

Making Pull-Based Graph Processing Performant

Samuel Grossman¹, Heiner Litz², and Christos Kozyrakis¹

¹Stanford University ²University of California, Santa Cruz

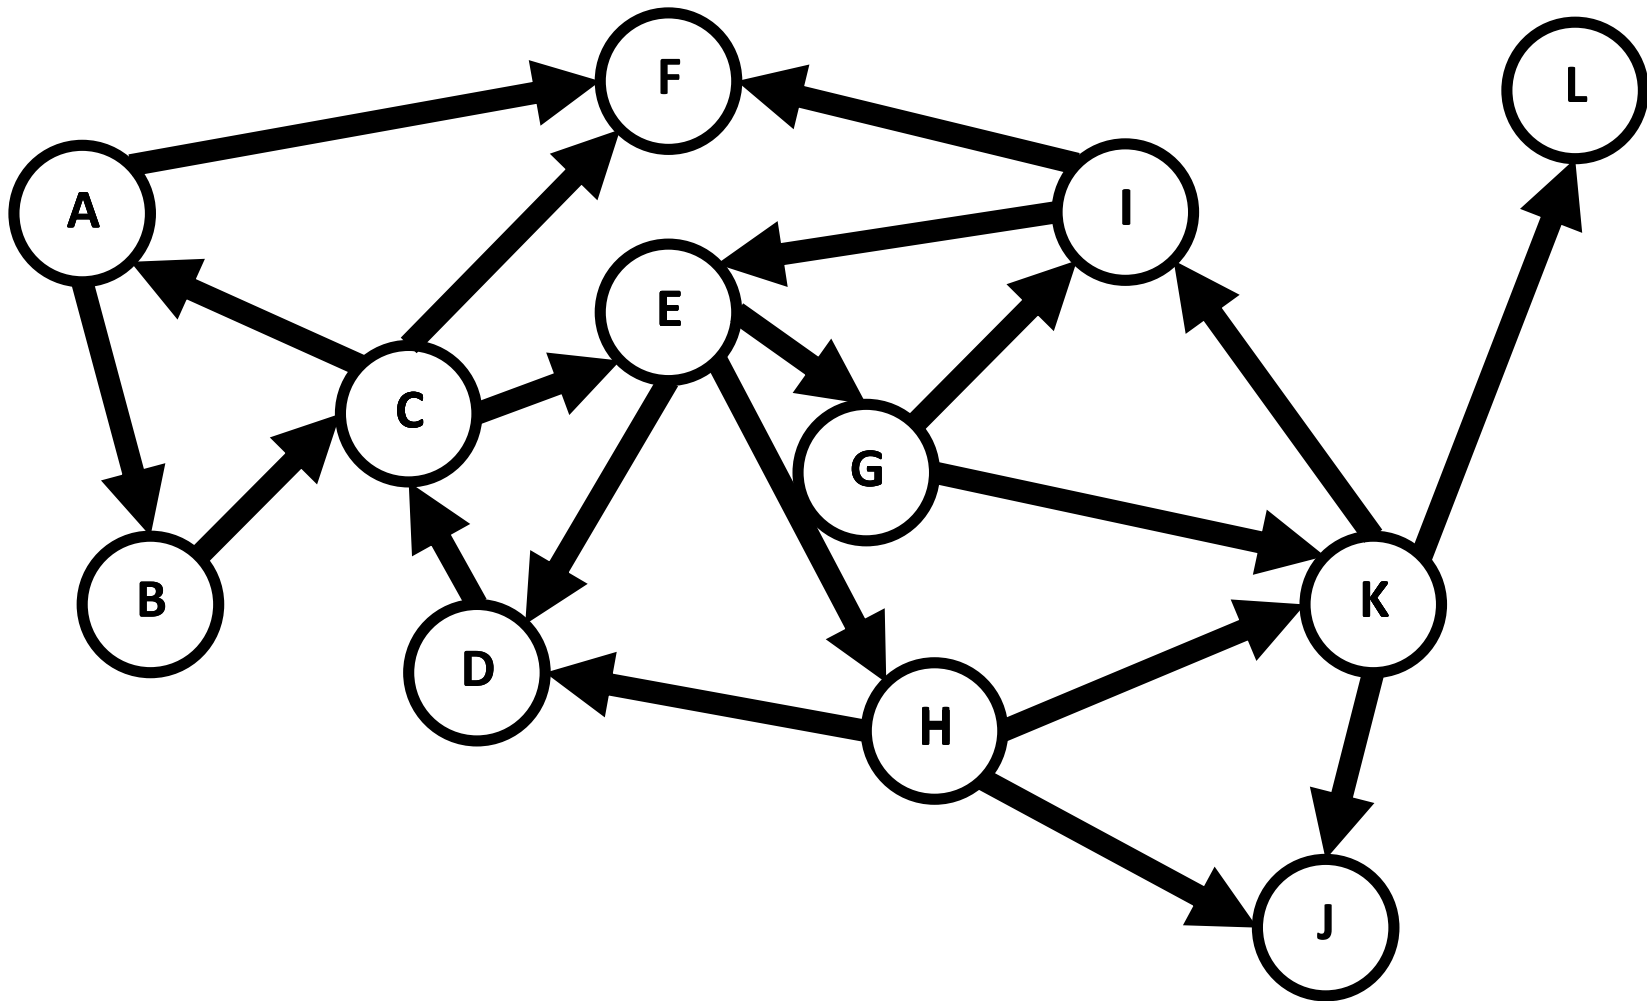
PPoPP 2018 · February 27, 2018



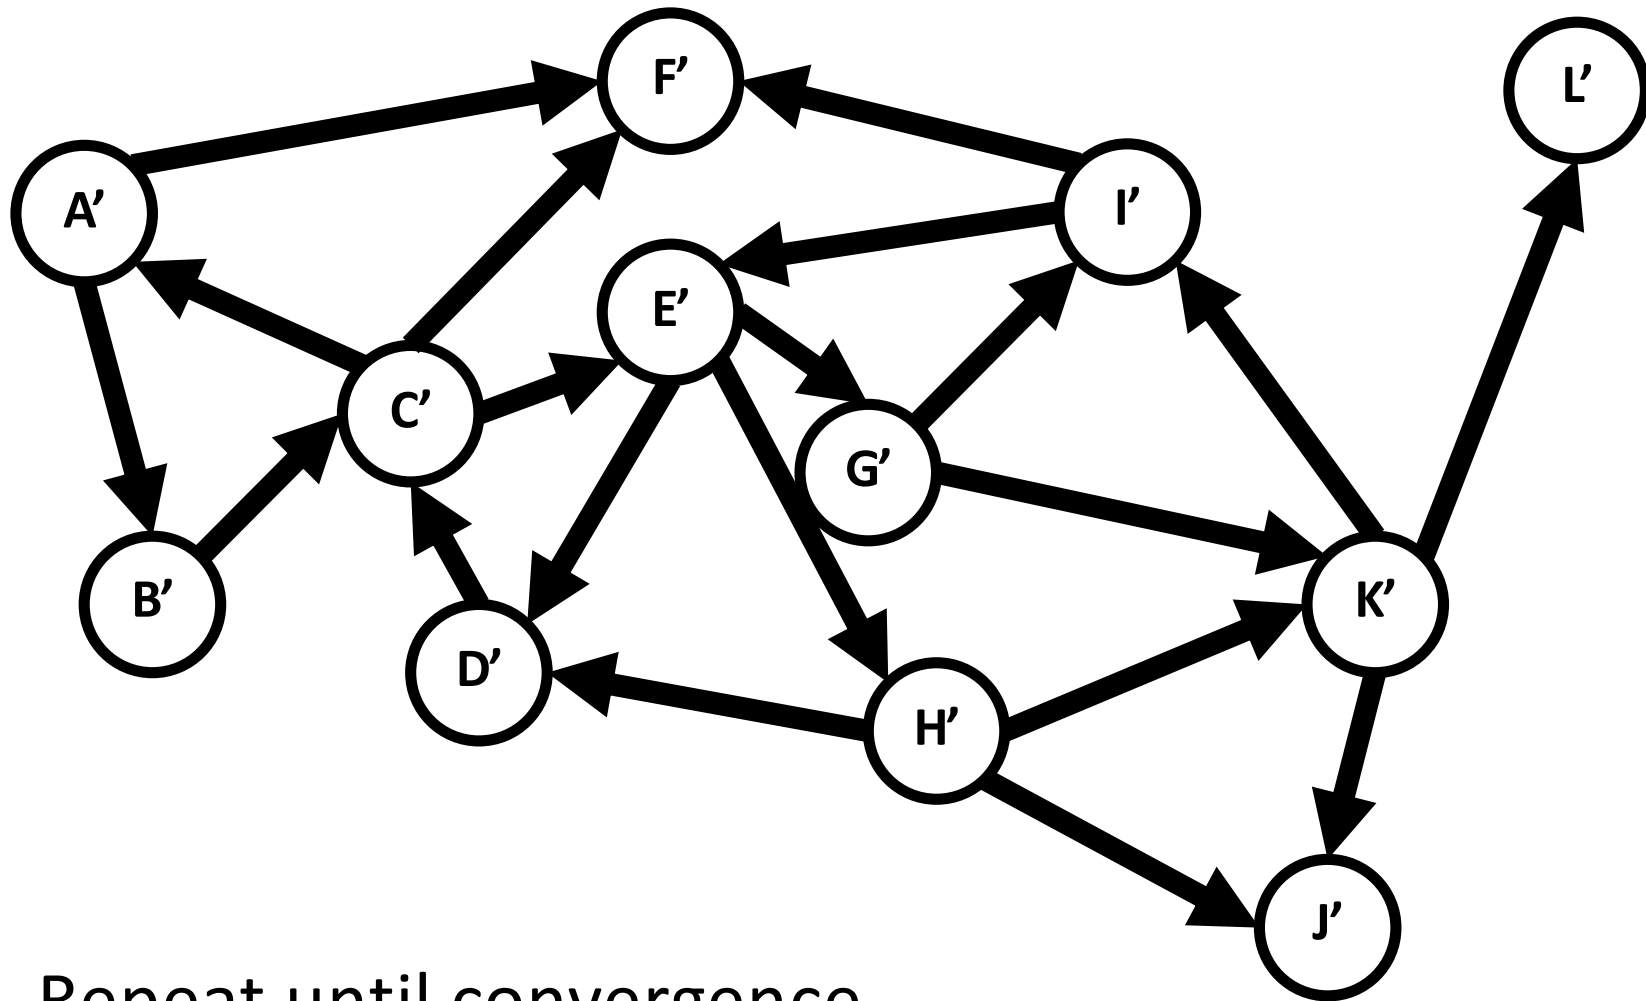
Graph Processing

- Problem modelled as **objects** (vertices) and **connections between them** (edges)
- Examples:
 - Internet (pages and hyperlinks)
 - Social network (people and friendships)
 - Roads and intersections
 - Products and ratings

