

# Decoupling Schedule, Topology Layout, and Algorithm to Easily Enlarge the Tuning Space of GPU Graph Processing

Shinnung Jeong<sup>1</sup>, Yongwoo Lee<sup>1</sup>, Jaeho Lee<sup>1</sup>, Heelim Choi<sup>1</sup>,

Seungbin Song<sup>1</sup>, Jinho Lee<sup>2</sup>, Yongsok Kim<sup>1</sup>, Hanjun Kim<sup>1</sup>

Yonsei University<sup>1</sup> Seoul National University<sup>2</sup>

October 10-12, 2022

# Graph Processing Is Important!



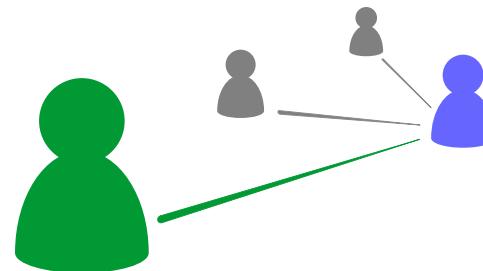
Google 검색 또는 URL 입력



Web Search



Neuroscience



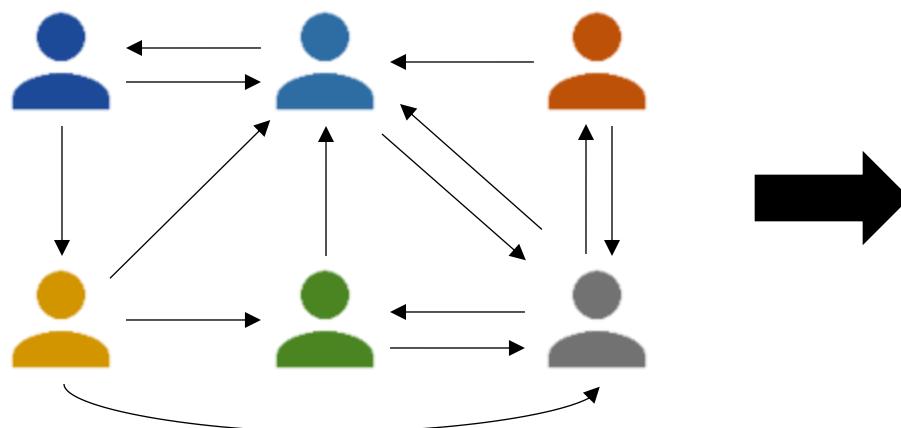
Social Network Analysis



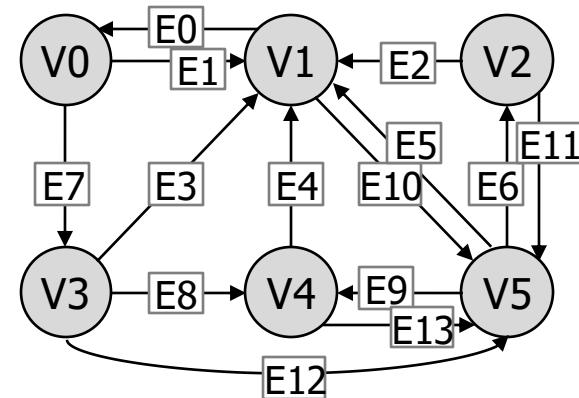
Single Source Shortest Path in Map

# Graph

- An abstract data structure with vertices and their pairs (edges)
- Graph = (Vertex, Edge)



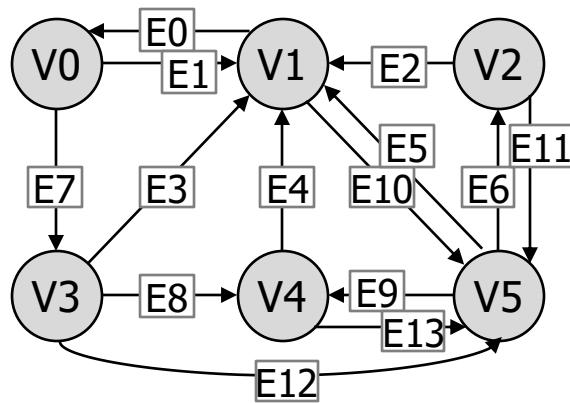
Friendship Relation



Graph

# Graph Processing

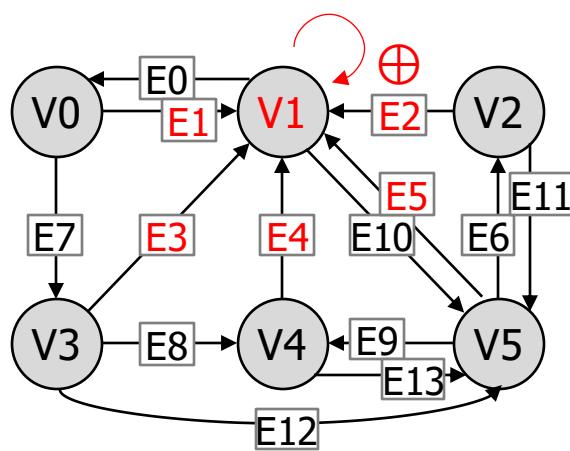
- Analyzing a graph with a given algorithm
- Computes each vertex value with its neighboring edges and vertices
- Graph processing = Algorithm + Schedule + Topology layout



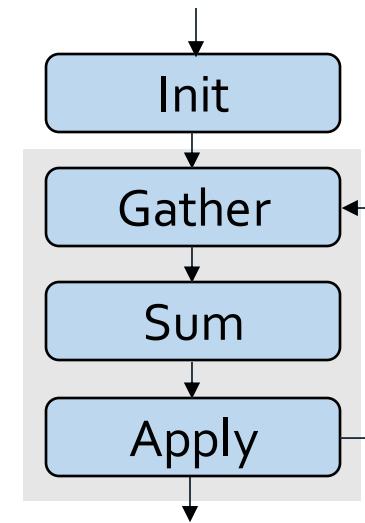
Graph

# Graph Processing : Algorithm

- Graph processing = **Algorithm** + Schedule + Topology layout
- Algorithm: **how to process a graph**
- Gather data from neighbors, Accumulate the data, and Update vertex



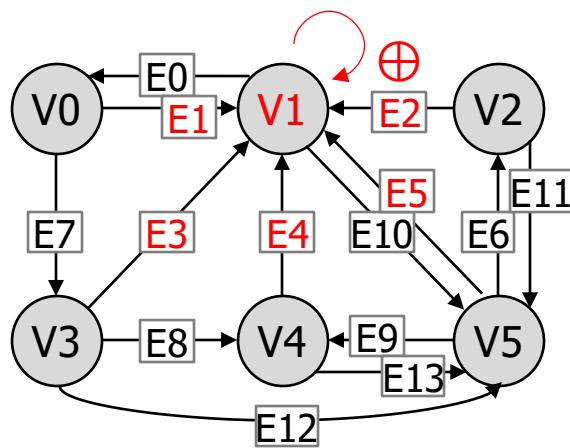
Graph



Algorithm

# Graph Processing : Schedule

- Graph processing = Algorithm + **Schedule** + Topology layout
- Schedule: **How to execute the algorithm**
- Determine which thread processes which parts in what order



Graph

T0	T1	T2	T3	T4	T5
E0	E1	E6	E7	E8	E11
	E2			E9	E12
	E3			E10	E13
	E4				
	E5				

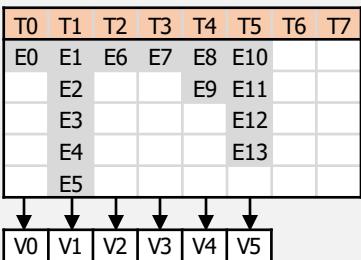
Time

Schedule (Vertex-Mapping)

