



H-INDEX: Hash-Indexing for Parallel Triangle Counting on GPUs

Santosh Pandey, Xiaoye Sherry Li, Aydin Buluc, Jiejun Xu, Hang Liu



Triangle Counting: Background



For edge (7,8):

Neighbor (7) = [1,2,3,4,5,6,9,10,11]

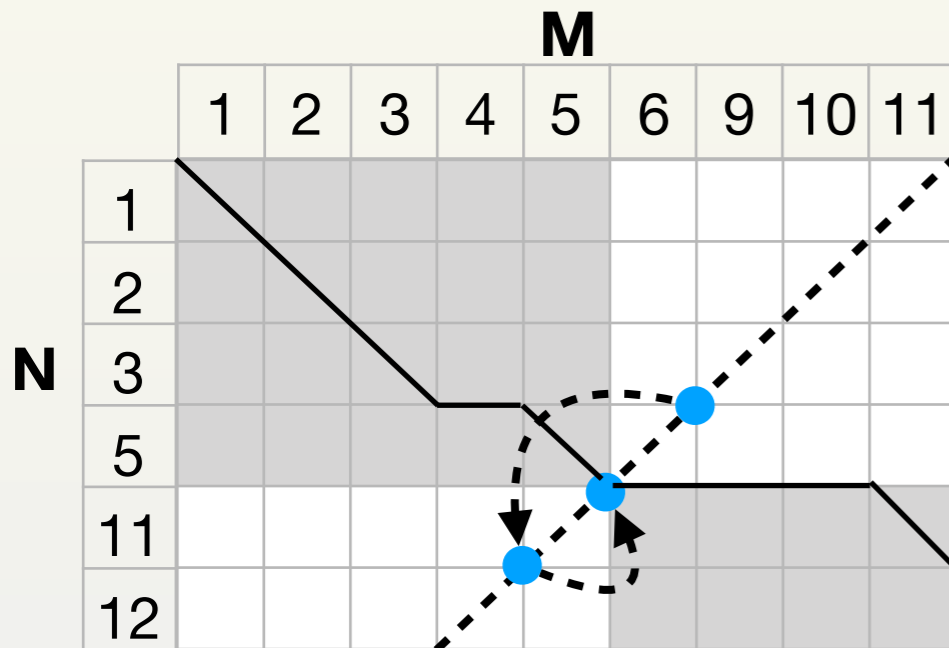
Neighbor (8) = [1,2,3,5,11,12]

Triangles = Intersection (Neighbor(7), Neighbor(8))


