zhang, et al.

Variety in Hardware

Universal Graph Framework (Under Submission)
Brahmakshatriya, Zhang, et al.

GPU vs CPU

GPU (Volta)



- Memory Bandwidth: ~900 GB/s
- Compute Power: ~10 TFLOPs
- Cache Size: ~3 MB Shared Cache, ~96KB L1 Cache per SM
- Memory Size: 32 GB
- Less Powerful Cores

CPU (Skylake)



- Memory Bandwidth: ~100 GB/s
- Compute Power: ~1 TFLOPs
- Cache Size: ~64MB Shared Cache, 256KB Private Cache per Core
- Memory Size: Up to 3 TB
- More Powerful Cores with Out-of-Order Execution, Branch predictor, ...

No Best Hardware

GPU (Volta)



CPU (Skylake)



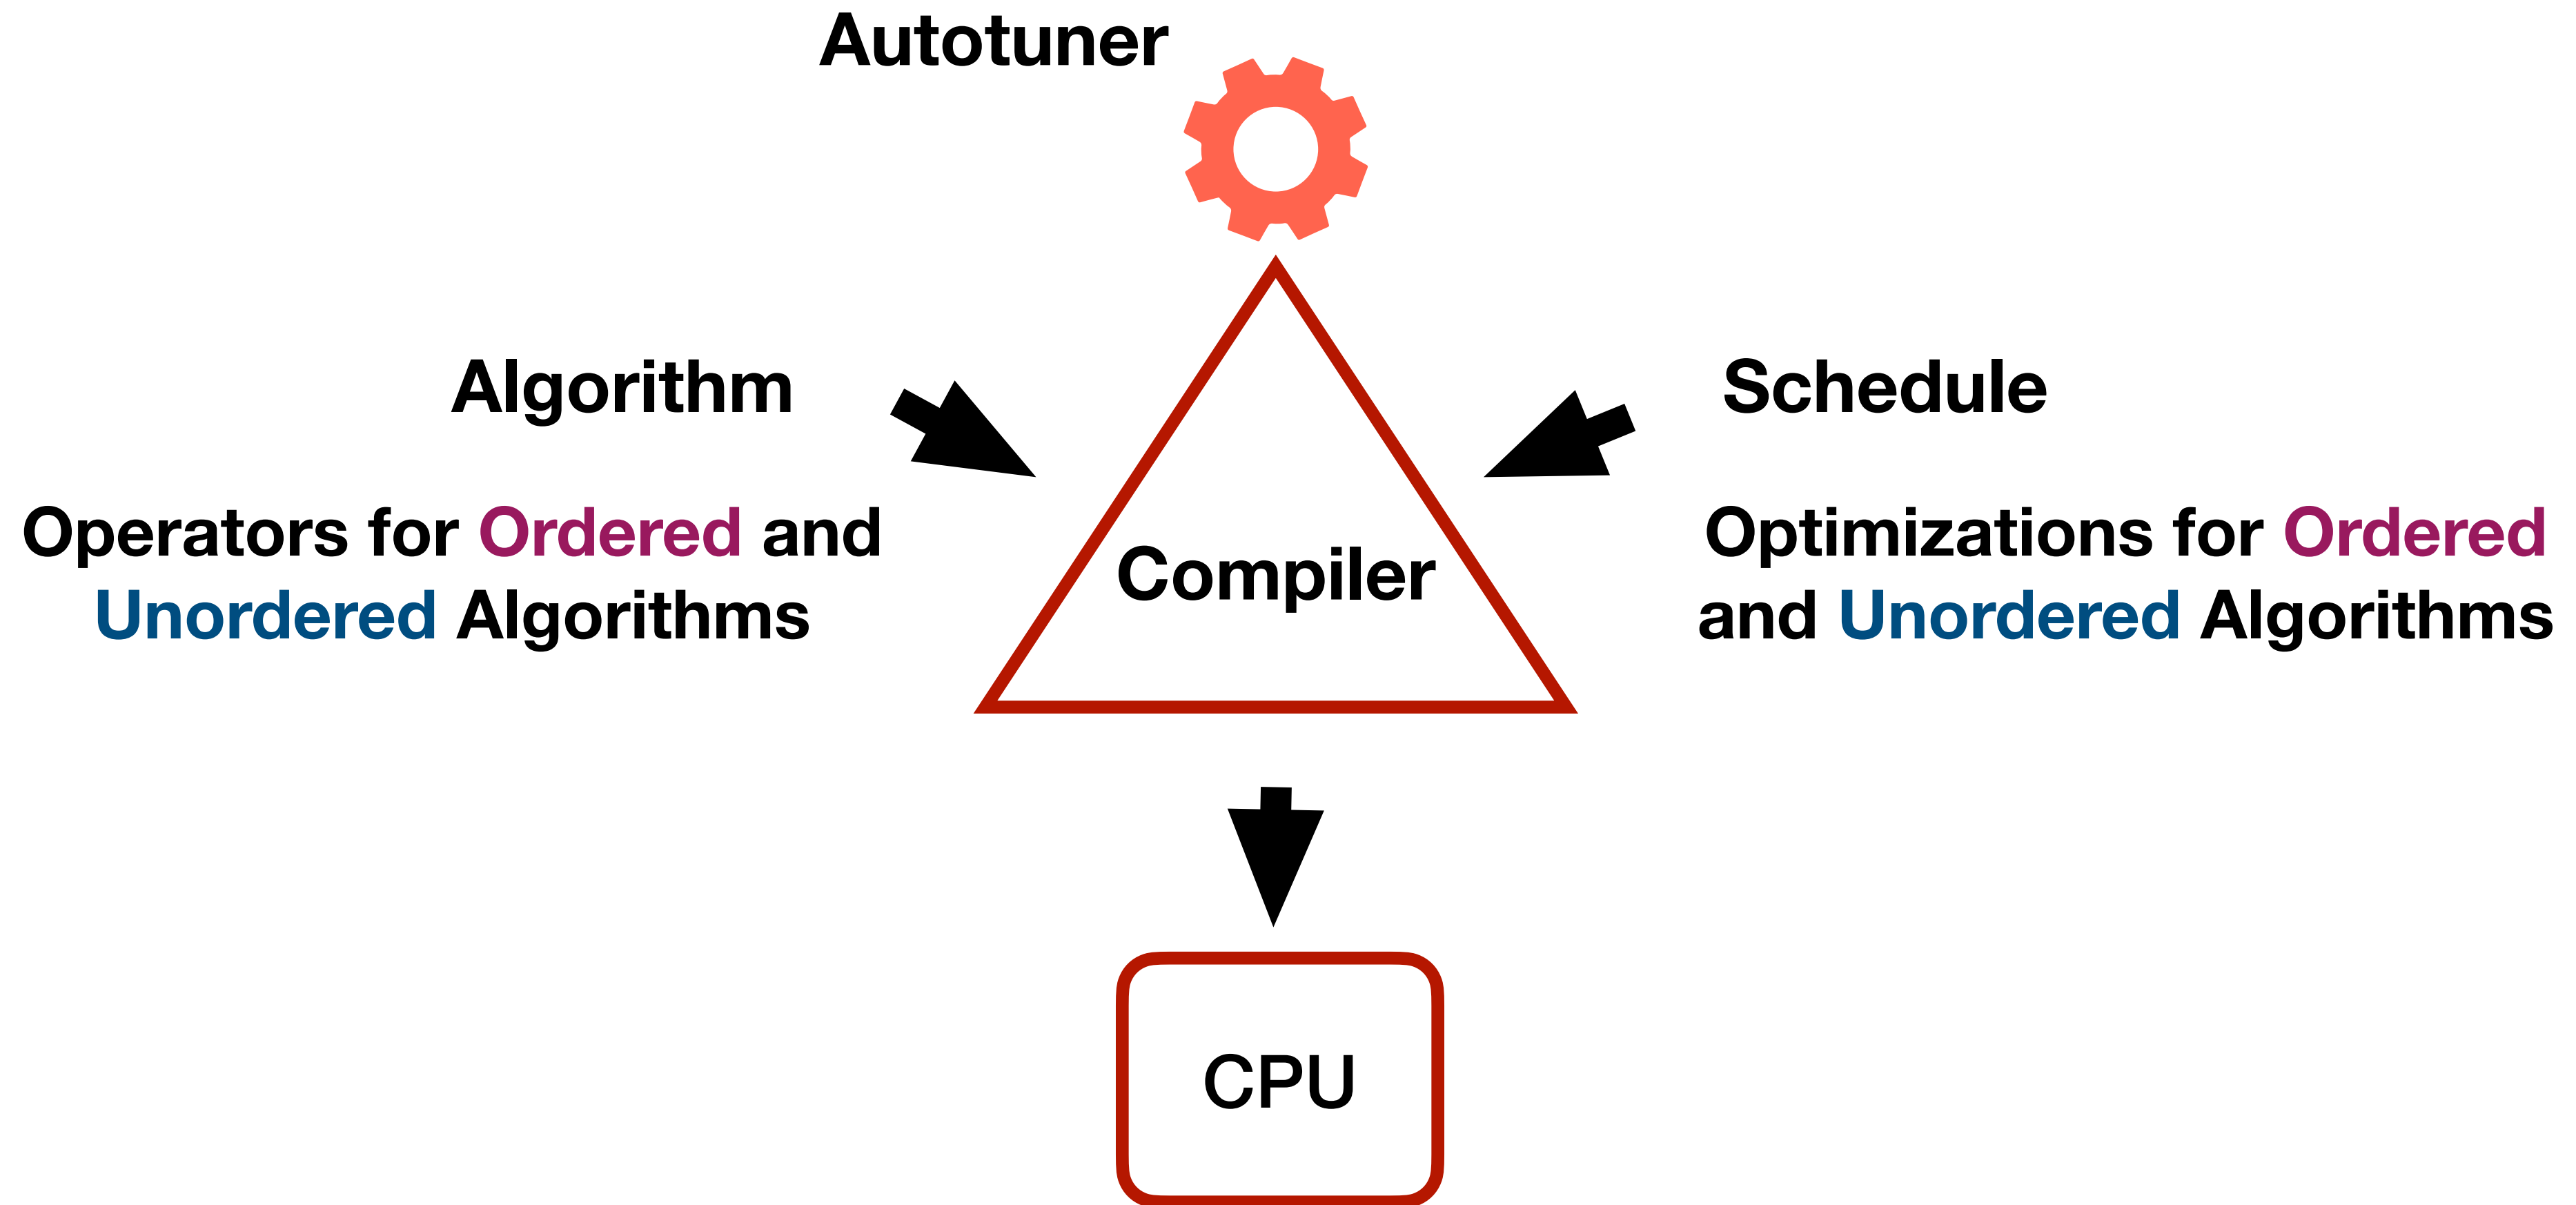
- Throughput-Oriented, Abundant Parallelism, Little Control Flow, Medium-Sized Graphs
- PageRank, Connected Components, Label Propagation on Social and Web Graphs ...