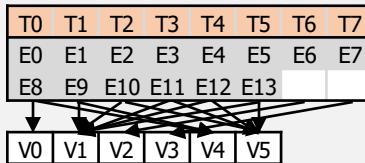
# Graph Processing : Schedule

- Depending on how to fetch edges, there are various schedules

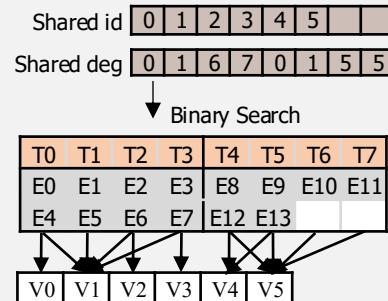
**VM (Vertex-Mapping)**



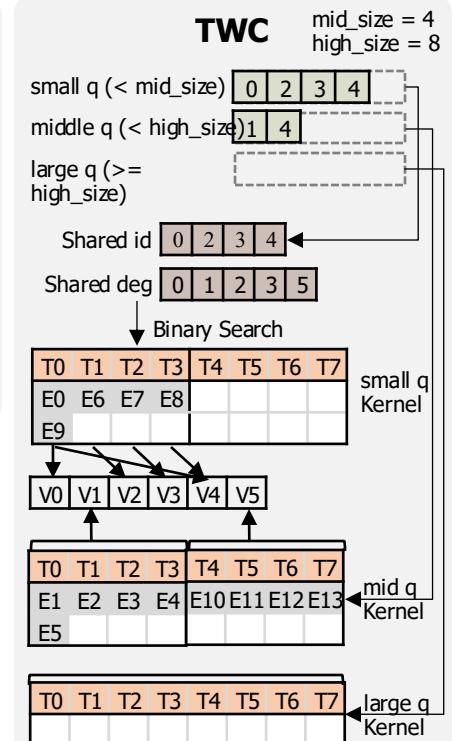
**EM (Edge-Mapping)**



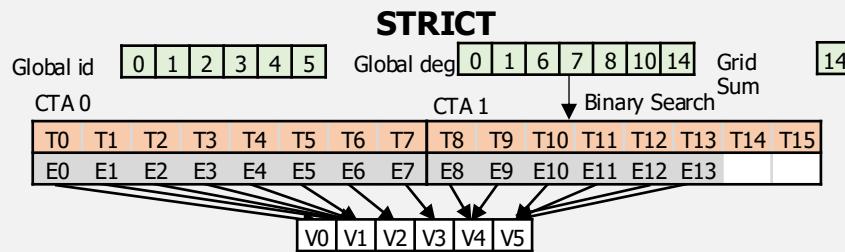
**WM / CM Ex) WM**



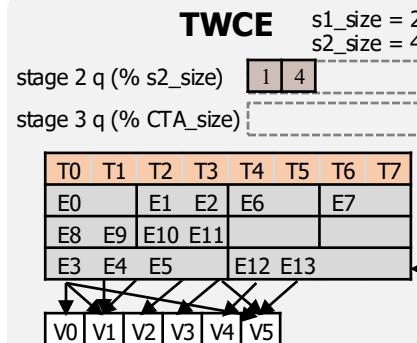
**TWC**



**STRICT**

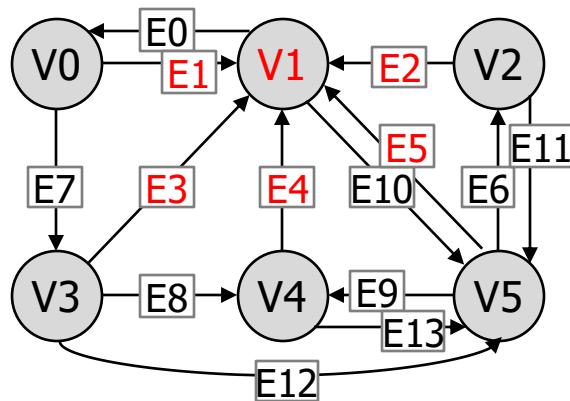


**TWCE**



# Graph Processing : Topology Layout

- Graph processing = Algorithm + Schedule + **Topology layout**
- Topology layout: **how to store a graph topology in GPU memory**
- Store edges (sources and destinations)



Graph

ptr	0	1	6	7	8	10	14
src	1	0	2	3	4	5	5
dest	0	1	1	1	1	1	2

Topology layout (COO)

# Graph Processing : Topology Layout

- Depending on how to store edges and how a vertex points its neighboring edges, there are various topology layouts

ptr	0	1	6	7	8	10	14
-----	---	---	---	---	---	----	----

src	1	0	2	3	4	5	5	0	3	5	1	2	3	4
-----	---	---	---	---	---	---	---	---	---	---	---	---	---	---

dest	0	1	1	1	1	1	2	3	4	4	5	5	5	5
------	---	---	---	---	---	---	---	---	---	---	---	---	---	---

COO

ptr	0	1	6	7	8	10	14
-----	---	---	---	---	---	----	----

list	1	0	2	3	4	5	5	0	3	5	1	2	3	4
------	---	---	---	---	---	---	---	---	---	---	---	---	---	---

CSR/CSC

plus	1	0	0	1	1	*
------	---	---	---	---	---	---

PLUS\_SIZE = 1

minus	*	1	1	0	0	1
-------	---	---	---	---	---	---

MINUS\_SIZE = 1

ptr	0	0	3	4	5	5	8
-----	---	---	---	---	---	---	---

src	3	4	5	5	1	1	2	3
-----	---	---	---	---	---	---	---	---

dest	1	1	1	2	3	5	5	5
------	---	---	---	---	---	---	---	---

DIA

Extra Data Ref = CSR

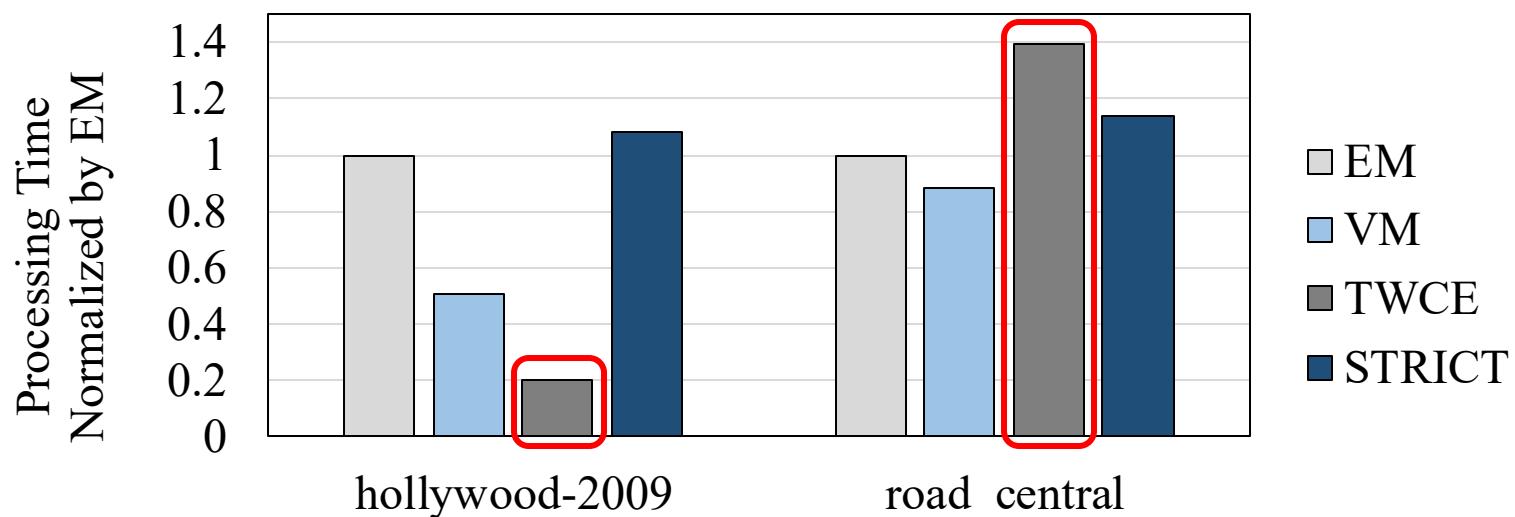
ptr	0	0	3	3	3	3	5
-----	---	---	---	---	---	---	---