Graph Processing



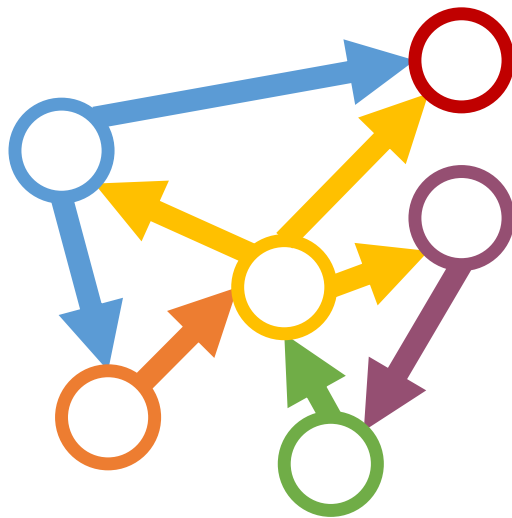
Graph Processing



Repeat until convergence

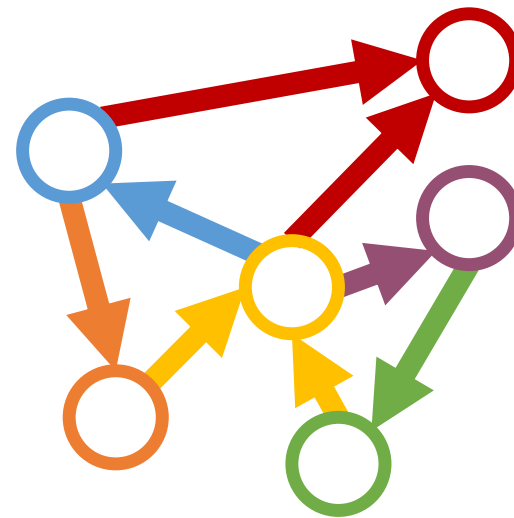
Graph Processing

Push



Group by **source vertex**

Pull



Group by **destination vertex**

Hybrid: dynamically select push or pull for each iteration

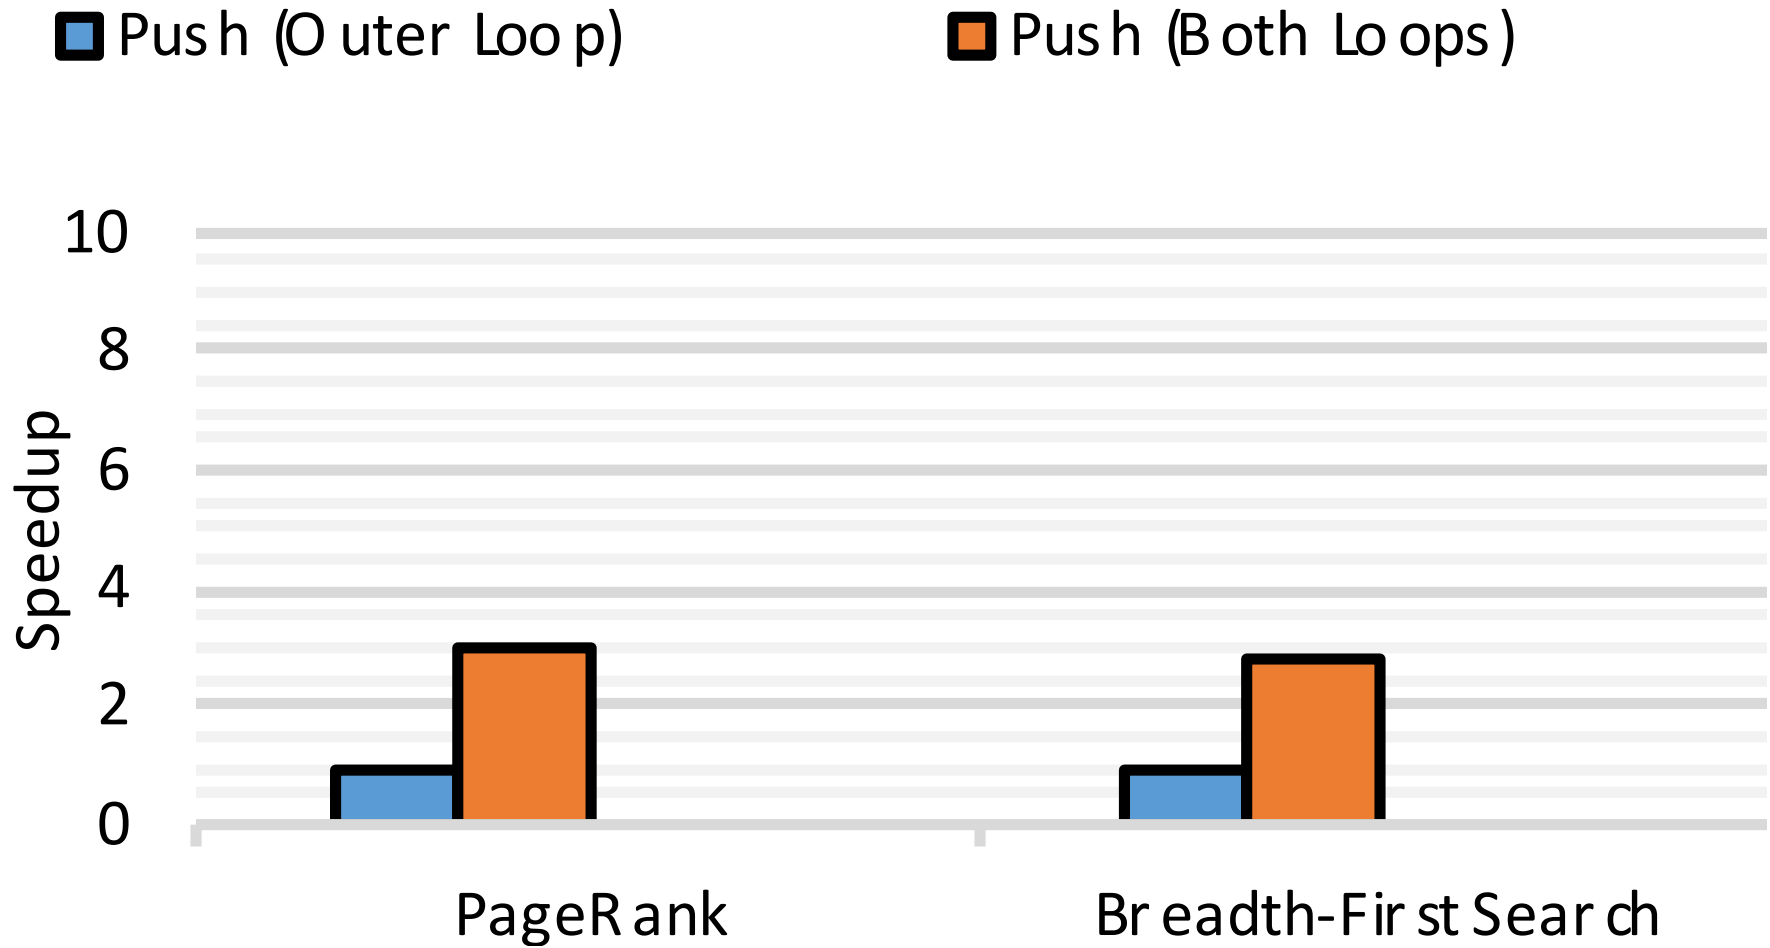
Graph Processing

```
foreach vertex v in graph.vertices  
    foreach edge e in v.(in|out)edges  
        // process the edge  
        ...
```

Parallelizing Graph Processing

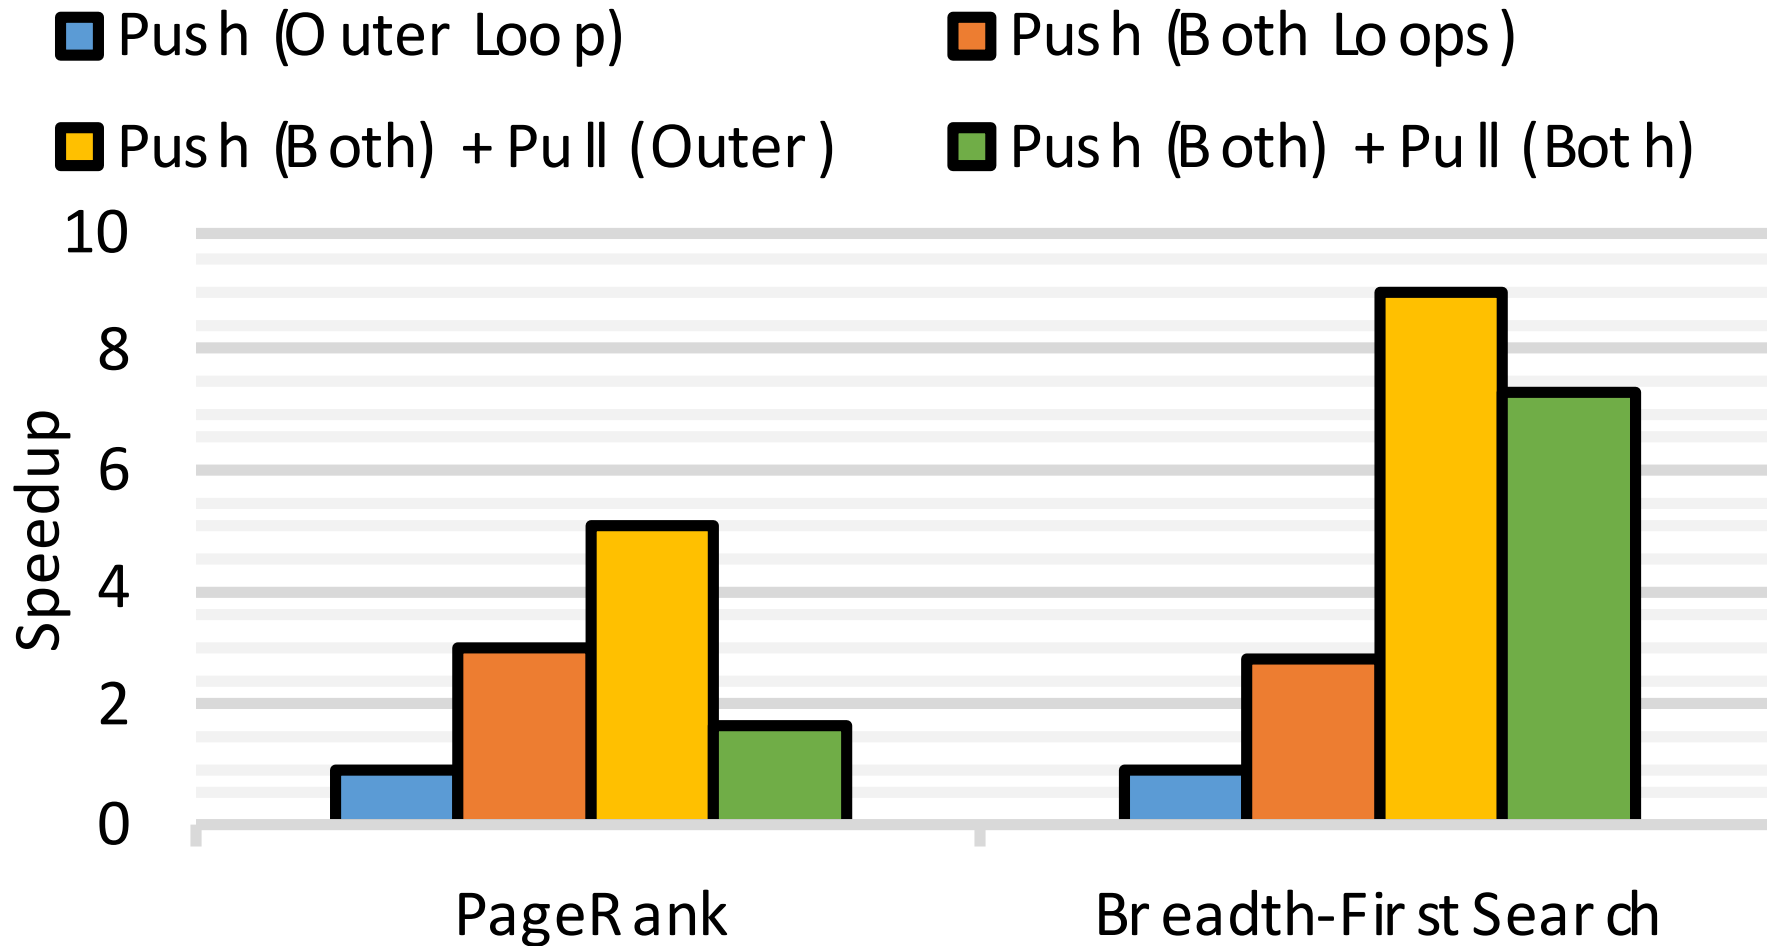
- Outer loop parallelization
 - **Between cores:** assign *entire vertices* to threads
- Inner loop parallelization
 - **Between cores:** subdivide the edges within each vertex
 - **Within one core:** vectorize the loop