List Intersection: Related Methods

M	1	2	3	4	5	6	9	10	11
N	1	2	3	5	11	12			

(a) Input list



(b) Merge-path: $O(M+N)$



(c) Binary-search: $O(N * \log M)$

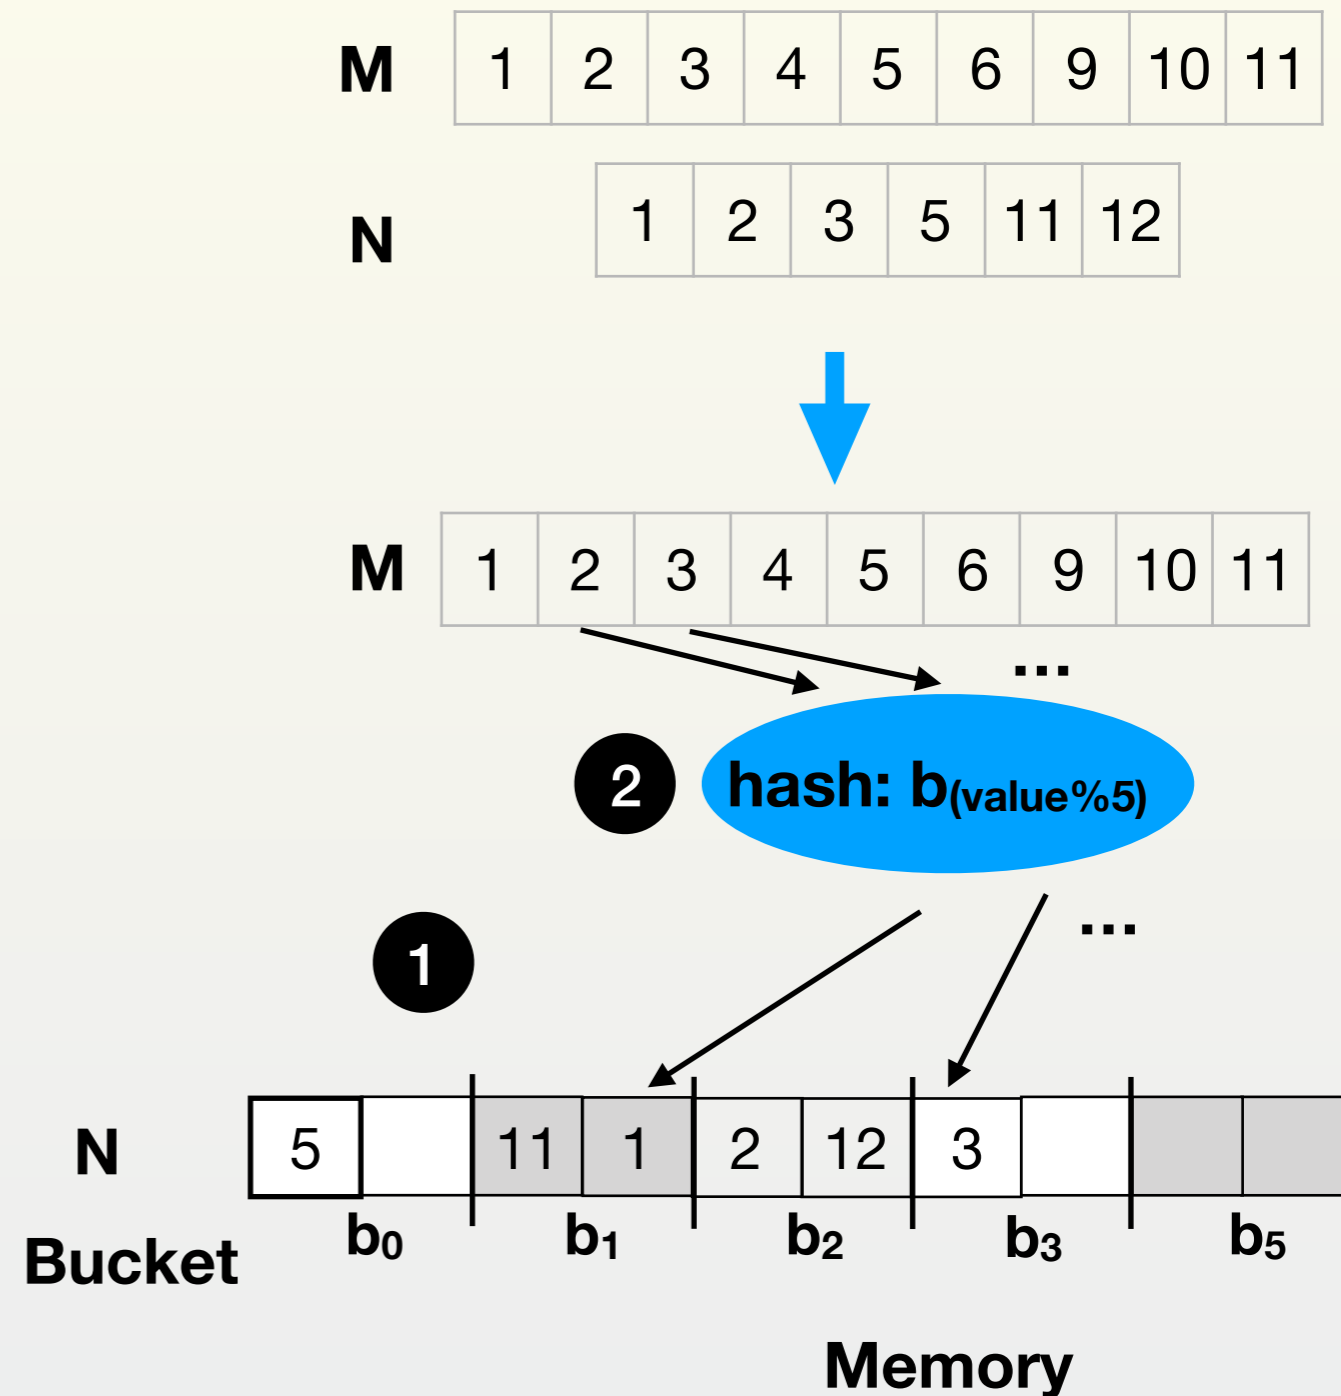
Partitioning Overhead
 Strided Memory Access
 Both Sorted Lists