list	3	4	5	3	4
------	---	---	---	---	---

ELL

# Tuning Schedule & Topology Layout

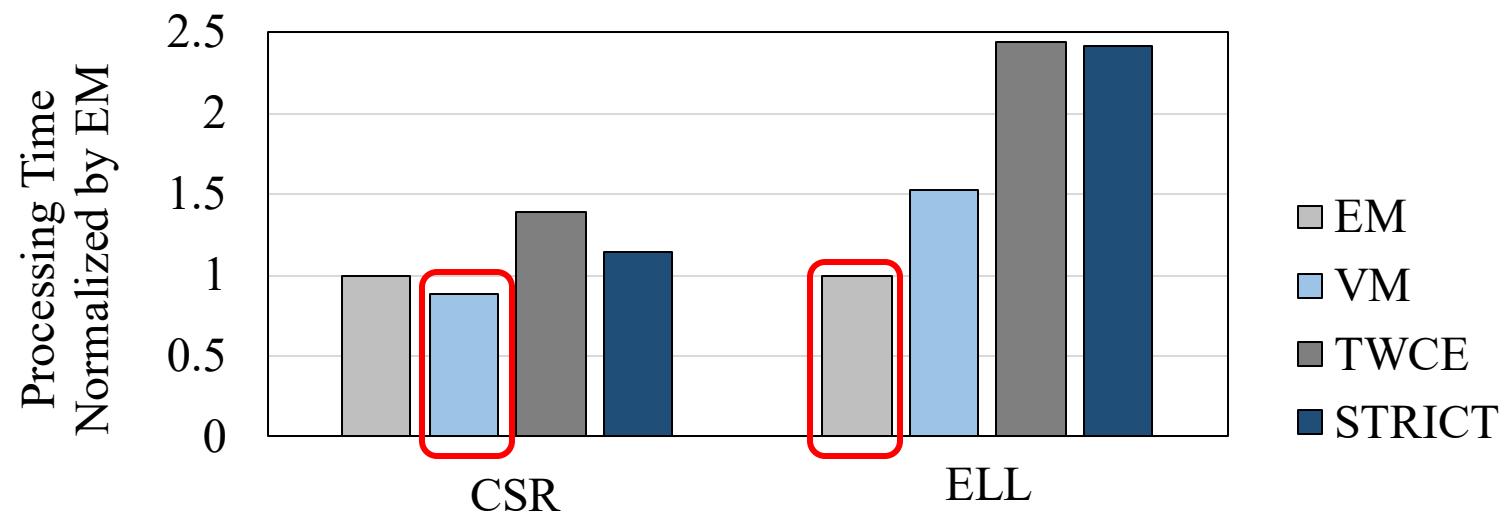
- Dataset affects the optimal schedule
- TWCE is **best for Hollywood-2009**, but **worst for Road\_centeral**



**Tuning a schedule is required for each dataset!**

# Tuning Schedule & Topology Layout

- Topology layout affects the optimal schedule
- VM is best for CSR, but EM is best for ELL

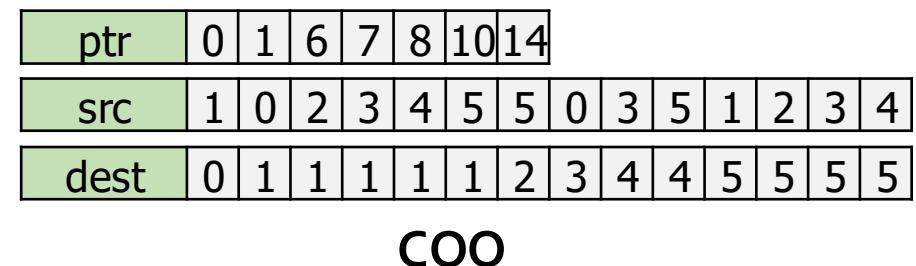
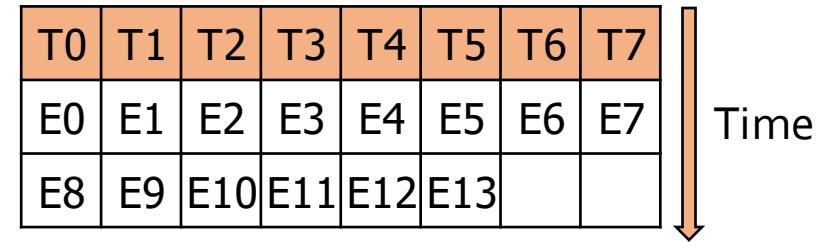


Schedule and topology layout should be considered together in tuning!

# Existing Graph Processing Model

- Only **algorithm** is decoupled (**gather, sum, apply**)
- **Schedule** and **Topology layout** are tightly **coupled** to the model

```
1 foreach tid in (0, #Thread){  
2   for eid in (tid, #Edge, #Thread){  
3     src = coo.srclist[eid]  
4     dest = coo.destlist[eid]  
5     if sum(dest, gather(src, dest, eid))  
6       activate(dest)  
7   }  
8   foreach vid in (0, #Vertex){  
9     apply(vid)  
10 }
```



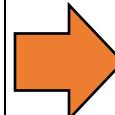
Schedule : Edge Mapping

Topology Layout : COO

Algorithm : Abstract API (gather,sum,apply)

# Existing Work : When to change Topology Layout

```
1 foreach tid in (0, #Thread){  
2   for eid in (tid, #Edge, #Thread){  
3     src = coo.srclist[eid]  
4     dest = coo.destlist[eid]  
5     if sum(dest, gather(src, dest, eid))  
6       activate(dest)  
7   } }  
8 foreach vid in (0, #Vertex){  
9   apply(vid) }
```

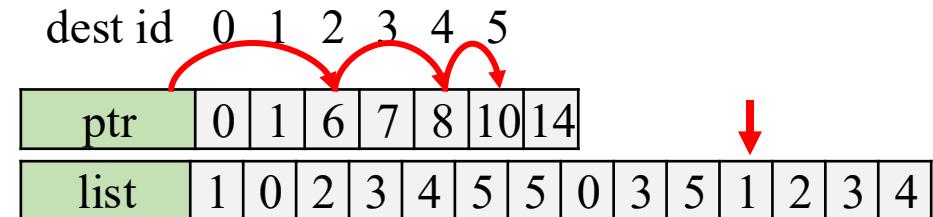


```
1 foreach tid in (0, #Thread){  
2   for eid in (tid, #Edge, #Thread){  
3     vid = binarySearch(csr.ptr, eid)  
4     if(vid == -1) continue  
5     dest = vid  
6     src = csr.list[eid]  
7     if sum(dest, gather(src, dest, eid))  
8       activate(dest)  
9   } }  
10 foreach vid in (0, #Vertex){  
11   apply(vid)}
```

Schedule : Edge Mapping  
Topology Layout : COO

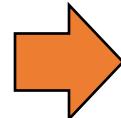
ptr	0	1	6	7	8	10	14
src	1	0	2	3	4	5	5
dest	0	1	1	1	1	2	3

Schedule : Edge Mapping  
Topology Layout : CSR



# Existing Work : When to change Schedule

```
1 foreach tid in (0, #Thread){  
2   for eid in (tid, #Edge, #Thread){  
3     vid = binarySearch(csr.ptr, eid)  
4     if(vid == -1) continue  
5     dest = vid  
6     src = csr.list[eid]  
7     if sum(dest, gather(src, dest, eid))  
8       activate(dest)  
9   }  
10  foreach vid in (0, #Vertex){  
11    apply(vid)  
12  }
```



```
1 foreach tid in (0, #Thread){  
2   for vid in (tid, #Vertex, #Thread){  
3     begin = csr.offset[vid]  
4     end = csr.offset[vid + 1]  
5     for eid in(begin, end):  
6       dest = vid  
7       src = csr.list[eid]  
8       if sum(dest, gather(src, dest, eid)):  
9         activate(dest)  
10    }  
11    foreach vid in (0, #Vertex){  
12      apply(vid)  
13    }
```

Schedule : Edge Mapping  
Topology Layout : CSR

Schedule : Vertex Mapping  
Topology Layout : CSR

Changing a schedule and layout requires to change the entire program!

