Lehrstuhl für Datenbanksysteme
Fakultät für Informatik
Technische Universität München

Breadth-First-Search

Seminar: Implementierungtechniken für Hauptspeicherdatenbanksysteme

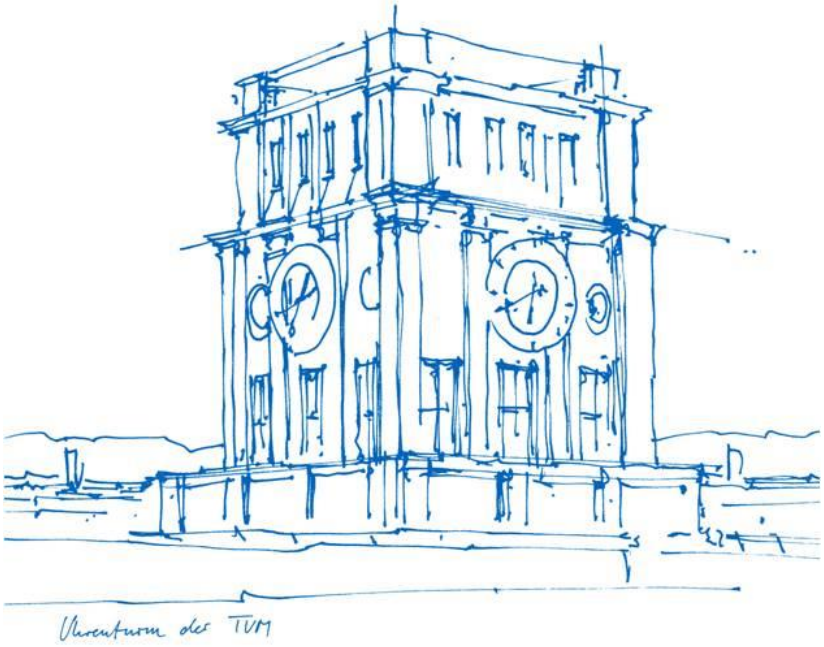

Mahrad Zoonemat Kermani

München, 14. November 2018

Motivation

- Finding neighbors
- Distance calculations
- Shortest path / Minimum spanning tree
  - All-pair shortest path (APSP)

Small-world networks

- One big connected graph
- Most nodes are not neighbors
- Distance between all the pairs is small (degrees of separation)
  - Six degrees of Kevin Bacon

