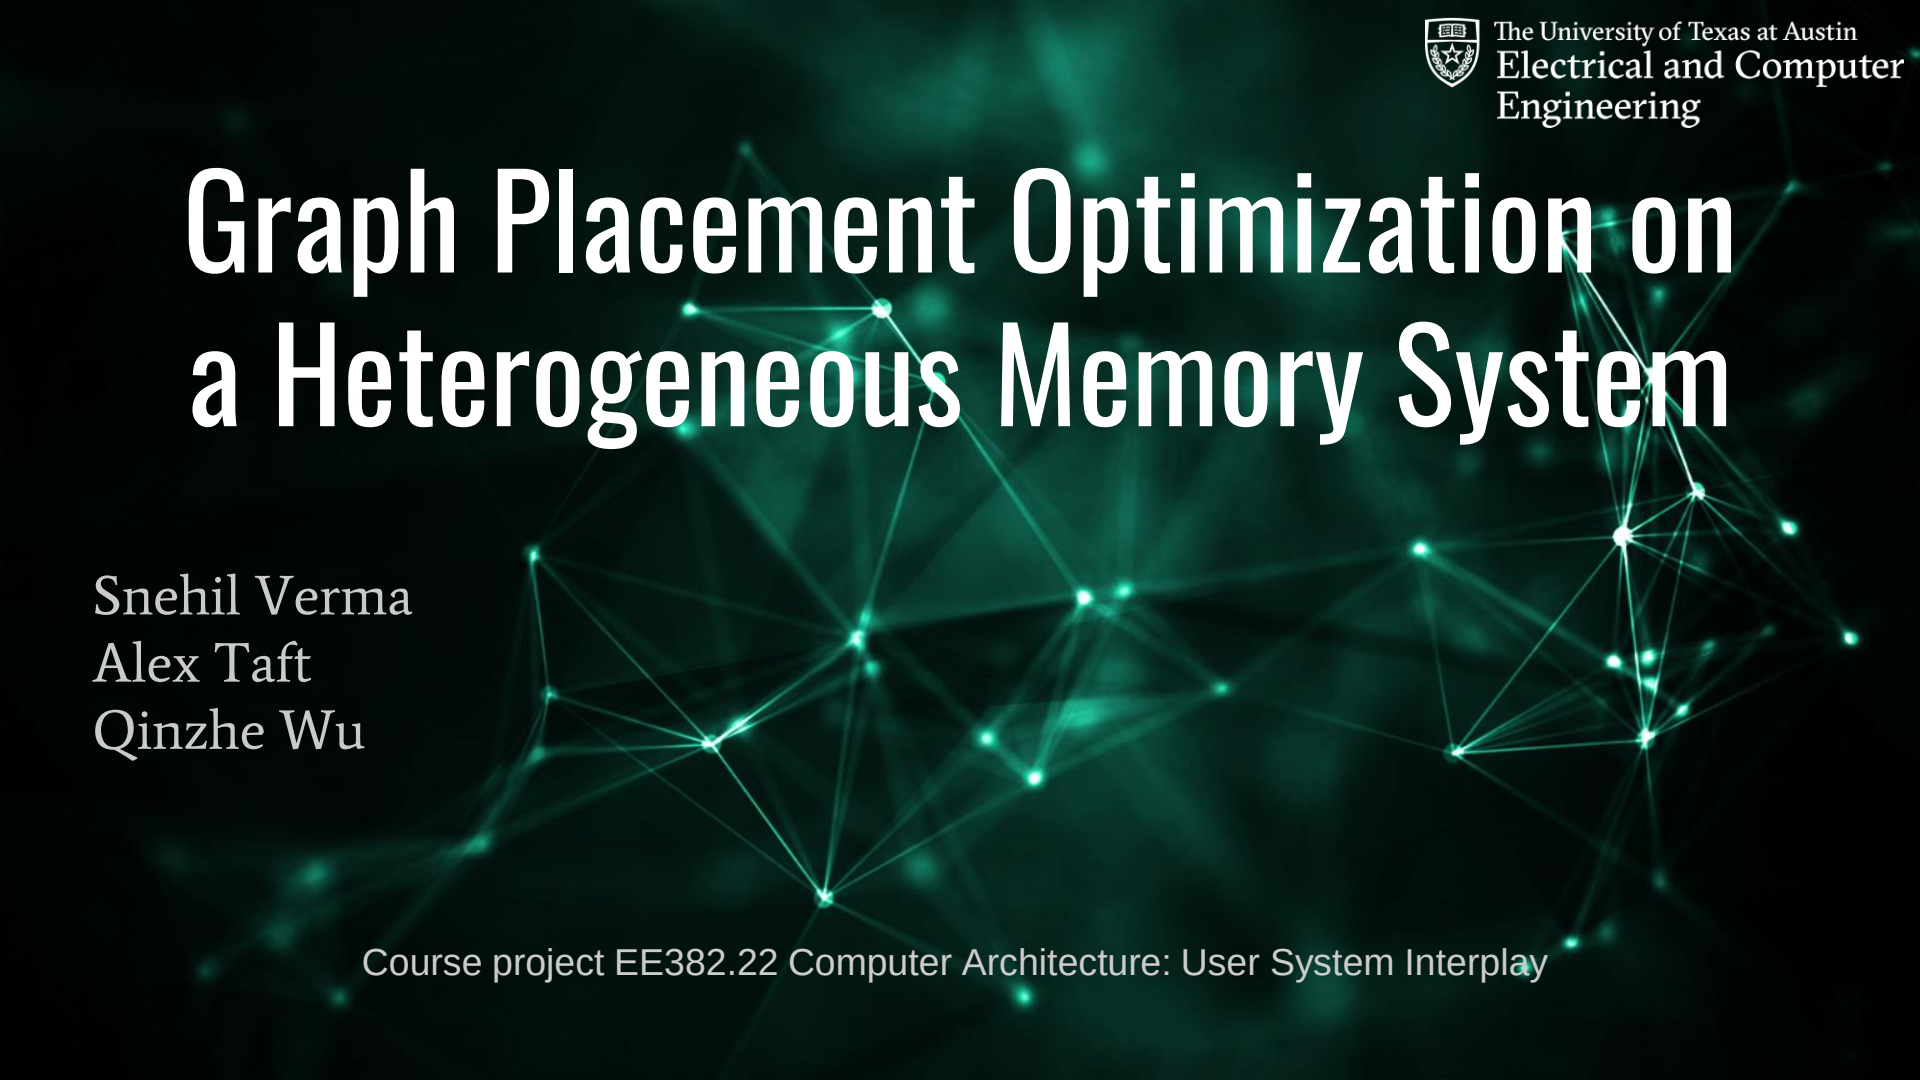
The University of Texas at Austin
Electrical and Computer Engineering
Graph Placement Optimization on a Heterogeneous Memory System
Snehil Verma
Alex Taft
Qinzhe Wu
Course project EE382.22 Computer Architecture: User System Interplay

