

THE MORE THE MERRIER: EFFICIENT MULTI-SOURCE GRAPH TRAVERSAL

Manuel Then*, Moritz Kaufmann*, Fernando Chirigati†, Tuan-Anh Hoang-Vu† ,
Kien Pham†, Huy T. Vo†, Alfons Kemper*, Thomas Neumann*

*Technische Universität München, †New York University

Published at VLDB 2015

Presented by Victor Ying

6.886 – February 23, 2021

Background

- Some applications do many BFS traversals (from different starting nodes) on one graph
 - *E.g., compute centrality metrics across graph*
- Prior work: parallel BFS, direction-optimizing BFS
 - *Speed up a single BFS traversal*

- Graph traversals have poor cache behavior

```
-   while visit ≠ ∅  
      for each v ∈ visit do  
        for each n ∈ neighborsv do  
          if n ∉ seen then  
            seen ← seen ∪ {n}  
            visitNext ← visitNext ∪ {n}  
            do BFS computation on n  
          visit ← visitNext  
          visitNext ← ∅
```

Finding neighbors in graph
is non-trivial expense

Read a single random bit in *seen* bitset
when traversing each edge

- Small-world phenomenon: **some** graphs have low diameter


Multi-Source BFS



- Concurrently **run many independent BFS traversals** on the same graph
 - 100s of BFSs on a single CPU core


Multi-Source BFS

- Concurrently **run many independent BFS traversals** on the same graph
 - 100s of BFSs on a single CPU core



Multi-Source BFS

- Concurrently **run many independent BFS traversals** on the same graph
 - 100s of BFSs on a single CPU core
- Store concurrent **BFSs state as 3 bitsets** per vertex



- Represent **BFS traversal as SIMD bit operations** on these bitsets

Multi-Source BFS (MS-BFS)


One round:

```
for  $i = 1, \dots, N$   
  if  $visit[v_i] = \mathbb{B}_\emptyset$ : skip  
  for each  $n \in neighbors[v_i]$   
     $\mathbb{D} \leftarrow visit[v_i] \ \& \ \sim seen[n]$   
    if  $\mathbb{D} \neq \mathbb{B}_\emptyset$   
       $visitNext[n] \leftarrow visitNext[n] \mid \mathbb{D}$   
       $seen[n] \leftarrow seen[n] \mid \mathbb{D}$ 
```

- Given frontiers, compute next frontiers
- By traversing **every** edge in the graph once.
- Cost of finding neighbors is amortized over several traversals