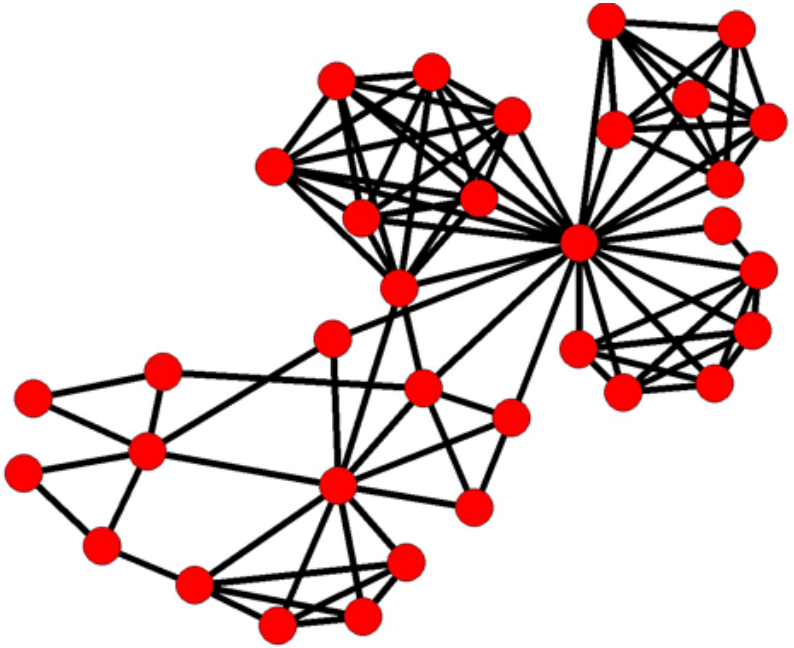

http://archive.news.ku.edu/2008/february/5/language.shtml

Naïve BFS on subsection of Small-World

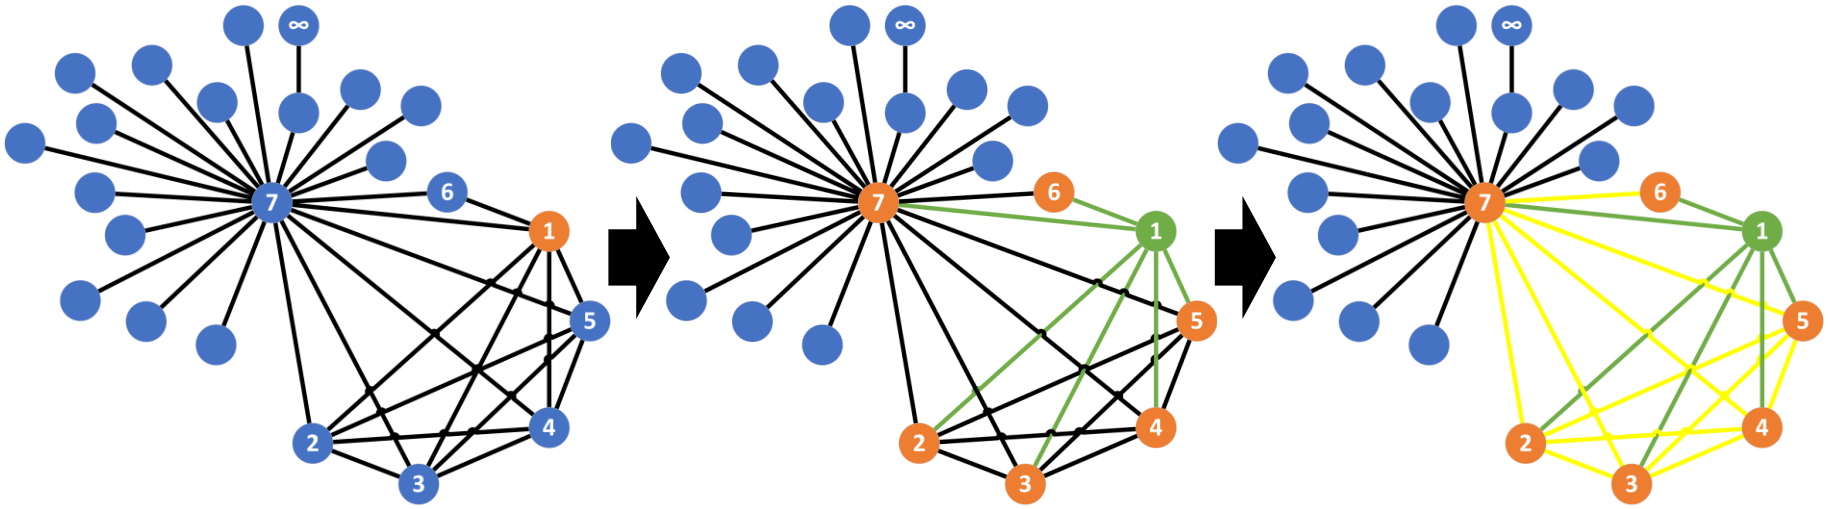

Pop back

Push front

Not visited

Not visited

MST

Ignored

Frontier out edges

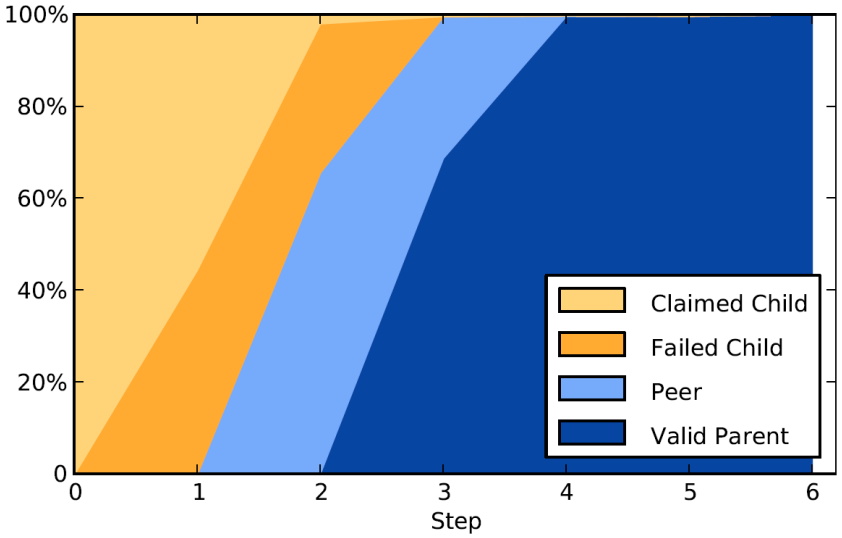