Motivation

Dynamic Access Partitioning (DAP) [HPCA 2017]

- Sacrifice memory side cache hit rate
- Online profiling

Lightweight Profiler (ProfDP) [ICS 2018]

- Profile data structures offline
  - Size, Importance, Bandwidth and Latency sensitivity
- Assumes optimal placement is when all data is in fast memory

Goals

Analyze the topology of graph and make choice if vertex stored in HBM or DRAM

- Relabel Tool
  - Use existing shared memory framework (Ligra {PPoPP 13})
- No programmer code change
- No new hardware or online profiling

Users

- Graph Application on Heterogeneous Memory Systems
- Those that don’t want to modify existing code
- Still be able to tune performance if desired

- Additional users:
  - more complex placement strategies on top of what we propose

Placement Strategies

- Degree
  - Outgoing and incoming edges
- Frontier
  - Size of the frontiers in BFS
- Interleave
  - Alternate every BFS frontier
- Intraleave
  - Alternate vertices within a frontier
- Random
  - Skewed and unskewed

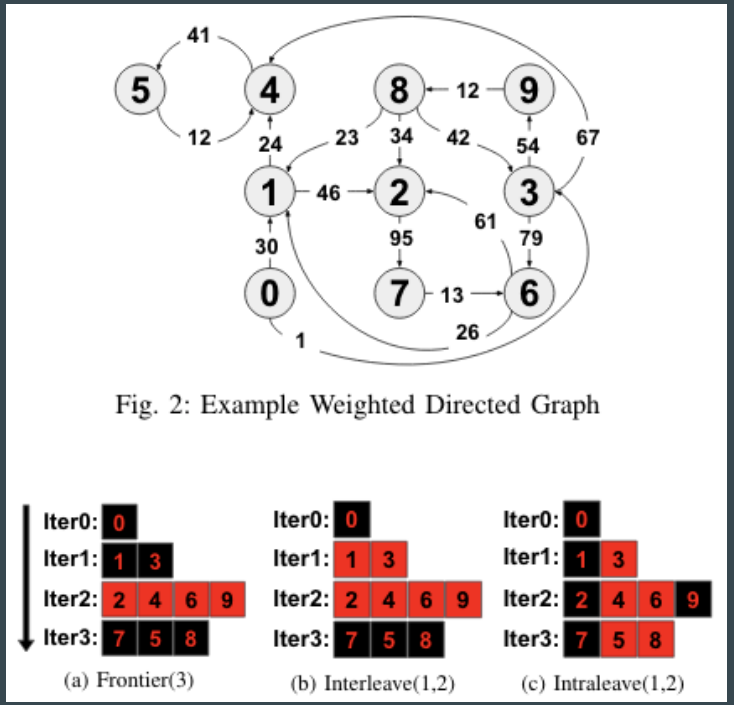

Fig. 2: Example Weighted Directed Graph

(a) Frontier(3) (b) Interleave(1,2) (c) Intraleave(1,2)

Evaluation

- Ligra Framework
  - Measure execution time
- 4 Graphs of varying Size/Connectivity
- 3 Applications
  - Breadth First Search (BFS)
  - Single Source Shortest Path (SSSP)
  - Connected Component (CC)
- All placement strategies tried with varied parameters
- Used Intel KNL at TACC
- Baselines were all in HBM, DRAM, or HBM as a Cache
- Each experiment repeated 50 times and reported the geometric mean

Abbr.	Name	$\|V\|$	$\|E\|$
TX	Road network Texas	1,393,383	3,843,320
OK	Orkut on-line social network	3,072,627	117,185,083
TW	Twitter (MPI)	52,579,683	1,963,263,821
RMAT	Synthesized RMAT graph	400,000,000	2,000,000,000

	Small Size (<= 16 GB)	**Large Size** (> 16 GB)
Small Degree (<= 10)	Road TX	RMAT2B
Large Degree (> 10)	Orkut	Twitter