Multi-Source BFS - Example

Initial




	b_1	b_2
1	X	
2		X
3		
4		
5		
6		

visit

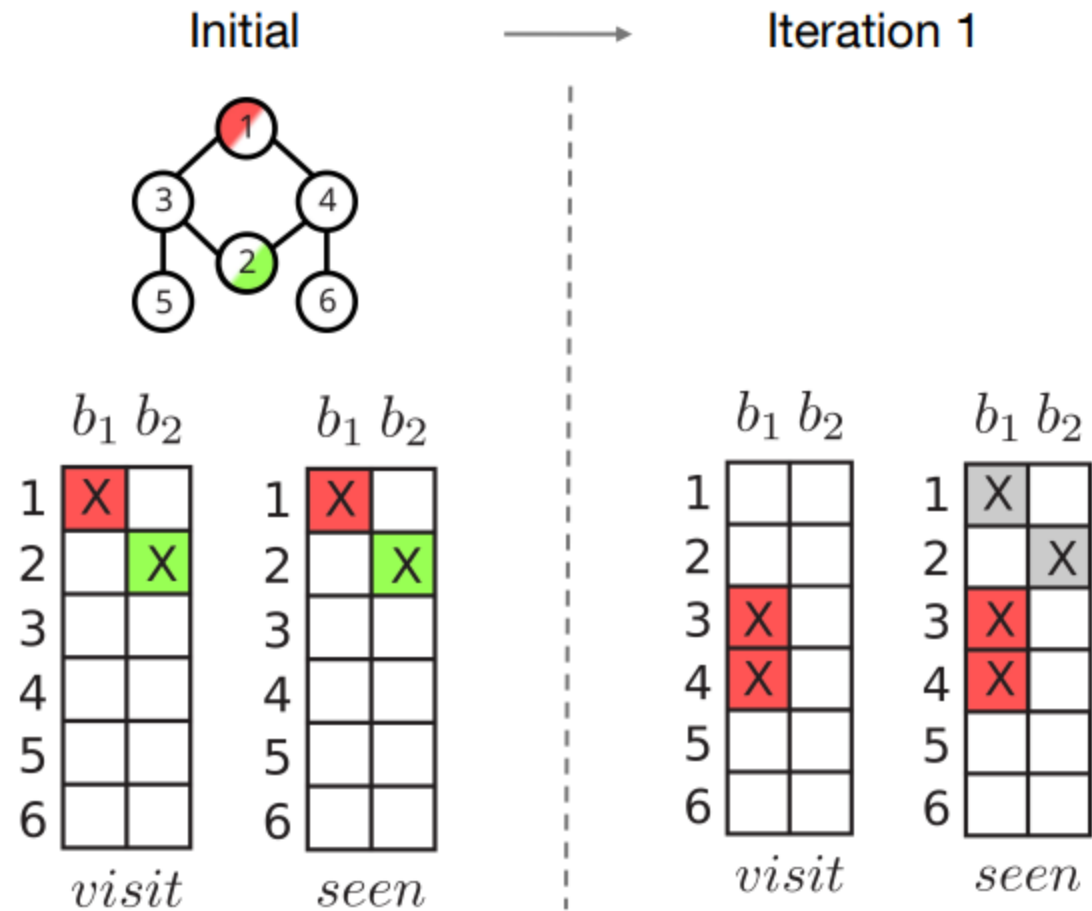
	b_1	b_2
1	X	
2		X
3		
4		
5		
6		

seen


Multi-Source BFS - Example



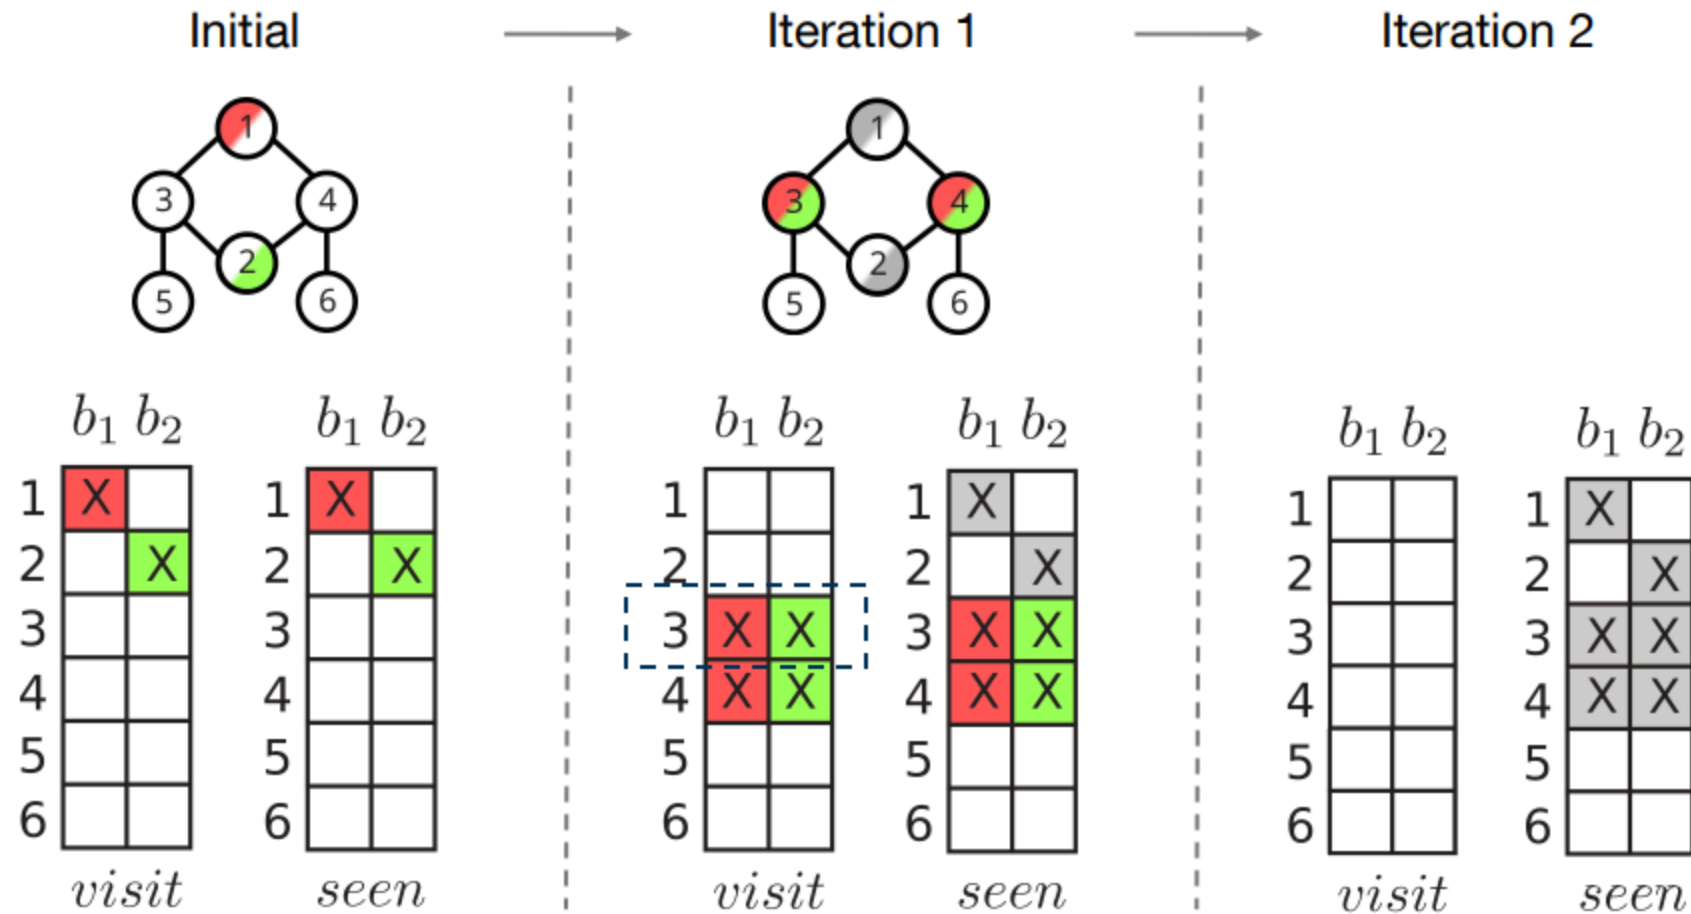
Multi-Source BFS - Example




Multi-Source BFS - Example



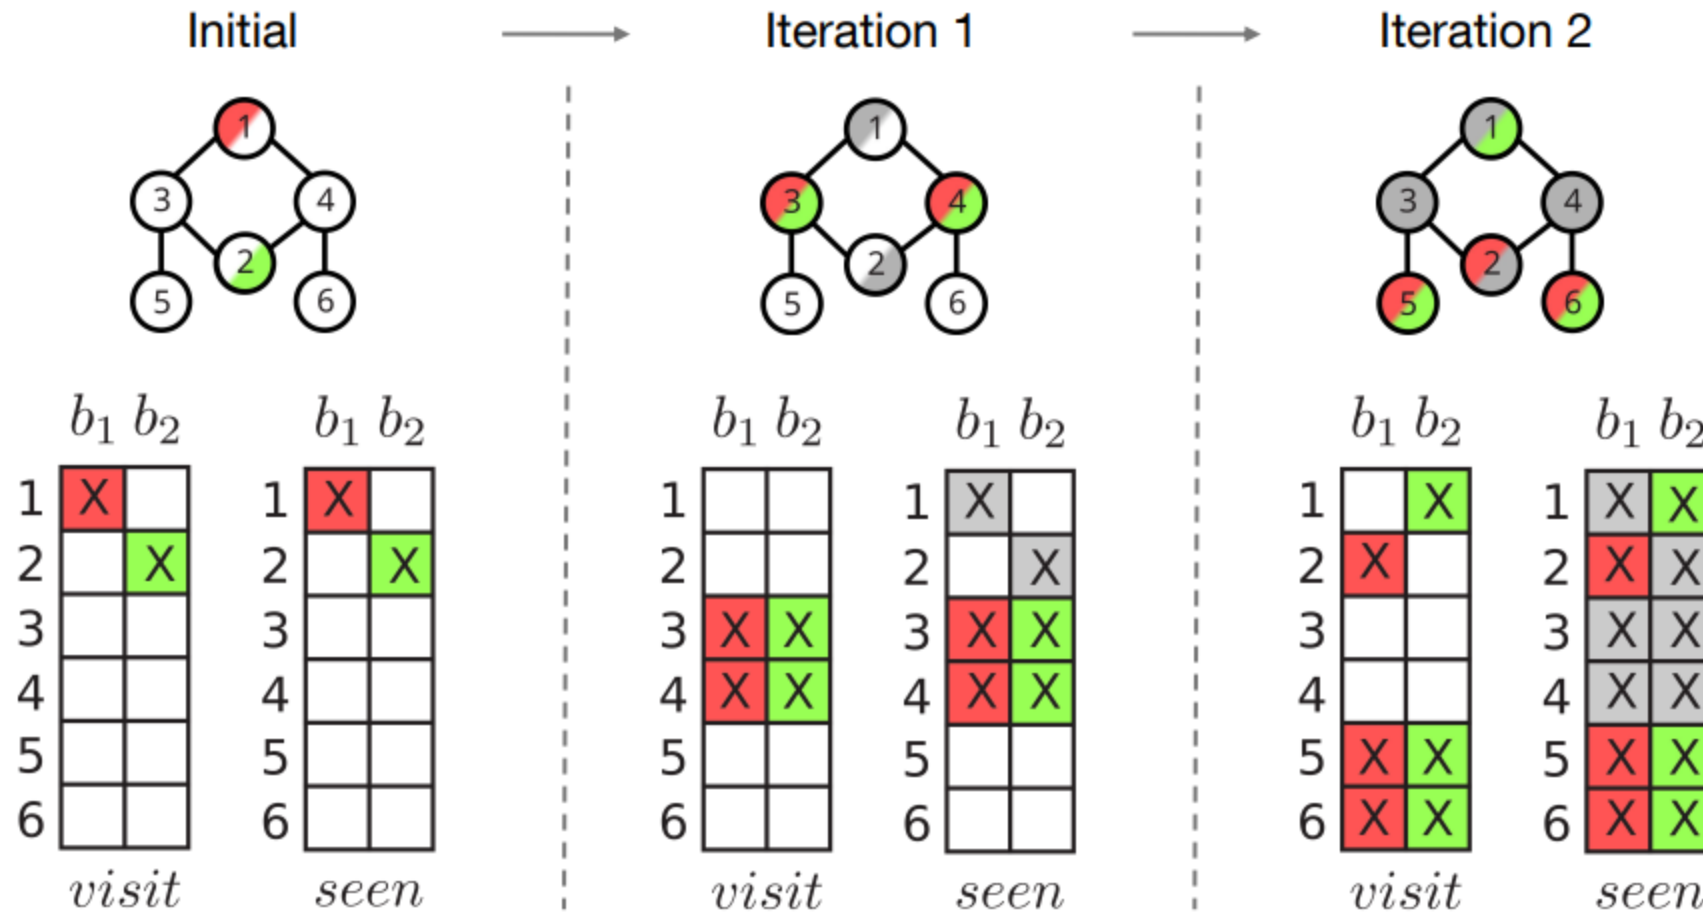
Multi-Source BFS - Example



Multi-Source BFS - Example



Multi-Source BFS - Example



MS-BFS work analysis


```
while  $visit \neq \emptyset$ 
  for  $i = 1, \dots, N$ 
    if  $visit[v_i] = \mathbb{B}_\emptyset$ , skip
    for each  $n \in neighbors[v_i]$ 
       $\mathbb{D} \leftarrow visit[v_i] \ \& \ \sim seen[n]$ 
      if  $\mathbb{D} \neq \mathbb{B}_\emptyset$ 
         $visitNext[n] \leftarrow visitNext[n] \mid \mathbb{D}$ 
         $seen[n] \leftarrow seen[n] \mid \mathbb{D}$ 
        do BFS computation on  $n$ 
   $visit \leftarrow visitNext$ 
  reset  $visitNext$ 
```

- $O(n+m)$ work per round
- $O(\text{diameter})$ rounds needed.
- $O((n+m) \times \text{diameter})$ total work for ω traversals.
- Textbook BFS takes $O(n+m)$ for one traversal

How wide to make the bitvectors?