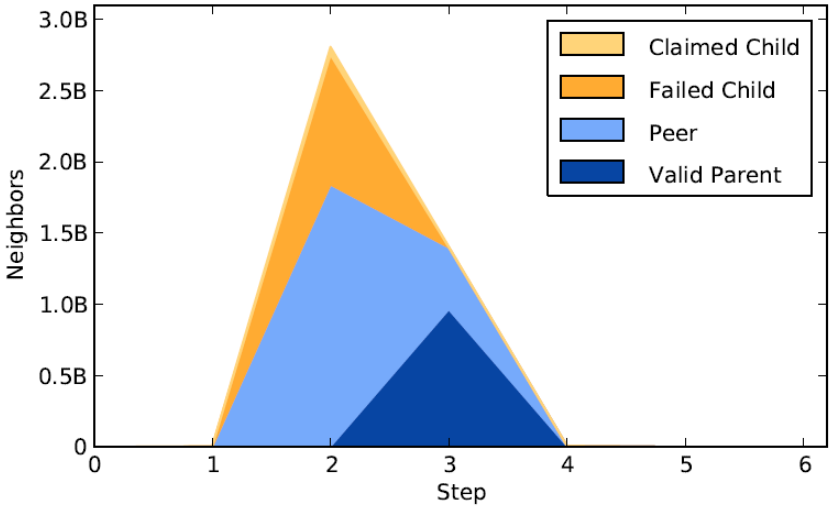

Beamer et al. [1]

Direction-Optimized BFS

- Traverse bottom-up once a threshold is hit
  - Frontier edges > remaining edges / α
- Bottom-up?
  - Iterate through the unvisited
  - Add nodes that are neighbors of frontiers

