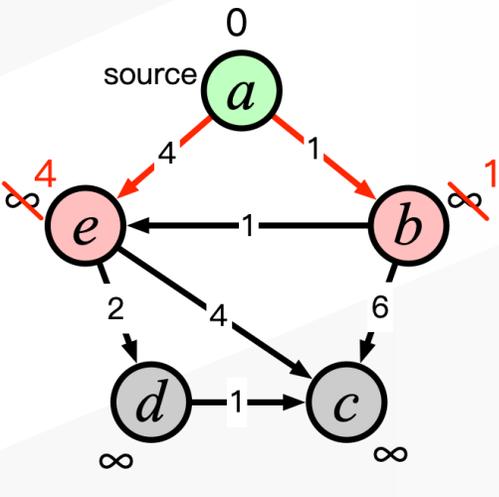


Fast Iterative Graph Computing with Updated Neighbor States

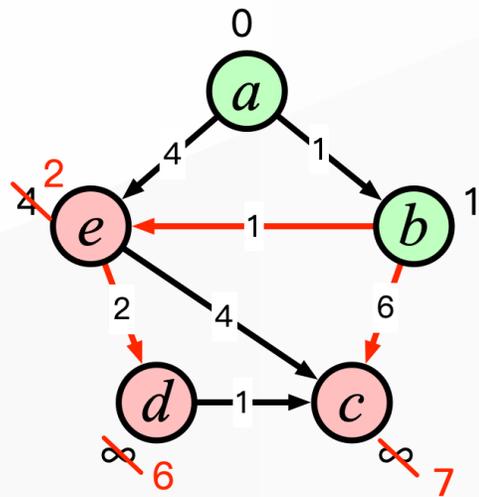
Yijie Zhou¹, **Shufeng Gong**¹, Feng Yao¹, Hanzhang Chen¹, Song Yu¹,
Pengxi Liu¹, Yanfeng Zhang¹, Ge Yu¹, Jeffrey Xu Yu²
Northeastern University¹, The Chinese University of Hong Kong²

Iterative Computation

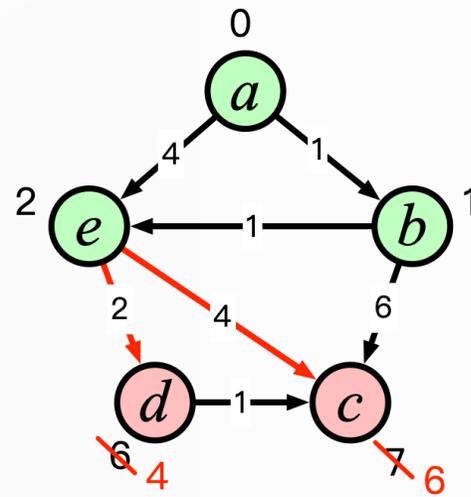
- PageRank, SSSP, BFS,
- Traversing the entire graph multiple times



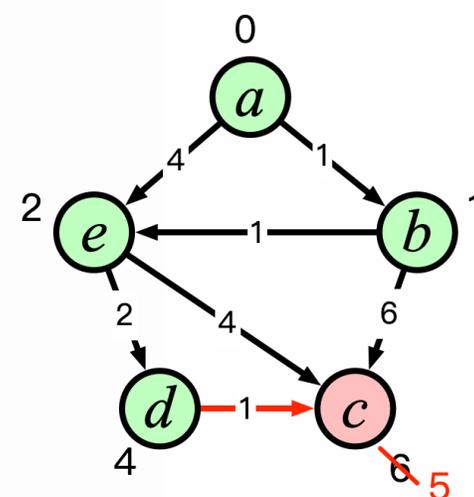
Iter. 1



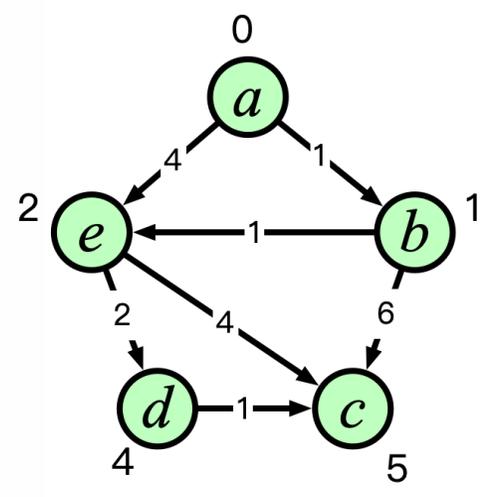
Iter. 2



Iter. 3



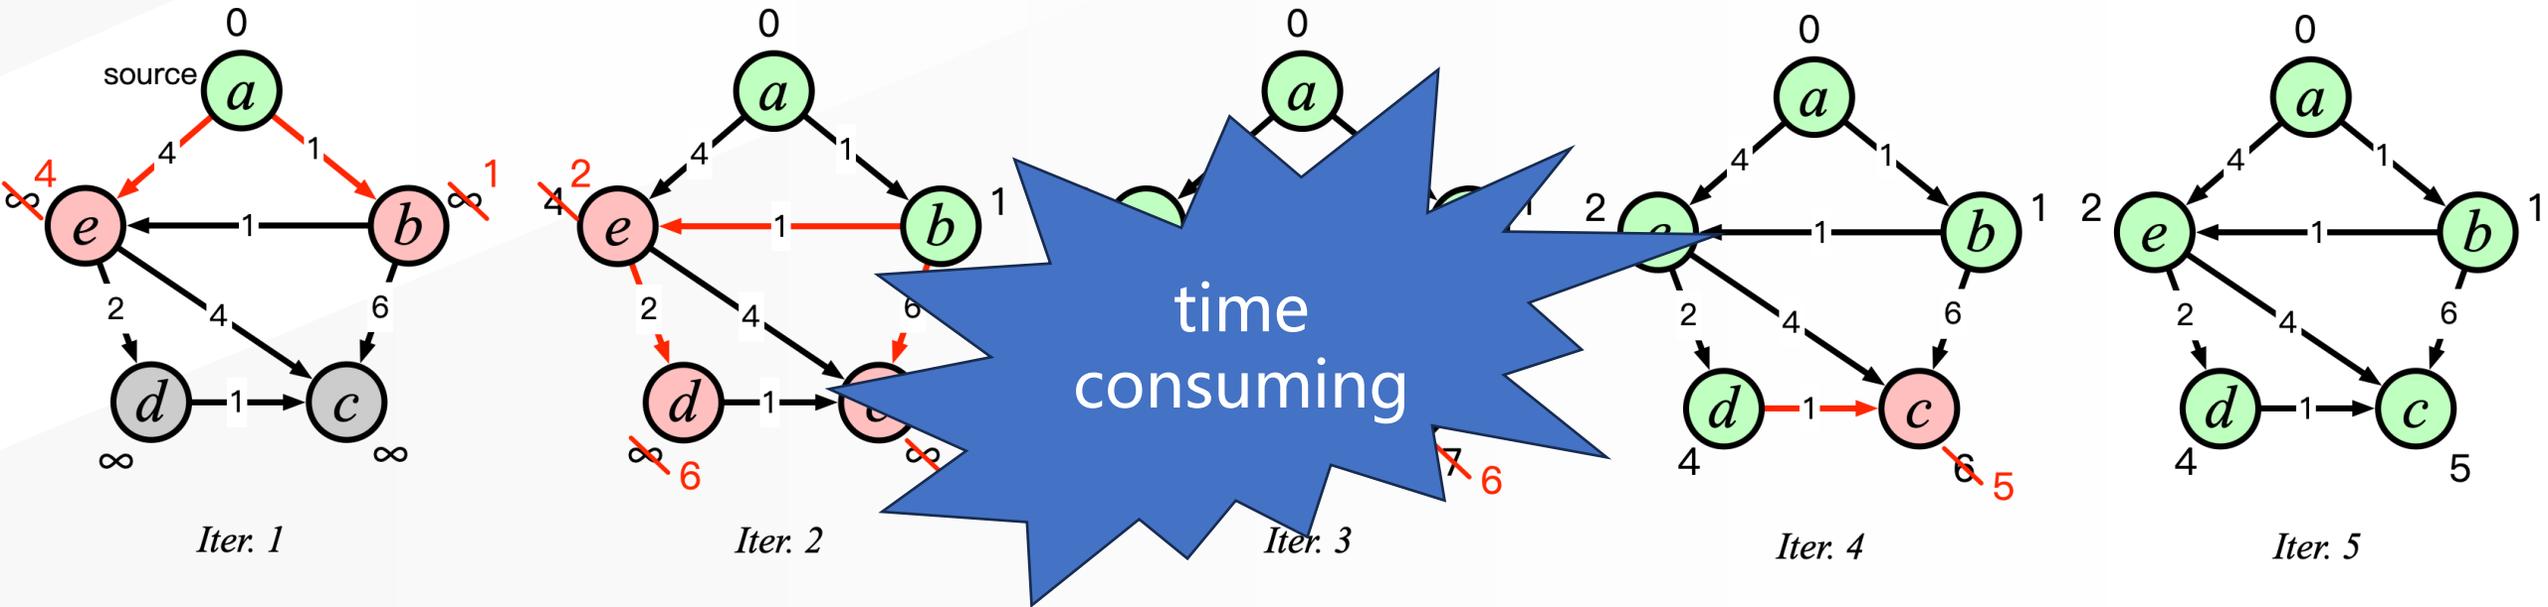
Iter. 4



Iter. 5

Iterative Computation

- PageRank, SSSP, BFS,
- Traversing the entire graph multiple times



Accelerate Iterative Computation

- How to accelerate iterative computation?
 - Reduce the time per iteration
 - Decrease the number of iterations

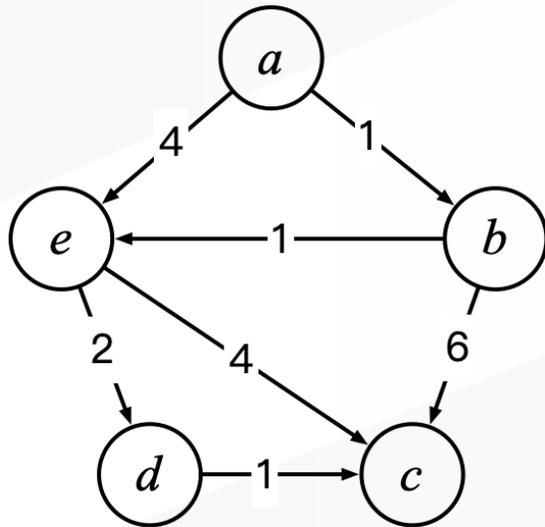
Vertex State Changes in Iterative Rounds