- Latency-Sensitive, Limited parallelism, or Large Graphs
- SSSP, PPSP, AStar on Road Networks and Smaller social networks, ...

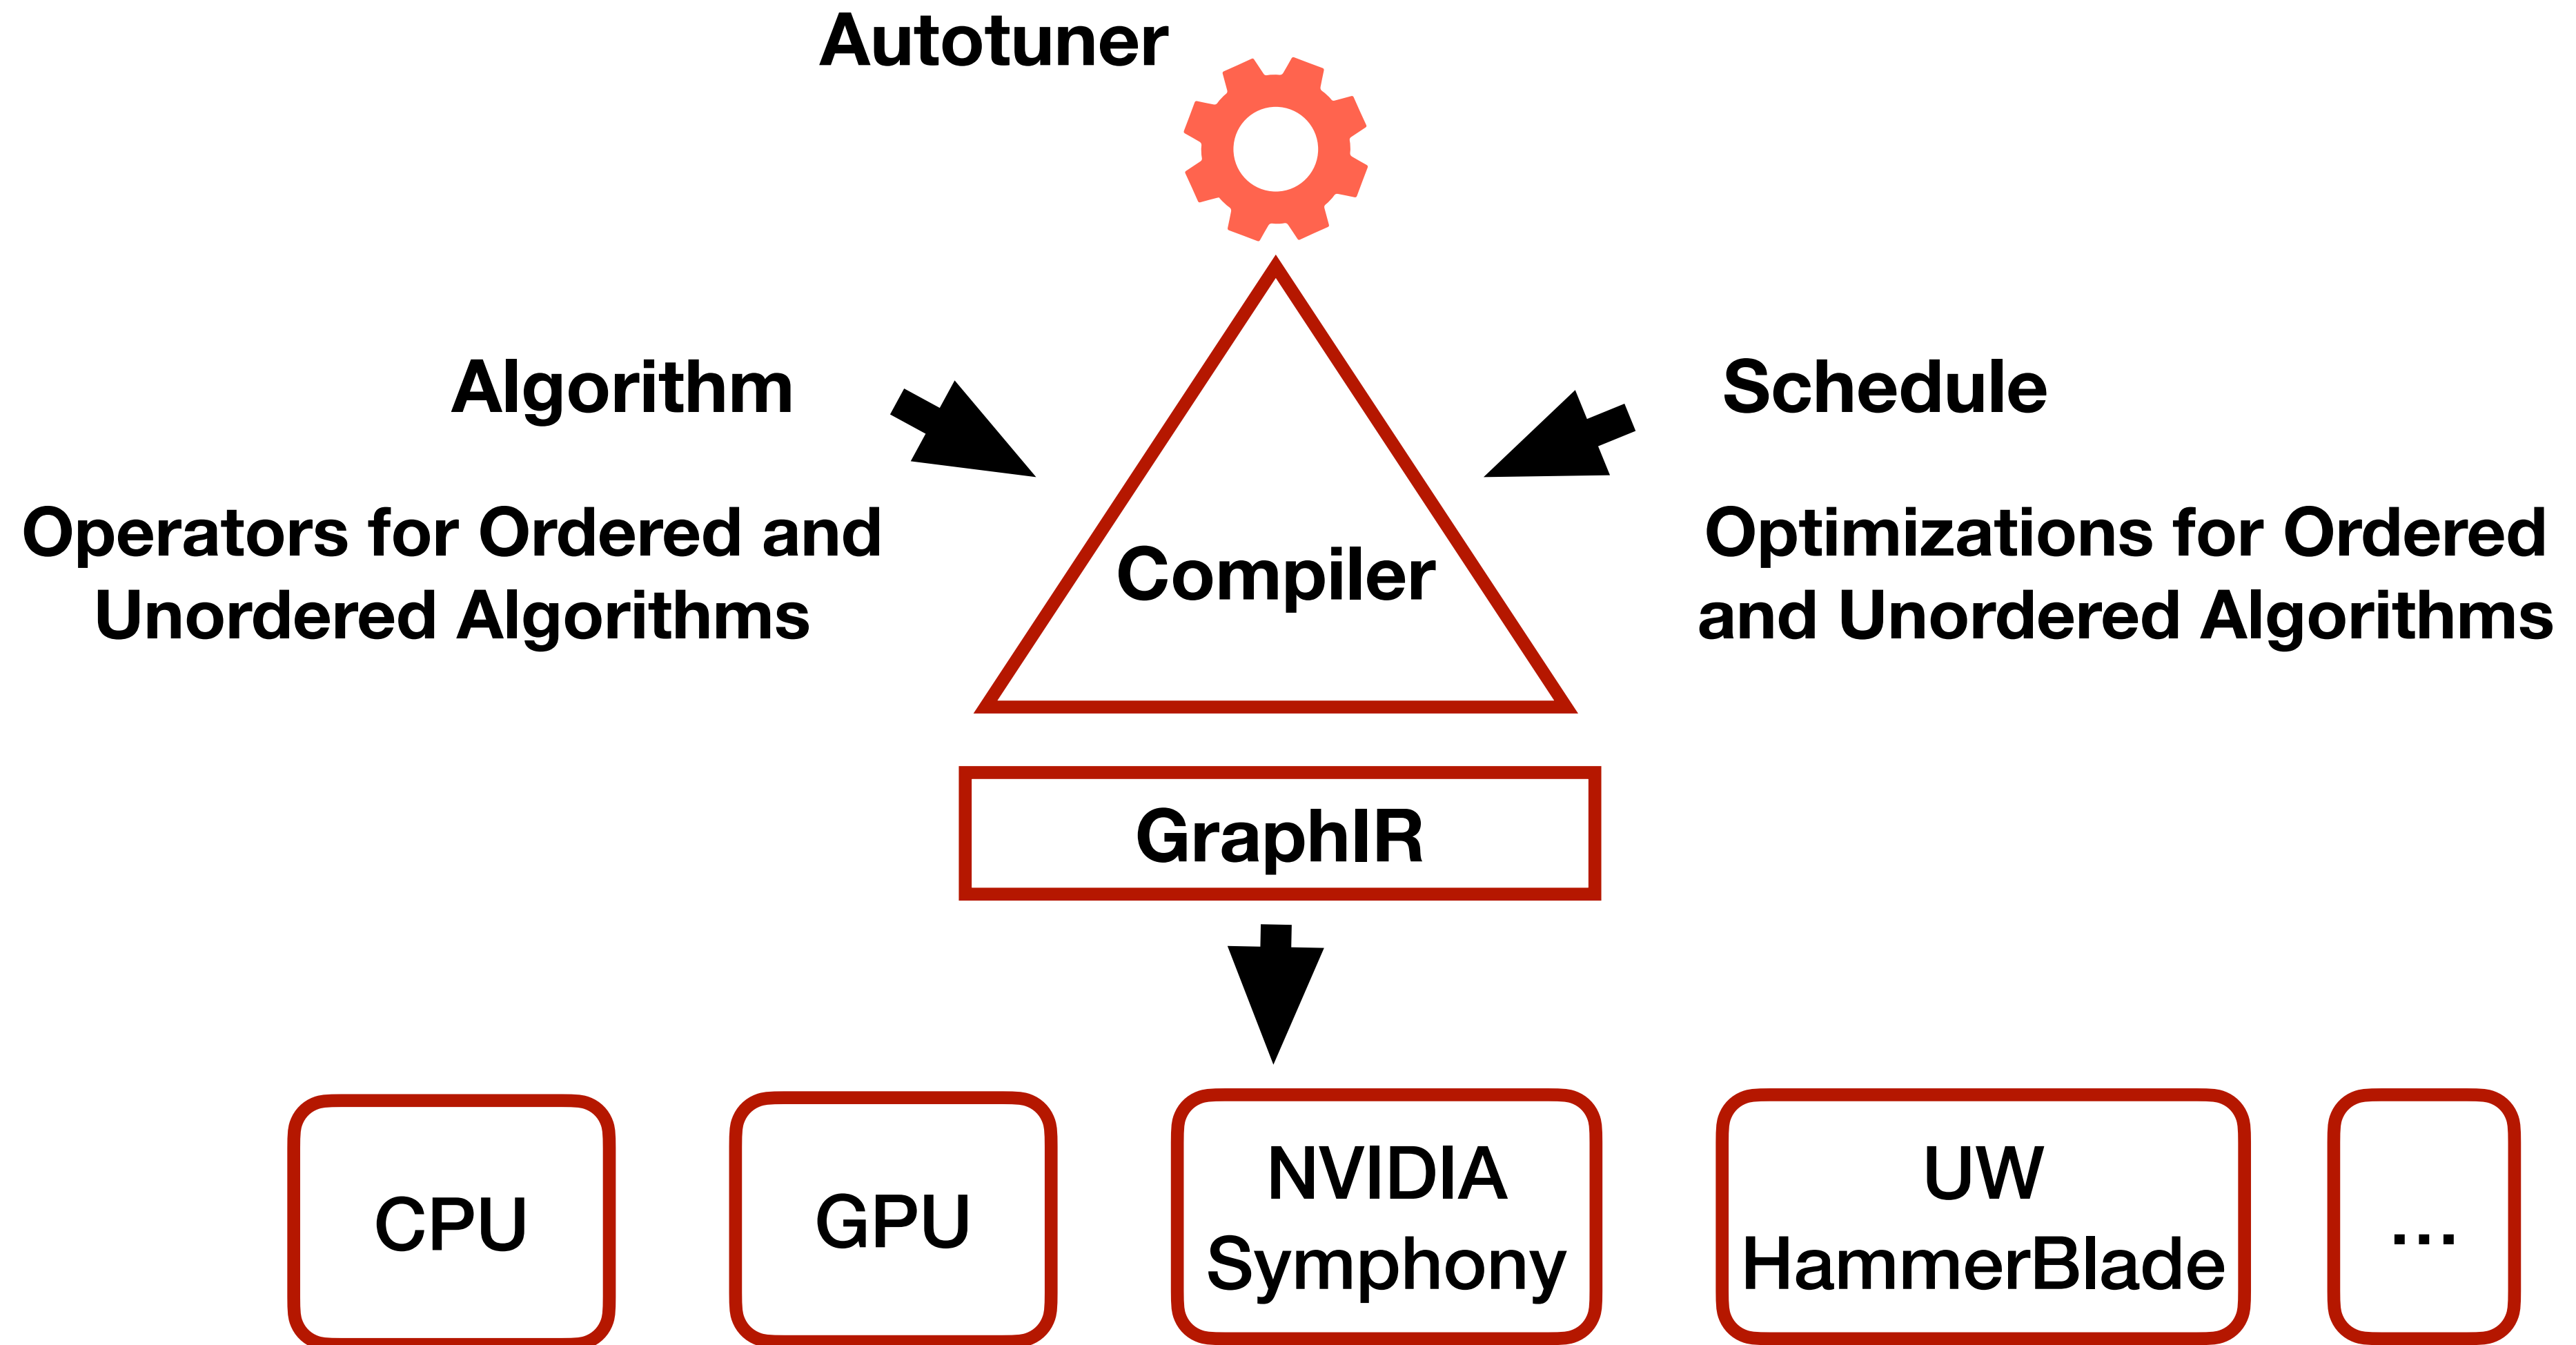
Key Optimizations for GPU

- Direction Optimization
- Cache Optimization
- Load Balance
- Kernel Fusion across Iterations
- Active Vertexset Creation
- Active Vertexset Deduplication
- Active Vertexset Processing Ordering (Ordered Algorithms)