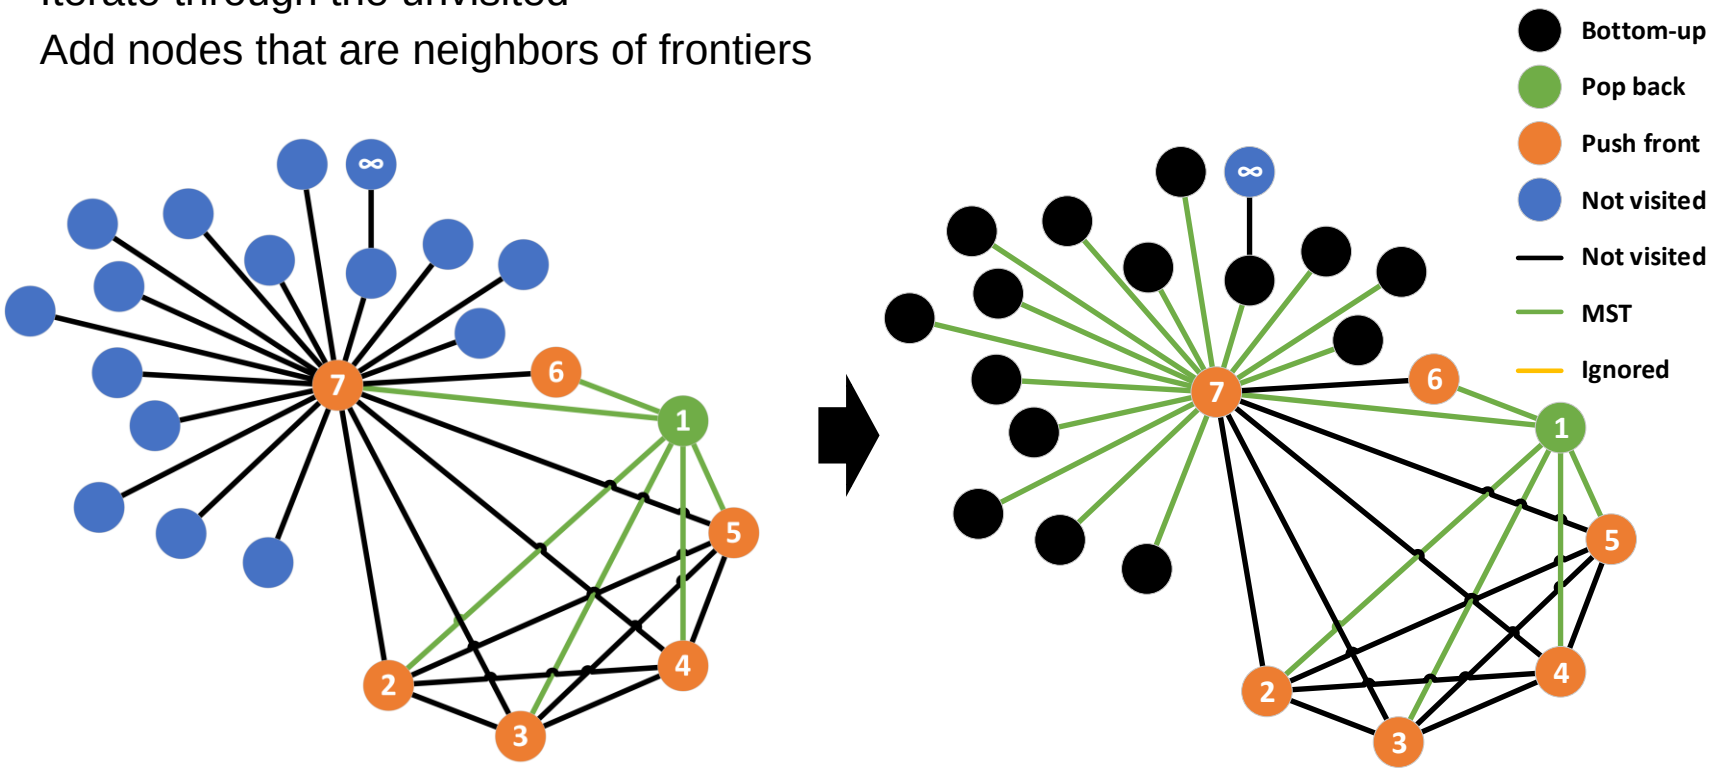

MS-BFS

- Maintain three n×k bitarrays
  - next (Black)
  - seen (~Blue)
  - frontier (Orange)
  - k: number of BFS instances
  - n: vertices
- Top-down:
  - Iterate frontier
  - Mark unseen neighbors as next
  - Mark next set bits as seen
- Bottom-up:
  - Iterate seen (unseen)
  - Mark next if unseen and neighbor is frontier
  - Mark next set bits as seen