Some Interesting Observations

BFS on RoadTX using Degree

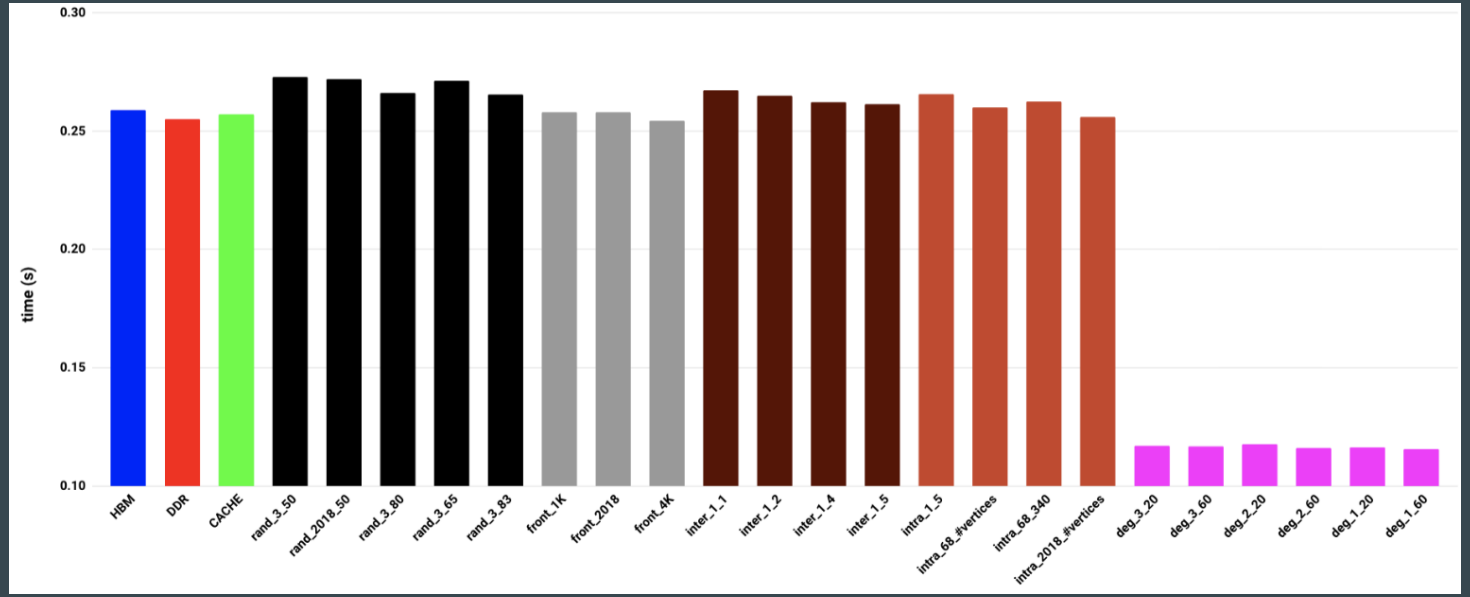

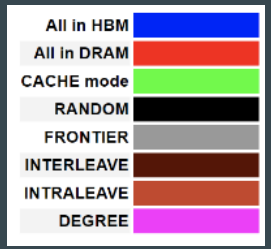

- 2x improvement over any other strategy
- but...

Degree (cont.)