Parallelizing Graph Processing



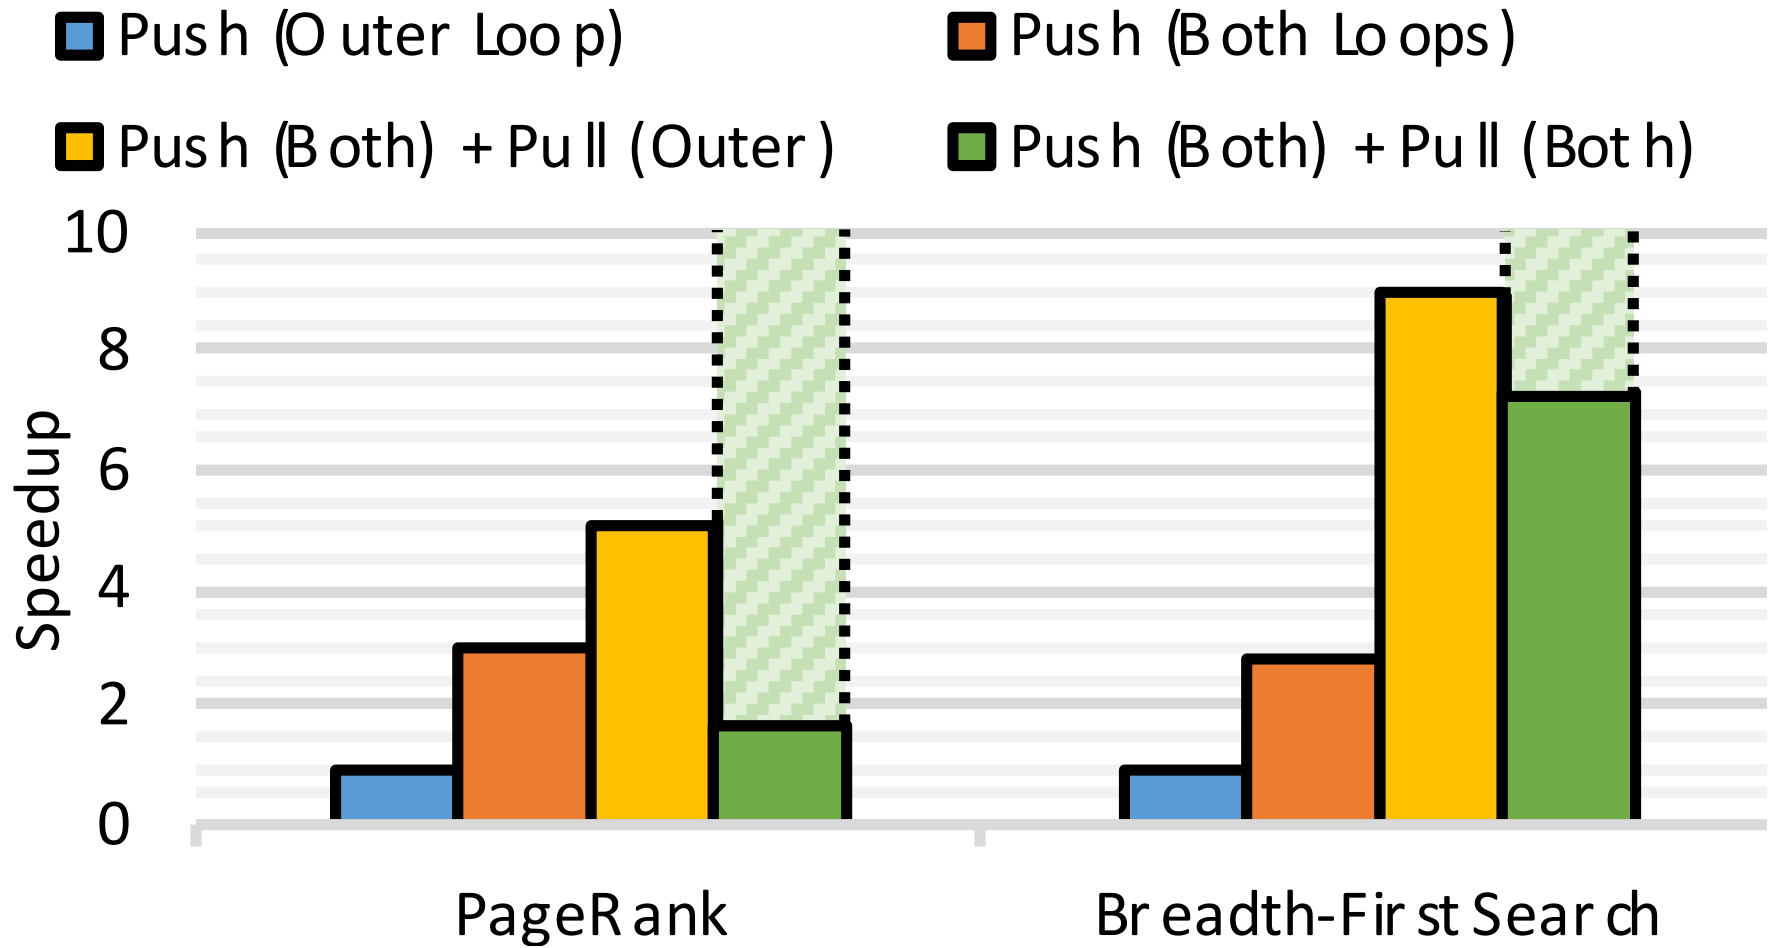
Running Ligra on *twitter-2010* graph

Parallelizing Graph Processing



Running Ligra on *twitter-2010* graph

Parallelizing Graph Processing



Running Ligra on *twitter-2010* graph

Pull's Performance Challenge

Serial Inner Loop

Parallel Inner Loop

Contribution #1: "Scheduler Awareness"

A technique that can be applied to the inner loop of a pull engine to parallelize it without introducing conflicts.

- One thread per vertex
- Updates are thread-private
- Multiple threads per vertex
- Updates will conflict

Pull's Performance Opportunity

- Further gains possible using SIMD vectorization

- Improve parallelism of the computation

- Improve memory bandwidth utilization

Contribution #2: “Vector-Sparse”

A low-level modification to a data structure commonly used to represent graphs, intended to enhance vectorization.

Data structure layout issues impede effective vectorization in existing work

Grazelle

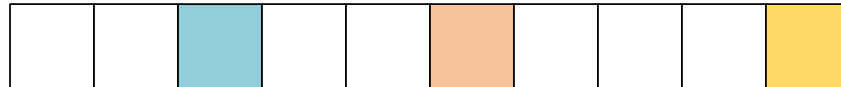
- A hybrid graph processing framework that embodies both of our contributions
- Outperforms the state-of-the-art by over 10× in some cases
- Available for download at <https://github.com/stanford-mast/Grazelle-PPoPP18>

Scheduler Awareness

Contribution #1

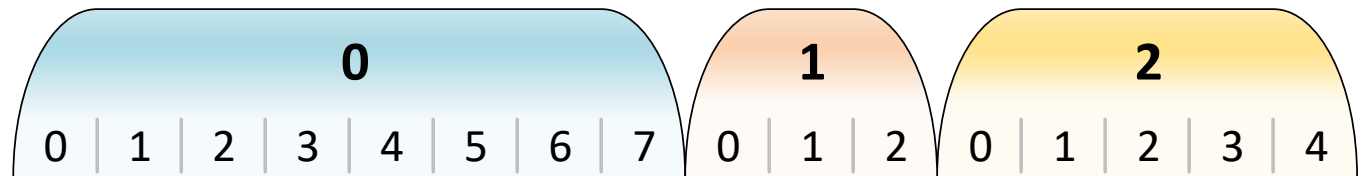
Serial Inner Loop

Vertex Data

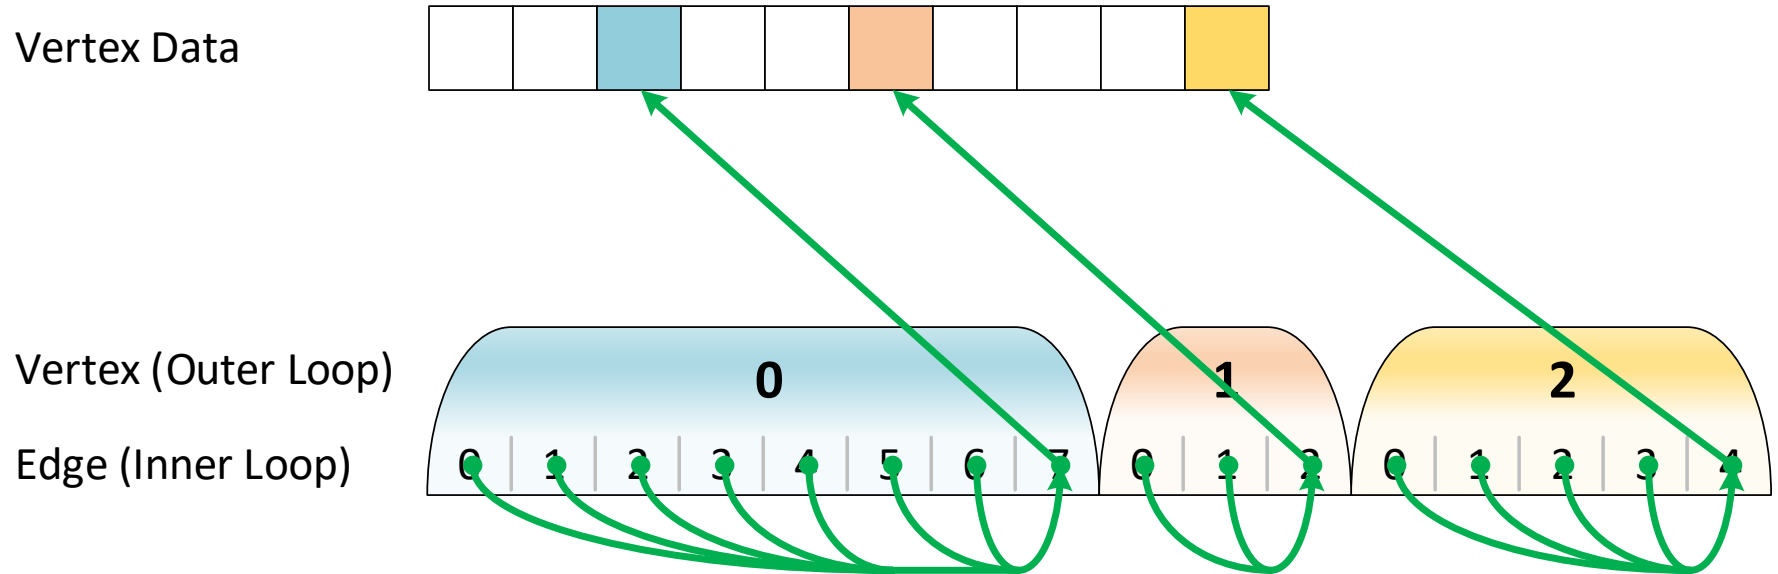


Vertex (Outer Loop)

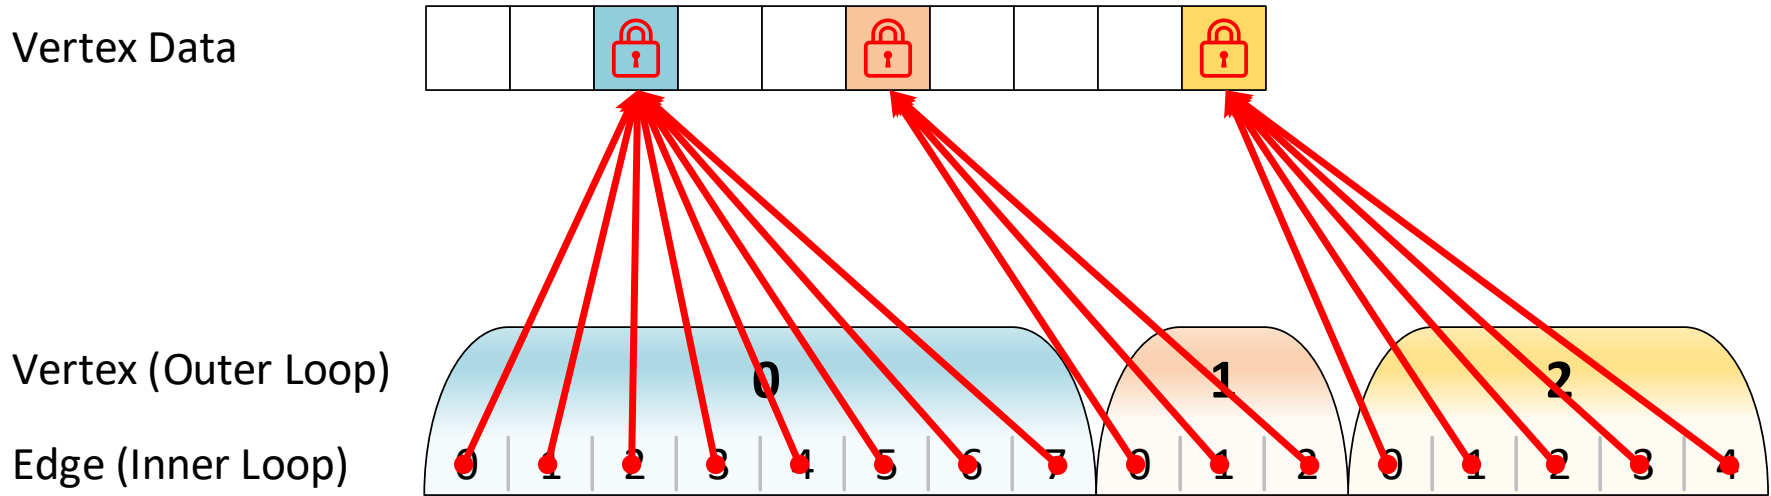
Edge (Inner Loop)