# Observation 1: Edge Scheduling

- Processing models **schedule edge**, and execute gather and sum

```
1 foreach tid in (0, #Thread){  
2   for eid in (tid, #Edge, #Thread){  
3     vid = binarySearch(csr.ptr, eid)  
4     if(vid == -1) continue  
5     dest = vid  
6     src = csr.list[eid]  
7     if sum(dest, gather(src, dest, eid))  
8       activate(dest)  
9   }  
10  foreach vid in (0, #Vertex){  
11    apply(vid)  
12  }
```

Schedule : Edge Mapping  
Topology Layout : CSR

```
1 foreach tid in (0, #Thread){  
2   for vid in (tid, #Vertex, #Thread){  
3     begin = csr.offset[vid]  
4     end = csr.offset[vid + 1]  
5     for eid in(begin, end):  
6       dest = vid  
7       src = csr.list[eid]  
8       if sum(dest, gather(src, dest, eid)):  
9         activate(dest)  
10    }  
11  foreach vid in (0, #Vertex){  
12    apply(vid)  
13  }
```

Schedule : Vertex Mapping  
Topology Layout : CSR

# Observation 1: Edge Scheduling

- Decouple the edge schedule step from the processing model

```
1 foreach tid in (0, #Thread){  
2  
3     for edge (edge schedule)  
4  
5         dest = vid  
6         src = csr.list[eid]  
7         if sum(dest, gather(src, dest, eid))  
8             activate(dest)  
9     }  
10    foreach vid in (0, #Vertex){  
11        apply(vid)  
12    }
```

Schedule : Edge Mapping  
Topology Layout : CSR

```
1 foreach tid in (0, #Thread){  
2  
3     for edge (edge schedule)  
4  
5  
6         dest = vid  
7         src = csr.list[eid]  
8         if sum(dest, gather(src, dest, eid)):  
9             activate(dest)  
10        }  
11        foreach vid in (0, #Vertex){  
12            apply(vid)  
13        }
```

Schedule : Vertex Mapping  
Topology Layout : CSR

# Observation 2: Topology Layout Access

- After scheduling edge, access source and destination for each edge

```
1 foreach tid in (0, #Thread){  
2   for edge (edge schedule)  
3     dest = vid  
4     src = csr.list[eid]  
5     if sum(dest, gather(src, dest, eid))  
6       activate(dest)  
7   }  
8   foreach vid in (0, #Vertex){  
9     apply(vid)  
10  }
```

Schedule : Edge Mapping  
Topology Layout : CSR

```
1 foreach tid in (0, #Thread){  
2   for edge (edge schedule)  
3     dest = vid  
4     src = csr.list[eid]  
5     if sum(dest, gather(src, dest, eid)):  
6       activate(dest)  
7     }  
8   foreach vid in (0, #Vertex){  
9     apply(vid)  
10  }
```

Schedule : Vertex Mapping  
Topology Layout : CSR

# Observation 2: Topology Layout Access

- Decouple edge data access from the model

```
1 foreach tid in (0, #Thread){  
2     for edge (edge schedule)  
3         Get Edge Info.  
4             (Topology Layout Access)  
5         if sum(dest, gather(src, dest, eid))  
6             activate(dest)  
7     }  
8     foreach vid in (0, #Vertex){  
9         apply(vid)  
10    }
```

Schedule : Edge Mapping  
Topology Layout : CSR

```
1 foreach tid in (0, #Thread){  
2     for edge (edge schedule)  
3         Get Edge Info.  
4             (Topology Layout Access)  
5         if sum(dest, gather(src, dest, eid)): 6             activate(dest)  
7     }  
8     foreach vid in (0, #Vertex){  
9         apply(vid)  
10    }
```

Schedule : Vertex Mapping  
Topology Layout : CSR

# Observation 3: Vertex Scheduling

- Apply is executed for each vertex (Vertex Scheduling)

```
1 foreach tid in (0, #Thread){  
2     for edge (edge schedule)  
3         Get Edge Info.  
4             (Topology Layout Access)  
5         if sum(dest, gather(src, dest, eid))  
6             activate(dest)  
7     }  
8     foreach vid in (0, #Vertex){  
9         apply(vid)  
10    }
```

Schedule : Edge Mapping  
Topology Layout : CSR

```
1 foreach tid in (0, #Thread){  
2     for edge (edge schedule)  
3         Get Edge Info.  
4             (Topology Layout Access)  
5         if sum(dest, gather(src, dest, eid)):  
6             activate(dest)  
7     }  
8     foreach vid in (0, #Vertex){  
9         apply(vid)  
10    }
```

Schedule : Vertex Mapping  
Topology Layout : CSR

# Observation 3: Vertex Scheduling

- Decouple the vertex schedule step from the processing model

```
1 foreach tid in (0, #Thread){  
2     for edge (edge schedule)  
3         Get Edge Info.  
4             (Topology Layout Access)  
5         if sum(dest, gather(src, dest, eid))  
6             activate(dest)  
7     }  
8     foreach vertex (vertex schedule)  
9         apply(vid)  
10    }
```

Schedule : Edge Mapping  
Topology Layout : CSR

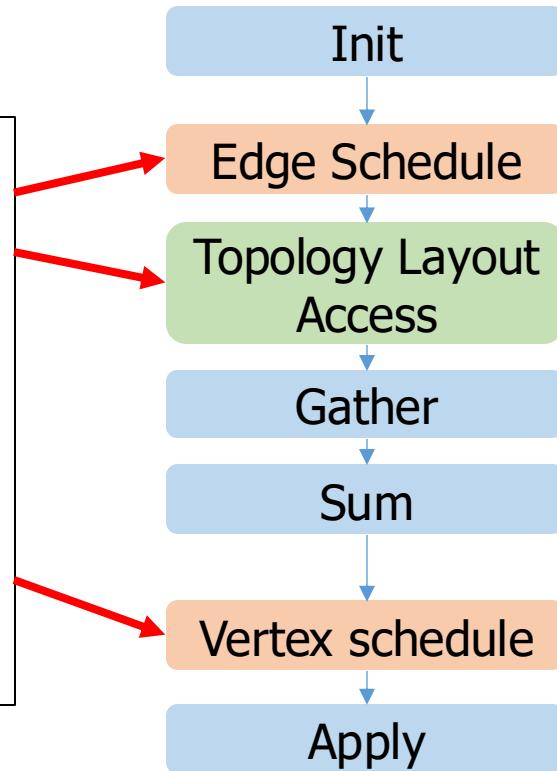
```
1 foreach tid in (0, #Thread){  
2     for edge (edge schedule)  
3         Get Edge Info.  
4             (Topology Layout Access)  
5         if sum(dest, gather(src, dest, eid)):  
6             activate(dest)  
7     }  
8     foreach vertex (vertex schedule)  
9         apply(vid)  
10    }
```