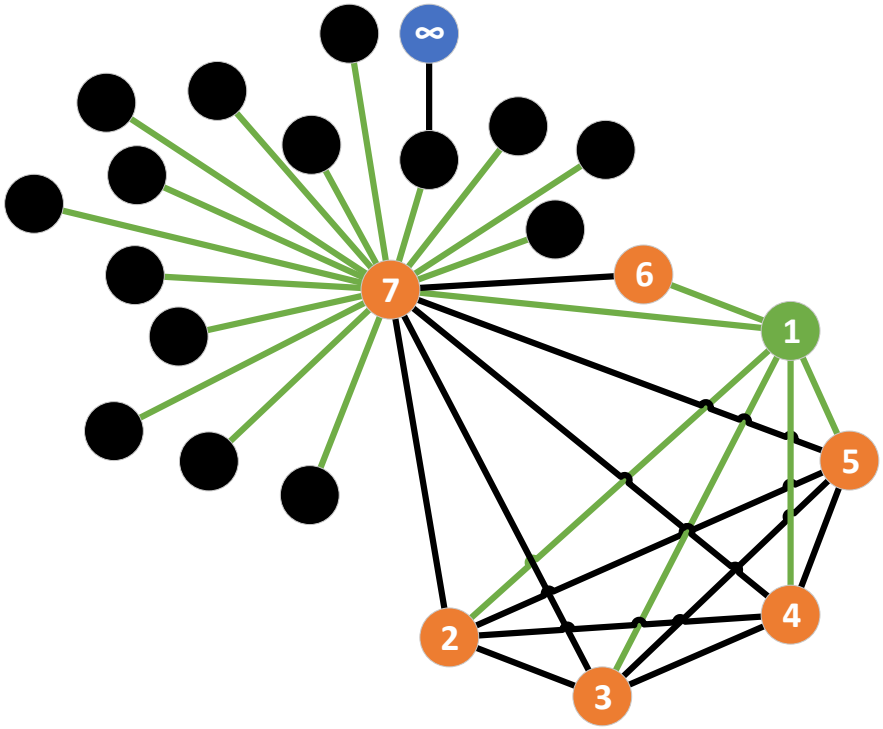

MS-BFS – Example

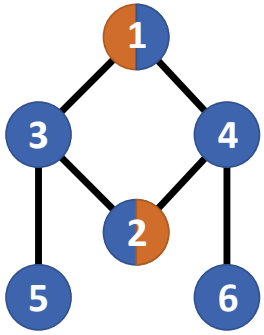

init

frontier		seen		next	
X		X			
	X		X		

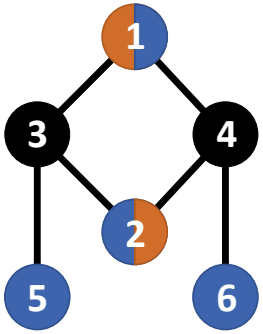

step1

frontier		seen		next	
X		X			
	X		X		
		X	X	X	X
		X	X	X	X

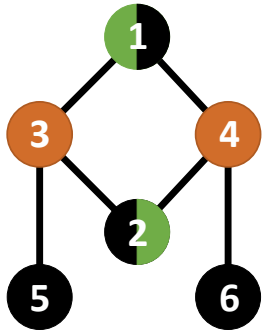

step2

frontier		seen		next	
		X	X		X
		X	X	X	
X	X	X	X		
X	X	X	X		
		X	X	X	X
		X	X	X	X

Then et al. [2]

MS-BFS Memory Usage

N	w	P	BFSs	MS-BFS Memory	Edges	Graph Memory
2^{20}	64	1	64	24 MB	16	128 MB
2^{20}	64	6	384	144 MB	16	128 MB
2^{20}	64	10	640	240 MB	48	384 MB
2^{20}	64	16	1024	384 MB	48	384 MB
2^{20}	512	10	5120	1.9 GB	48	384 MB
$50x2^{20}$	512	10	5120	93.8 GB	48	6.3 GB

N	number of vertices
w	word size
P	parallel runs (threads)
BFSs	simultaneous BFS runs
Edges	number of edges per vertex (Graph500 benchmark, Kronecker graph)

MS-PBFS

- Partition the vertices into subsets
- Split operations into random access and sequential access
  - section 1: two iterators
  - section 2: one iterator (no sync required)
- No synchronization for sequential
- Synchronize the random access
- Lock-free array updates for random access