GraphIt

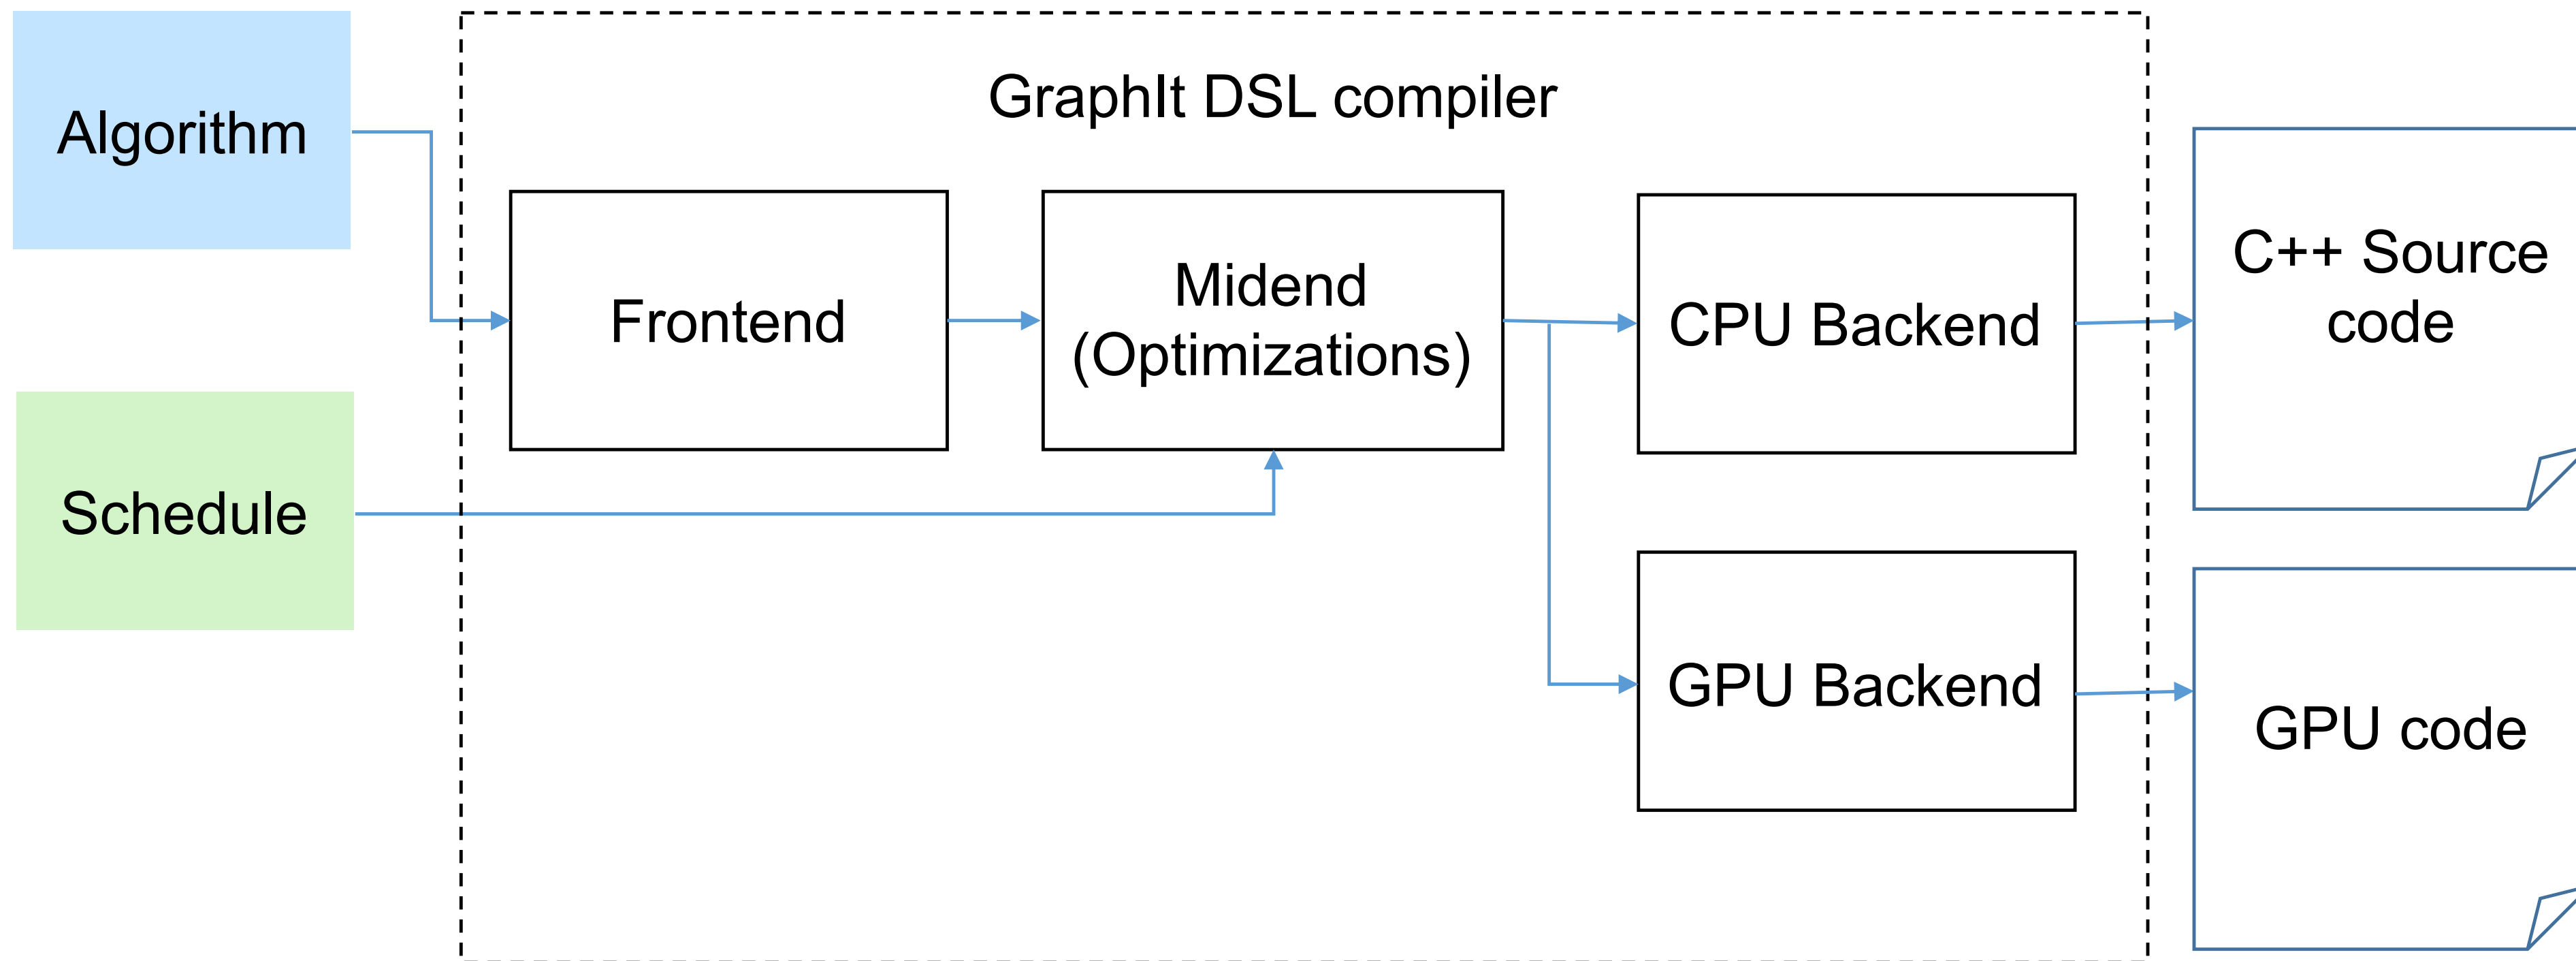


Portability

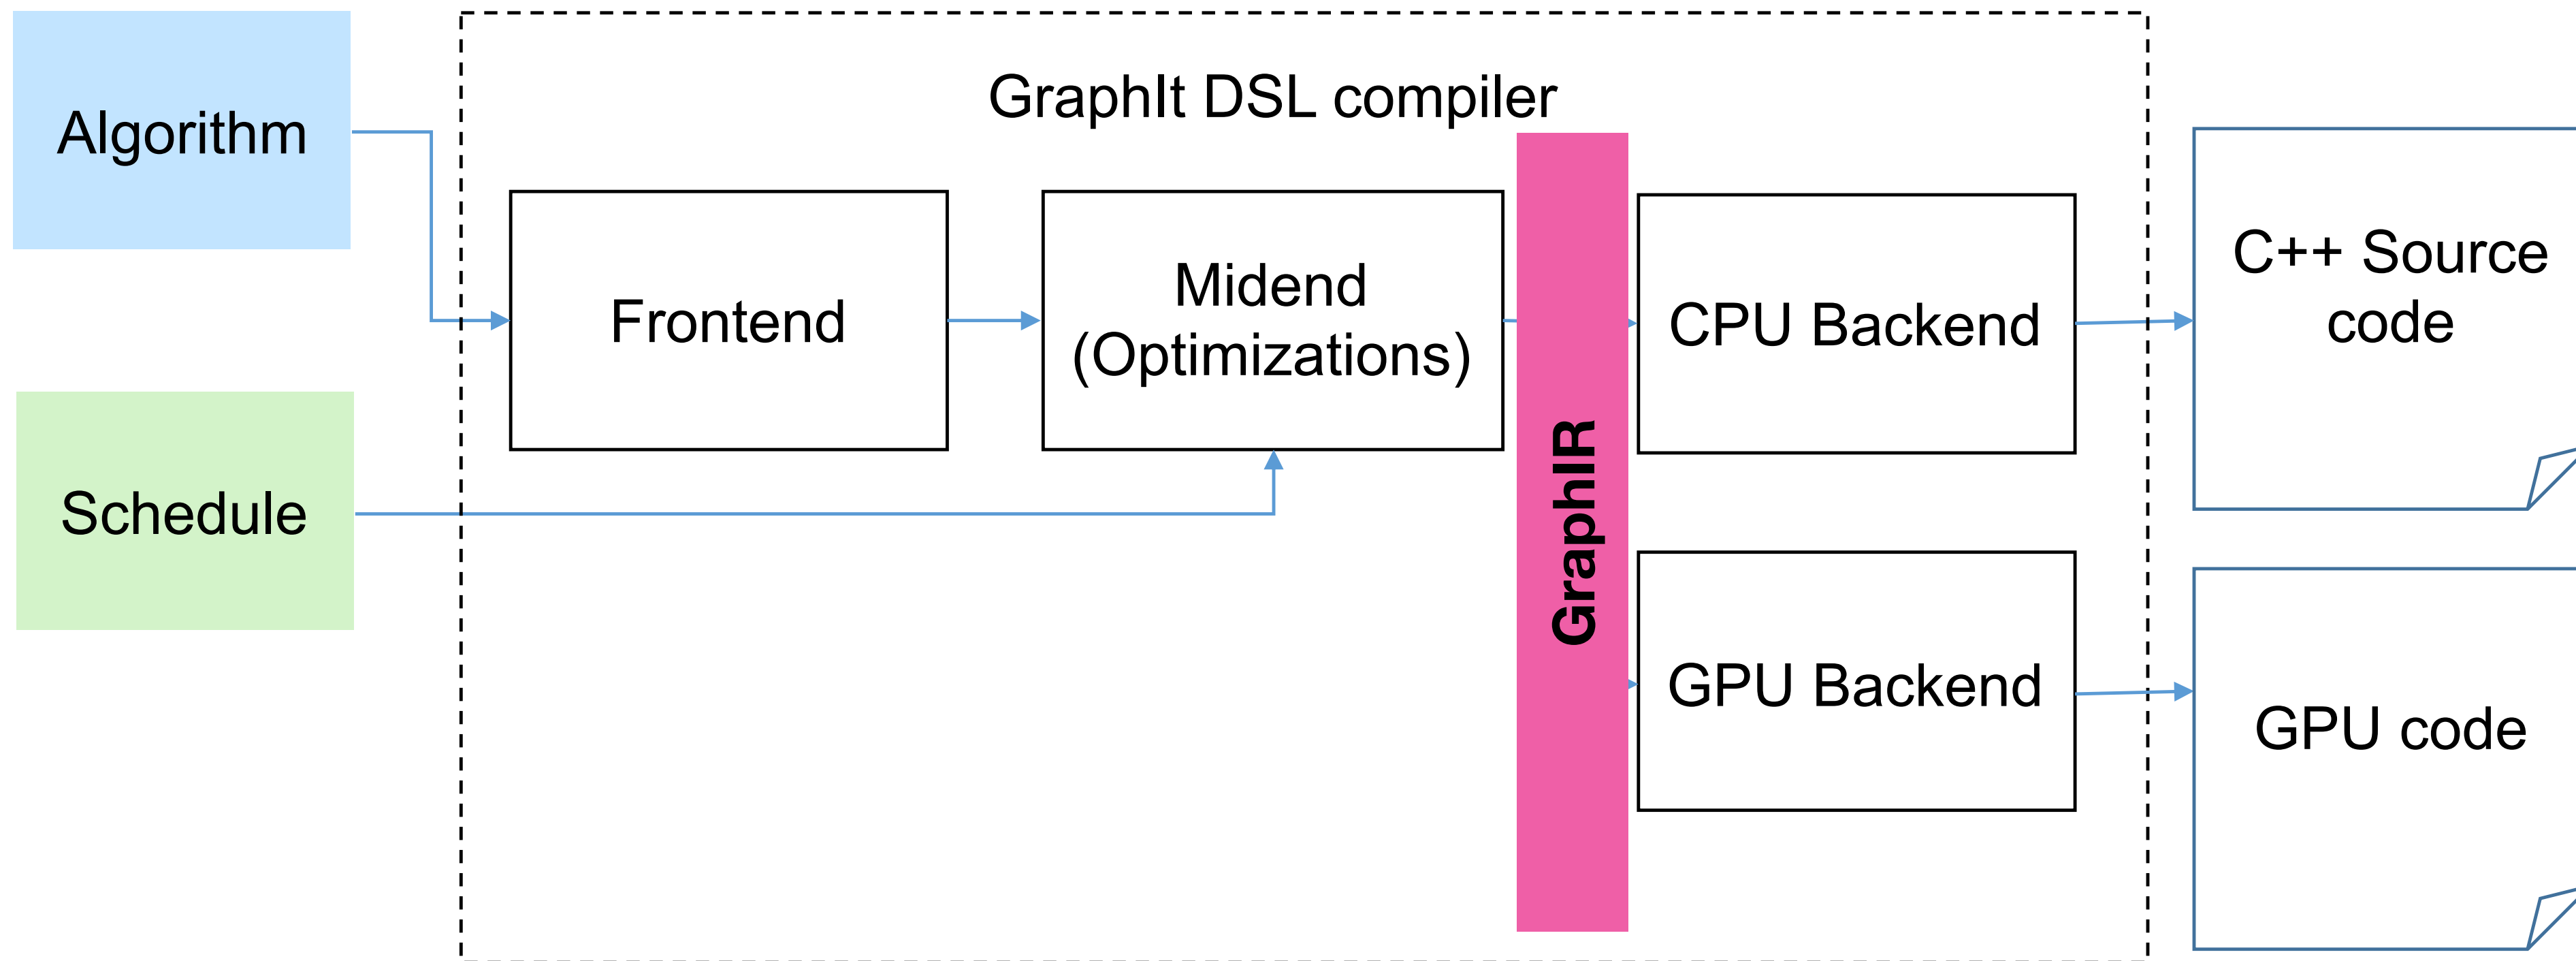


Work In Progress Brahmakshatriya, Zhang, Hong, et al.


GraphIR



GraphIR




GPU Evaluation



**Achieves Consistent
High Performance over
Different Algorithms and
Graphs on GPU**

GPU Evaluation



Related Work

DSLs that separates algorithm from schedules and data layouts:

Halide [PLDI13], Tiramisu [CGO19], Taco [CGO19], TVM [OSDI18], Taichi [SIGGRAPH Asia19] ...

Graph frameworks, optimizations, and architectures that support unordered parallelism:

Ligra [PPoPP13], Galois [SOSP13], GraphBLAS[HPEC 2016], GunRock[PPoPP16], Propagation Blocking [IPDPS17], Cagra [BigData17], Grazelle [PPoPP18], GSwitch[PPoPP19], Sep-Graph[PPoPP19], PHI [MICRO19]
...