- Bitvector width (ω) = number of concurrent BFSs (per thread)
- Maximize SIMD parallelism by matching the width of largest registers?
- Wider, by using multiple registers?

Evaluation - The More the Merrier



MS-BFS improves cache performance

```
while  $visit \neq \emptyset$ 
  for  $i = 1, \dots, N$ 
    if  $visit[v_i] = \mathbb{B}_\emptyset$ , skip
    for each  $n \in neighbors[v_i]$ 
       $\mathbb{D} \leftarrow visit[v_i] \ \& \ \sim seen[n]$ 
      if  $\mathbb{D} \neq \mathbb{B}_\emptyset$ 
         $visitNext[n] \leftarrow visitNext[n] \ | \ \mathbb{D}$ 
         $seen[n] \leftarrow seen[n] \ | \ \mathbb{D}$ 
        do BFS computation on  $n$ 
   $visit \leftarrow visitNext$ 
  reset  $visitNext$ 
```

- Each cache line in $seen[]$ accessed once per adjacent edge
- Many concurrent BFSs amortize cost of cache line movement.
- Most expensive line is still this random access to $seen[]$

MS-BFS with “aggregated neighbor processing”

```
while  $visit \neq \emptyset$   
  for  $i = 1, \dots, N$   
    if  $visit[v_i] = \mathbb{B}_\emptyset$ , skip  
    for each  $n \in neighbors[v_i]$   
       $visitNext[n] \leftarrow visitNext[n] \mid visit[v_i]$ 
```

```
for  $i = 1, \dots, N$   
  if  $visitNext[v_i] = \mathbb{B}_\emptyset$ , skip  
   $visitNext[v_i] \leftarrow visitNext[v_i] \& \sim seen[v_i]$   
   $seen[v_i] \leftarrow seen[v_i] \mid visitNext[v_i]$   
  if  $visitNext[v_i] \neq \mathbb{B}_\emptyset$   
    do BFS computation on  $v_i$ 
```

```
 $visit \leftarrow visitNext$   
reset  $visitNext$ 
```