- Not so great on some other graphs/applications

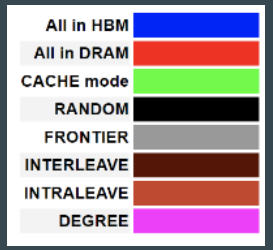

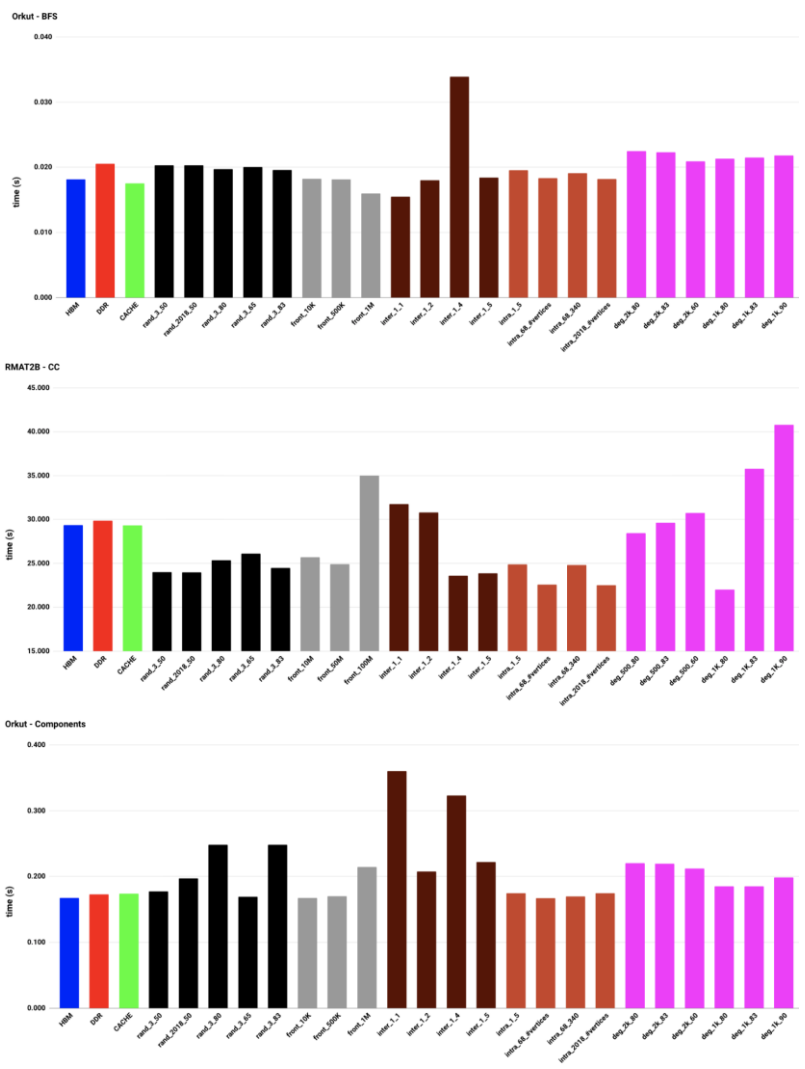

Random

- Randomly selecting provides improvement about as good as any intelligently chosen strategy
  - For graphs with large size