Synchronous



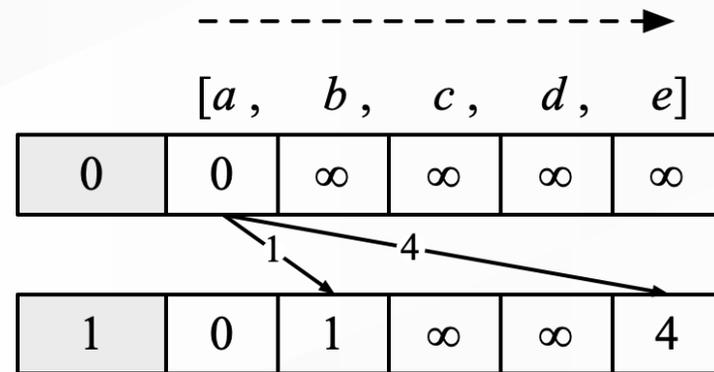
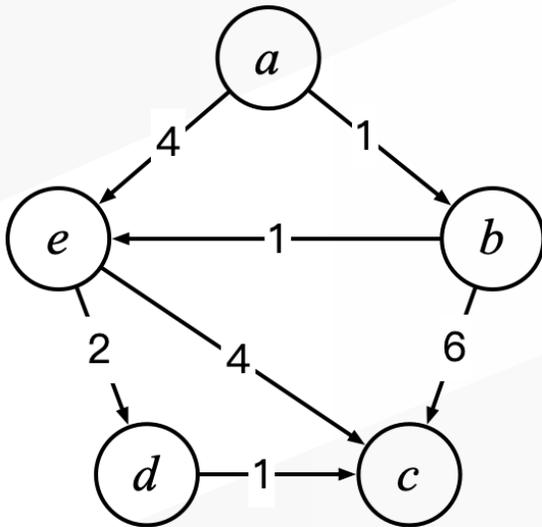
[*a*, *b*, *c*, *d*, *e*]

0	0	∞	∞	∞	∞
---	---	----------	----------	----------	----------



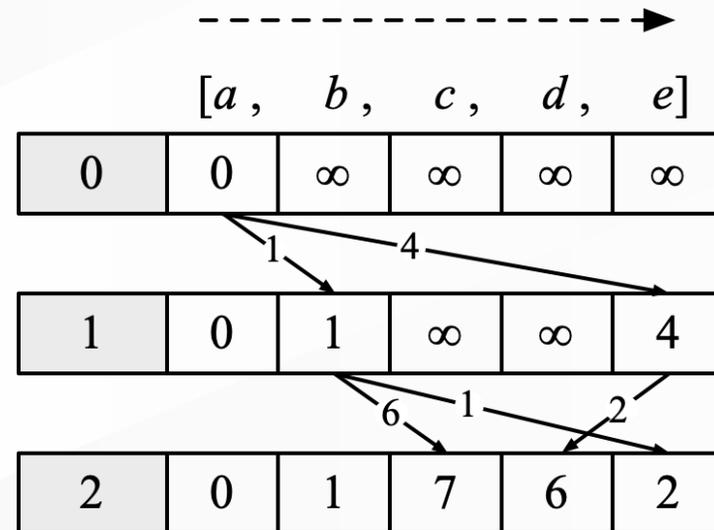
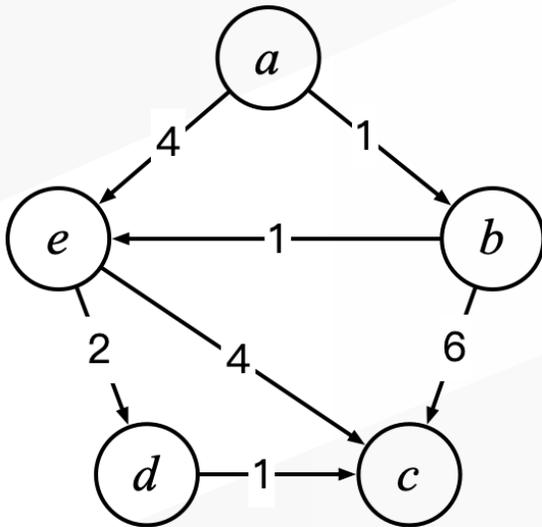
Vertex State Changes in Iterative Rounds

Synchronous



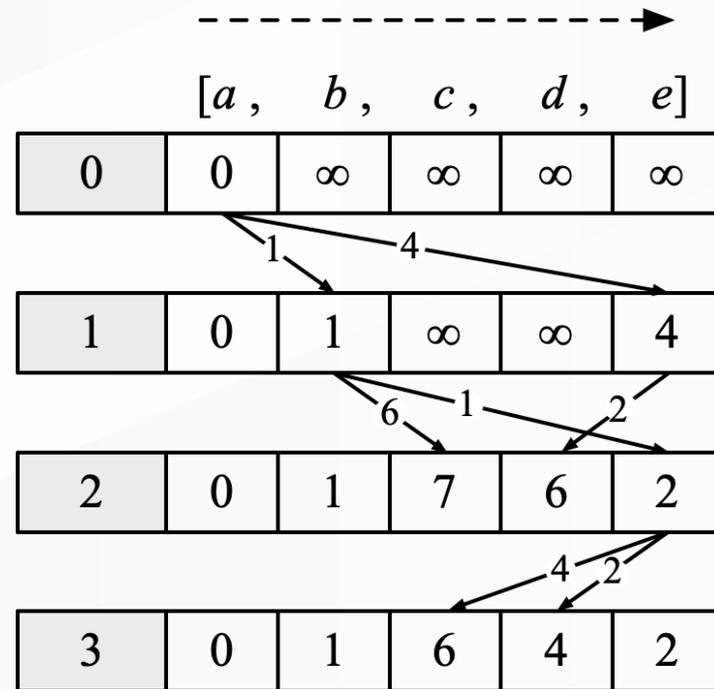
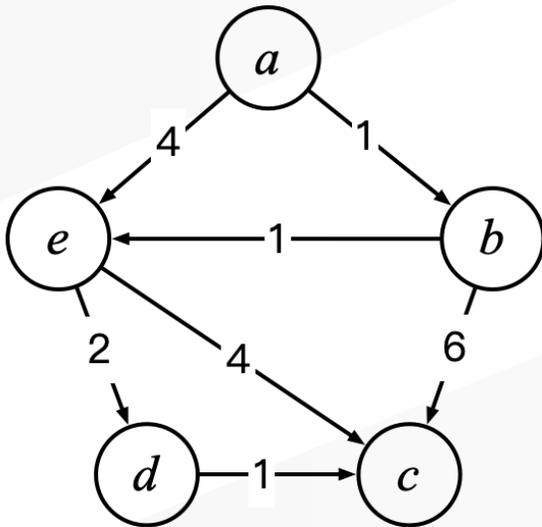
Vertex State Changes in Iterative Rounds

Synchronous

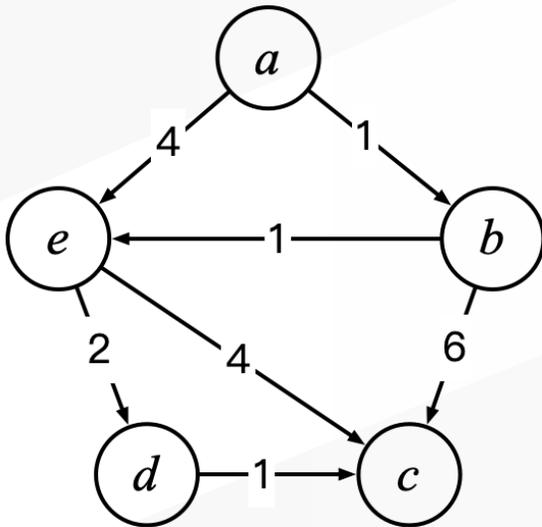


Vertex State Changes in Iterative Rounds

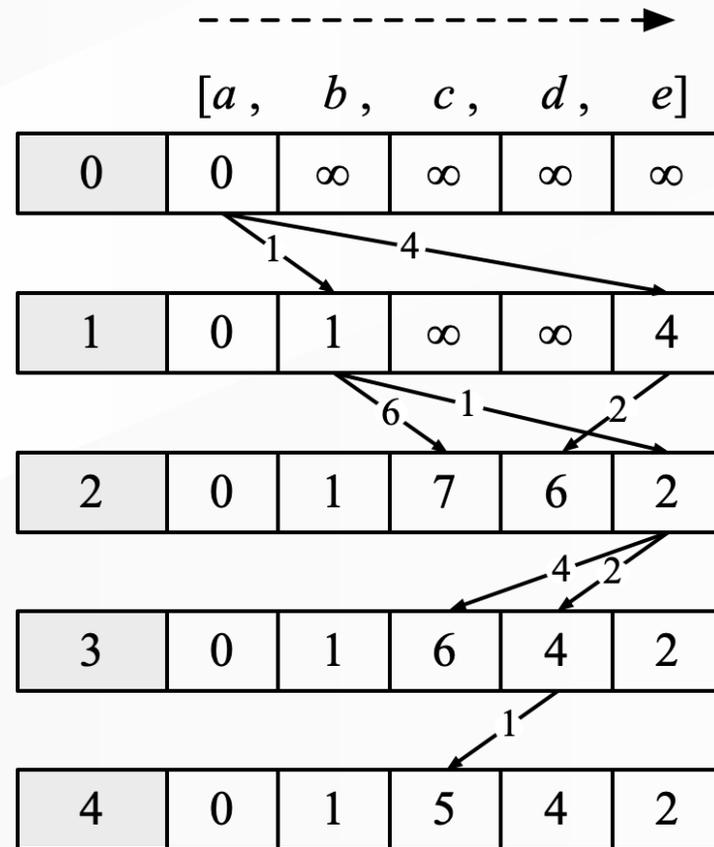
Synchronous



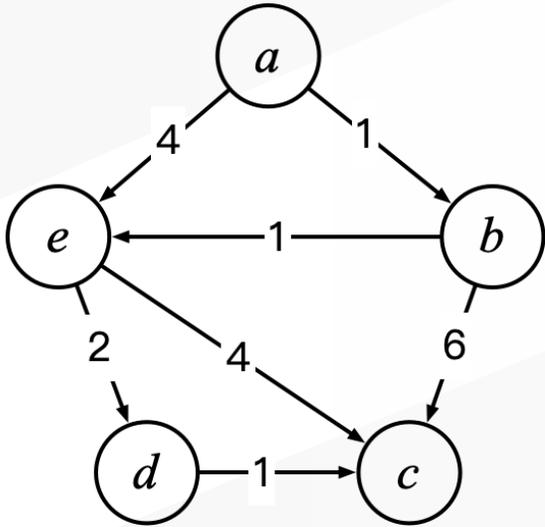
Vertex State Changes in Iterative Rounds