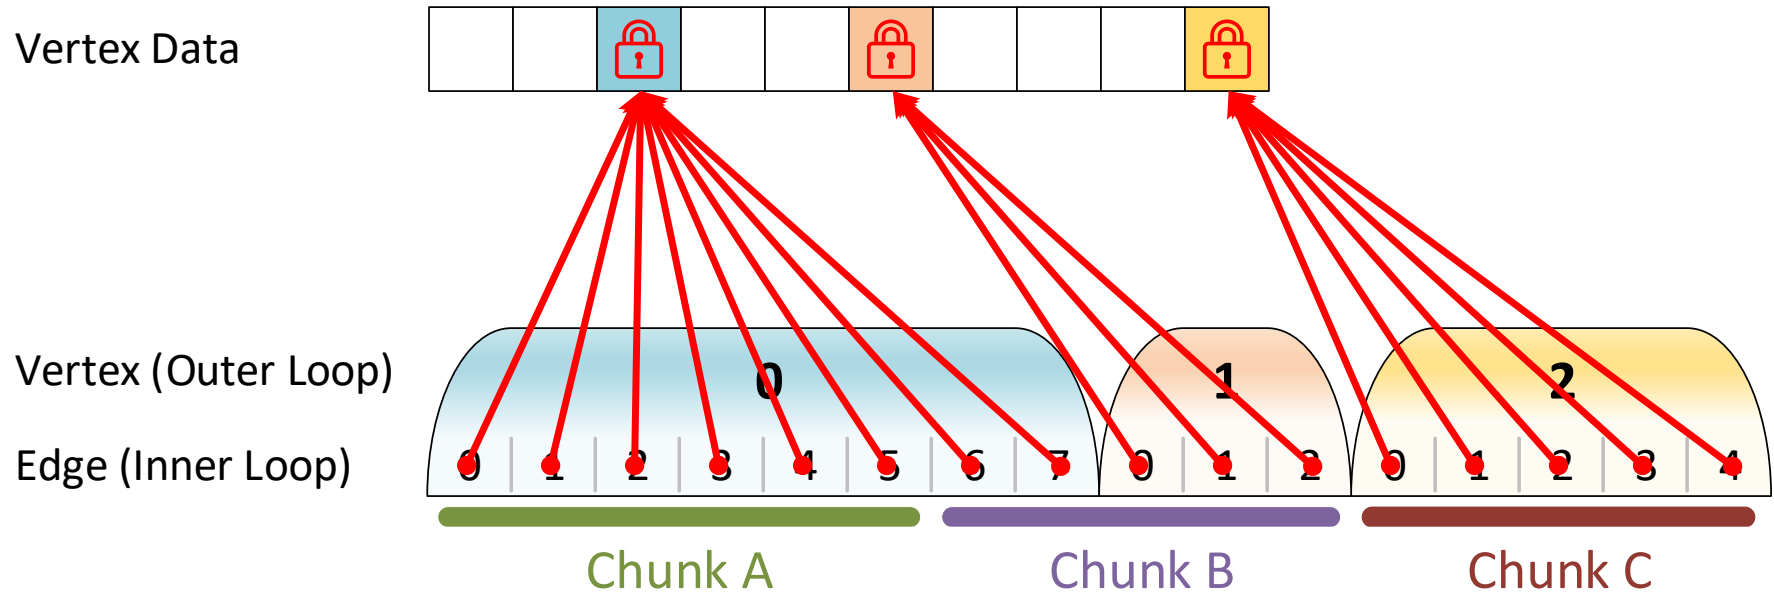
Serial Inner Loop



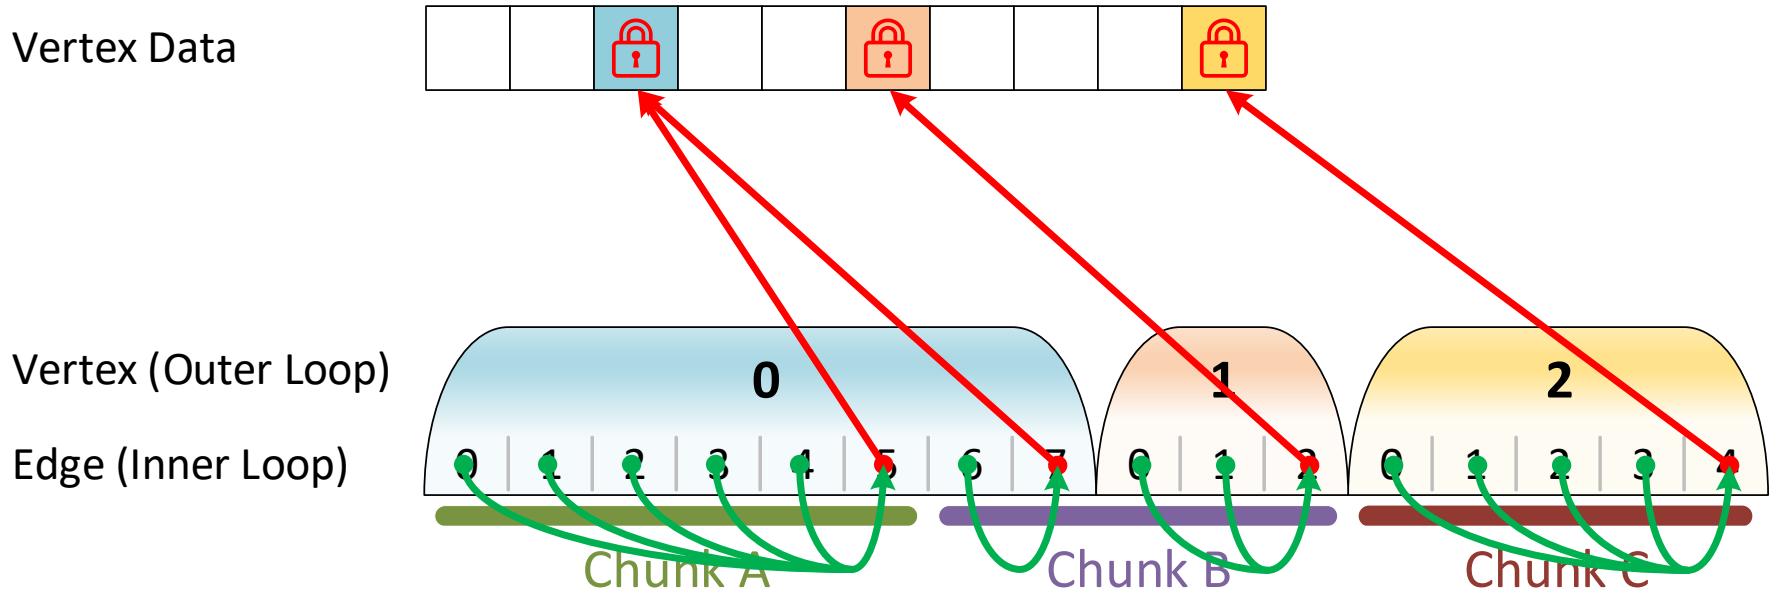
Scheduler Un-Awareness



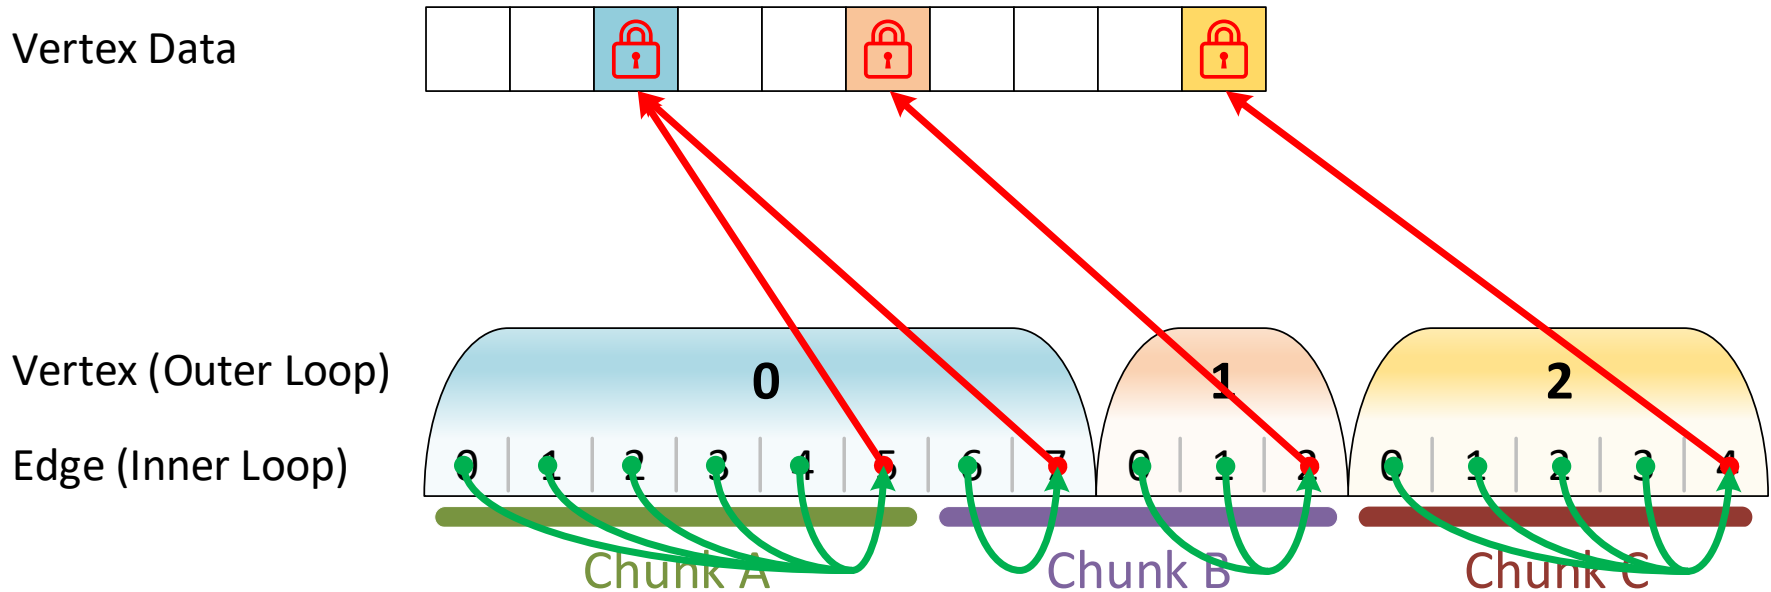
Scheduler Un-Awareness



Scheduler Awareness

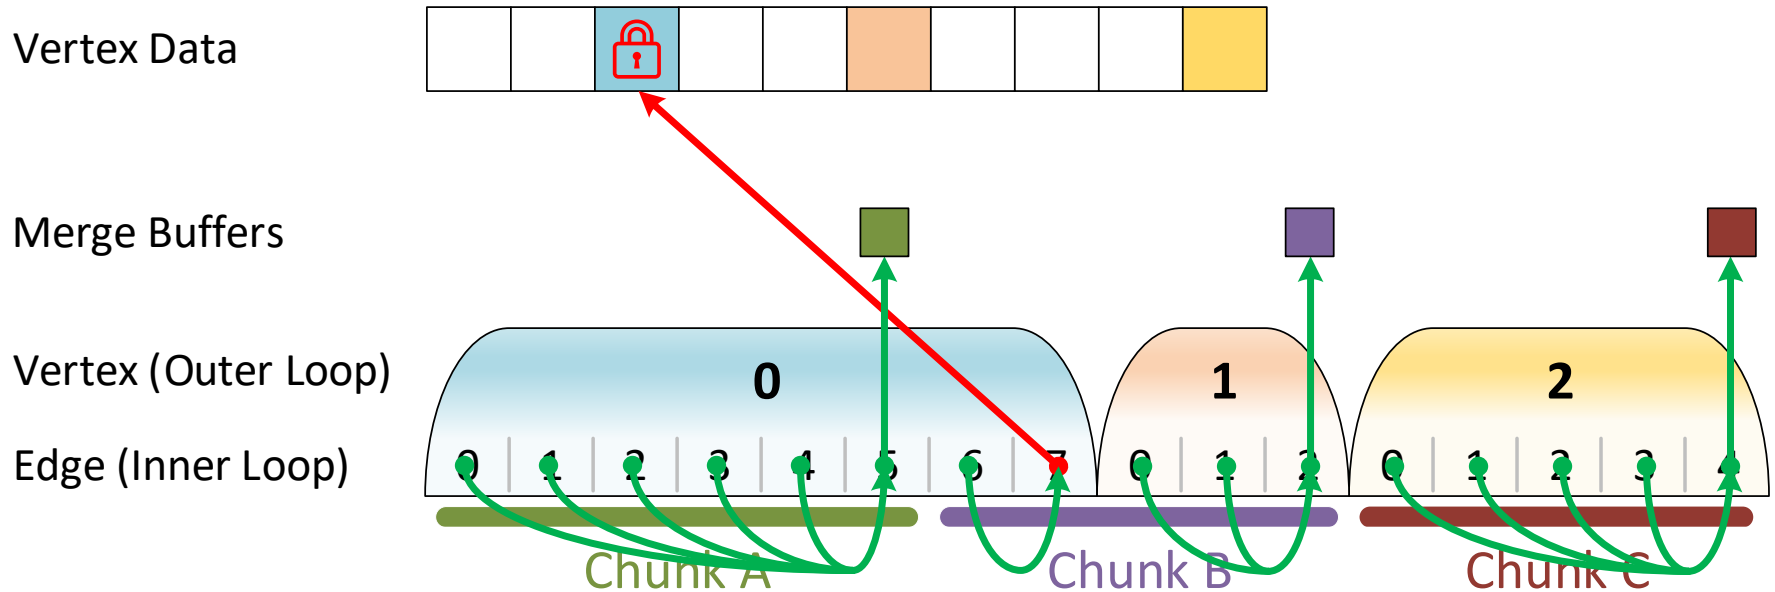


Scheduler Awareness



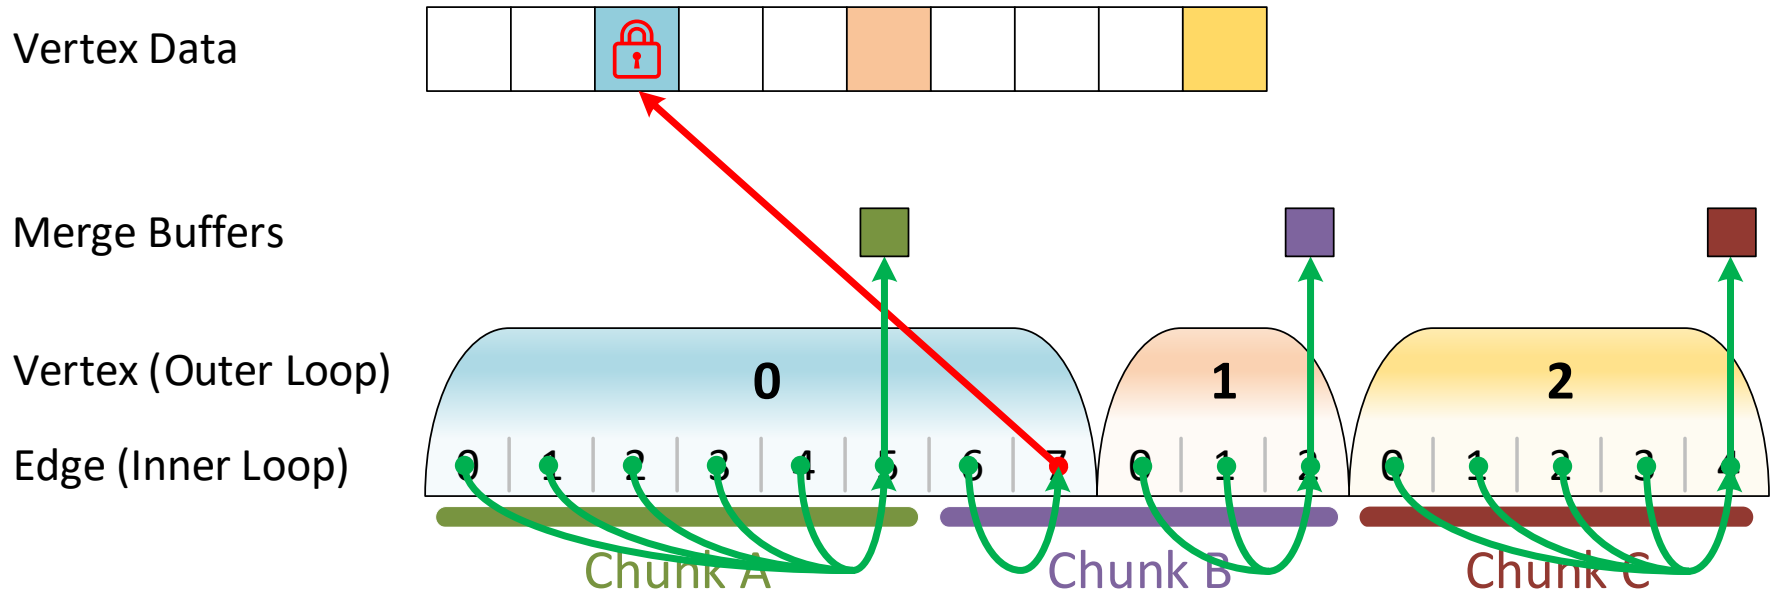
1. Writes **at the end of a chunk** are redirected to a private per-chunk merge buffer.