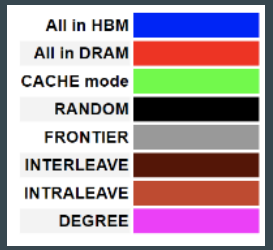

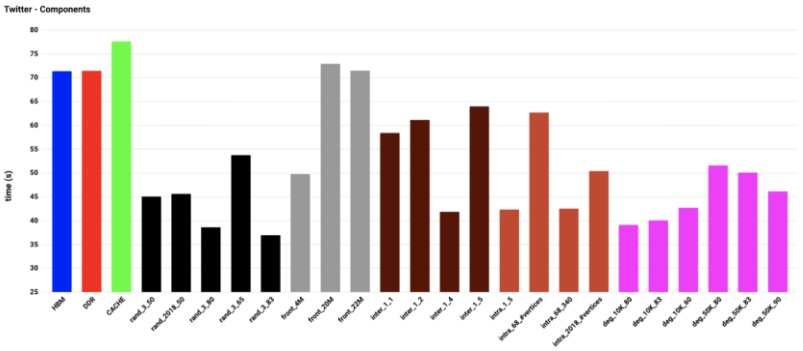

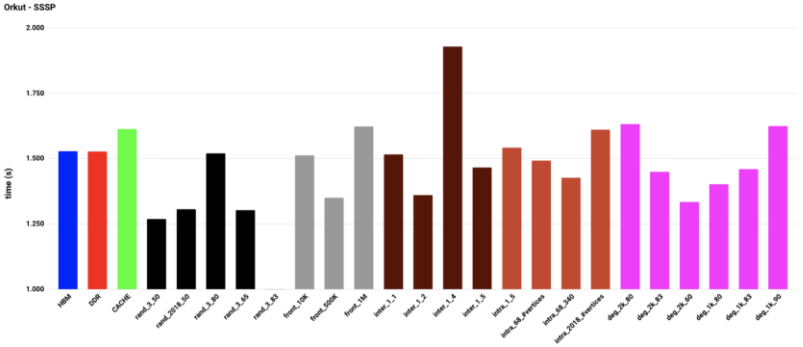

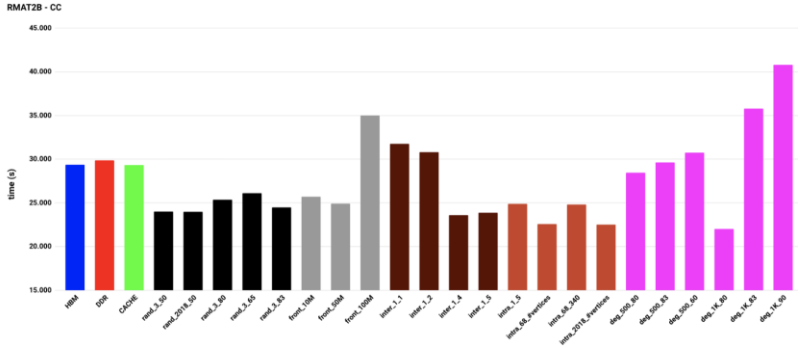

Limitations and Future Work

- Memory Hierarchy not taken in consideration
- Capacity Constraints of HBM (16GB)

- Use as building block for more involved strategy
- Comparison with other works

Conclusions

- Performance improves via data partitioning
  - Achieves better speedup than placing data only in one type of the memory
- Provided a tool to relabel graph in heterogeneous memory system
  - Integrated into Ligra framework
  - No changes in code required
  - Knobs for the programmers
- No one strategy performs the best across all graphs/applications

Thank You

Questions?