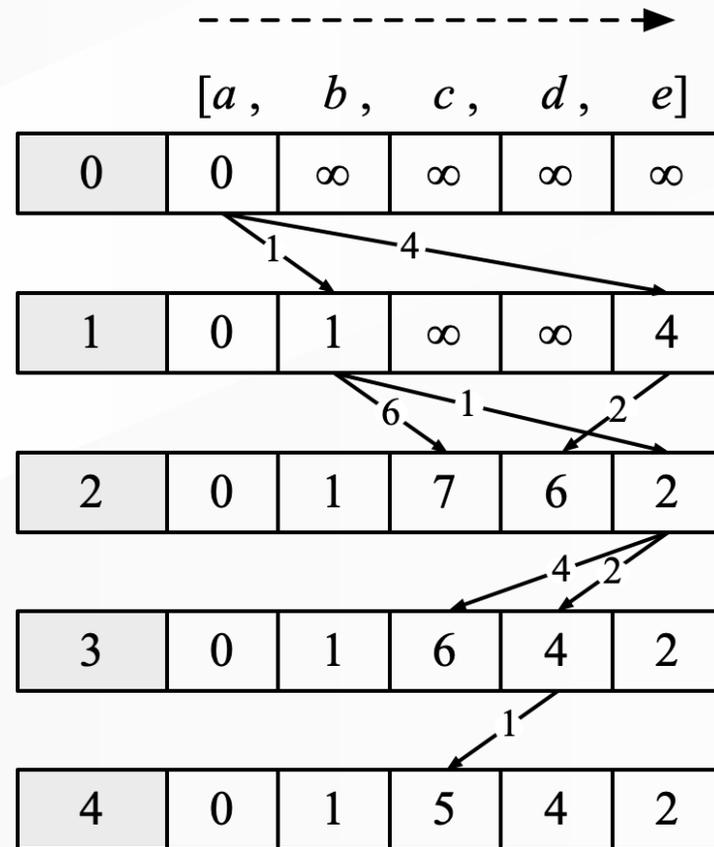
Synchronous



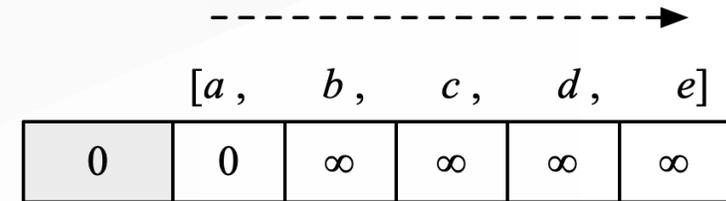
Vertex State Changes in Iterative Rounds



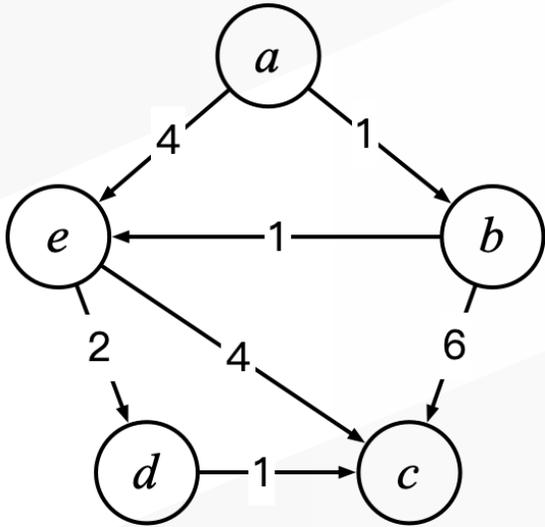
Synchronous



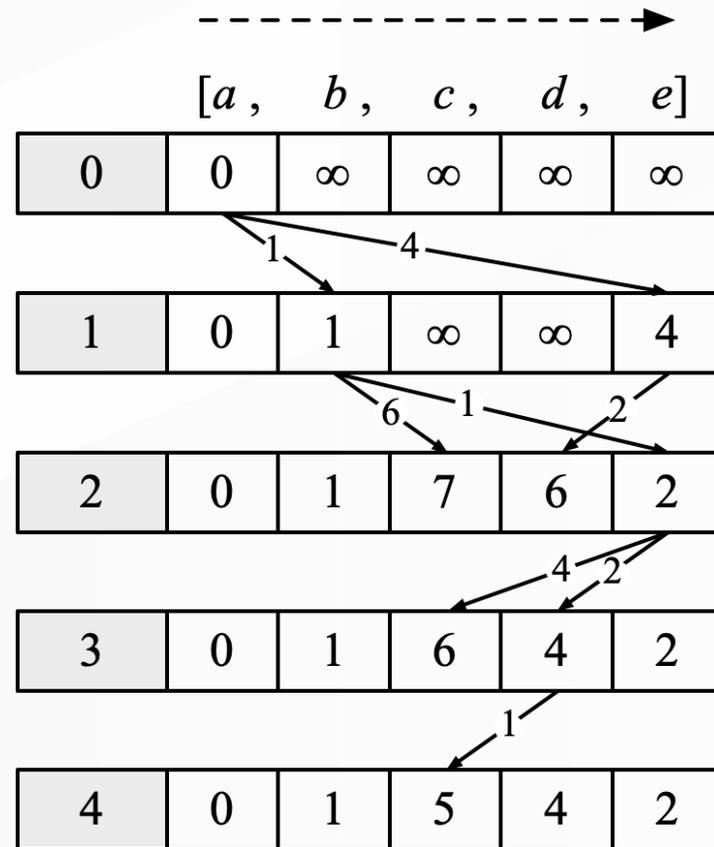
Asynchronous



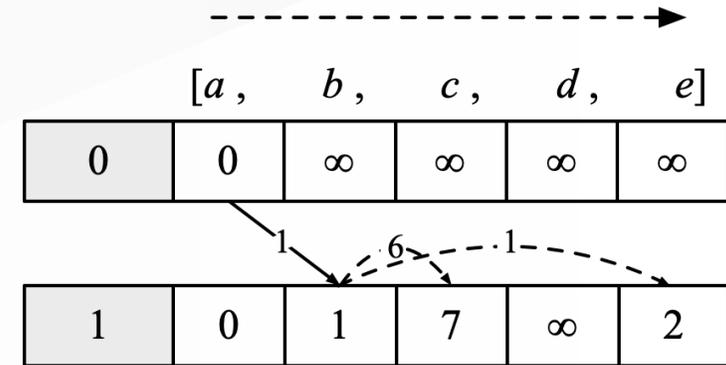
Vertex State Changes in Iterative Rounds



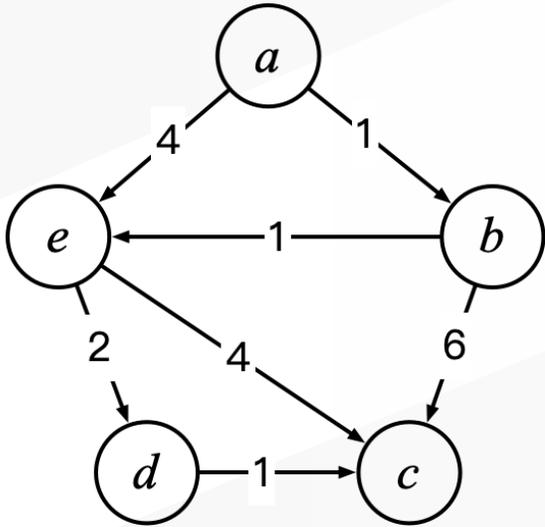
Synchronous



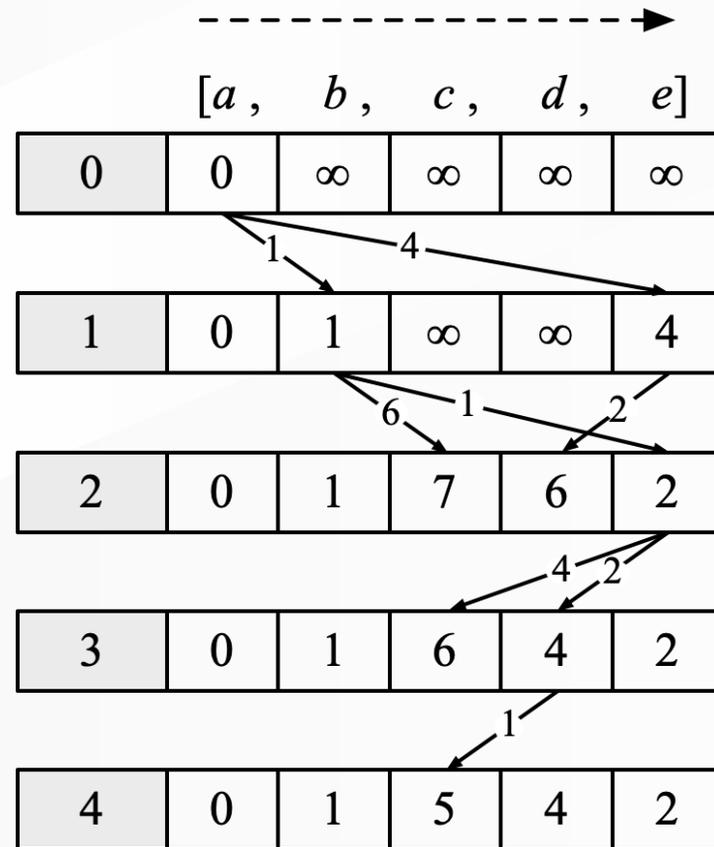
Asynchronous



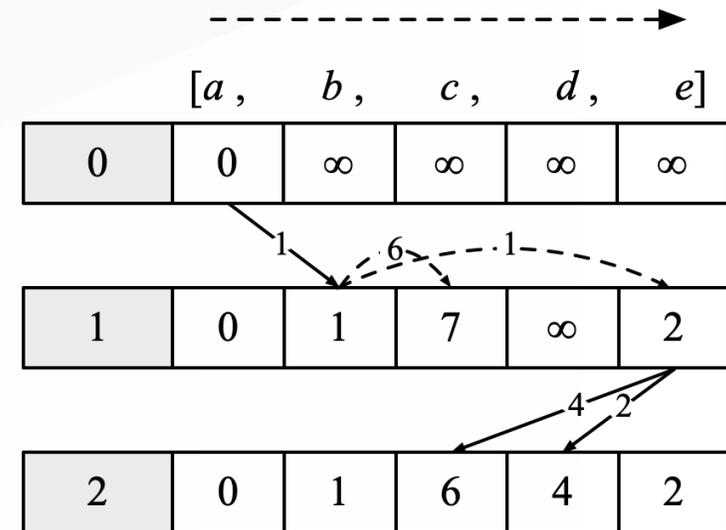
Vertex State Changes in Iterative Rounds