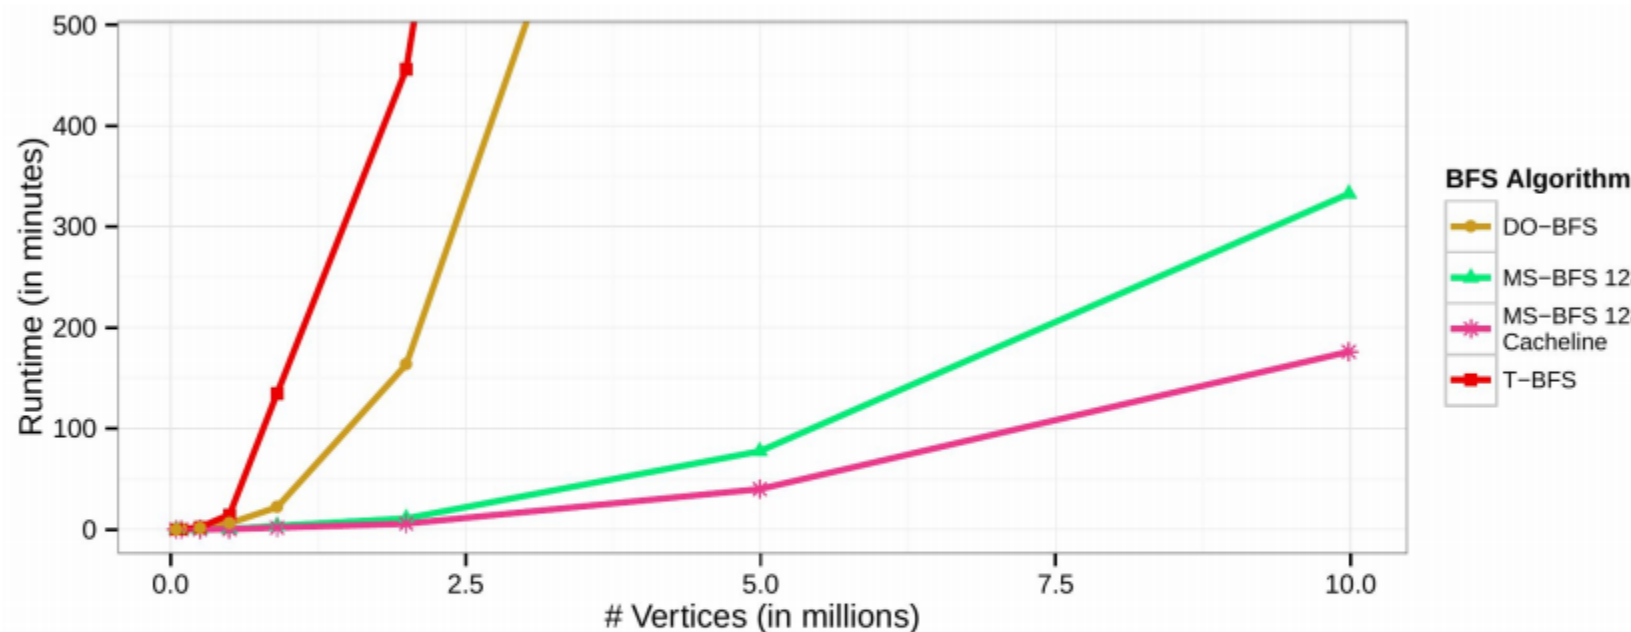
- Defer accesses to $seen[]$, and then do the accesses in scanning fashion, so each entry in $seen[]$ is accessed at most once

MS-BFS: further improvements

- Direction-optimizing
- Explicit prefetching
- Heuristics to decide what groups of BFS transversals to run together

Evaluation

- MS-BFS-based closeness centrality. 4x Intel Xeon E7-4870v2, 1TB



Graph	MS-BFS	Speedup over	
		T-BFS	DO-BFS
LDBC 1M	0:02h	73.8x	12.1x
LDBC 10M	2:56h	88.5x	28.7x
Wikipedia	0:26h	75.4x	29.5x
Twitter (1M)	2:52h	54.6x	12.7x

Multi-core scalability?!