Scheduler Awareness



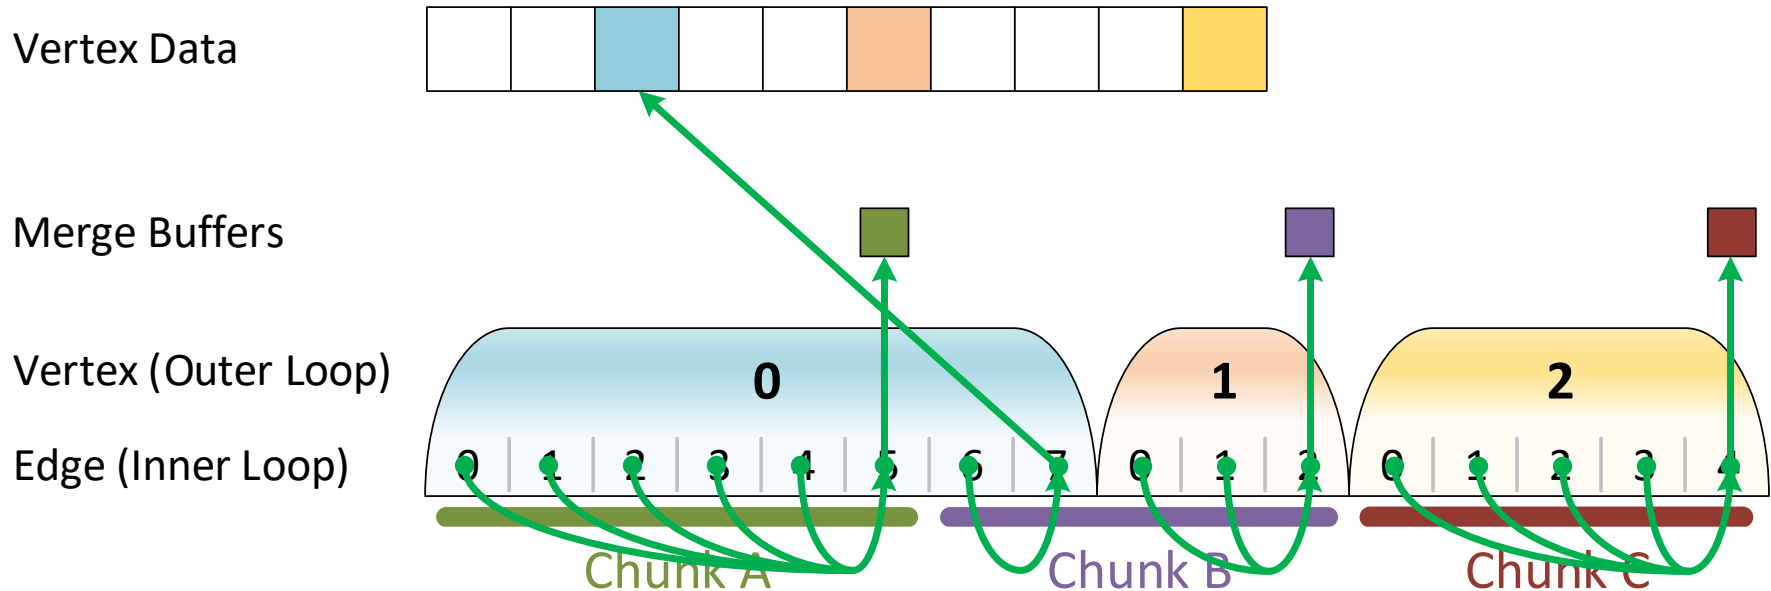
1. Writes **at the end of a chunk** are redirected to a private per-chunk merge buffer.

Scheduler Awareness



1. Writes **at the end of a chunk** are redirected to a private per-chunk merge buffer.
2. Writes **in the middle of a chunk** can be committed to shared state without synchronization.

Scheduler Awareness



1. Writes **at the end of a chunk** are redirected to a private per-chunk merge buffer.
2. Writes **in the middle of a chunk** can be committed to shared state without synchronization.

Analyzing Scheduler Awareness

- Performance impact depends primarily on the **scheduling granularity**
- **Scheduler Un-Awareness:** trade-off between load balance and probability of write conflicts
- **Scheduler Awareness:** finer granularity leads to increased merge operation overhead

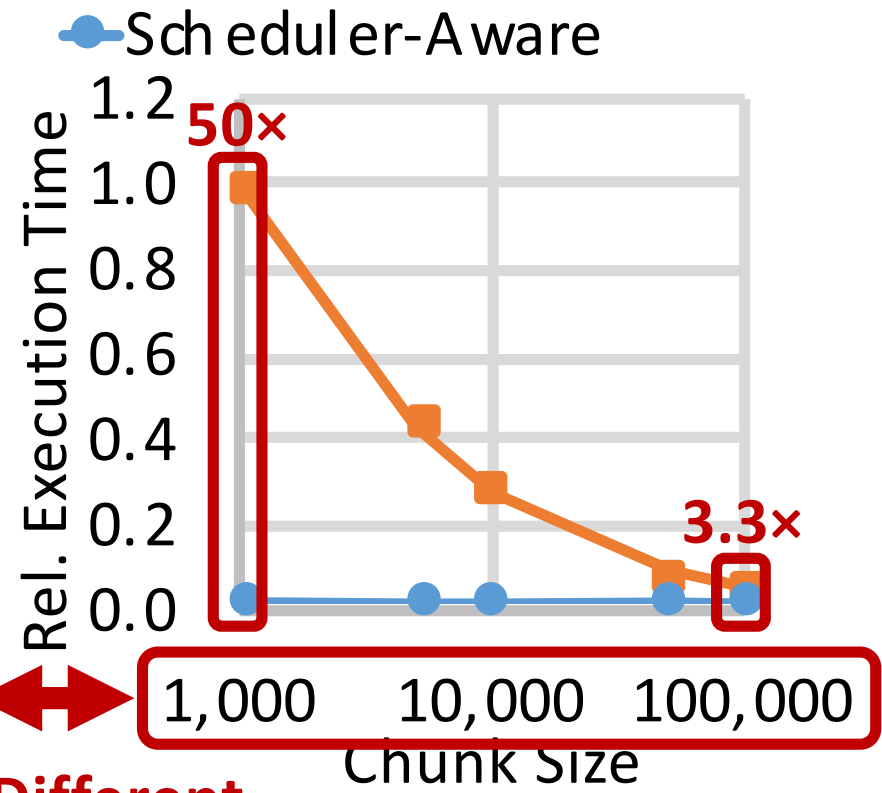
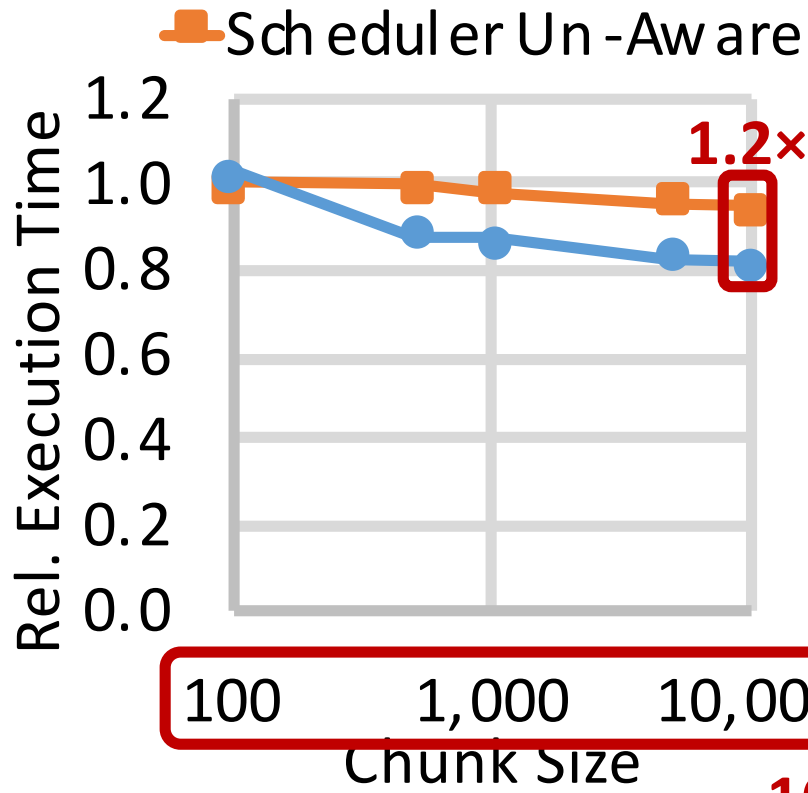
PageRank: Performance vs. Scheduling Granularity

dimacs-usa

uk-2007

(low, even degree distribution)

(extremely skewed)



10x Different

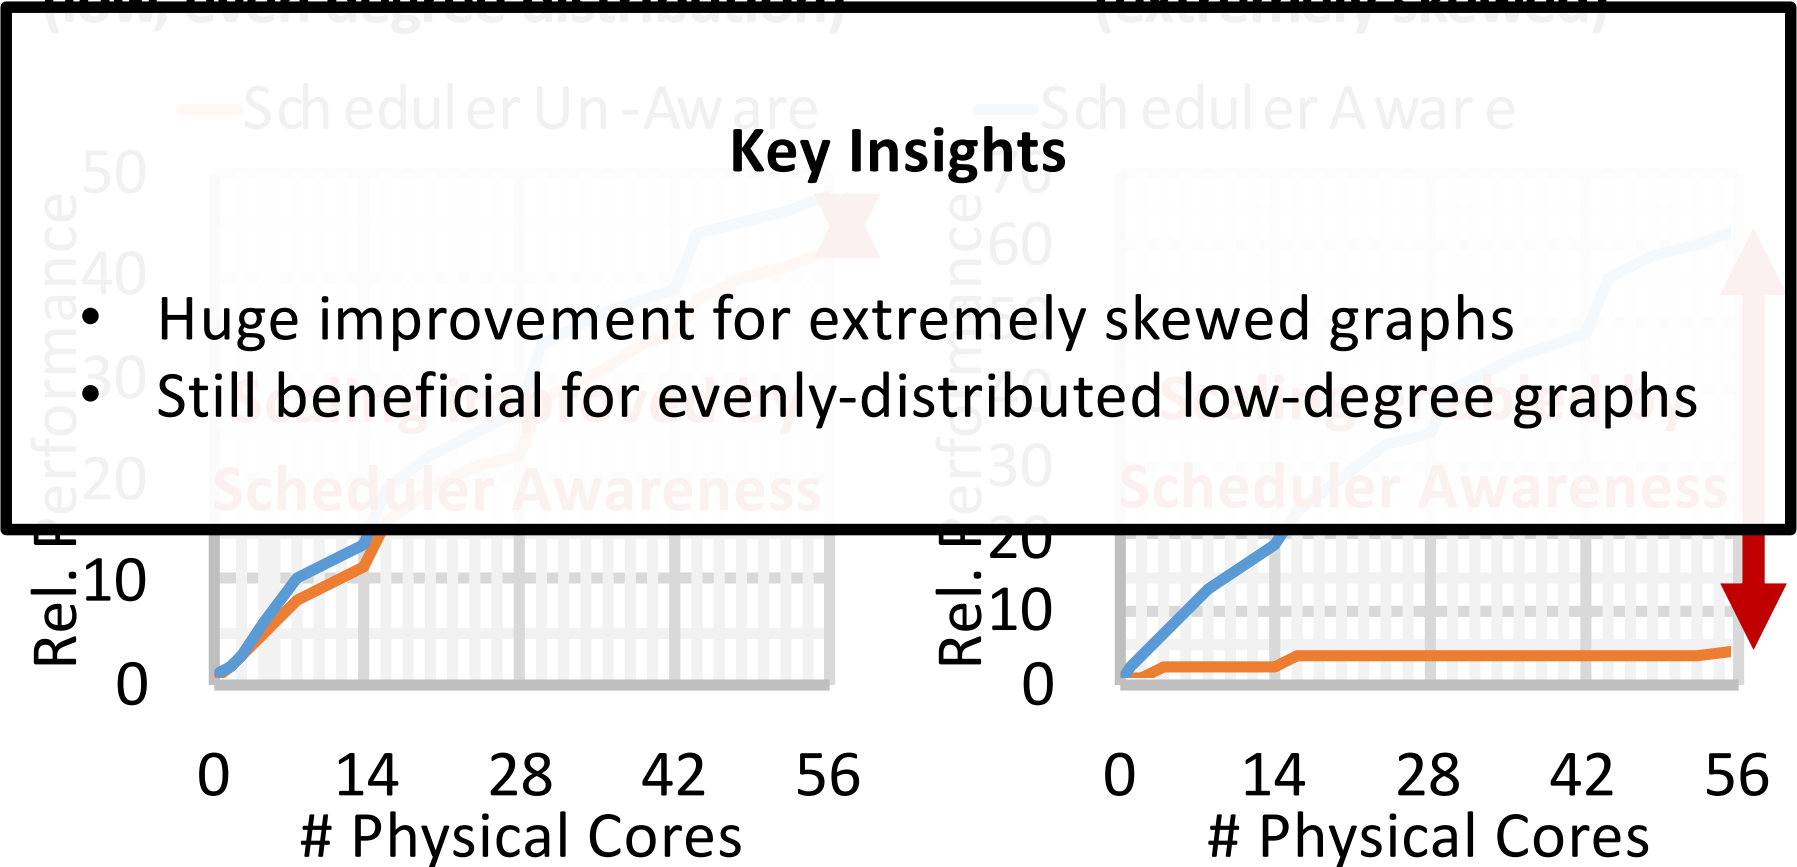
PageRank: Performance vs. Number of Cores