Drawbacks

Higher Complexity
 Strided Memory Access
 One Sorted List

Hashed Based Intersection

- Hash based work partition.
 - One list to hash;
 - Another list as search key.
- Complexity $\sim O(M + N)$.
- No sorting required.




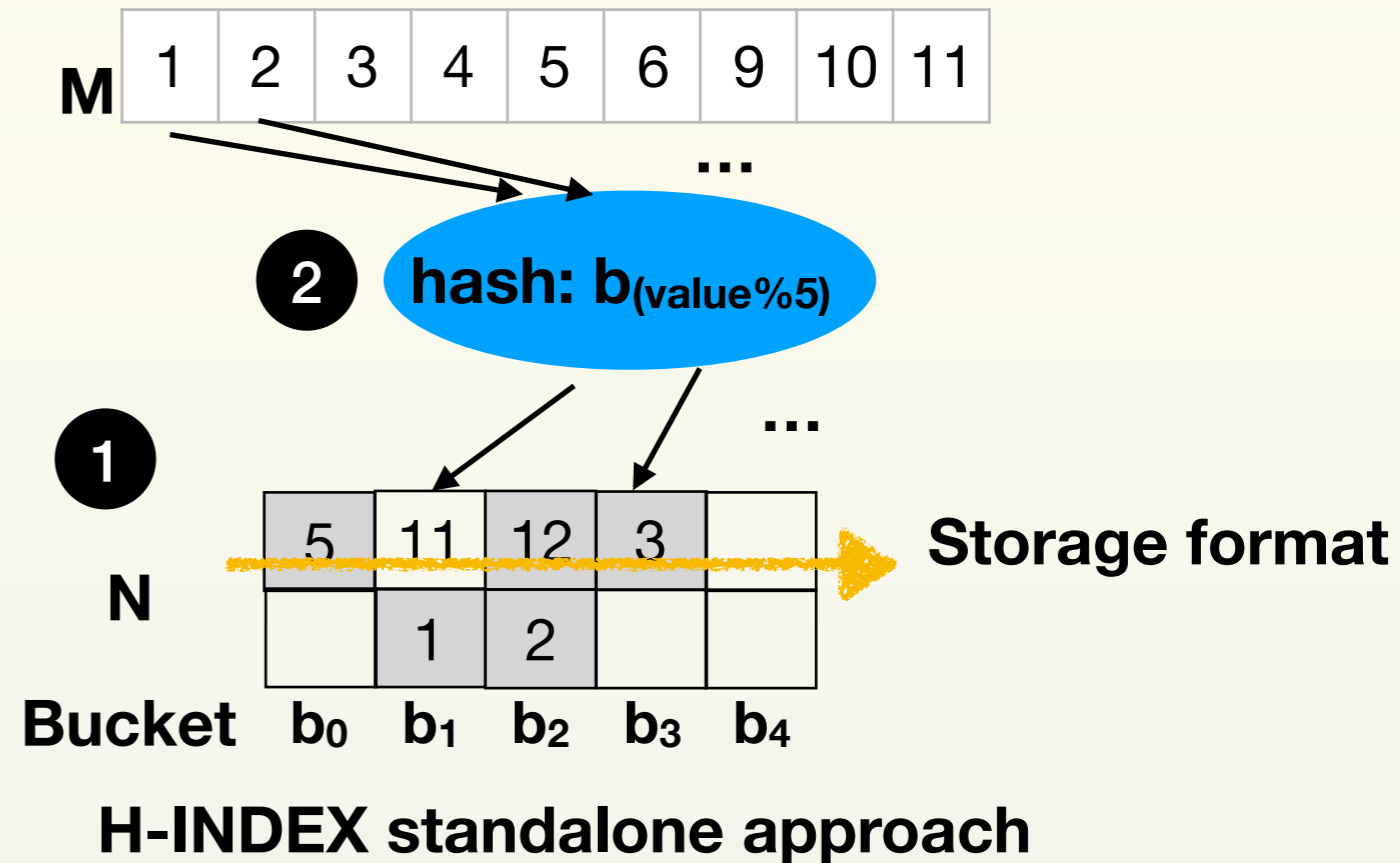
Drawbacks

Strided Memory Access

Hash Collisions

Contribution 1: Coalesced Memory Access

- GPU friendly bucket placement.
- Coalesced memory access for each depth of hash bins.
- Use bin width equal to warp size.