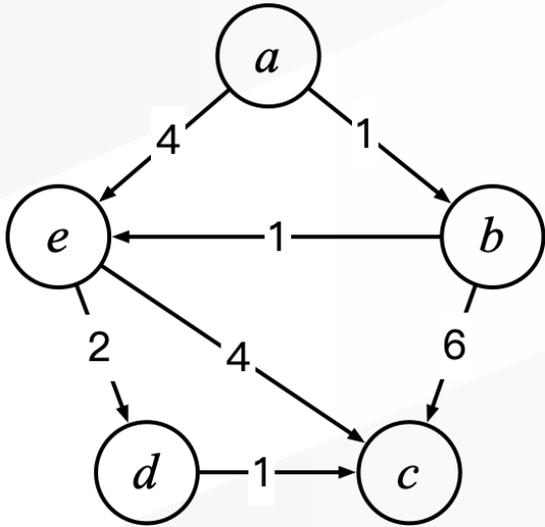
Synchronous



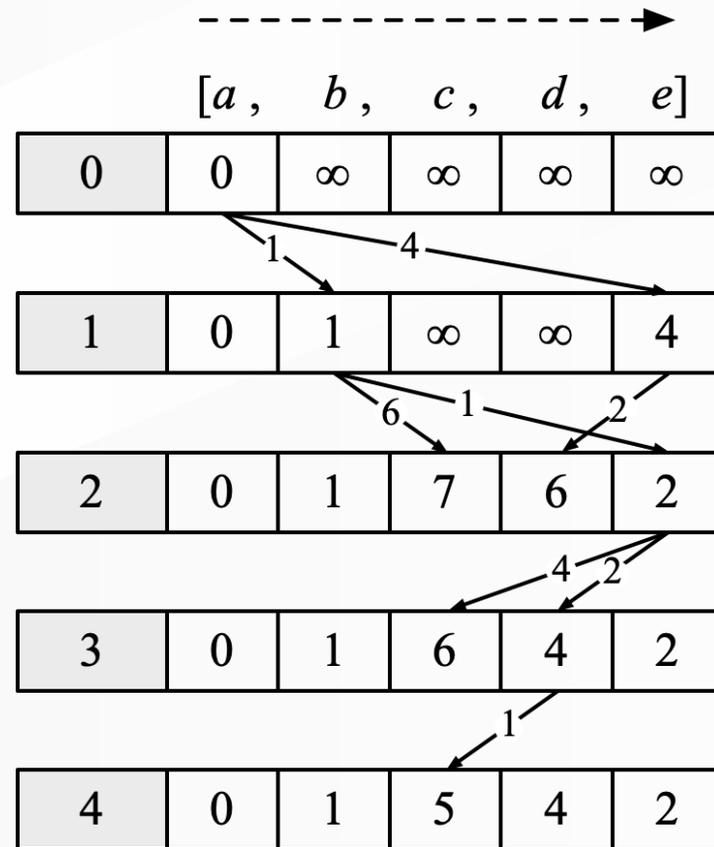
Asynchronous



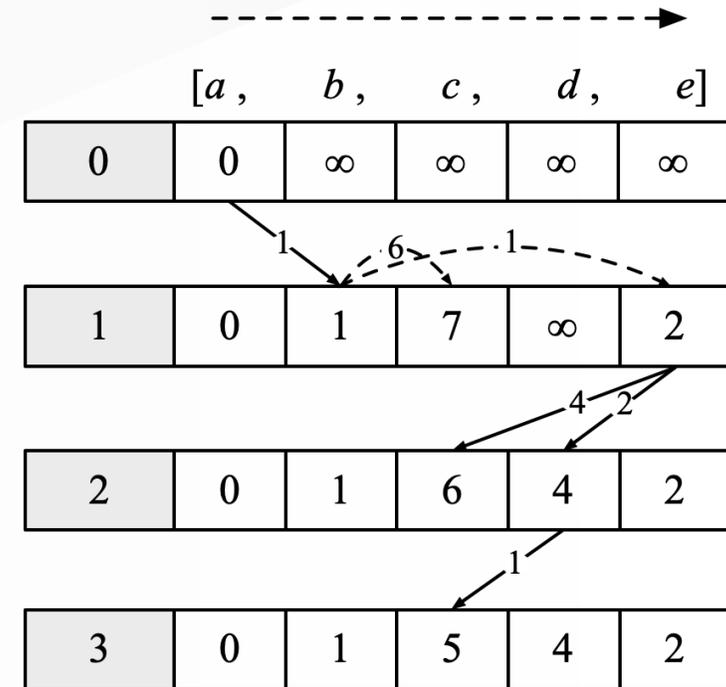
Vertex State Changes in Iterative Rounds



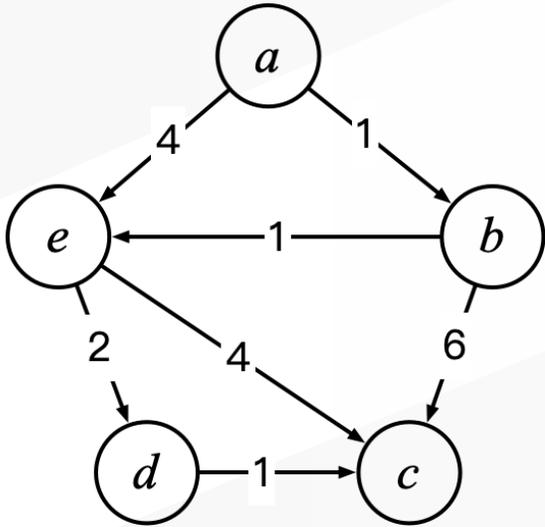
Synchronous



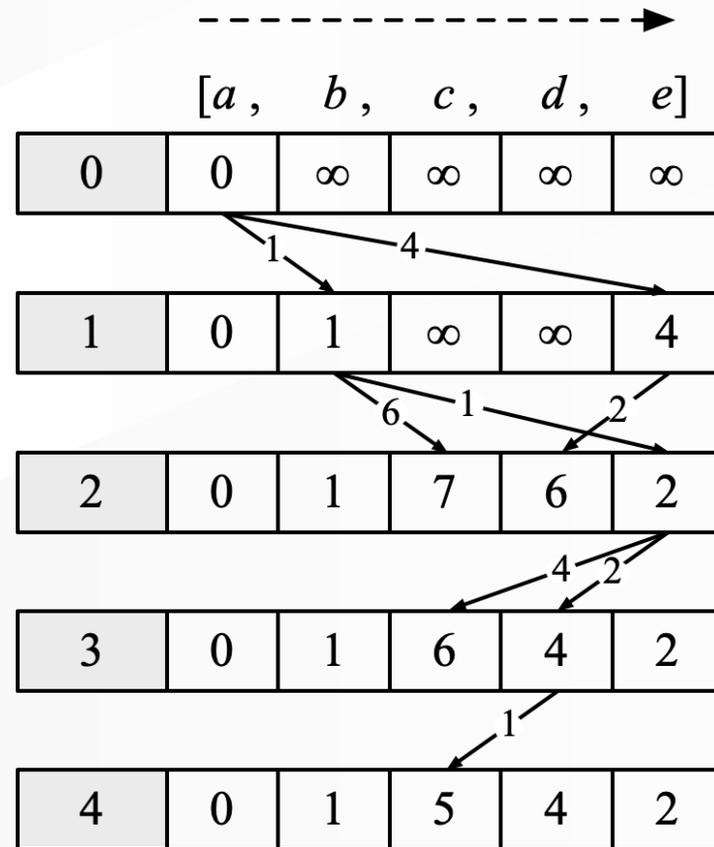
Asynchronous



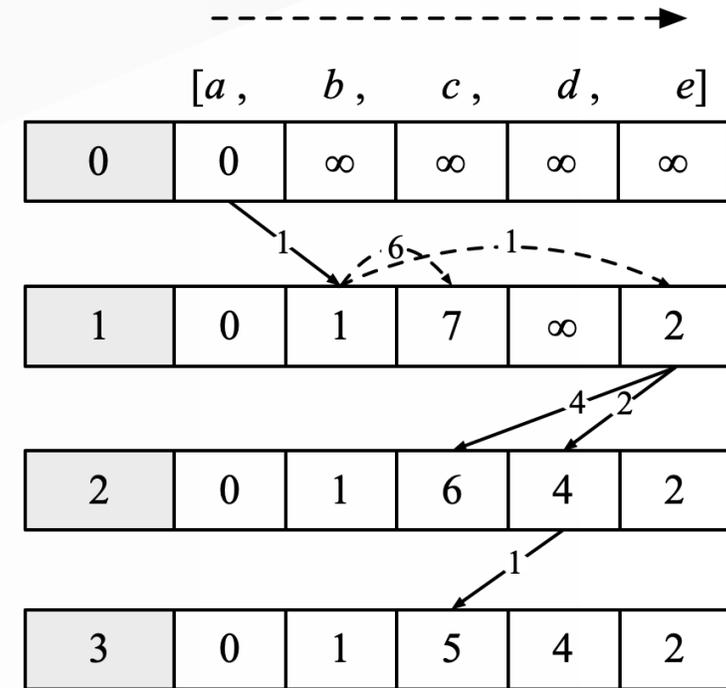
Vertex State Changes in Iterative Rounds



Synchronous

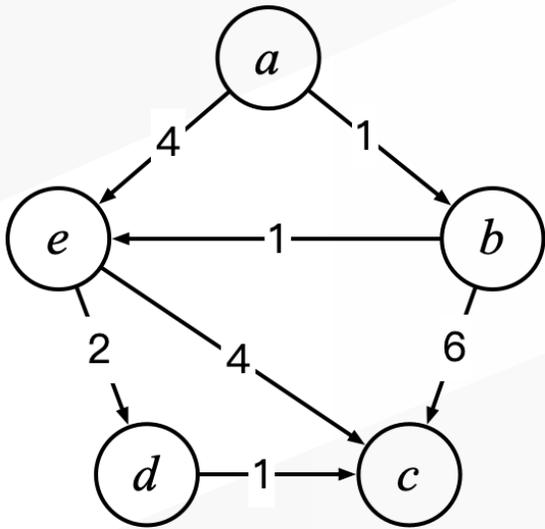


Asynchronous



Asynchronous iteration reduces the number of rounds because each vertex can immediately use the latest state value.