Schedule : Vertex Mapping  
Topology Layout : CSR

# New processing model

- Schedule and topology layout are decoupled from processing model

```
1 foreach tid in (0, #Workload){  
2   while(getNextEdgeID(tid, data, vid, eid)){  
3     getEdge(vid, eid, src, dest)  
5     if sum(dest, gather(src, dest, eid))  
6       activate(dest)  
7     }  
8   }  
9 }  
10 foreach tid in (0, #Vertex){  
11   apply(tid)  
12 }
```



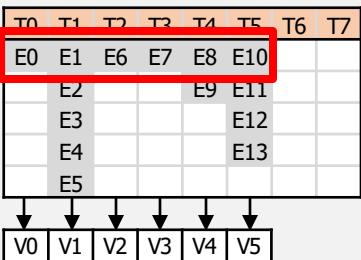
New Processing Model

New Processing Steps

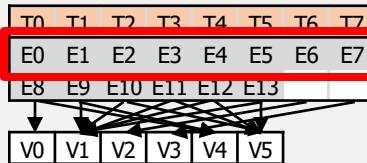
# Schedule Abstraction

- Edge scheduling Abstraction: ***getNextEdgeID***
  - Return edge id for each iteration for each thread

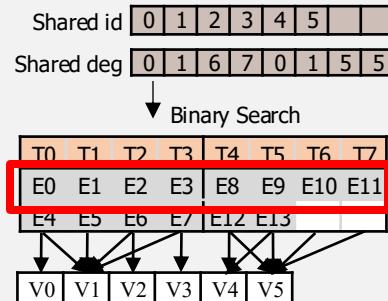
**VM (Vertex-Mapping)**



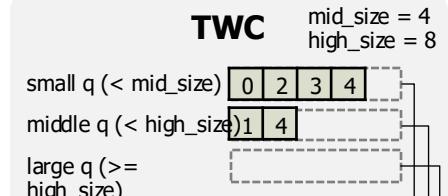
**EM (Edge-Mapping)**



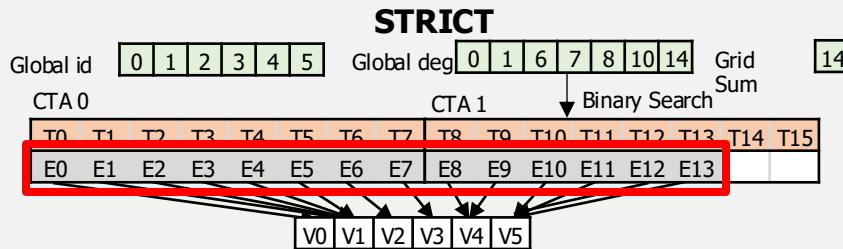
**WM / CM Ex) WM**



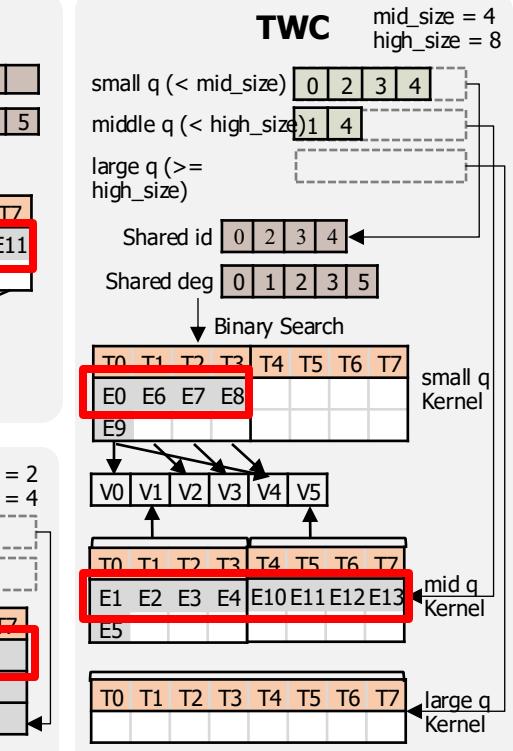
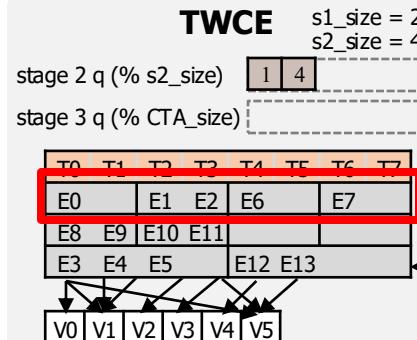
**TWC**



**STRICT**



**TWCE**



# Schedule Abstraction

- Some schemes support load balancing: *initShared*, *initGlobal*
  - Setting **shared** or **global memory** before edge scheduling

**VM (Vertex-Mapping)**

T0	T1	T2	T3	T4	T5	T6	T7
E0	E1	E6	E7	E8	E10		
E2				E9	E11		
E3					E12		
E4						E13	
E5							

V0 V1 V2 V3 V4 V5

**EM (Edge-Mapping)**

T0	T1	T2	T3	T4	T5	T6	T7
E0	E1	E2	E3	E4	E5	E6	E7
E8	E9	E10	E11	E12	E13		

V0 V1 V2 V3 V4 V5

**WM / CM Ex) WM**

Shared id	0	1	2	3	4	5	
Shared deg	0	1	6	7	0	1	5

Binary Search

T0	T1	T2	T3	T4	T5	T6	T7
E0	E1	E2	E3	E8	E9	E10	E11
E4	E5	E6	E7	E12	E13		

V0 V1 V2 V3 V4 V5

**TWC**

mid_size = 4	high_size = 8
small q (< mid_size)	0 2 3 4
middle q (< high_size)	1 4
large q (>= high_size)	

Shared id	0	2	3	4
Shared deg	0	1	2	3

Binary Search

T0	T1	T2	T3	T4	T5	T6	T7
E0	E6	E7	E8				
E9							

**STRICT**

Global id	0	1	2	3	4	5
Global deg	0	1	6	7	8	10
Grid Sum	14					
CTA 0						
CTA 1						

Binary Search

T0	T1	T2	T3	T4	T5	T6	T7
E0	E1	E2	E3	E4	E5	E6	E7
E8	E9	E10	E11	E12	E13		

V0 V1 V2 V3 V4 V5

**TWCE**

s1_size = 2	s2_size = 4
stage 2 q (% s2_size)	1 4

stage 3 q (% CTA_size)	
------------------------	--

T0	T1	T2	T3	T4	T5	T6	T7
E0							
E8							

T0	T1	T2	T3	T4	T5	T6	T7
E0							
E8							

T0	T1	T2	T3	T4	T5	T6	T7
E0							
E8							

# Schedule Abstraction

- Support different schedules w/o changing the processing model

```
1 foreach tid in (0, #Workload){  
2   initShared(tid, data)  
3   while(getNextEdgeID(tid, data, vid, eid)){  
4     getEdge(vid, eid, src, dest)  
5     if sum(dest, gather(src, dest, eid))  
6       activate(dest)  
7   }  
8 }  
9 }  
10 foreach tid in (0, #Vertex){  
11   apply(tid)  
12 }
```