Contribution 2: Reducing Collision

M

1	2	3	4	5	6	9	10	12	15	20	22
---	---	---	---	---	---	---	----	----	----	----	----

N

1	2	3	5	11	12
---	---	---	---	----	----

Input list

Hashing longer neighbor list:

5	6	12	3	4
10	1	2		9
15	6	22		
20				

↓
Depth

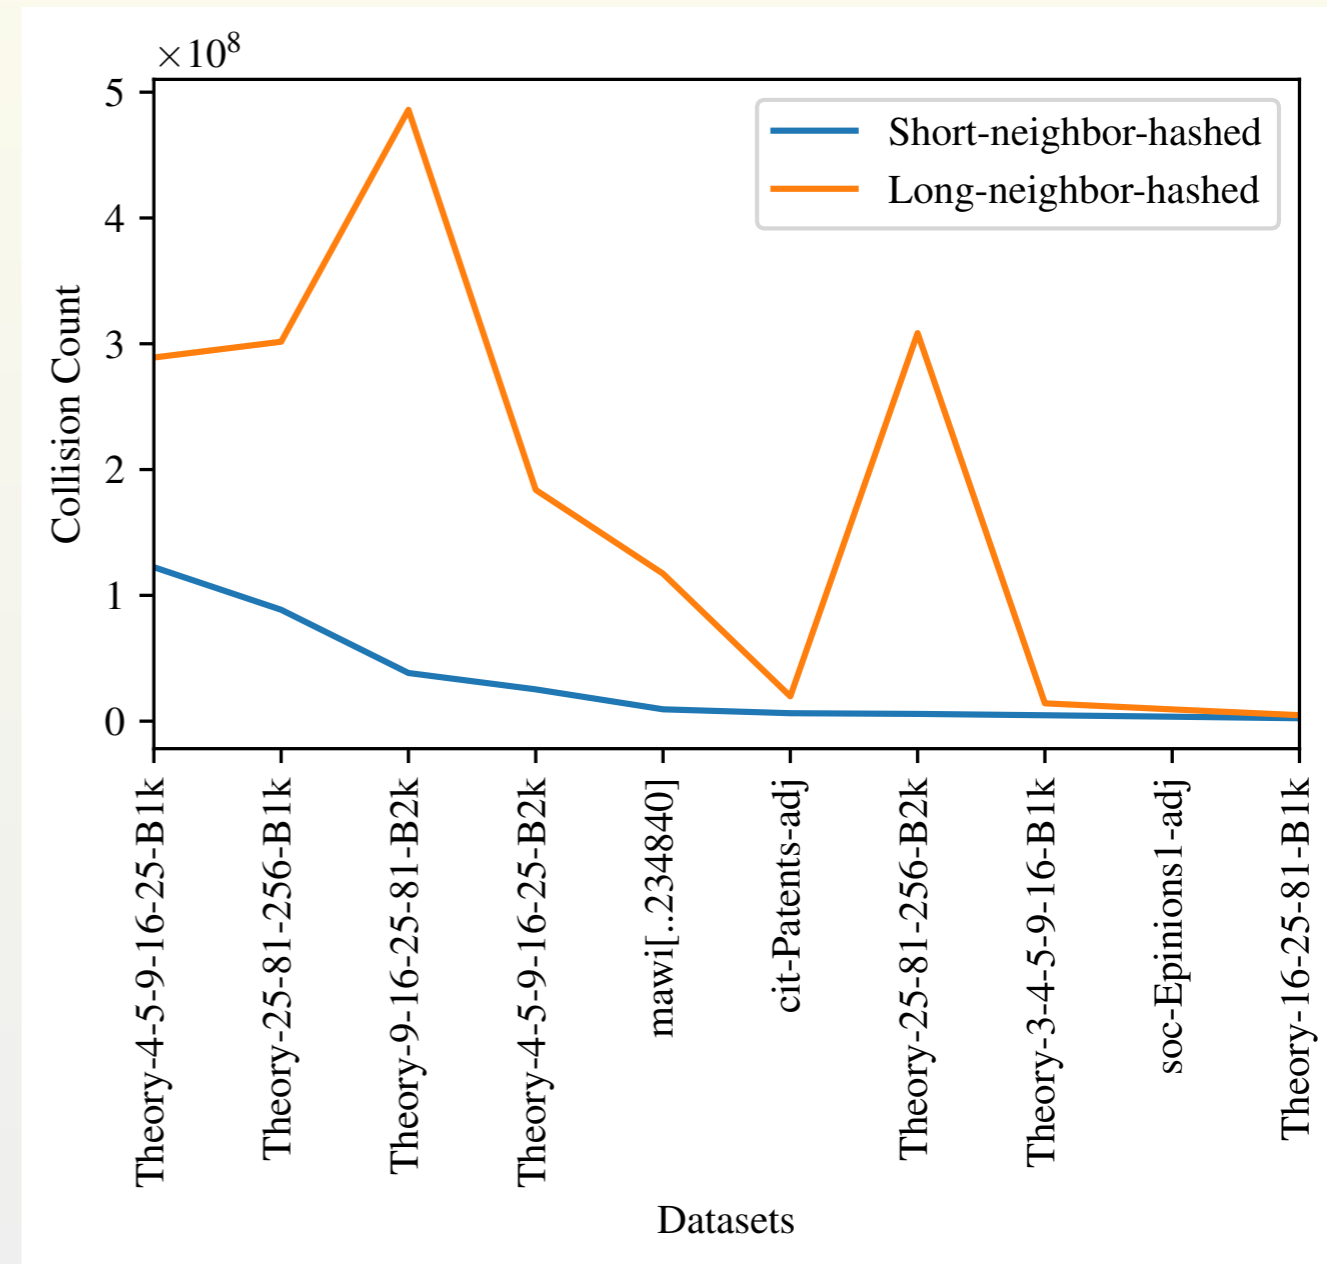
Bucket b_0 b_1 b_2 b_3 b_4

Hashing shorter neighbor list:

5	11	12	3	
	1	2		

↓
Depth

Bucket b_0 b_1 b_2 b_3 b_4




Less collision



Faster intersection

Scaling Triangle Counting



Edge List

(8,1) (8,2)
(7,3) (7,5)
(7,6) (7,8)



GPU 0

(9,7) (7,10)
(7,4) (11,8)
(7,6) (12,5)