Vertex State Changes in Iterative Rounds



Asynchronous

----->

[a, b, c, d, e]

0	0	∞	∞	∞	∞
1	0	1	7	∞	2
2	0	1	6	4	2
3	0	1	5	4	2

Diagram illustrating the asynchronous update process. The table shows the state of vertices a, b, c, d, and e over four iterations. The first column (shaded) indicates the iteration number. The second column (shaded) is the source vertex's state. The third column (shaded) is the state of vertex 'a'. The fourth column (shaded) is the state of vertex 'b'. The fifth column (shaded) is the state of vertex 'c'. The sixth column (shaded) is the state of vertex 'd'. The seventh column (shaded) is the state of vertex 'e'. Arrows indicate updates: from iteration 0 to 1, 'a' updates 'b' (1) and 'e' (4). From iteration 1 to 2, 'b' updates 'c' (6) and 'e' (2). From iteration 2 to 3, 'e' updates 'd' (2) and 'c' (4). Dashed arrows show that 'c' and 'd' do not update 'a' because their current values (1 and 2) are greater than 'a's value (0).

Asynchronous
with reordered order

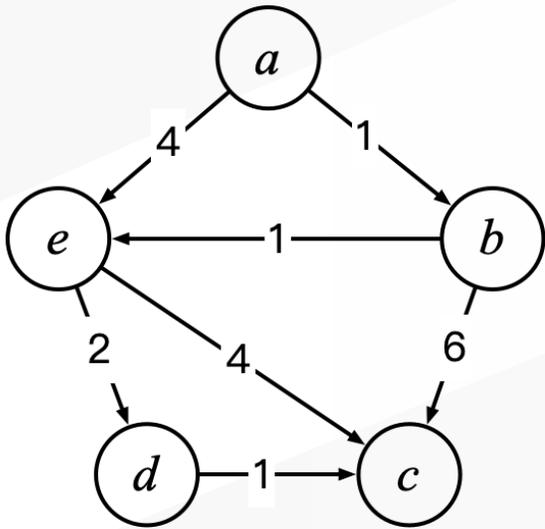
----->

[a, b, e, d, c]

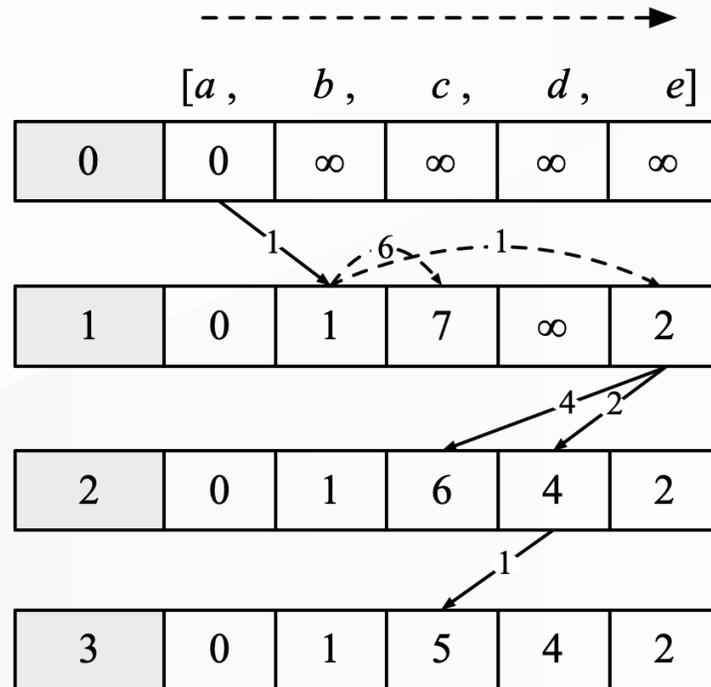
0	0	∞	∞	∞	∞
---	---	----------	----------	----------	----------

Diagram illustrating the asynchronous update process with a reordered order. The table shows the state of vertices a, b, e, d, and c over four iterations. The first column (shaded) indicates the iteration number. The second column (shaded) is the source vertex's state. The third column (shaded) is the state of vertex 'a'. The fourth column (shaded) is the state of vertex 'b'. The fifth column (shaded) is the state of vertex 'e'. The sixth column (shaded) is the state of vertex 'd'. The seventh column (shaded) is the state of vertex 'c'. The eighth column (shaded) is the state of vertex 'c' (repeated). Arrows indicate updates: from iteration 0 to 1, 'a' updates 'b' (1) and 'e' (4). From iteration 1 to 2, 'b' updates 'c' (6) and 'e' (2). From iteration 2 to 3, 'e' updates 'd' (2) and 'c' (4). Dashed arrows show that 'c' and 'd' do not update 'a' because their current values (1 and 2) are greater than 'a's value (0).

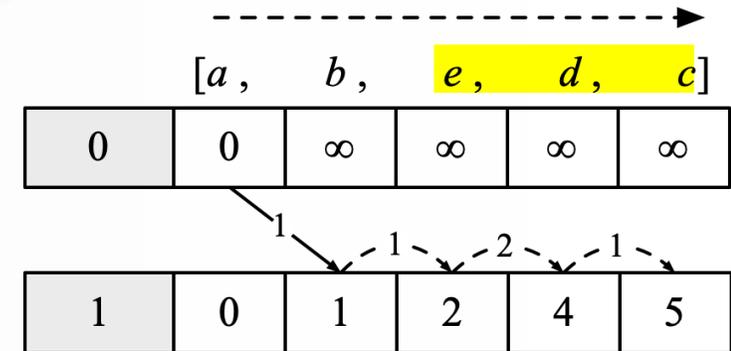
Vertex State Changes in Iterative Rounds