## Proposed Processing Model

### Vertex Mapping

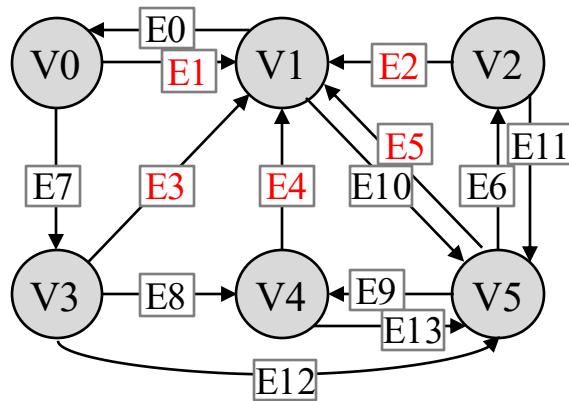
```
1 initShared (tid, data){  
2   getNeighbor(tid, start, end)  
3   data[0] = start  
4   data[1] = end}  
5 getNextEdgeID (tid, data, &vid, &eid){  
6   if(data[0] >= data[1]) return false  
7   vid = tid  
8   eid = data[0]++  
9   return true}
```

### Edge Mapping

```
1 initShared (tid, data){  
2   data[0] = 0}  
3 getNextEdgeID (tid, data, &vid, &eid){  
4   if((data[0]++) = 1) return false  
5   if(searchVID(tid, vid)) return false  
6   eid = tid  
7   return true}
```

# Topology Layout Abstraction

- Topology layout abstraction interfaces
  - ***getEdge*** : Return source and destination vertex for a given edge
  - ***getNeighbor*** : Return incoming and outgoing edges for a given vertex

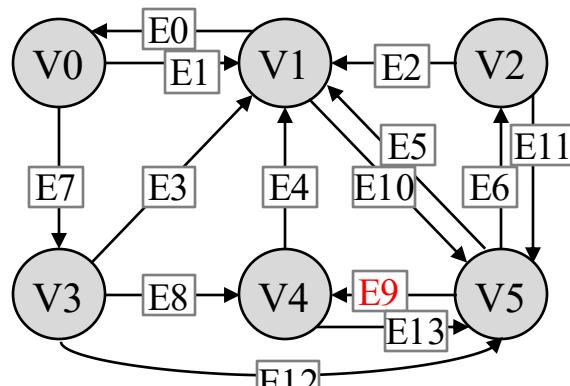


Graph

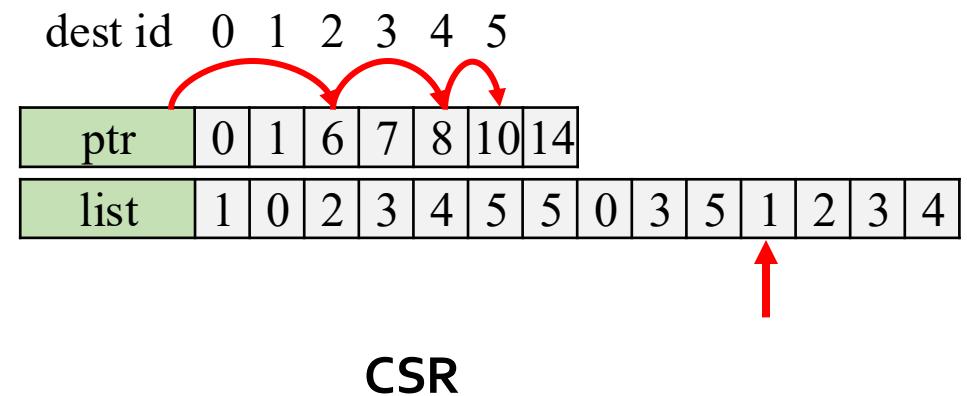
$\text{getNeighbor}(v_1) \rightarrow e_1, e_2, e_3, e_4, e_5$   
 $\text{getEdge}(e_1) \rightarrow (v_0, v_1)$

# Topology Layout Abstraction

- Some schemes (CSR) store only one of source or destination ID
- For fast data access
  - Split ***getEdge*** into two steps: ***searchVID***, ***getEdge***



Graph



# Topology Layout Abstraction

- Support different schedules w/o changing the processing model

```
1 foreach tid in (0, #Workload){  
2   initShared(tid, data)  
3   while(getNextEdgeID(tid, data, vid, eid)){  
4     getEdge(vid, eid, src, dest)  
5     if sum(dest, gather(src, dest, eid))  
6       activate(dest)  
7     }  
8   }  
9 }  
10 foreach tid in (0, #Vertex){  
11   apply(tid)  
12 }
```

## Proposed Processing Model

### COO

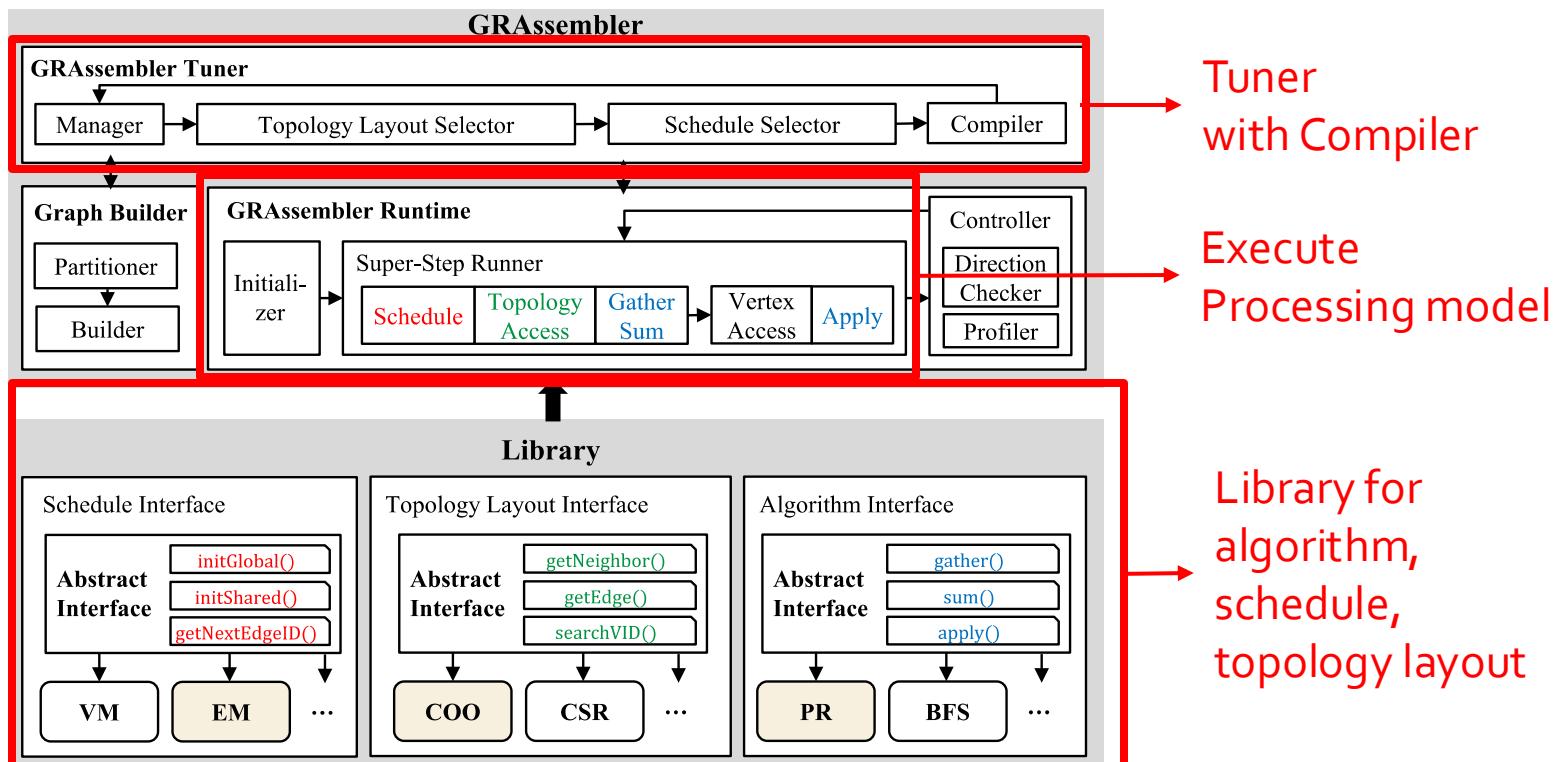
```
1 getNeighbor (vid, &start, &end){  
2   src = ptr[vid]  
3   dest = ptr[vid + 1]}  
4 searchVID (eid, &vid, &cond){  
5   vid = srclist[eid]}  
6 getEdge (vid, eid, &src, &dest){  
7   src = srclist[eid]  
8   dest = destlist[eid]}
```