GPU 1

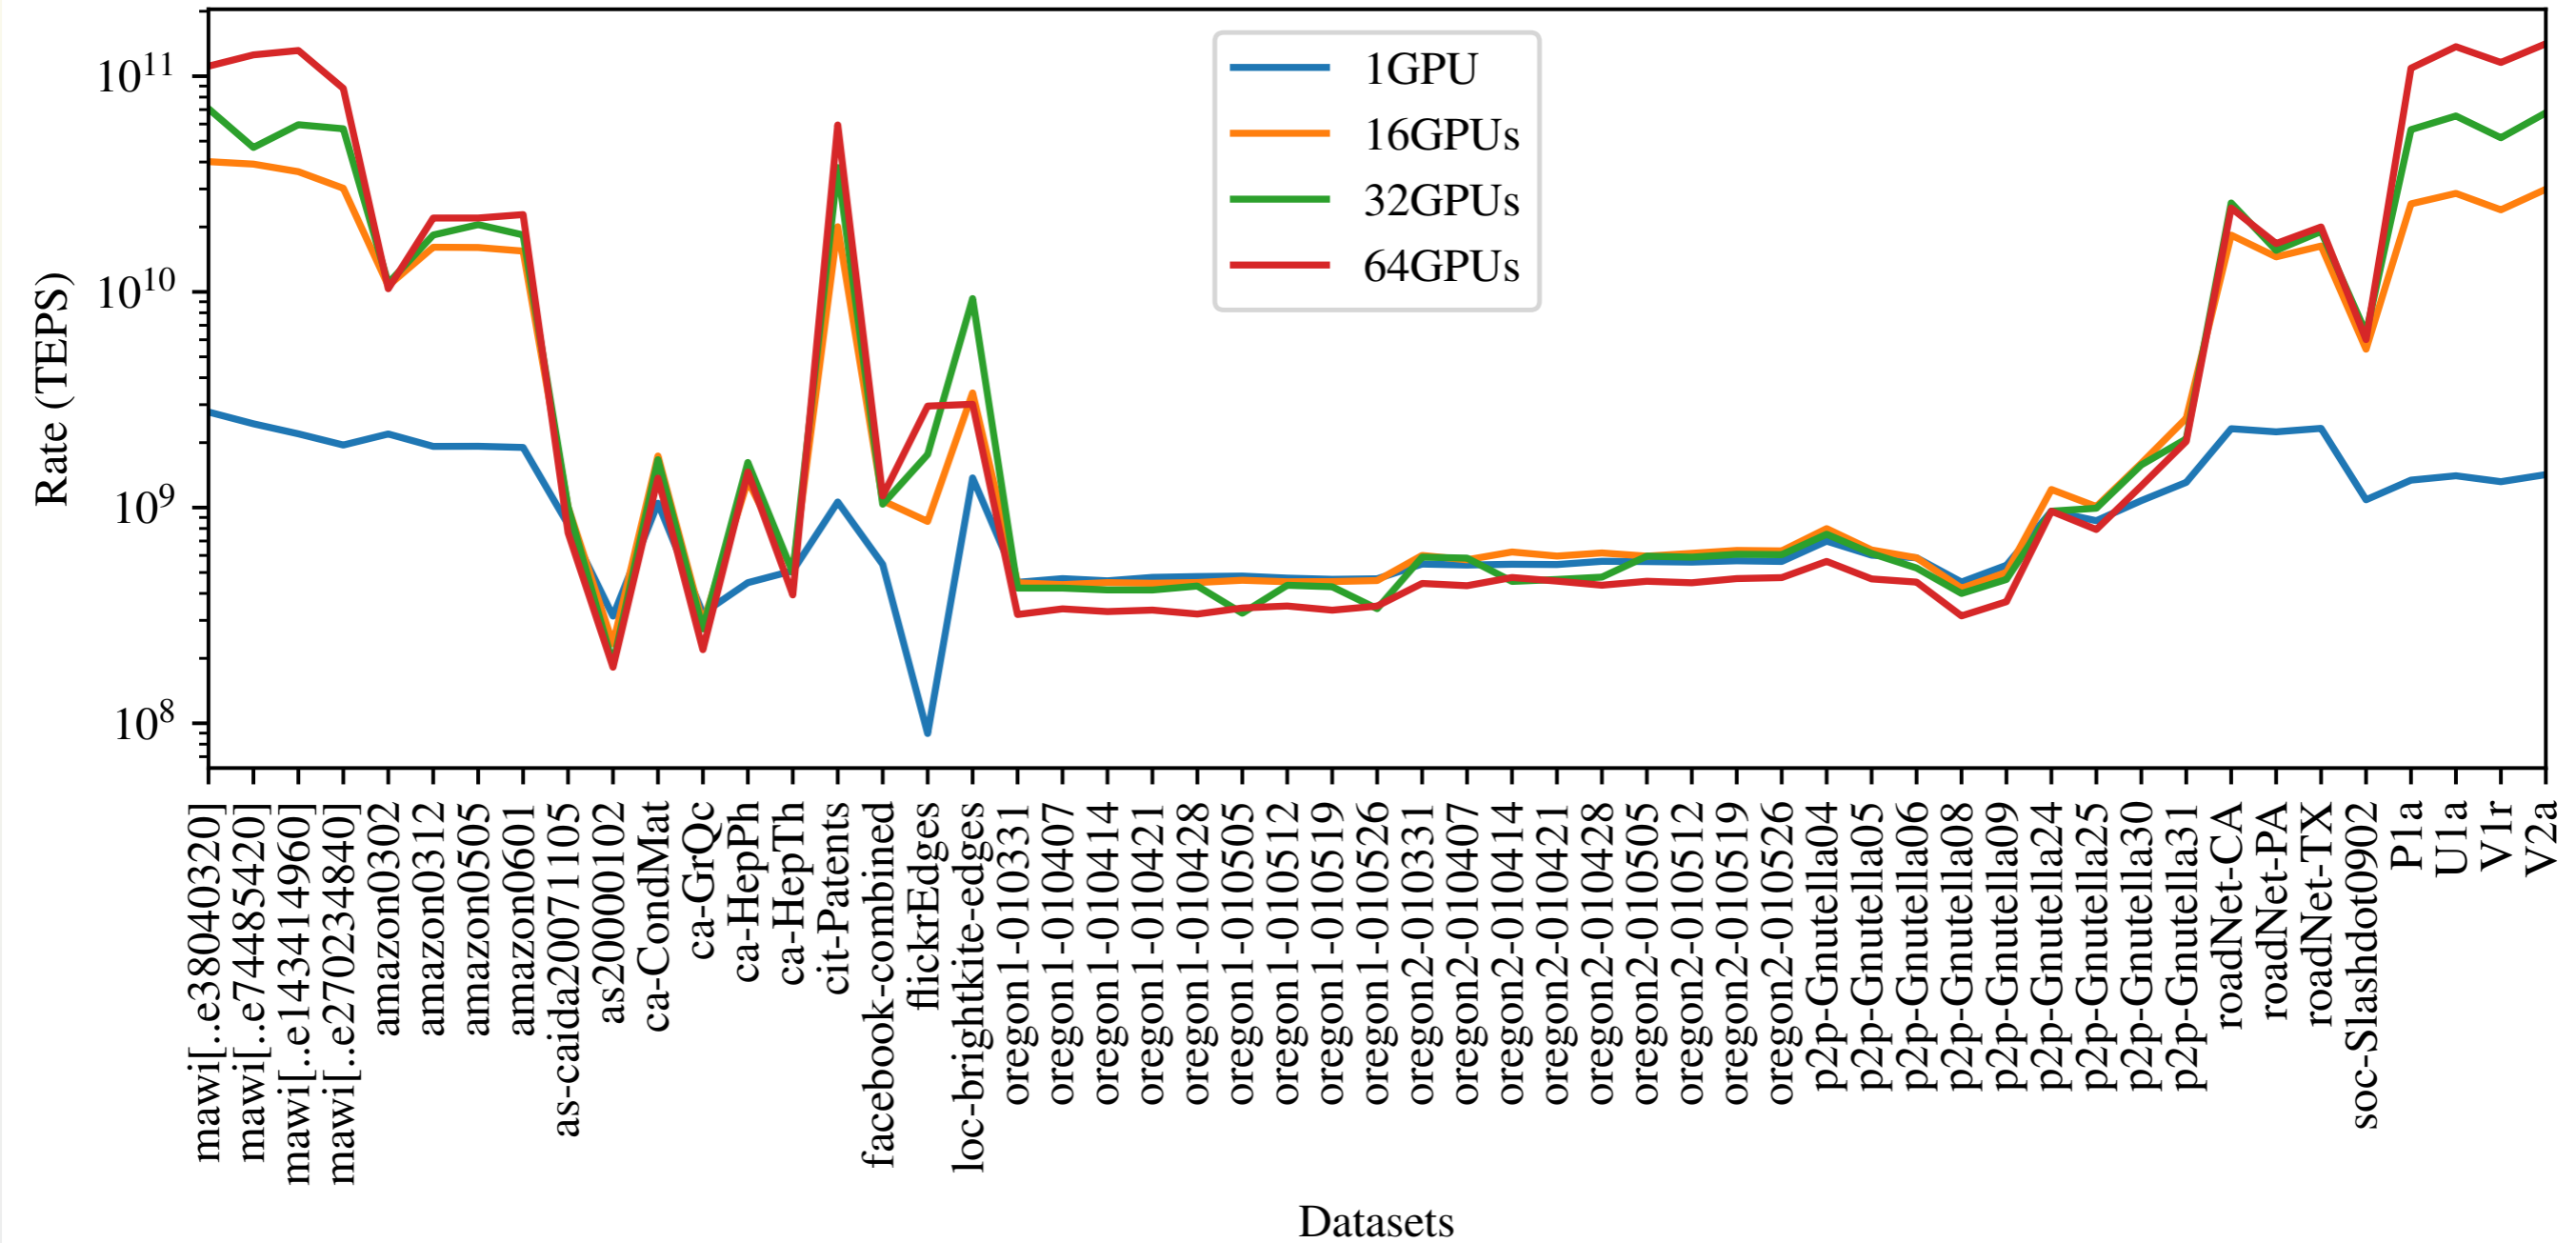
(8,12) (7,1)
(8,3) (5,8)
(6,9) (11,8)