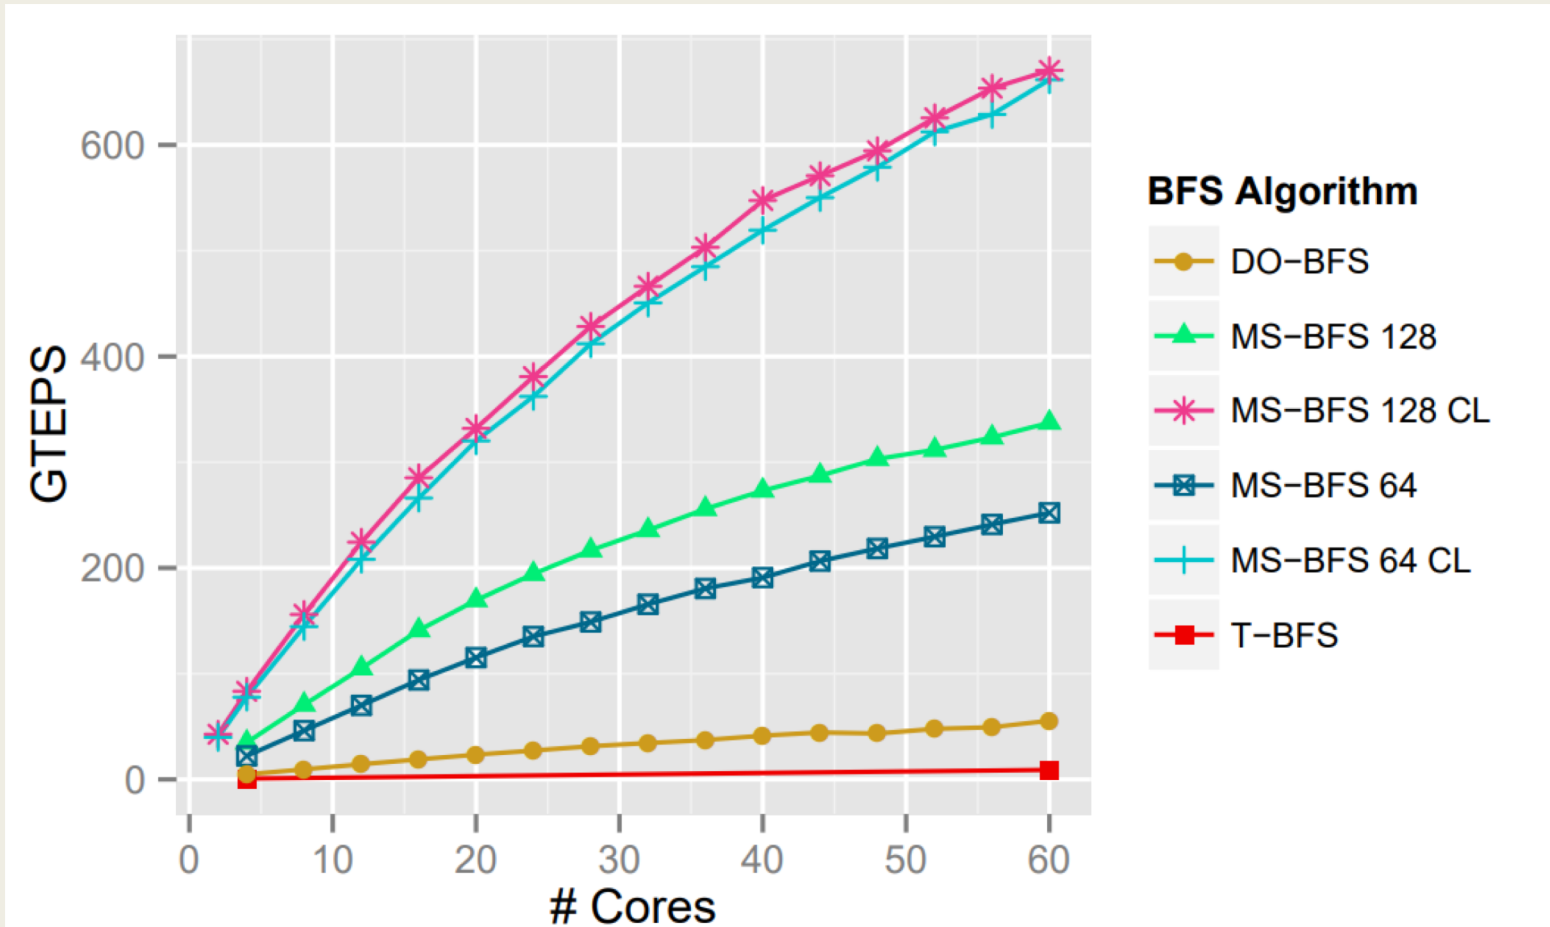


Figure 5: Multi-core scalability results.

Results only for LDBC
1M-vertex graph,
which is ~314 MB

Conclusions

- MS-BFS runs multiple BFSs
 - *On the same graph*
 - *Within a single thread*
 - *Amortizes cache line movement cost*
- For low-diameter graphs, a large fraction of vertices are visited each round, so you can amortize cost of traversing graph over many concurrent traversals.
- >10x speedup over direction-optimizing BFS
- Changing random accesses to predictable array scans improves efficiency.

Future work

- Combining parallelism across traversals with parallelism within traversals.
- Alternative architectures:
 - *GPUs should be good at exploiting SIMD-style parallelism?*
- Applications beyond closeness centrality.
- Other graphs. Is there a hybrid approach that works if graphs have moderate diameter?
- Other types of traversals besides BFS. Does it make sense to do multi-source SSSP/“weighted BFS” traversals on weighted graphs?
- Integrating into a graph analytics framework or a graph processing benchmark set?