### CSR

```
1 getNeighbor (vid, &start, &end){  
2   src = offset[vid]  
3   dest = offset[vid + 1]}  
4 searchVID (eid, &vid, &cond){  
5   vid = binarySearch(offset, eid)}  
6 getEdge (vid, eid, &src, &dest){  
7   src = vid  
8   dest = list[eid]}
```

# GRAssembler : Graph Processing Framework

- Execute the proposed processing model
  - Integrating algorithm, schedule, and topology layout
- Optimize integrated graph processing program



# Optimization: Dead Code Elimination

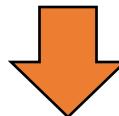
- Some algorithm and schedule skip some processing steps
  - Vertex-Mapping: Do nothing for initGlobal
  - BFS algorithm: Do nothing for apply
- We can eliminate processing step functions if not used

```
1 initGlobal()
2 Foreach Thread tid in (0, #workload){
3     initShared(tid, data)
4     while(getNextEdgeID(tid, data, vid, eid)){
5         getEdge(vid, eid, src, dest)
6         if sum(dest, gather(src, dest, eid))
7             activate(dest)
8     } }
9 Foreach Thread tid in (0, #Vertex){
10    apply(tid)
11 }
```

# Optimization: Atomic Operation Reduction

- If each vertex value is updated by only a thread, the atomic operation is not necessary
- Remove the unnecessary synchronization

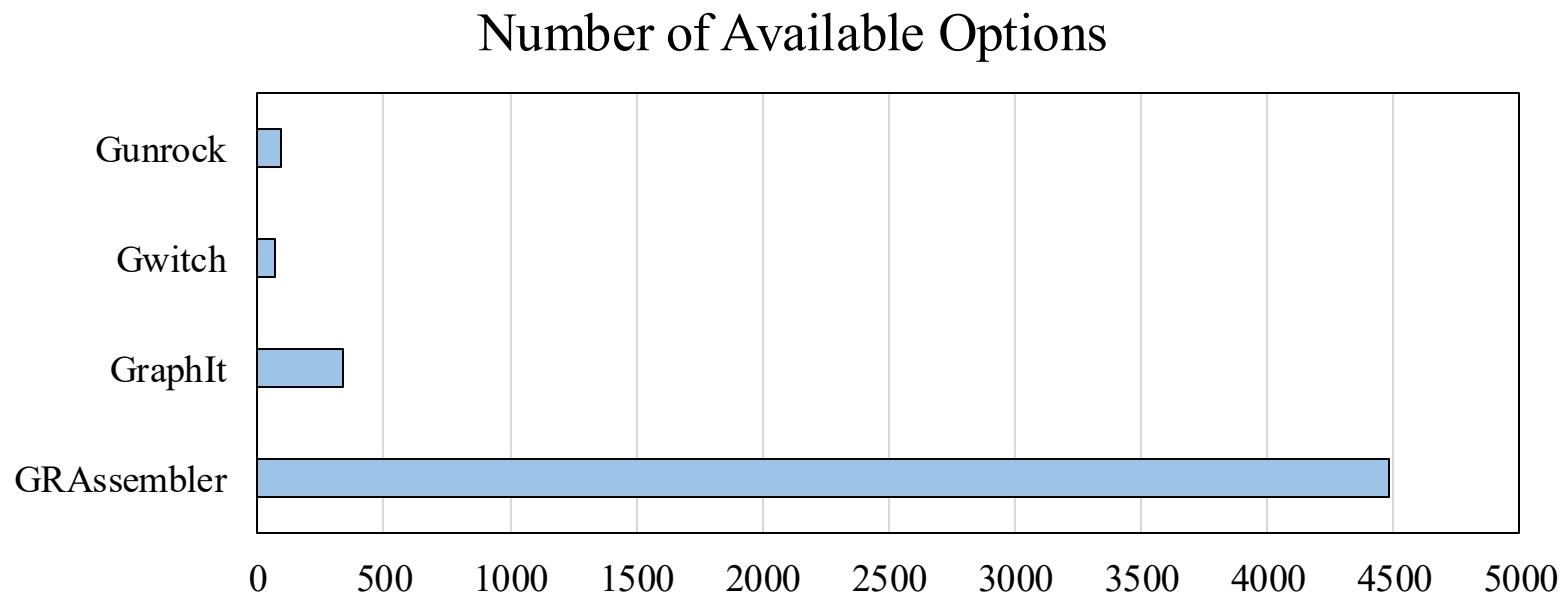
```
1 sum(dest, data){  
2     CAS(new_rank[dest], -1, data);  
3 }
```



```
1 sum(dest, data){  
2     if(new_rank[dest] != -1)  
3         new_rank[dest] = data;  
4 }
```

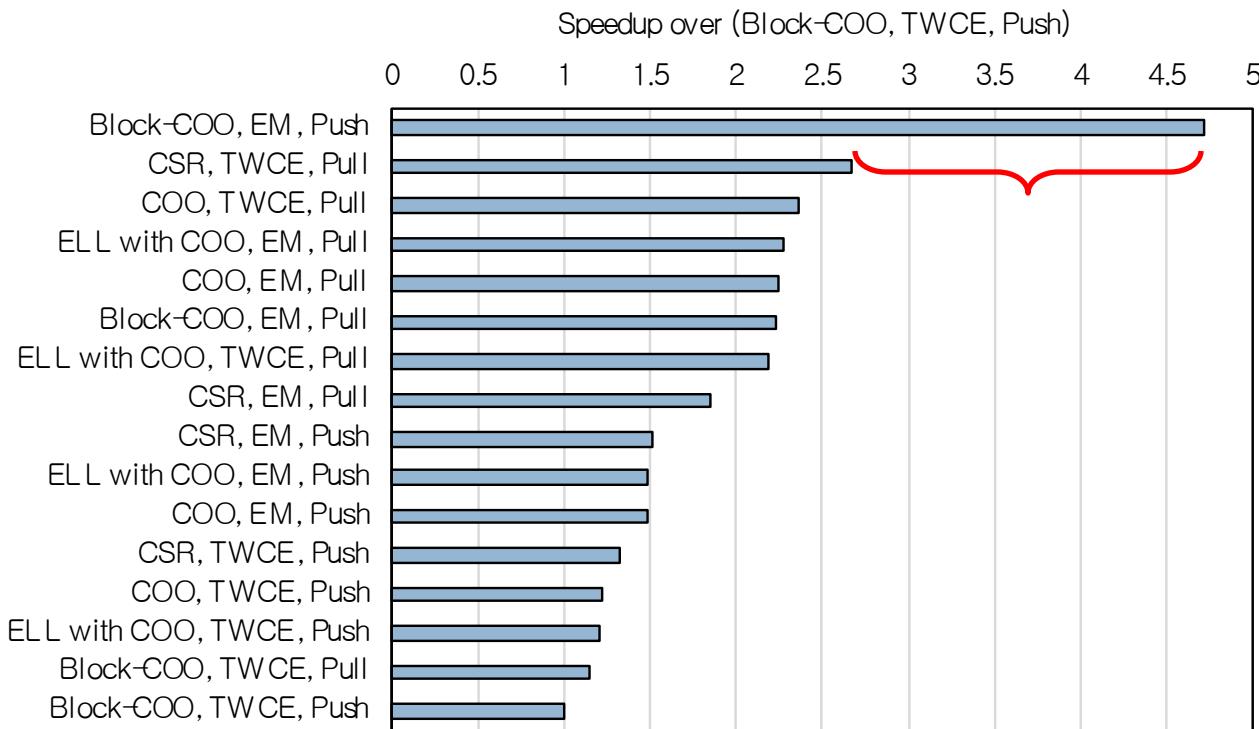
# Evaluation: Tuning Space

- Decoupling schedule and topology layout increases tuning space from 336 (GraphIt) to 4480 (GRAssembler)
- GRAssembler supports additional tuning options
  - CTA Size, Blocking, Active Dataset Structure