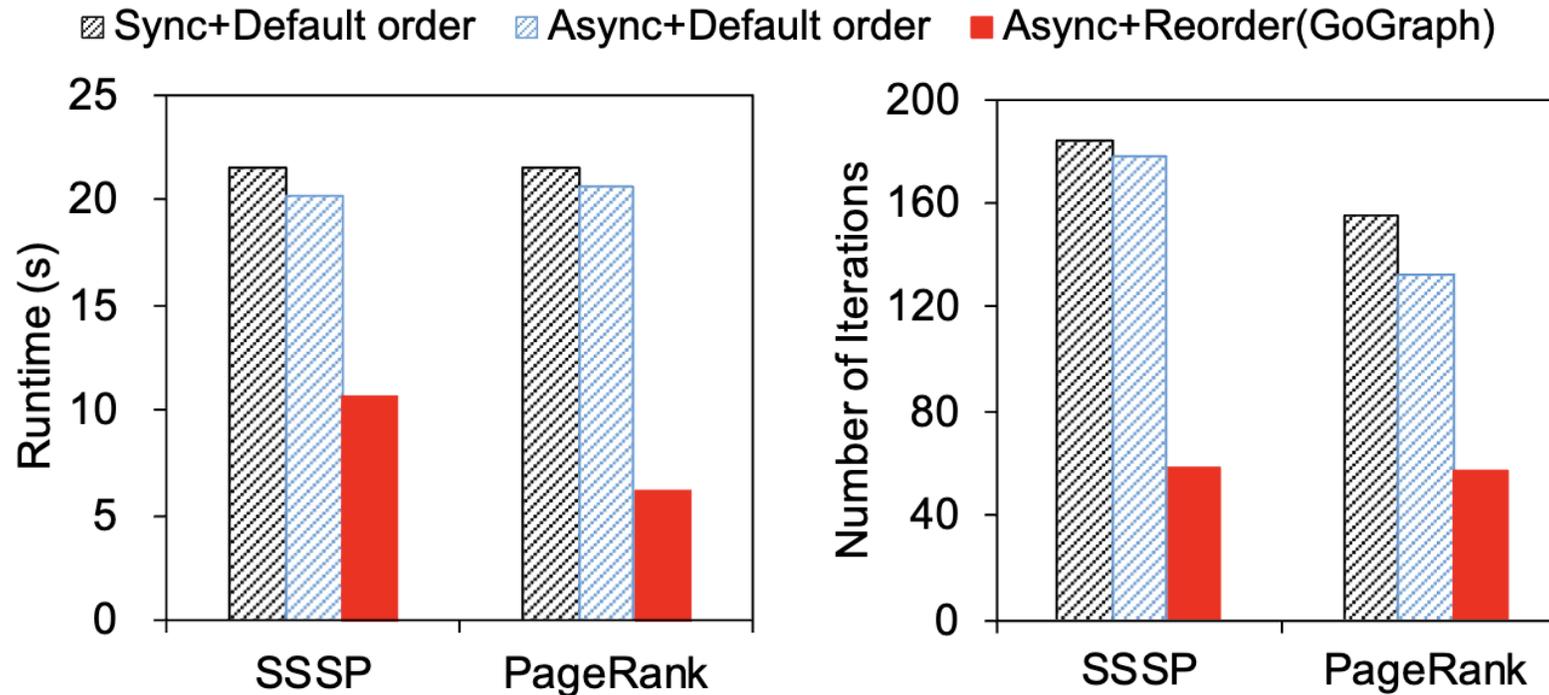
Asynchronous



Asynchronous
with reordered order



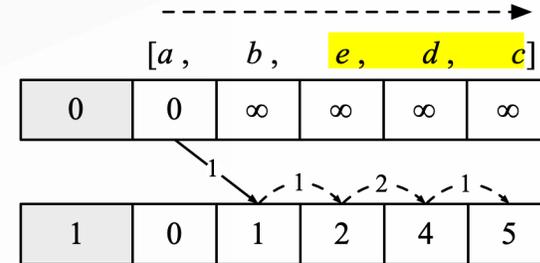
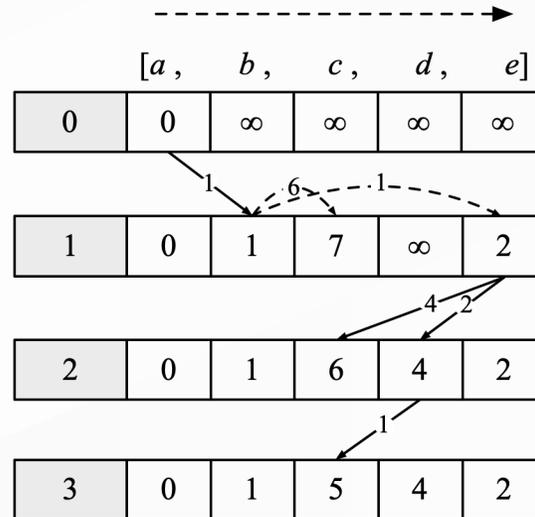
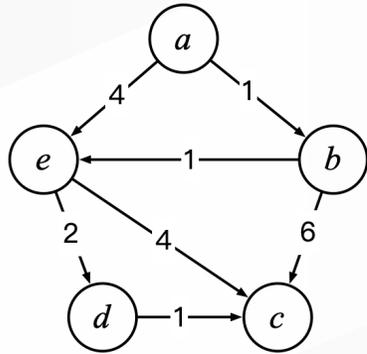
Converging Quicker After Reordering



(a) Runtime

(b) Iteration rounds

Goal

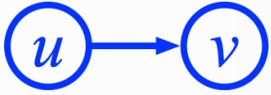


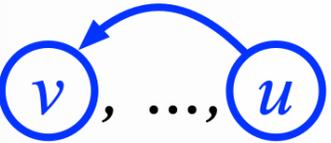
- Construct an efficient vertex processing order to accelerate the iterative computation.

Challenges

- Which processing order is better?
 - Challenge 1. Design a metric to measure the quality of the processing order.
- How to reorder the vertex to make the iterations converge faster?
 - Challenge 2. Design a vertex reordering method.

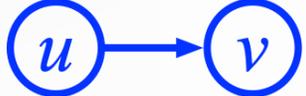
Positive/Negative Edge

For an existing edge: 

- Processing order 1: $O_V = [a, b, \dots, u, \dots, v, \dots, z]$
 - $\langle u, v \rangle$ is a **positive edge**, since v can use u 's latest state in the same round.
- Processing order 2: $O_V = [a, b, \dots, v, \dots, u, \dots, z]$
 - $\langle u, v \rangle$ is a **negative edge**, since v can only use u 's latest state in the next round.

Metric Function

Intuition: as mentioned before, the more positive edges, the more vertex state values can be utilized per round, speeding up convergence.

- Counts the number of positive edges ( , u is processed prior to v)

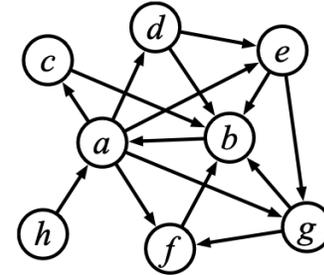
$$M(O_V) = n(e_{pos}) = \sum_{(u,v) \in E} \chi(u,v)$$

$$\chi(u,v) = \begin{cases} 1, & \text{if } (u,v) \text{ is positive,} \\ 0, & \text{if } (u,v) \text{ is negative.} \end{cases}$$

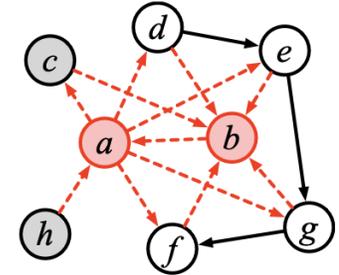
- The goal of our reordering method: maximize the quantity of positive edges

GoGraph

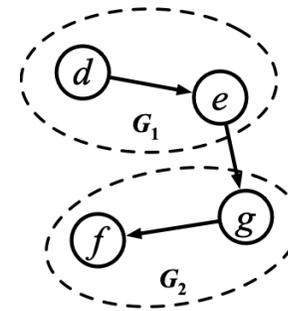
- 1) Extract high-degree and isolated vertices
- 2) Divide other vertices
- 3) Reorder vertices within subgraphs
- 4) Reorder subgraphs
- 5) Insert high-degree and isolated vertices



(a) The initial graph.



(b) Extract high-degree vertices.



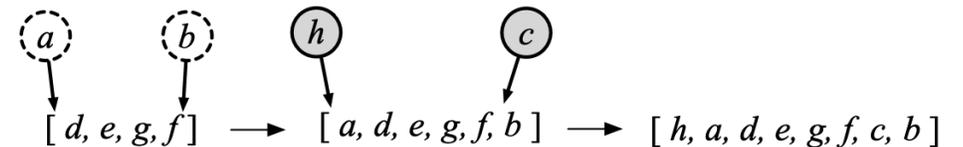
(c) Divide the remaining vertices

$$O_{V_1} = [d, e] \quad O_{V_2} = [g, f]$$

(d) Reordering vertices intra-subgraphs.

$$[[d, e], [g, f]] \rightarrow O_V = [d, e, g, f]$$

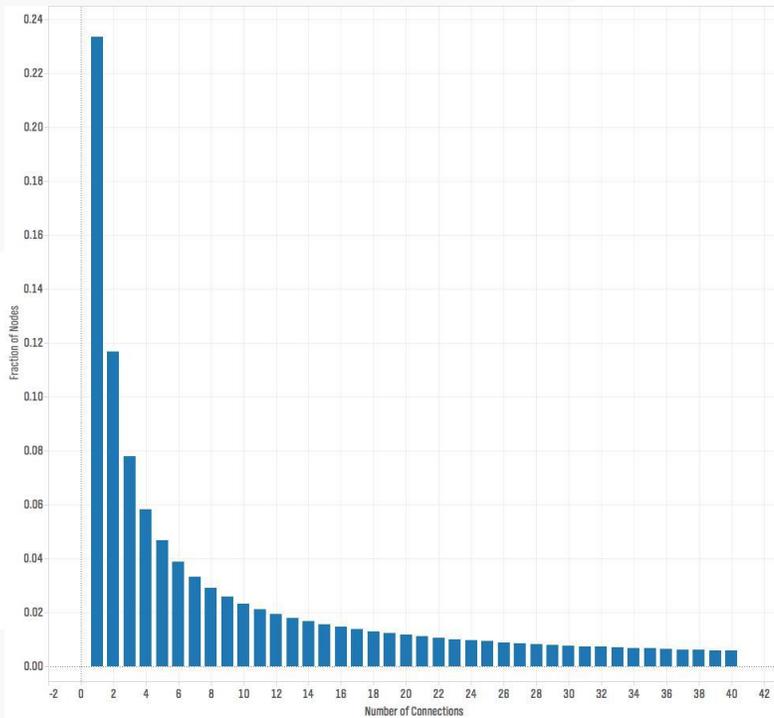
(e) Reordering vertices inter-subgraphs.



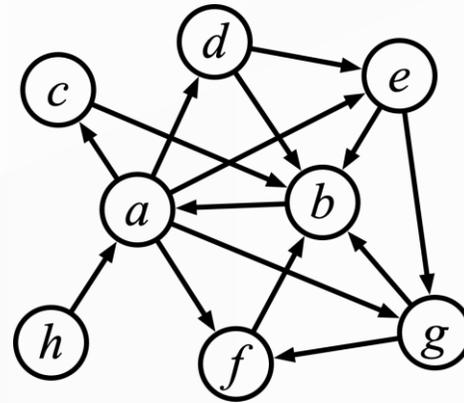
(f) Inserting high-degree/isolated vertices into the processing order.

GoGraph

1) Extract high-degree and isolated vertices

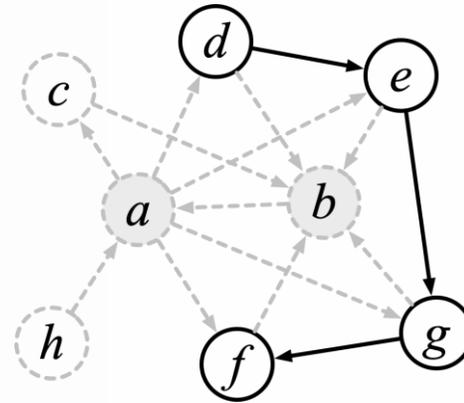


Distribution of vertices of different degrees for scale-free graphs



$$O_V^1 = [d, e, c, b, h, a, g, f]$$

$$M(O_V^1) = 10$$



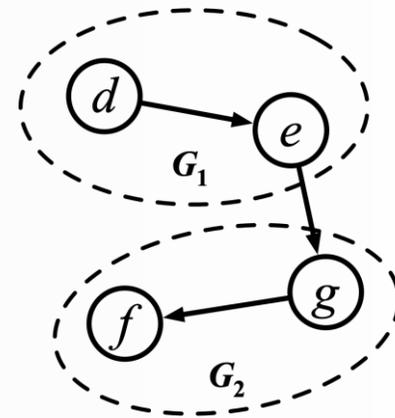
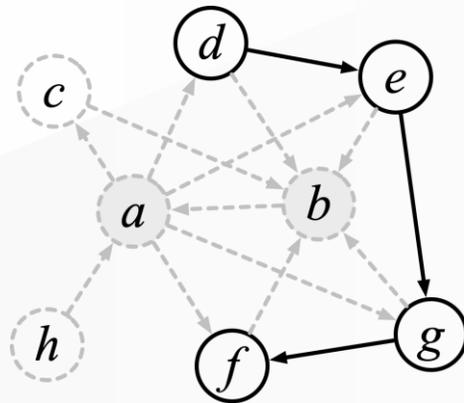
$$O_V^2 = [h, a, c, d, e, g, f, b]$$

$$M(O_V^2) = 14$$

GoGraph

2) Divide other vertices

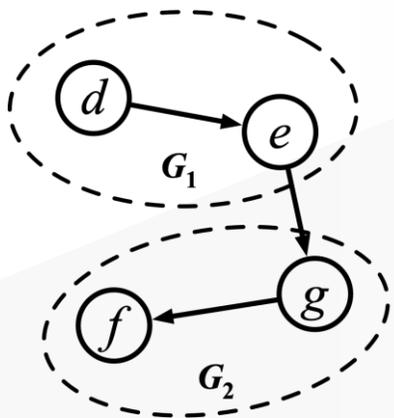
- Louvain, Metis, ...