Graph frameworks, optimizations, and architectures that support ordered parallelism

Julienne [SPAA17], Galois [SOSP13, PPoPP11], GAPBS [arxiv, IISWC15], Swarm [MICRO15, MICRO16, ISCA17, MICRO18]

Graph DSLs:

GreenMarl [ASPLOS12], EmptyHeaded [SIGMOD16], Elixir [OOPSLA12], Gluon [PLDI18], Abelian [EuroPar18]...

Related Work

DSLs that separates algorithm from schedules and data layouts:

Focus on the Graph Domain

Halide [PLDI13], Tiramisu [CGO19], Taco [CGO19], TVM [OSDI18], Taichi [SIGGRAPH Asia19] ...

Graph frameworks, optimizations, and architectures that support unordered parallelism:

Ligra [PPoPP13], Galois [SOSP13], GraphBLAS[HPEC 2016], GunRock[PPoPP16], Propagation Blocking [IPDPS17], Cagra [BigData17], Grazelle [PPoPP18], GSwitch[PPoPP19], Sep-Graph[PPoPP19], PHI [MICRO19]

...

Graph frameworks, optimizations, and architectures that support ordered parallelism

Focus on Cache and Composing Optimizations

Julienne [SPAA17], Galois [SOSP13, PPoPP11], GAPBS [arxiv, IISWC15], Swarm [MICRO15, MICRO16, ISCA17, MICRO18]

Graph DSLs:

GreenMarl [ASPLOS12], EmptyHeaded [SIGMOD16], Elixir [OOPSLA12], Gluon [PLDI18], Abelian [EuroPar18]...


Focus on Searching through Optimization Space

Lessons Learned

- No set of optimizations fits all graph applications well
- Select the best algorithm, optimization strategy, and hardware platform for each application
- Need software optimizations to fully utilize the underlying hardware (CPU or GPU)

Sparse Graph Computations

- Hardware Utilization
- Peak Performance (PageRank, SpMv)
- < 10% Peak of CPU and GPU