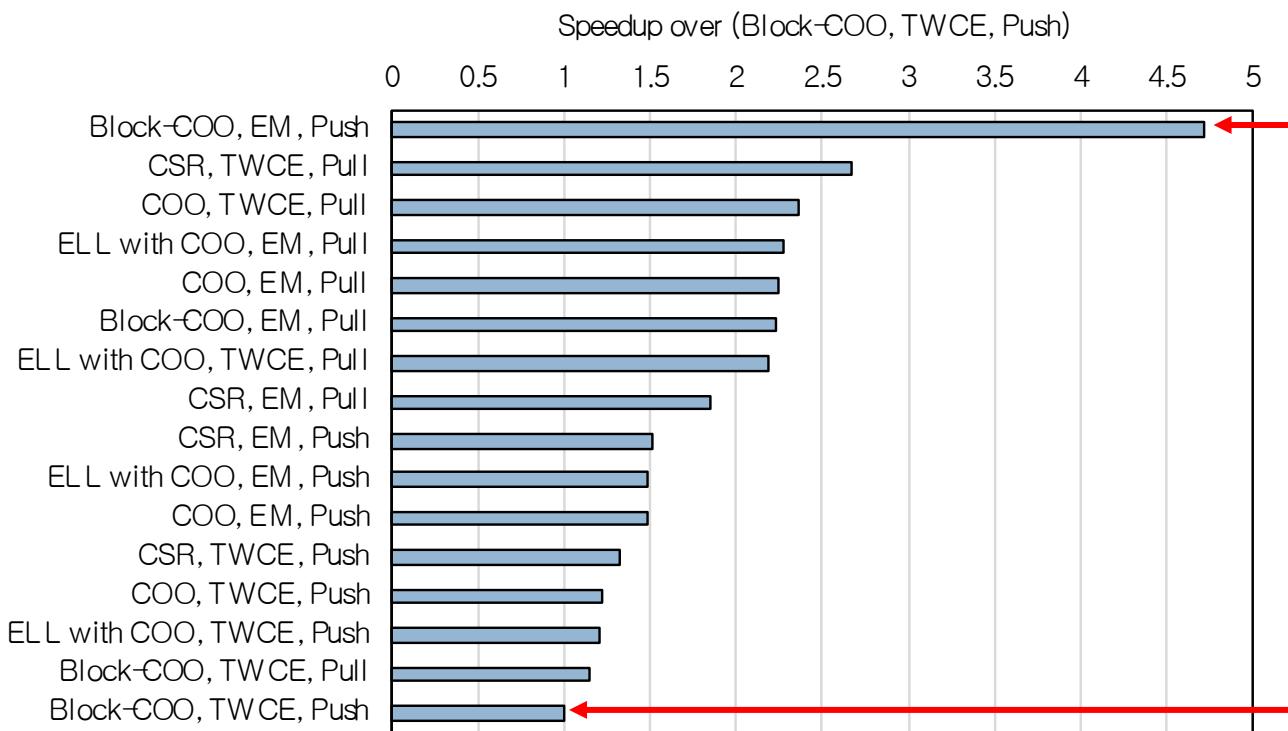
# Extending Tuning Space is Critical

- Almost two times better performance than second optimal solution
- The second optimal solution use different topology layout and schedule



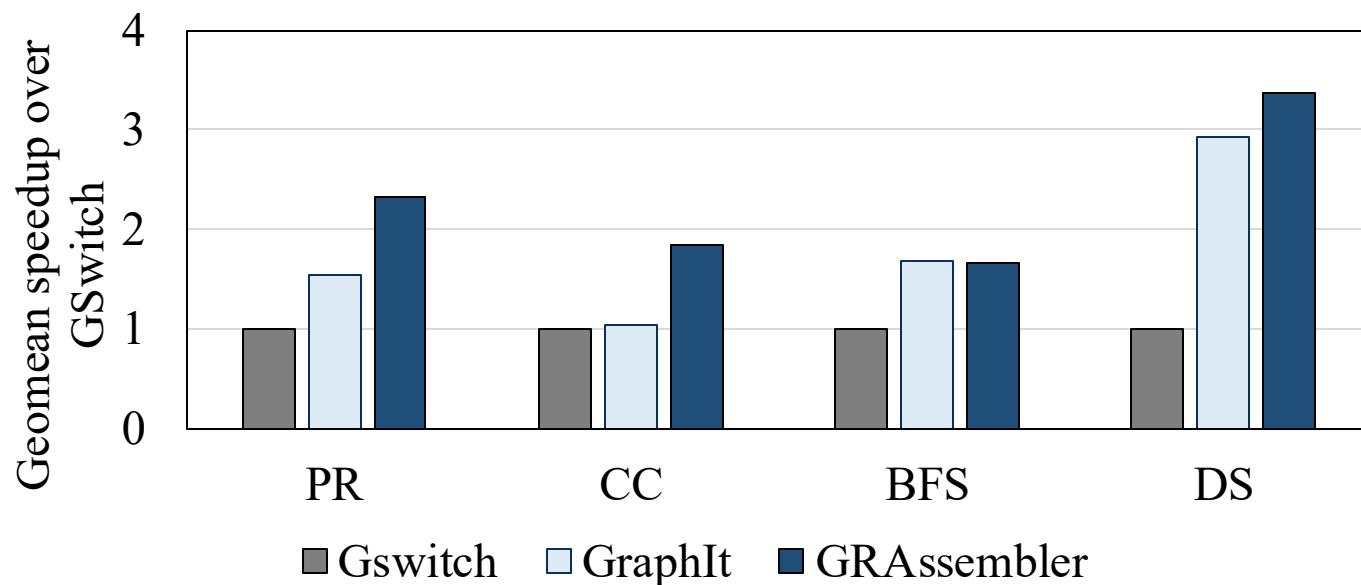
# Synergistic Performance Effect

- There exist synergistic performance effects between tuning options.
- Both the worst and the best use same topology layout.



# GRAssembler Performance Evaluation

- GRAssembler achieves 2.21x and 1.30x geomean speedup respectively over compared to Gswitch and GraphIt
  - GPU : NVIDIA GeForce RTX 3090



# Conclusion

- Decoupling schedule and topology layout from processing model
  - Makes changing a new schedule and topology layout easier
  - Allows new combinations of existing schedules and topology layouts
  - Increases tuning spaces : 336 (GraphIt) to 4480 (GRAssembler)
- GRAssembler: a new prototype graph processing framework
  - Supports the new processing model that decouples schedule and topology layout
  - Optimizes integrated graph processing program
- Increased tuning space yields 30% speedup compared to state-of-the-art framework



# Thank You!



Shinnung Jeong, shin0403@yonsei.ac.kr