MS-PBFS II

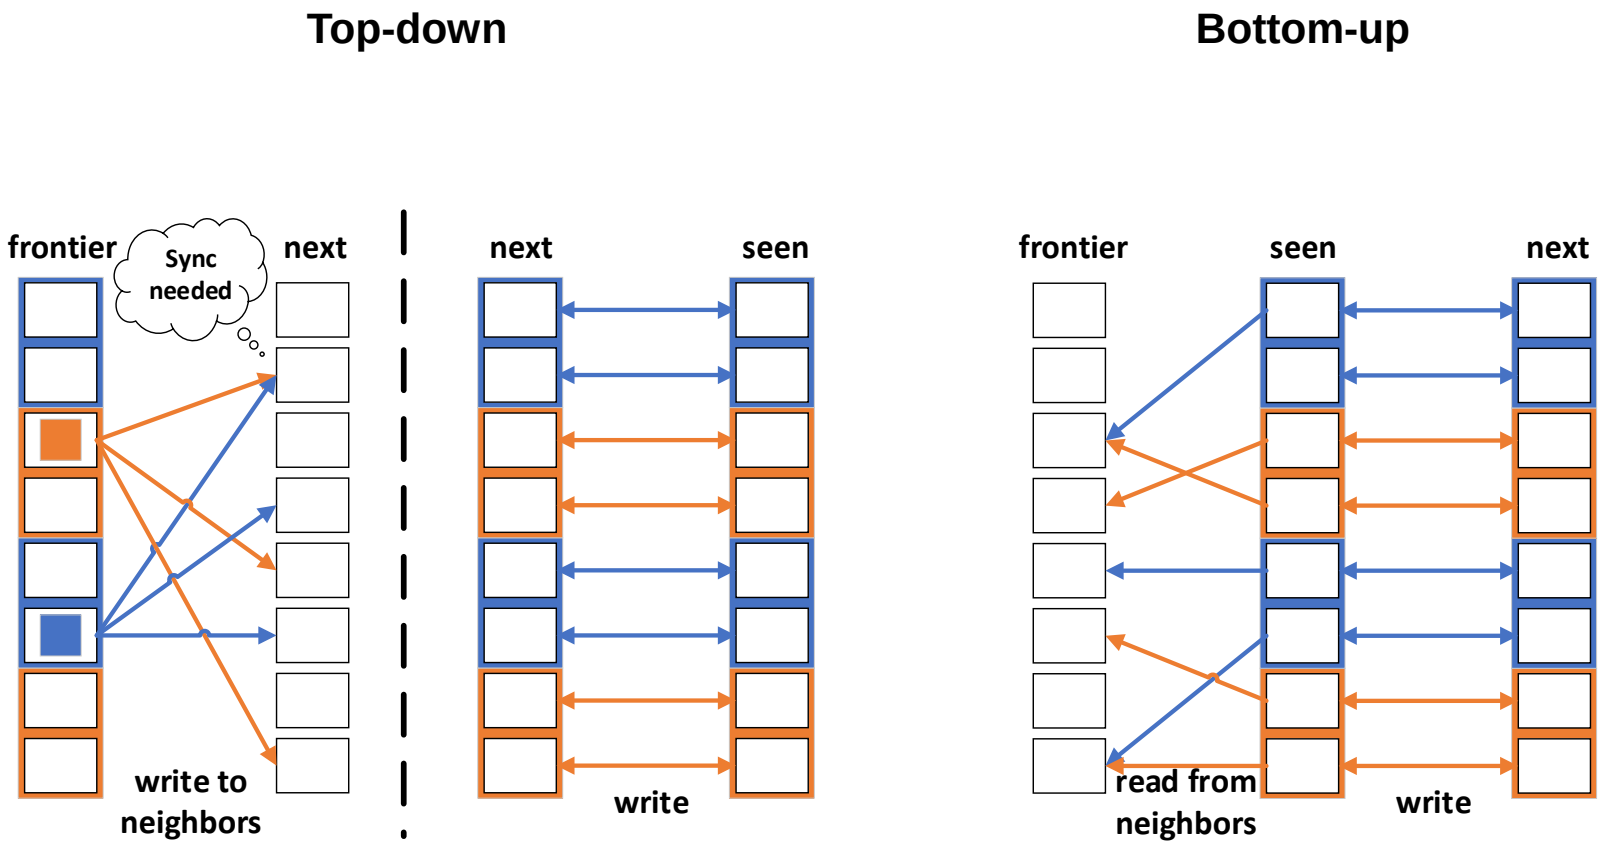

Kaufmann et al. [3]