- Programming System

```
template<typename APPLY_FUNC>
void edgeset_apply_pull_parallel(Graph &g, APPLY_FUNC apply_func) {
    int64_t numVertices = g.num_nodes(), numEdges = g.num_edges();
    parallel_for(int n = 0; n < numVertices; n++) {
        for (int socketId = 0; socketId < omp_get_num_places(); socketId++) {
            local_new_rank[socketId][n] = new_rank[n]; } }
    int numPlaces = omp_get_num_places();
    int numSegments = g.getNumSegments("s1");
    int segmentsPerSocket = (numSegments + numPlaces - 1) / numPlaces;
    #pragma omp parallel num_threads(numPlaces) proc_bind(spread){
    int socketId = omp_get_place_num();
    for (int i = 0; i < segmentsPerSocket; i++) {
        int segmentId = socketId + i * numPlaces;
        if (segmentId >= numSegments) break;
        auto sg = g.getSegmentedGraph(std::string("s1"), segmentId);
        #pragma omp parallel num_threads(omp_get_place_num_procs(socketId)) proc_bind(close){
        #pragma omp for schedule(dynamic, 1024)
        for (NodeID localId = 0; localId < sg->numVertices; localId++) {
            NodeID d = sg->graphId[localId];
            for (int64_t ngh = sg->vertexArray[localId]; ngh < sg->vertexArray[localId +
1]; ngh++) {
                NodeID s = sg->edgeArray[ngh];
                local_new_rank[socketId][d] += contrib[s]; }}}}
    parallel_for(int n = 0; n < numVertices; n++) {
        for (int socketId = 0; socketId < omp_get_num_places(); socketId++) {
            new_rank[n] += local_new_rank[socketId][n]; }}}}
    struct updateVertex {
    void operator() (NodeID v) {
        double old_score = old_rank[v];
        new_rank[v] = (base_score + (damp * new_rank[v]));
        error[v] = fabs((new_rank[v] - old_rank[v]));
        old_rank[v] = new_rank[v];
        new_rank[v] = ((float) 0); } };
    void pagerank(Graph &g, double *new_rank, double *old_rank, int *out_degree, int max_iter) {
        for (int i = (0); i < (max_iter); i++) {
            parallel_for(int v_iter = 0; v_iter < builtin_getVertices(edges); v_iter++) {
                contrib[v] = (old_rank[v] / out_degree[v]);};
            edgeset_apply_pull_parallel(edges, updateEdge());
            parallel_for(int v_iter = 0; v_iter < builtin_getVertices(edges); v_iter++) {
                updateVertex()(v_iter); } };
```

Sparse Graph Computations

- Hardware Utilization
- Peak Performance (PageRank, SpMv)
 - ~ 20% Peak of CPU and GPU with cache and other optimizations




- Programming System

```
template<typename APPLY_FUNC>
void edgeset_apply_pull_parallel(Graph &g, APPLY_FUNC apply_func) {
    int64_t numVertices = g.num_nodes(), numEdges = g.num_edges();
    parallel_for(int n = 0; n < numVertices; n++) {
        for (int socketId = 0; socketId < omp_get_num_places(); socketId++) {
            local_new_rank[socketId][n] = new_rank[n]; } }
    int numPlaces = omp_get_num_places();
    int numSegments = g.getNumSegments("s1");
    int segmentsPerSocket = (numSegments + numPlaces - 1) / numPlaces;
    #pragma omp parallel num_threads(numPlaces) proc_bind(spread){
    int socketId = omp_get_place_num();
    for (int i = 0; i < segmentsPerSocket; i++) {
        int segmentId = socketId + i * numPlaces;
        if (segmentId >= numSegments) break;
        auto sg = g.getSegmentedGraph(std::string("s1"), segmentId);
        #pragma omp parallel num_threads(omp_get_place_num_procs(socketId)) proc_bind(close){
        #pragma omp for schedule(dynamic, 1024)
        for (NodeID localId = 0; localId < sg->numVertices; localId++) {
            NodeID d = sg->graphId[localId];
            for (int64_t ngh = sg->vertexArray[localId]; ngh < sg->vertexArray[localId +
1]; ngh++) {
                NodeID s = sg->edgeArray[ngh];
                local_new_rank[socketId][d] += contrib[s]; }}}}
    parallel_for(int n = 0; n < numVertices; n++) {
        for (int socketId = 0; socketId < omp_get_num_places(); socketId++) {
            new_rank[n] += local_new_rank[socketId][n]; }}}}
struct updateVertex {
    void operator() (NodeID v) {
        double old_score = old_rank[v];
        new_rank[v] = (base_score + (damp * new_rank[v]));
        error[v] = fabs((new_rank[v] - old_rank[v]));
        old_rank[v] = new_rank[v];
        new_rank[v] = ((float) 0); } };
void pagerank(Graph &g, double *new_rank, double *old_rank, int *out_degree, int max_iter) {
    for (int i = (0); i < (max_iter); i++) {
        parallel_for(int v_iter = 0; v_iter < builtin_getVertices(edges); v_iter++) {
            contrib[v] = (old_rank[v] / out_degree[v]);};
        edgeset_apply_pull_parallel(edges, updateEdge());
        parallel_for(int v_iter = 0; v_iter < builtin_getVertices(edges); v_iter++) {
            updateVertex()(v_iter); } } }
```

Sparse Graph Computations

- Hardware Utilization
- Peak Performance (PageRank, SpMv)
 - ~ 20% Peak of CPU and GPU with cache and other optimizations



- Programming System

```
func updateEdge (src: Vertex, dst: Vertex)
  new_rank[dst] += old_rank[src] / out_degree[src]
end

func updateVertex (v: Vertex)
  new_rank[v] = base_score + 0.85*new_rank[v];
  old_rank[v] = new_rank[v];
  new_rank[v] = 0;
end

func main()
  for i in 1:11
    #s1# edges.apply(updateEdge);
    vertices.apply(updateVertex);
  end
end
```

schedule:


```
program->configApplyDirection("s1", "DensePull");
program->configApplyParallelization("s1", "dynamic-vertex-parallel");
program->configApplyNumSSG("s1", "fixed-vertex-count", 10);
```

- Open source (graphit-lang.org).



Sparse Graph Computations

- Hardware Utilization
- Peak Performance (PageRank, SpMv)
 - ~ 20% Peak of CPU and GPU with cache and other optimizations



- Programming System
 - Easy-to-Use
 - High-Performance across Different Algorithms and Data
 - Portable across Architectures

- Open source (graphit-lang.org).



Acknowledgements


- Professors and Postdocs
 - Saman Amarasinghe, Julian Shun, Shoaib Kamil, Matei Zaharia, Riyadh Baghdadi, Changwan Hong
- PhD Students
 - Ajay Brahmakshatriya, Vladimir Kiriansky, Laxman Dhulipala, Charith Mendis
- Master and Undergraduate Students
 - Mengjiao (Sherry) Yang, Tugsbayasgalan Manlaibaatar, Xinyi Chen, Claire Hsu, Haokuan Luo, Kenny Yang

This Work Supported By:



Summary

- Hardware Utilization
- Peak Performance (PageRank, SpMv)
 - ~ 20% Peak of CPU and GPU with cache and other optimizations



- Programming System
 - Easy-to-Use
 - High-Performance across Different Algorithms and Data
 - Portable across Architectures

- Open source (graphit-lang.org).