dimacs-usa

uk-2007

(low, even degree distribution)

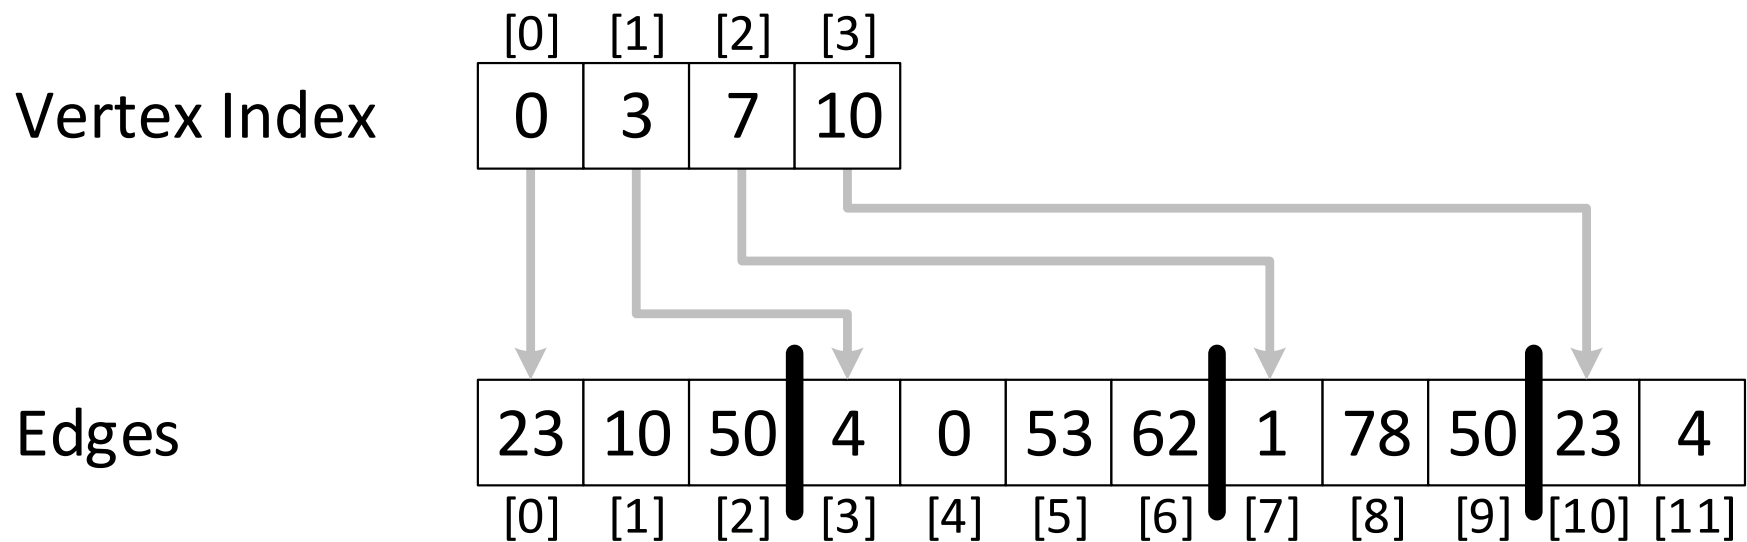
(extremely skewed)



Vector-Sparse

Contribution #2

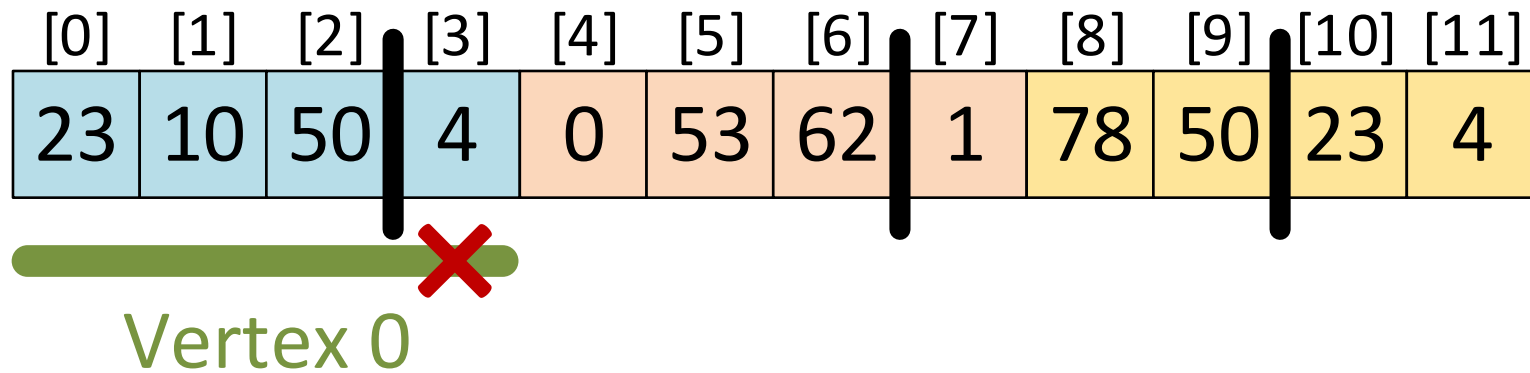
Compressed-Sparse



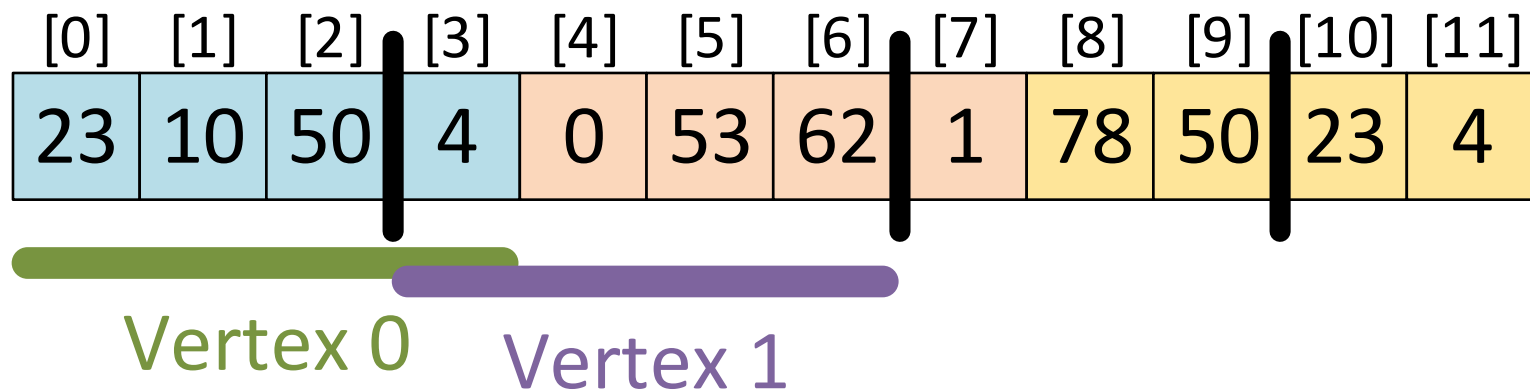
Vectorizing Compressed-Sparse

[0]	[1]	[2]	[3]	[4]	[5]	[6]	[7]	[8]	[9]	[10]	[11]
23	10	50	4	0	53	62	1	78	50	23	4