GPU 2




H-Index Scalability



H-Index vs. State-of-the-art

Datasets	Vertices	Edges	Triangles	Rate (billion TEPS)		
				Tricore (2018 Champion)	A. Yaşar et. al. (2018 Champion)	H-INDEX
				8 x P100 GPUs	Skylake CPU	1 x V100 GPU
Amazon0302	262,111	899,792	717,719	1.46	-	2.19
Amazon0312	400,727	2,349,869	3,686,467	2.64	0.387	1.922
roadNet-PA	1,088,092	1,541,898	67,150	1.73	-	2.245
roadNet-TX	1,379,917	1,921,660	82,869	2.03	-	2.33
soc-Slashdot0902	82,168	504,230	602,592	0.793	0.15	1.09

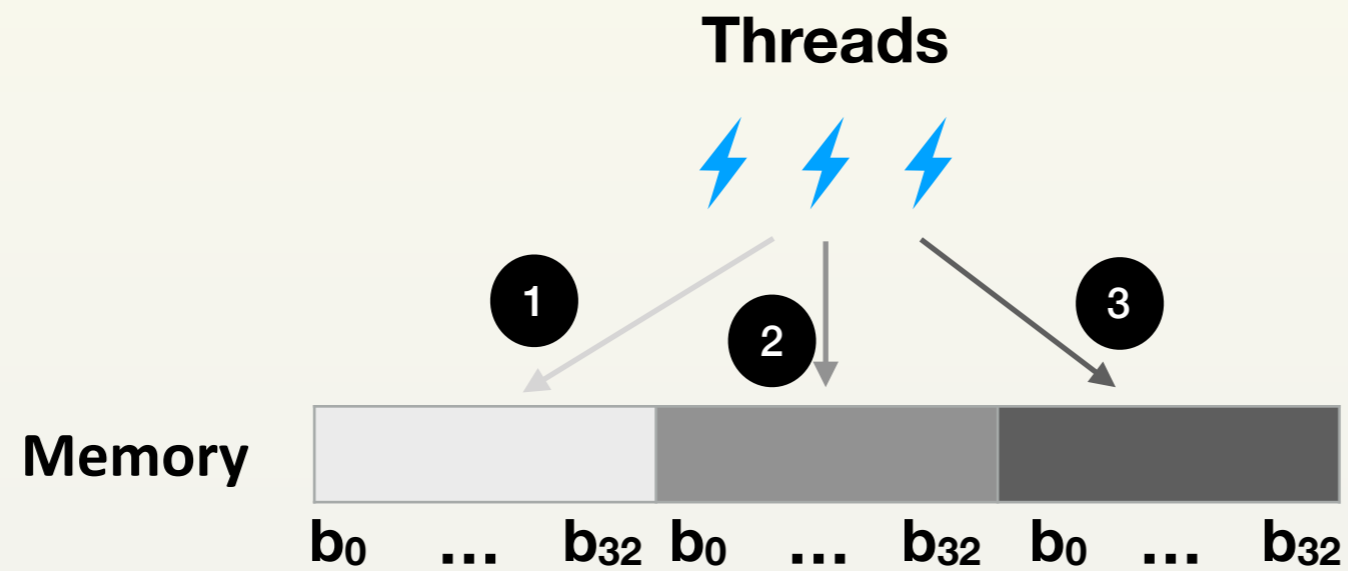
Acknowledgement




Thank you. Questions?


Reference Slides

More on coalesced memory access






Graph Orientation



Undirected Graph



Directed Graph