Evaluation – LDBC

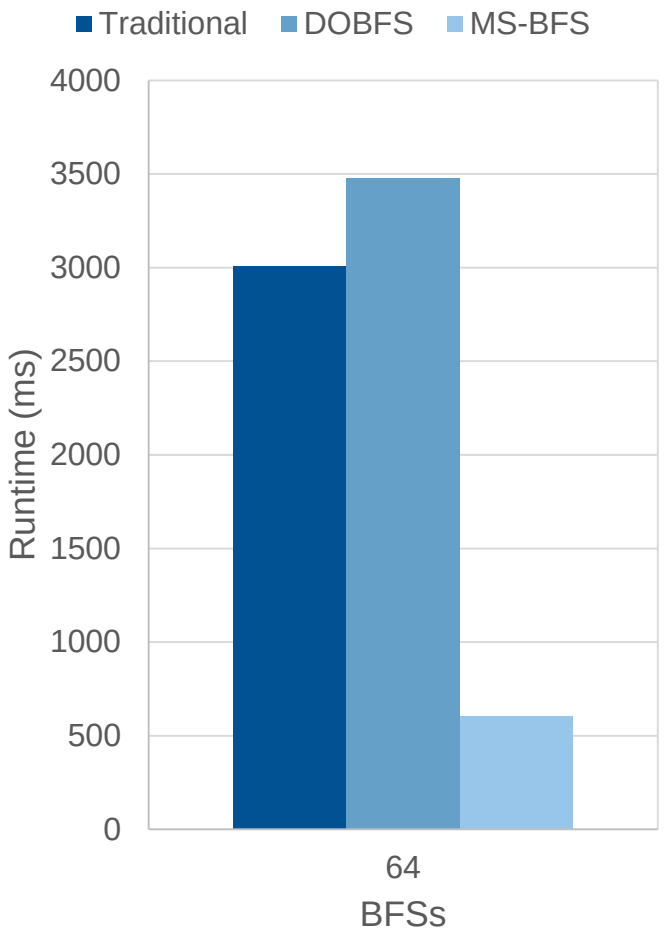

Parameter	LDBC*
Vertices	431157
Edge factor	46
Self-loops	No
Multiple-edges	No
Processor	Core^{2}Duo T9400
Cores	2
Memory	4 GB

Original Evaluation

Scale Factor	MS-PBFS (GTEPS)	MS-BFS (GTEPS)	MS-BFS 64 (GTEPS)
100	267	76.6	3.01
1000	118	45.5	1.30

Kaufmann et al. [3]

Sources

[1] Beamer, Scott; Asanovic, Krste; Patterson, David (2012). “Direction-Optimizing Breadth-First Search”

[2] Then, Manuel; Kaufmann, Moritz; Chirigati, Fernando; Hoang-Vu, Tuan-Anh; Pham, Kien; Kemper, Alfons; Neumann, Thomas; T. Vo, Huy (2015). “The More the Merrier: Efficient Multi-Source Graph Traversal”

[3] Kaufmann, Moritz; Then, Manuel; Kemper, Alfons; Neumann, Thomas (2017). “”Parallel Array-Based Single- and Multi-Source Breadth First Searches on Large Dense Graphs”

[5] https://github.com/sbeamer/gapbs/

[6] https://github.com/mtodat/ms-bfs/

[7] https://www.cc.gatech.edu/dimacs10/archive/kronecker.shtml

[8] https://github.com/ldbc/ldbc_snb_datagen