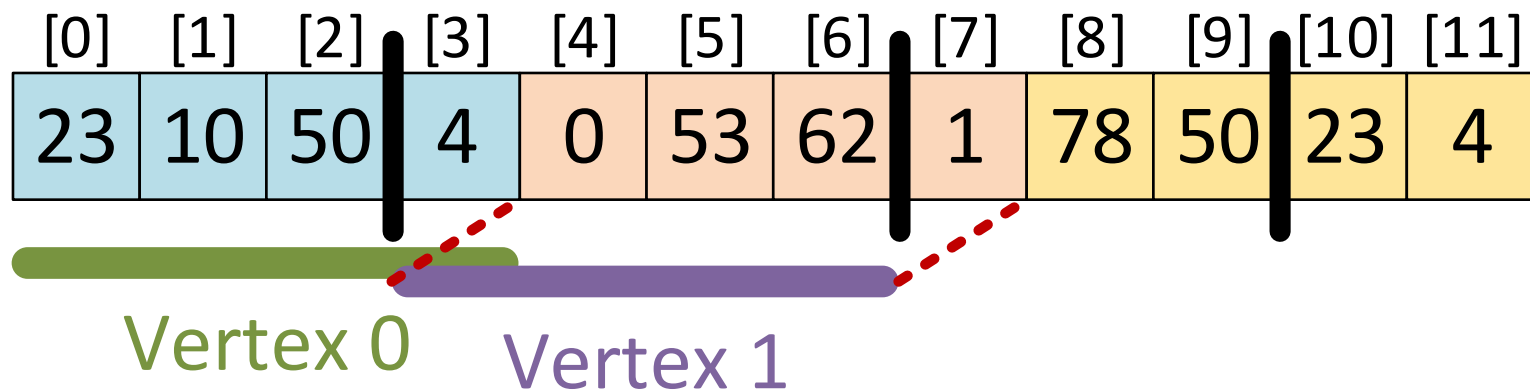
Vectorizing Compressed-Sparse



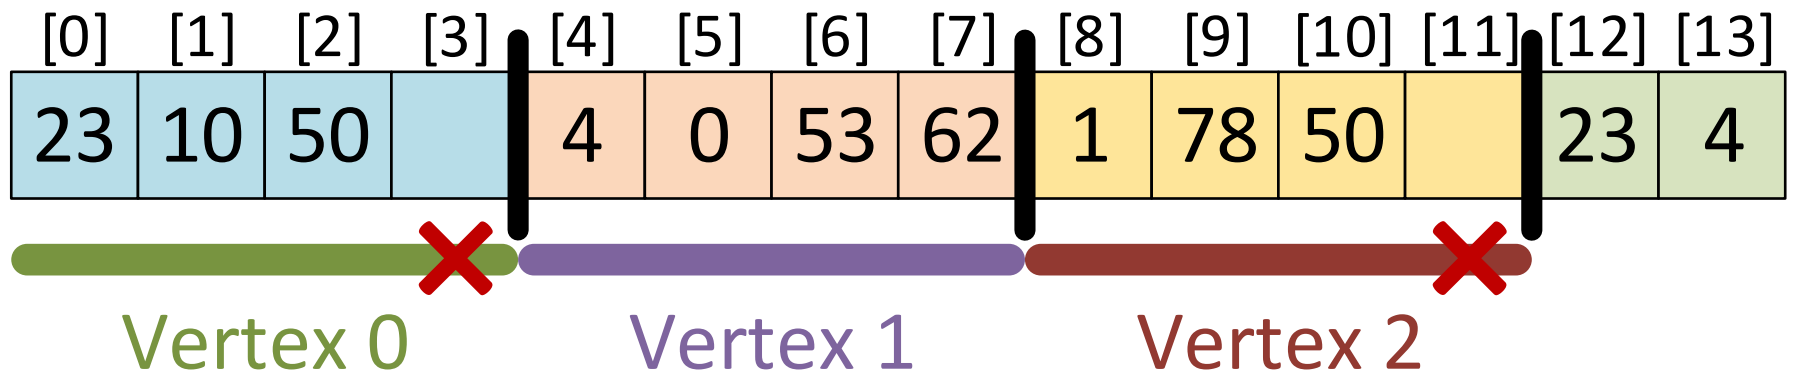
Vectorizing Compressed-Sparse



Vectorizing Compressed-Sparse

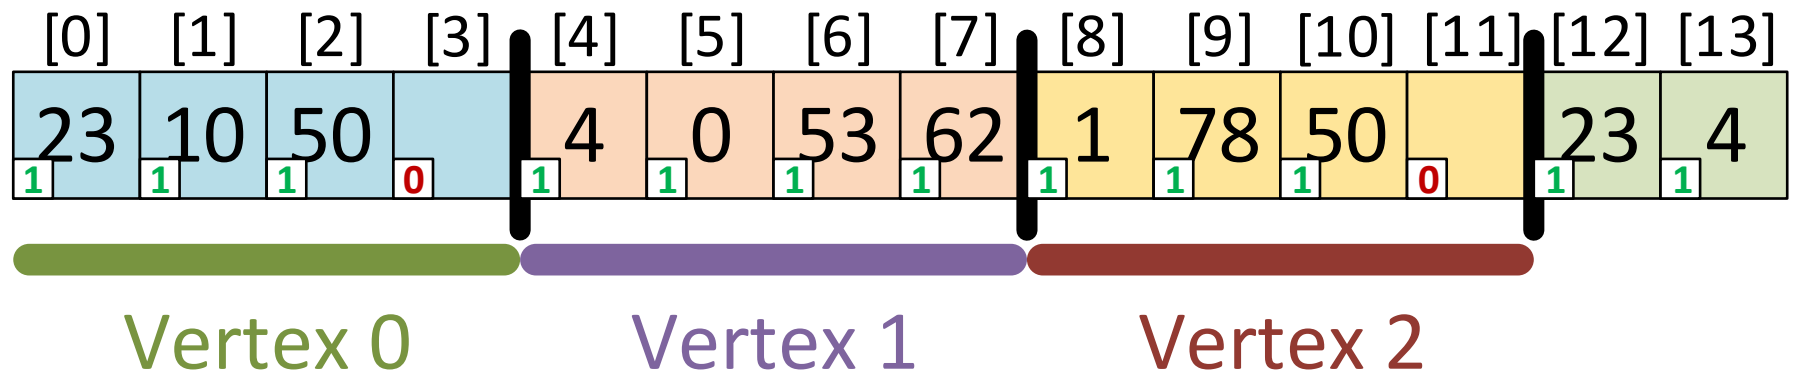


Vector-Sparse



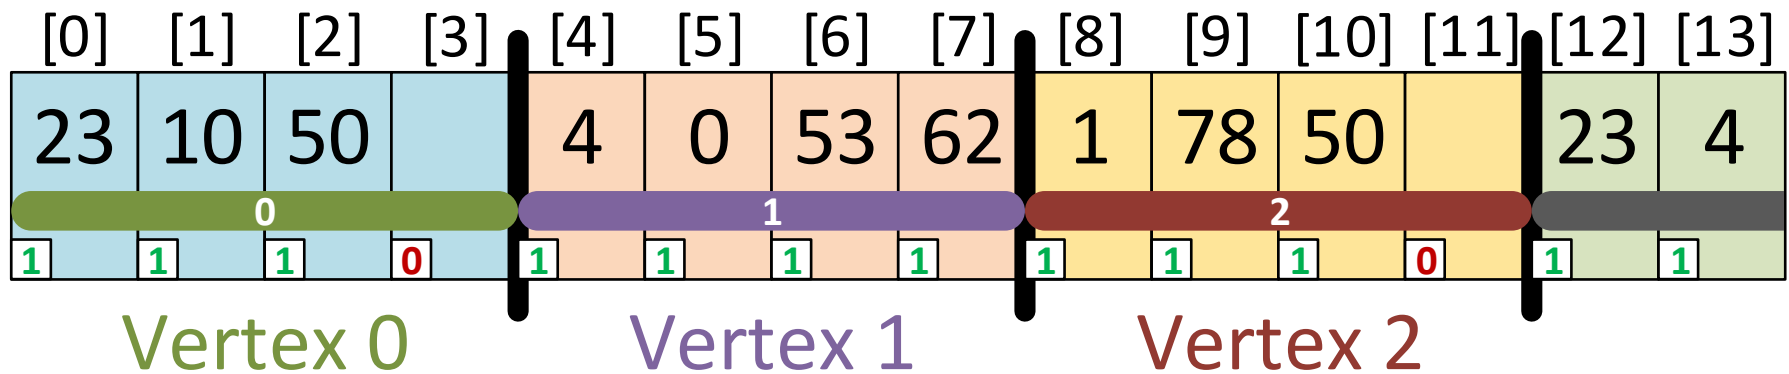
Padding

Vector-Sparse



Padding + "valid" bits

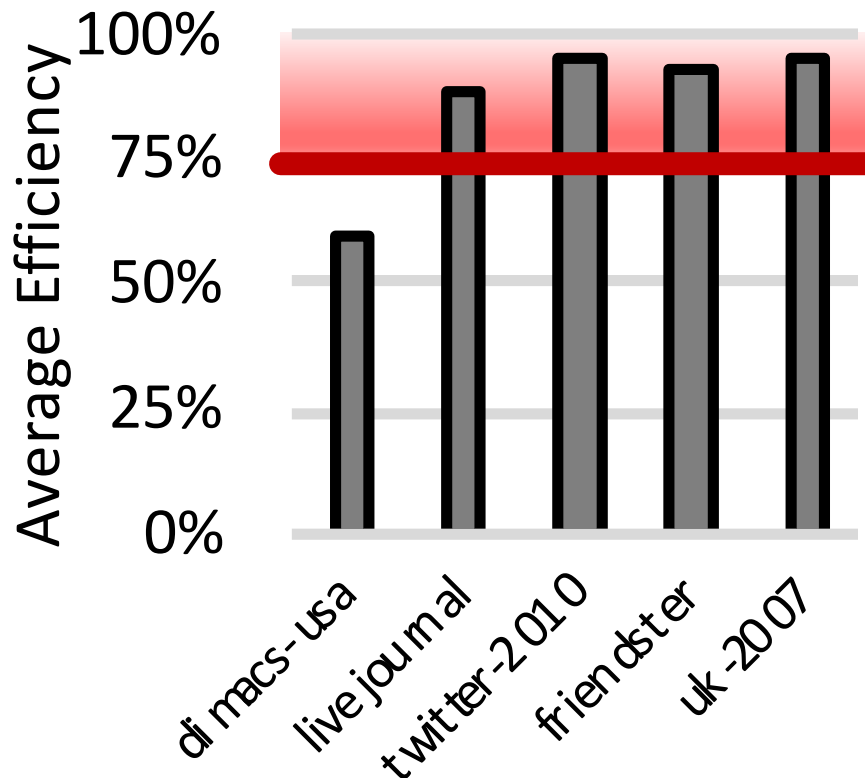
Vector-Sparse



Padding + “valid” bits + top-level vertex spread-encoding

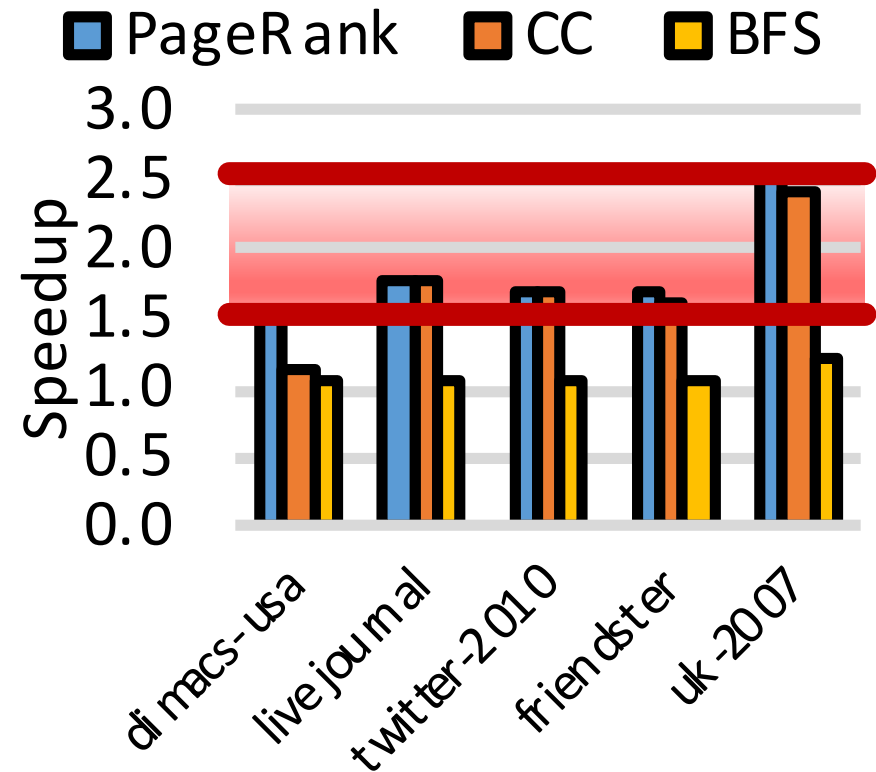
Analyzing Vector-Sparse

Packing Efficiency



Generally $\geq 75\%$

Performance Impact



1.5x to 2.5x

Performance Comparison

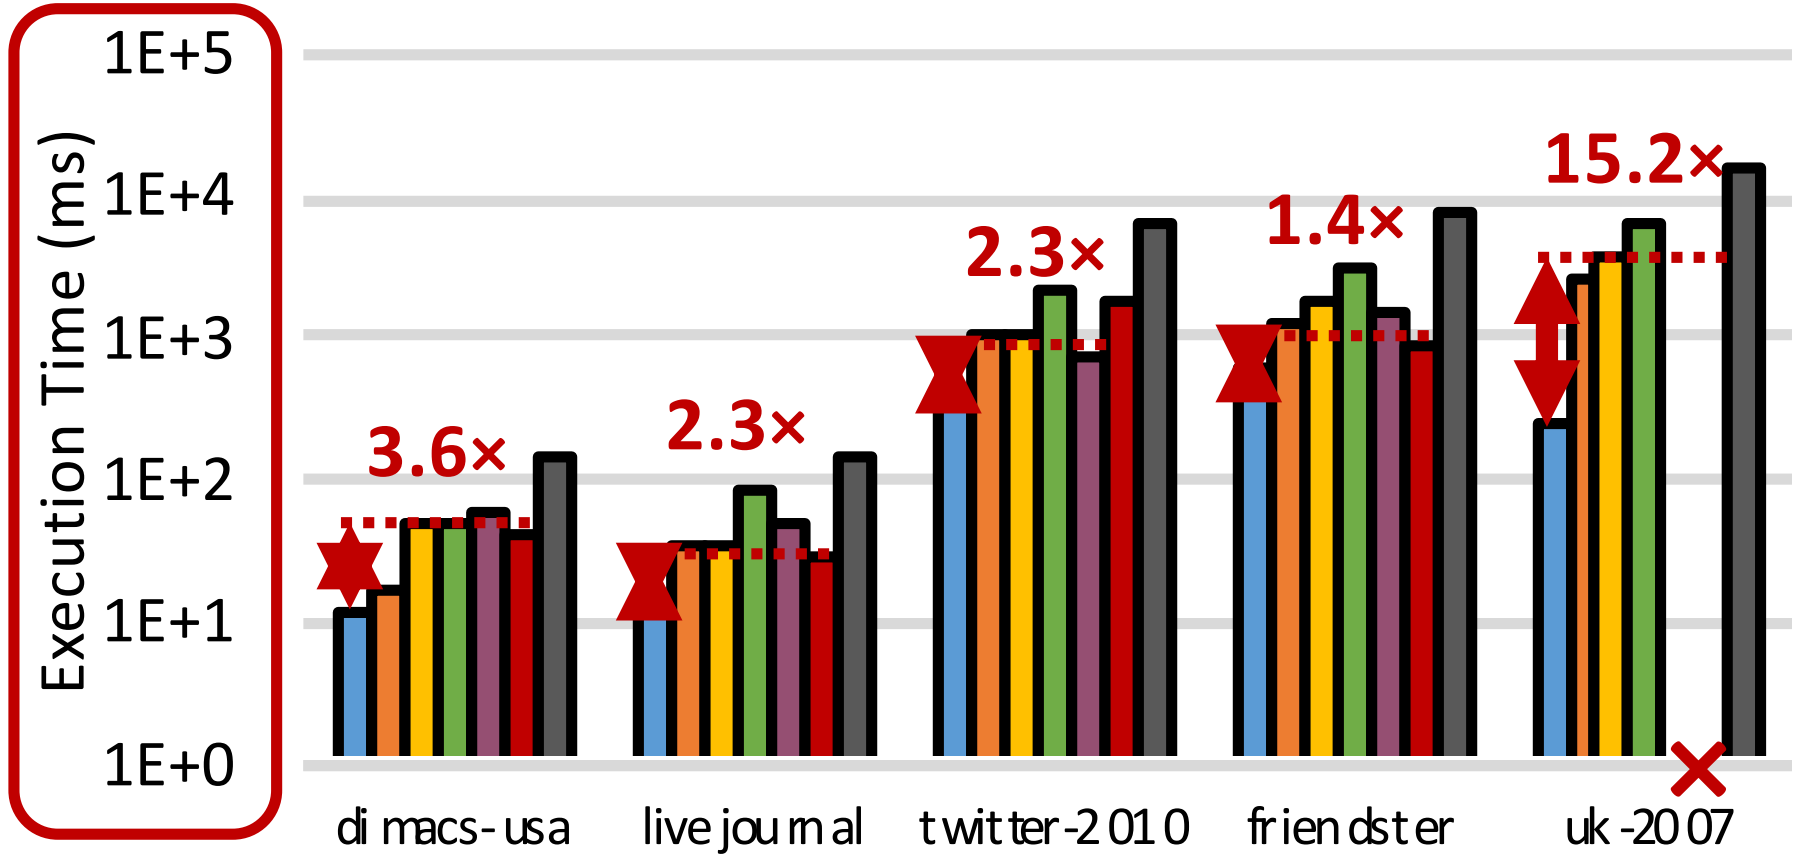
Putting it all together

Evaluation Scope

- Grazelle is compared with Ligra, Polymer, GraphMat, and X-Stream
- Three applications: PageRank, Connected Components, Breadth-First Search
- Running on a machine equipped with four Intel Xeon E7-4850 v3 processors
 - 14 physical cores / 28 logical cores per socket