Backup

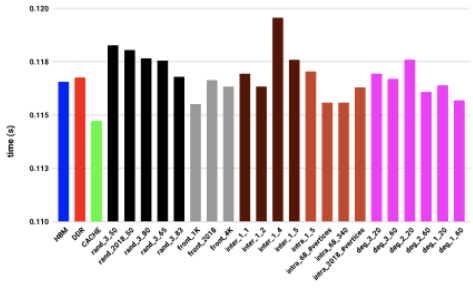

(a) Road TX BFS

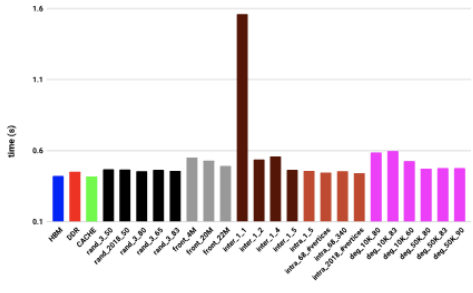

(b) Twitter BFS

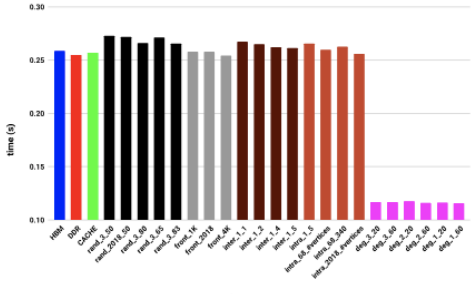

(c) Road TX SSSP

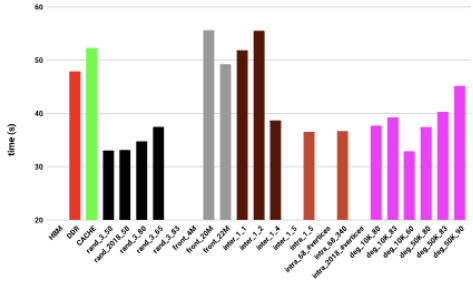

(d) Twitter SSSP

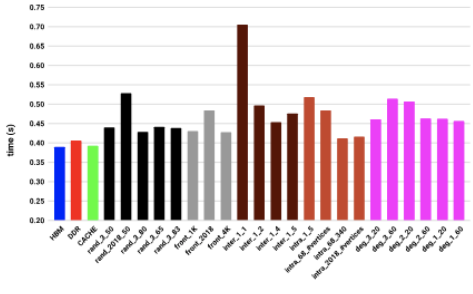

(e) Road TX CC

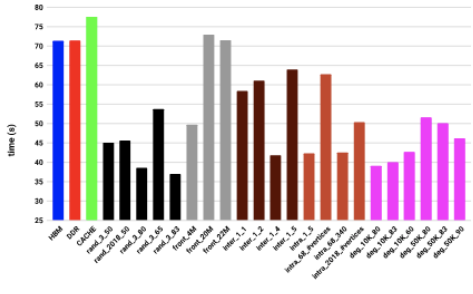

(f) Twitter CC

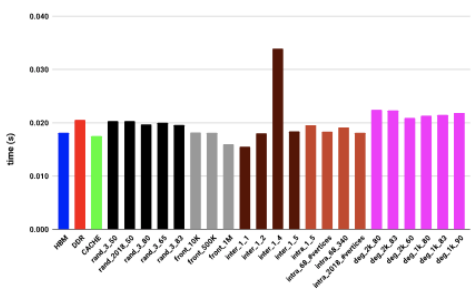

(a) Orkut BFS

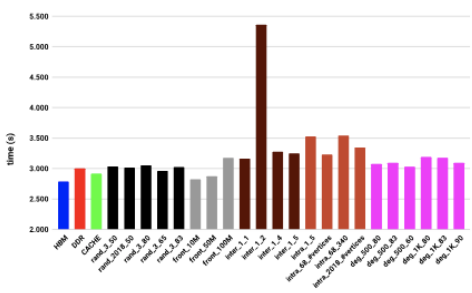

(b) RMAT2B BFS

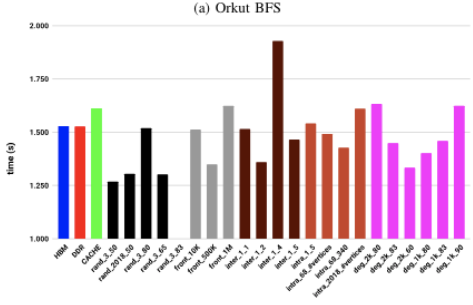

(c) Orkut SSSP

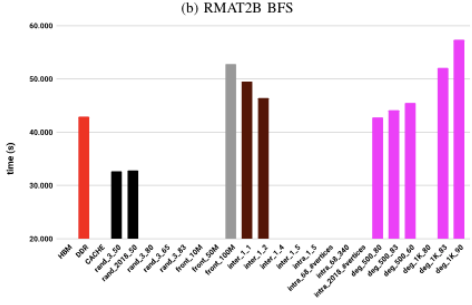

(d) RMAT2B SSSP

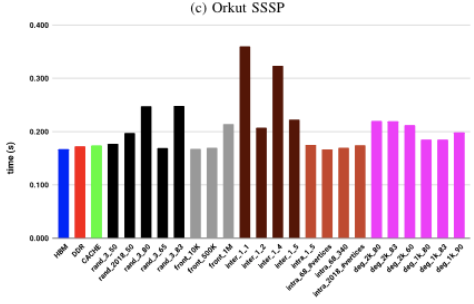

(e) Orkut CC

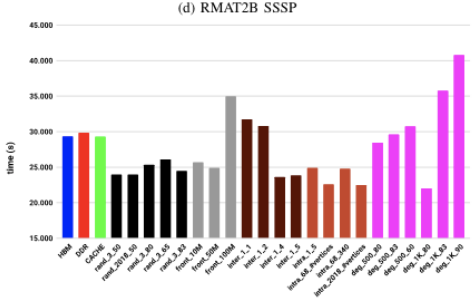

(f) RMAT2B CC