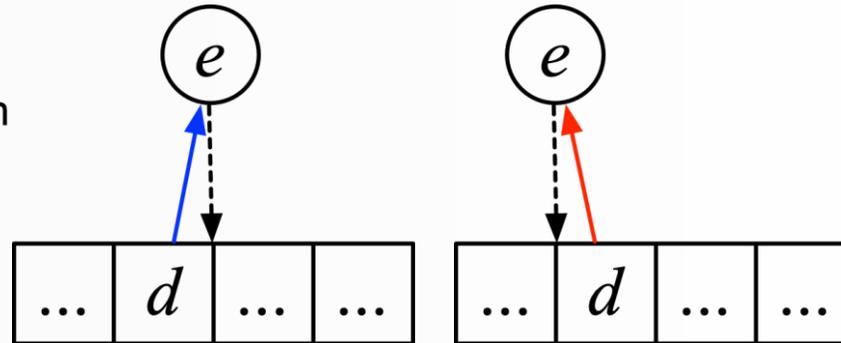
GoGraph

3) Reorder vertices within subgraphs

- Calculate the M value based on where the vertices are inserted;
- Find the position that maximizes M (maximize the number of positive edges and minimize the number of negative edges).



- > insert to a position
- > positive edge
- > negative edge



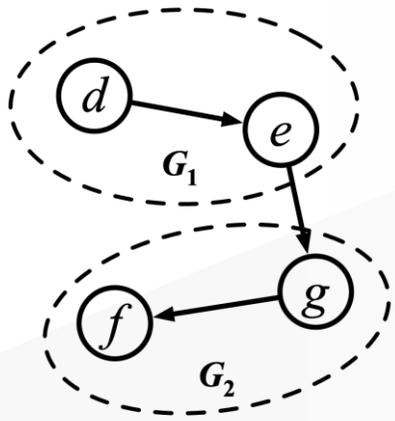
$$O_{G_1} = [d, e]$$

$$O_{G_2} = [g, f]$$

GoGraph

4) Reorder subgraphs

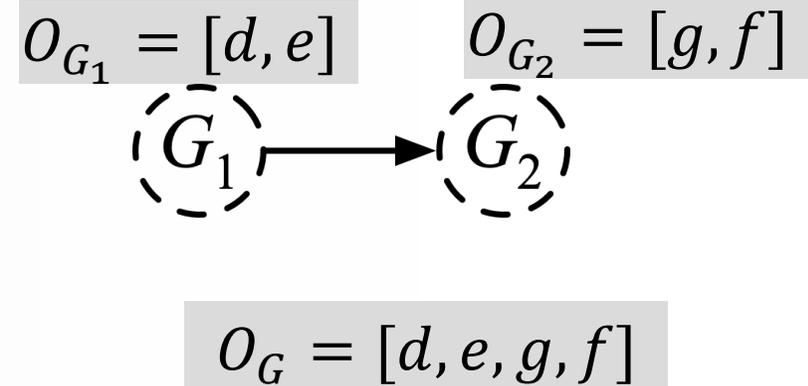
- Consider each subgraph as a super vertex;
- Perform the same operation in step 3.



$$w_{G_1, G_2} = |\{\{u, v\} | u \in G_i, v \in G_j\}|$$

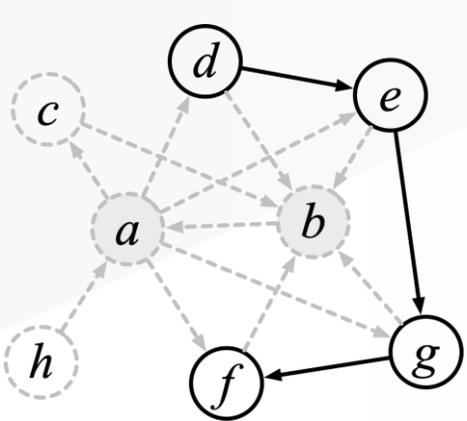
$$M(O_G) = \sum_{(G_i, G_j) \in P} \chi(G_i, G_j)$$

$$\chi(G_i, G_j) = \begin{cases} w_{G_1, G_2}, & \text{if } e_{(G_i, G_j)} \text{ is positive,} \\ 0, & \text{if } e_{(G_i, G_j)} \text{ is negative.} \end{cases}$$

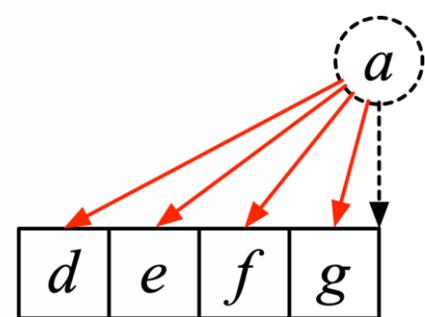
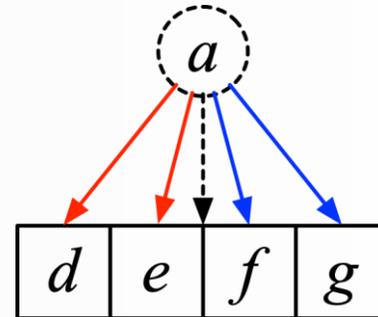
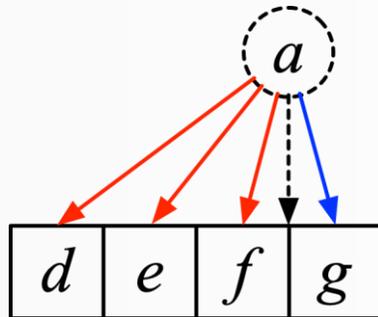
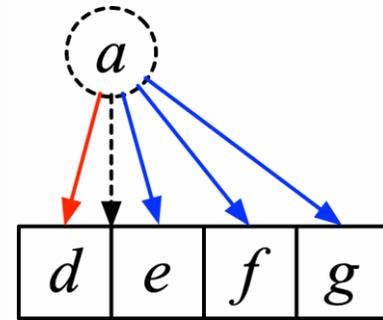
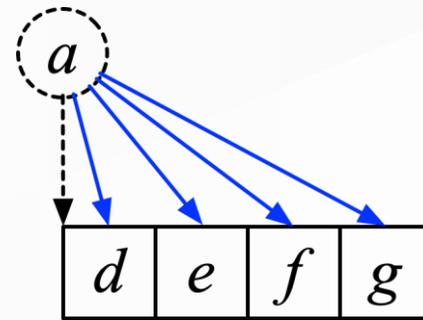


GoGraph

5) Insert high-degree and isolated vertices



- > insert to a position
- > positive edge
- > negative edge



[*h, a, d, e, g, f, c, b*]

Experiments

- **Competitors**
Degree Sorting, Hub Sorting, Hub Clustering, Rabbit, Gorder
- **Workloads**
PageRank, SSSP, BFS, PHP
- **Environment**
Linux server, 98 GB RAM, Ubuntu 22.04 (64-bit), GCC 7.5
- **Datasets**

Dataset	Vertices	Edges	Abbreviation
Indochina [16]	11,358	49,138	IC
SK-2005 [16]	121,422	36,7579	SK
Google [46]	875,713	5,241,298	GL
Wiki-2009 [16]	1,864,433	4,652,358	WK
Cit-Patents [47]	3,774,768	18,204,371	CP
LiveJournal [16]	4,033,137	27,972,078	LJ

Overall performance

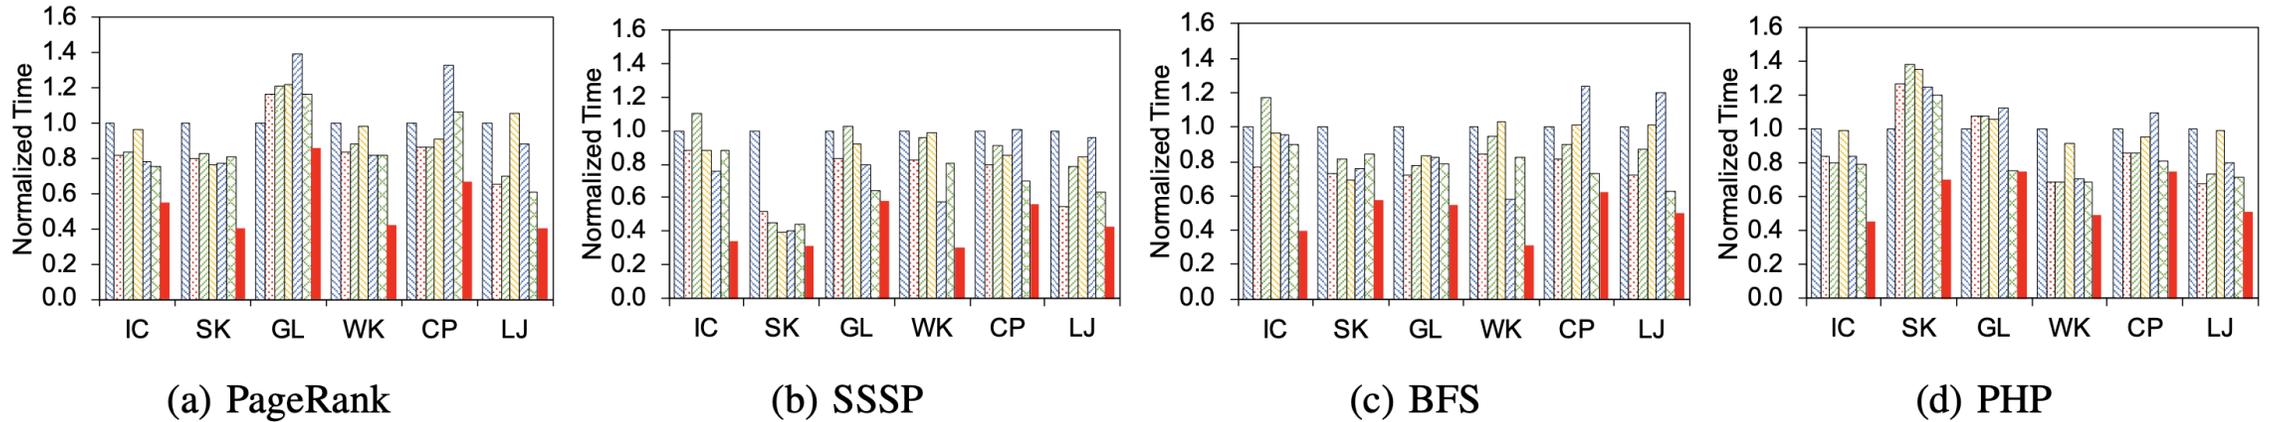


Fig. 1 The comparison of runtime

GoGraph outperforms competitors by an average of 1.83× in runtime.