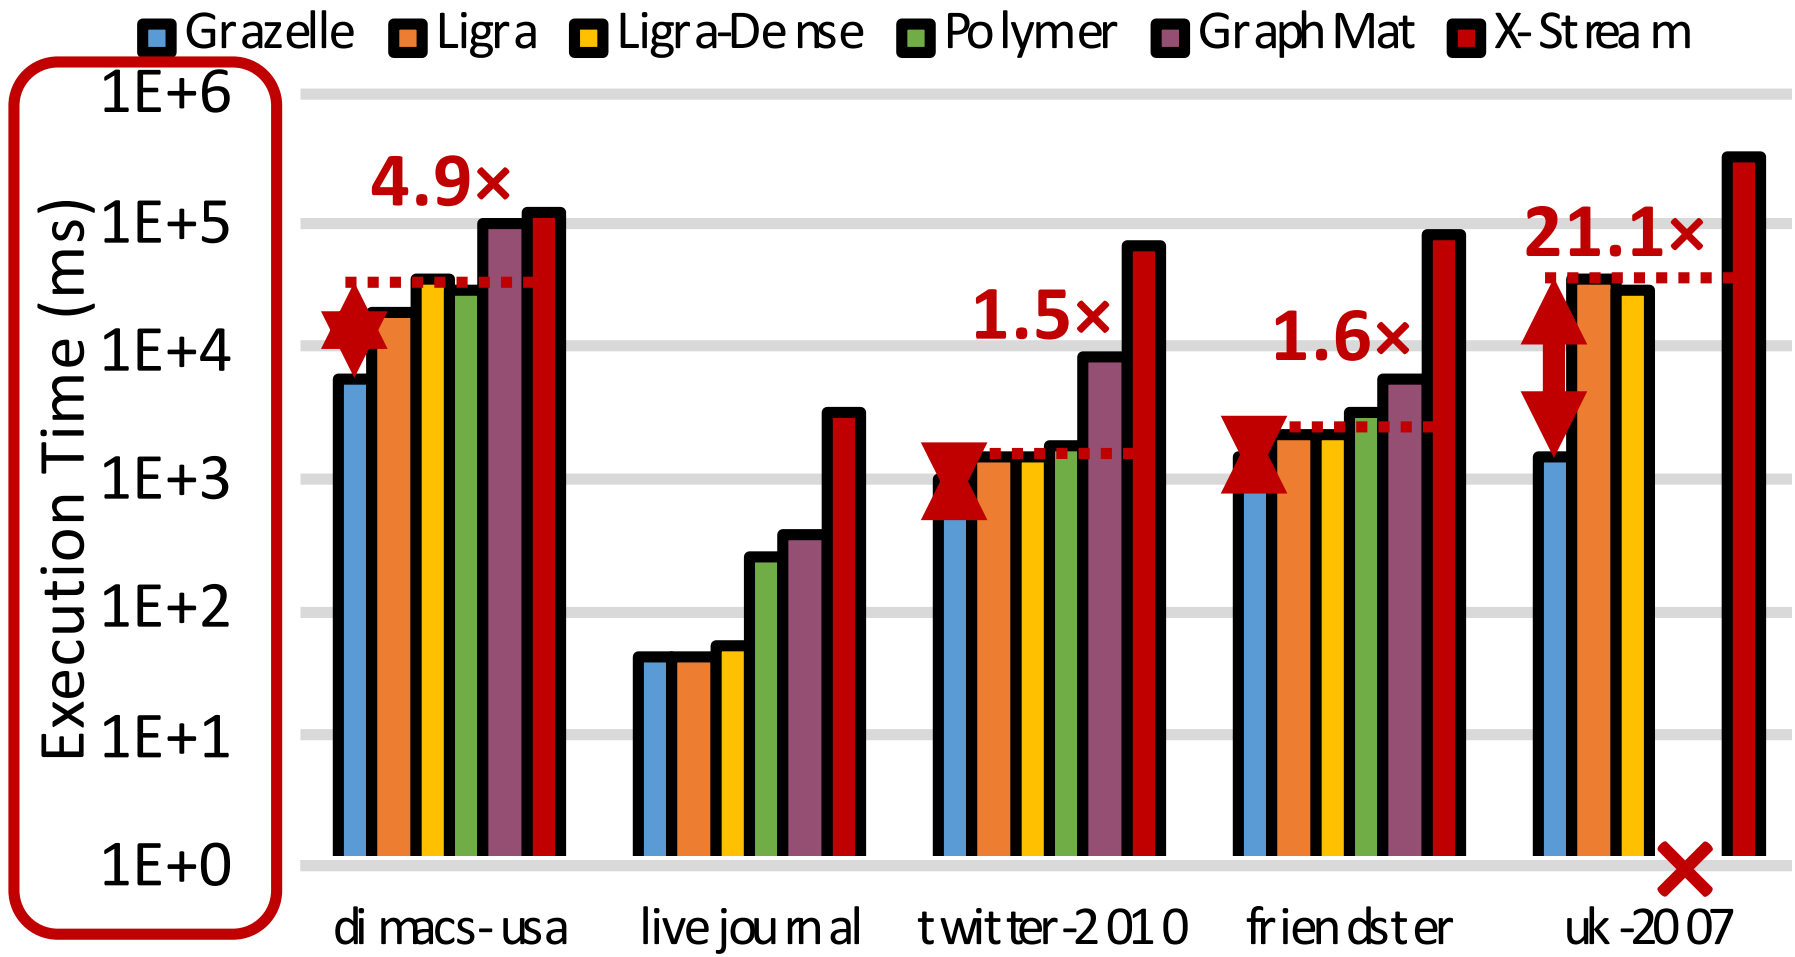
PageRank: Peak Processing Throughput

- Grazelle-Pull
- Grazelle-Push
- Ligma-Pull
- Ligma-Push
- Polymer
- Graph Mat
- X-Stream



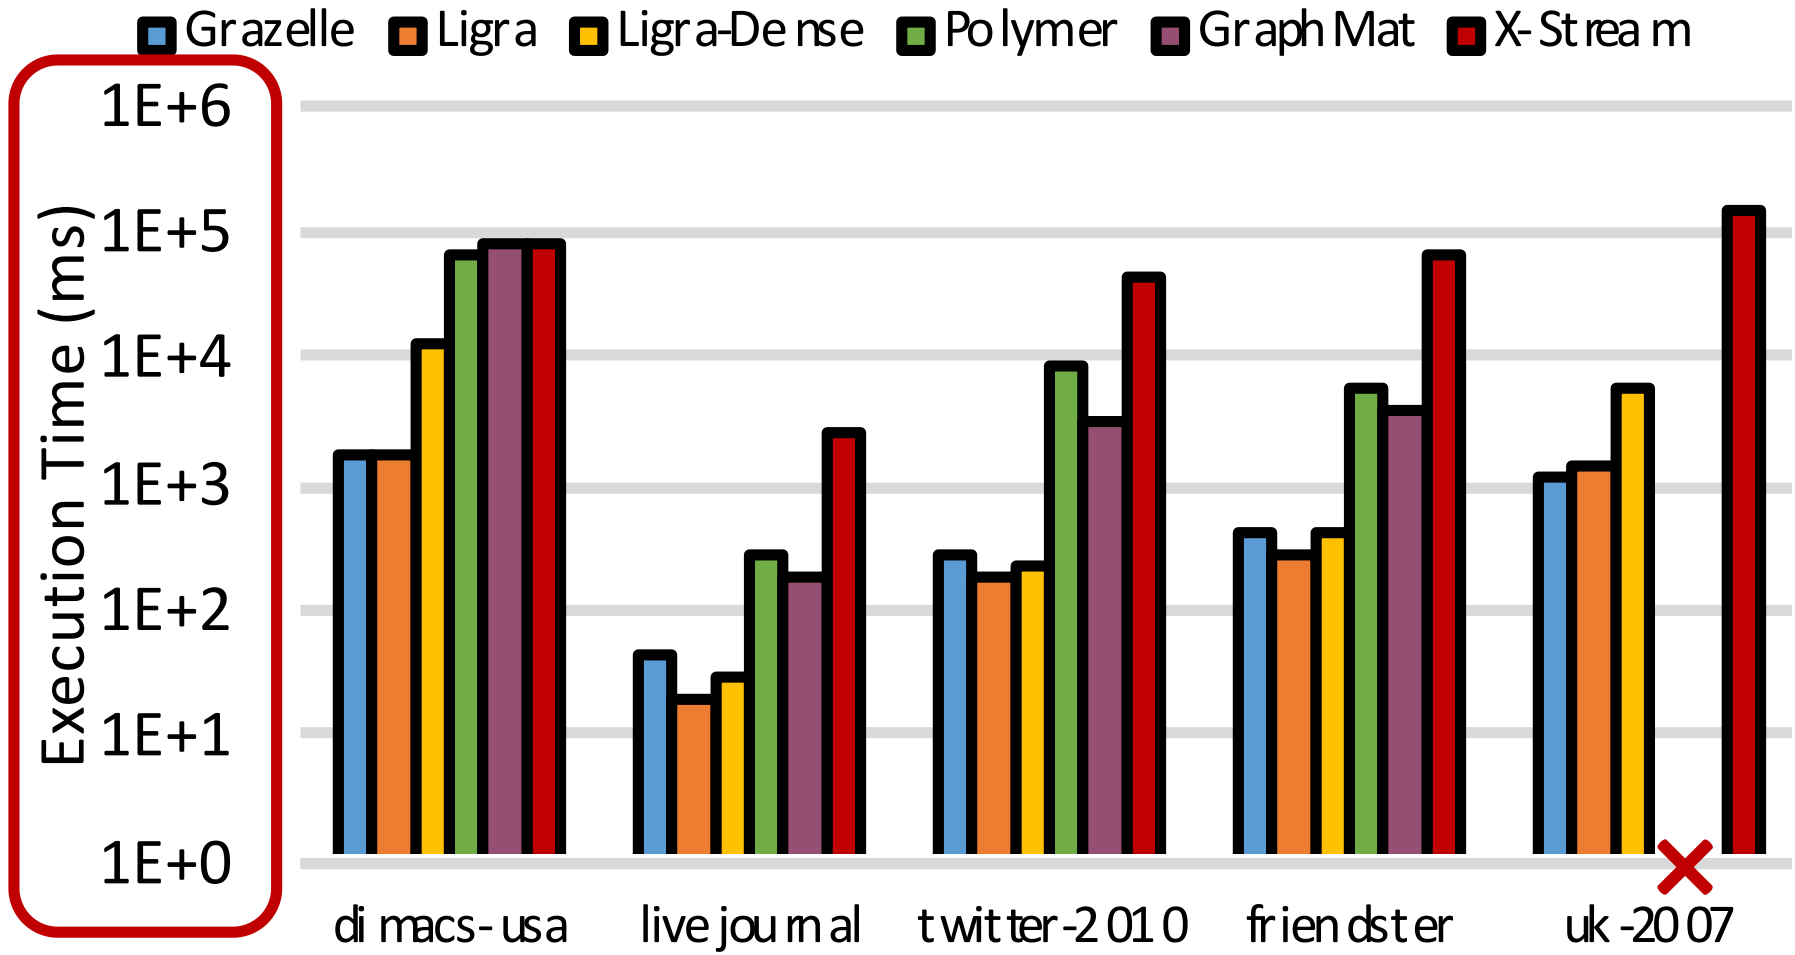
Logarithmic

Connected Components: Dynamic Control Flow



Logarithmic

Breadth-First Search: Compatibility of Optimizations



Logarithmic

Conclusion

- Two contributions to improve inner loop parallelization for pull-based graph processing
 - **Scheduler Awareness:** eliminate write conflicts
 - **Vector-Sparse:** enable SIMD vectorization
- Grazelle significantly out-performs state-of-the-art, in some cases by over 10×
- Grazelle is available for download at <https://github.com/stanford-mast/Grazelle-PPoPP18>

Thank You

Questions?