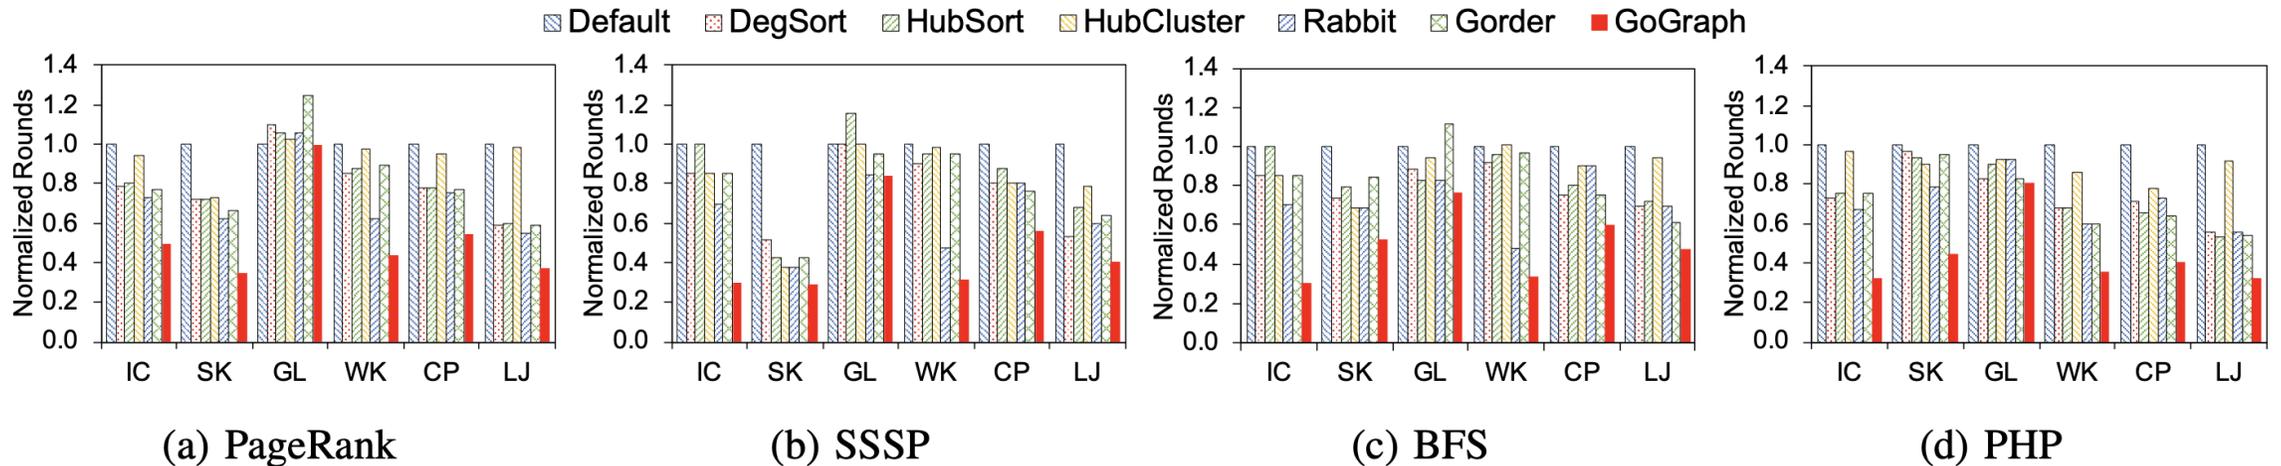


Fig. 2 The comparison of iteration rounds

GoGraph outperforms competitors by 41% on average reduction in iteration rounds.

Convergence comparison

- The distance from the state value at time t to the converged state value: $dist_t = \left| \sum_{v \in V} x^* - \sum_{v \in V} x_t \right|$

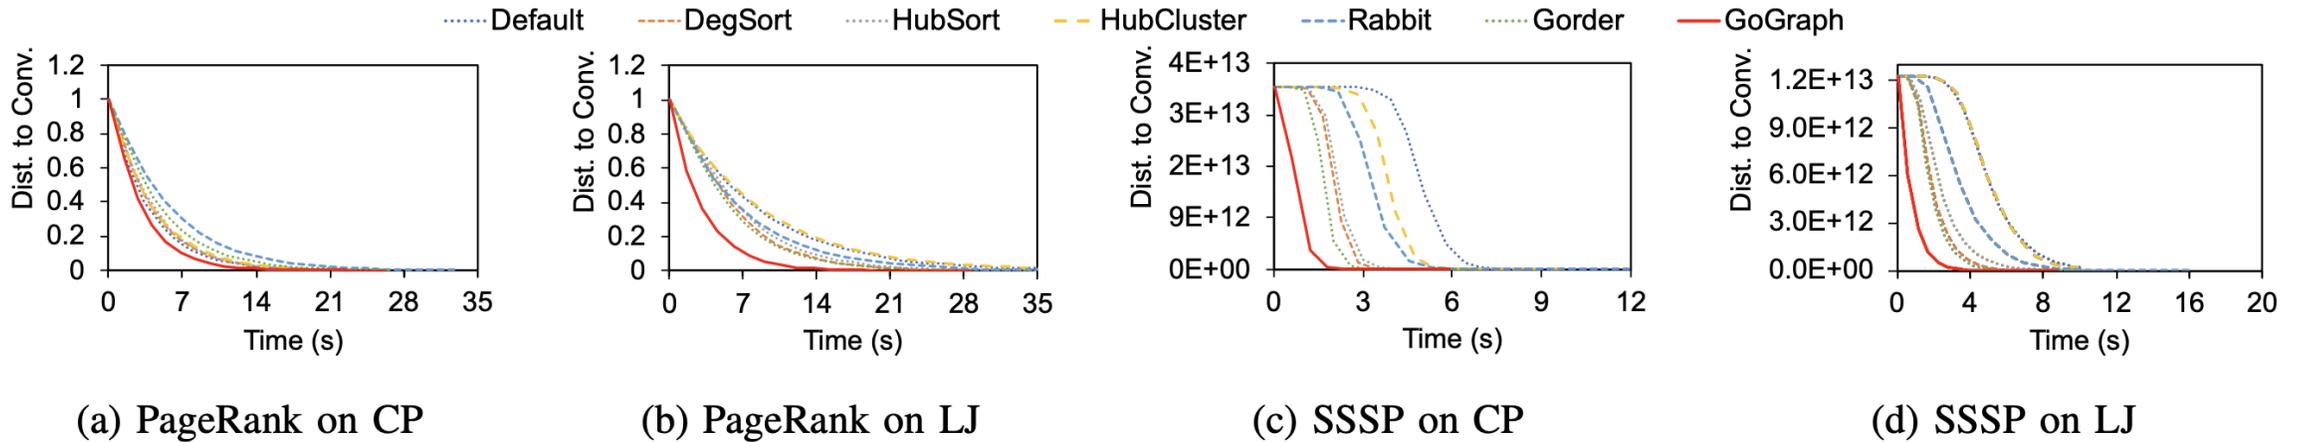


Fig. 3 The comparison of convergence speed

GoGraph algorithm consumes 59% of the average time used by competitors (with a minimum requirement of 37%) to reach convergence.

CPU cache miss

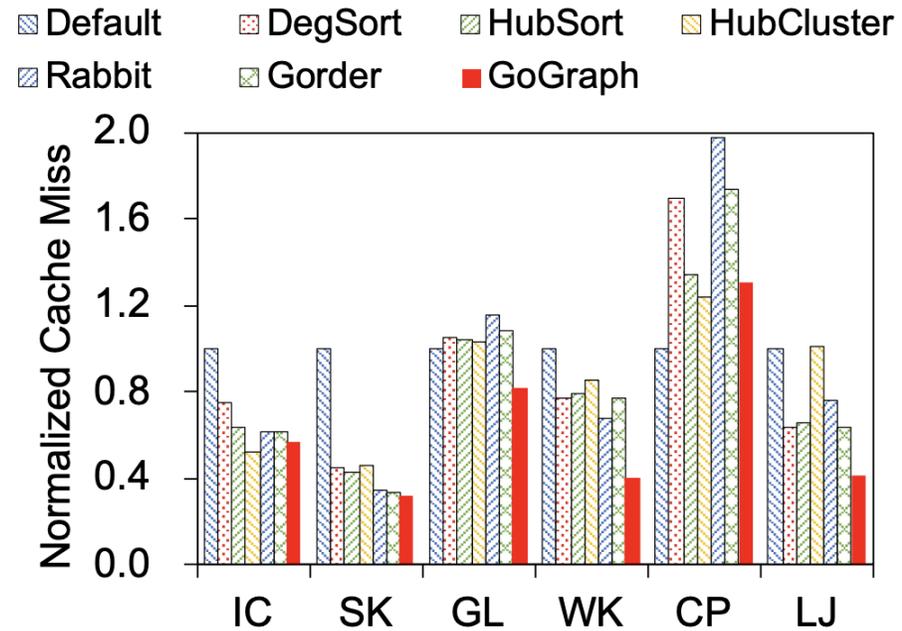


Fig. 4 The comparison of CPU cache miss

GoGraph can reduce the cache miss by 30% on average.

Conclusion

- We propose GoGraph, a graph reordering algorithm
- We propose a metric to measure the efficiency of the vertex processing order

Thank you for listening!