

THE MORE THE MERRIER

---

**EFFICIENT MULTI-SOURCE GRAPH TRAVERSAL**

# BACKGROUND

---

- ▶ Graph analytics
- ▶ Multi core machines
- ▶ Graph traversal on same graph from different sources
  - ▶ Calculating graph centralities
  - ▶ Enumerating neighborhoods for all vertices
  - ▶ All-pairs shortest distance problem

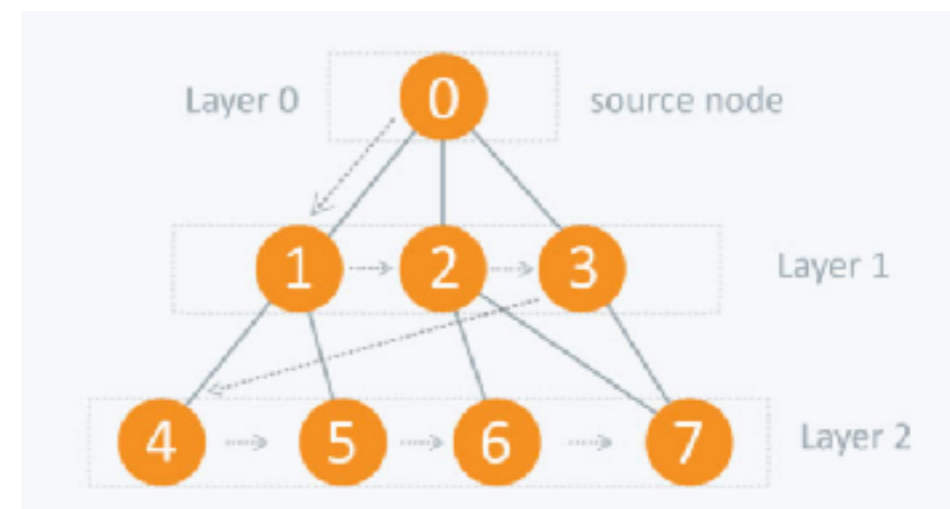


# BFS

- ▶ Textbook BFS
- ▶ Building block for other graph traversals
- ▶ Max levels - diameter(G)
- ▶ Random memory accesses every time it checks if a neighbor has been visited

**Listing 1: Textbook BFS algorithm.**

```
1 Input:  $G, s$ 
2  $seen \leftarrow \{s\}$ 
3  $visit \leftarrow \{s\}$ 
4  $visitNext \leftarrow \emptyset$ 
5
6 while  $visit \neq \emptyset$ 
7   for each  $v \in visit$ 
8     for each  $n \in neighbors_v$ 
9       if  $n \notin seen$ 
10          $seen \leftarrow seen \cup \{n\}$ 
11          $visitNext \leftarrow visitNext \cup \{n\}$ 
12         do BFS computation on  $n$ 
13    $visit \leftarrow visitNext$ 
14    $visitNext \leftarrow \emptyset$ 
```



# OPTIMIZING TEXTBOOK BFS

---

- ▶ Level by level parallelization

$$\textit{Work:} \quad T_1(n) = \Theta(m+n)$$

$$\textit{Span:} \quad T_\infty(n) = \Theta(d)$$

- ▶ Beamer et. All

- ▶ Bottom up approach - Explores based on unvisited nodes
- ▶ Hybrid approach - Uses bottom up for large frontiers, top up otherwise

Variable	Description
$\alpha$	Tuning parameter
$\beta$	Tuning parameter
$m_f$	# Edges in frontier
$m_u$	# Unexplored vertices
$n_f$	# Vertices in frontier
$n$	# Vertices

$$m_f > \frac{m_u}{\alpha} = C_{TB}$$

$$n_f < \frac{n}{\beta} = C_{BT}$$

# MOTIVATION

- ▶ Large graphs often must be searched from various starting nodes
- ▶ Lots of overlap when executing BFS from multiple nodes
- ▶ Small world graphs have even more overlap - large fanout, each level grows rapidly

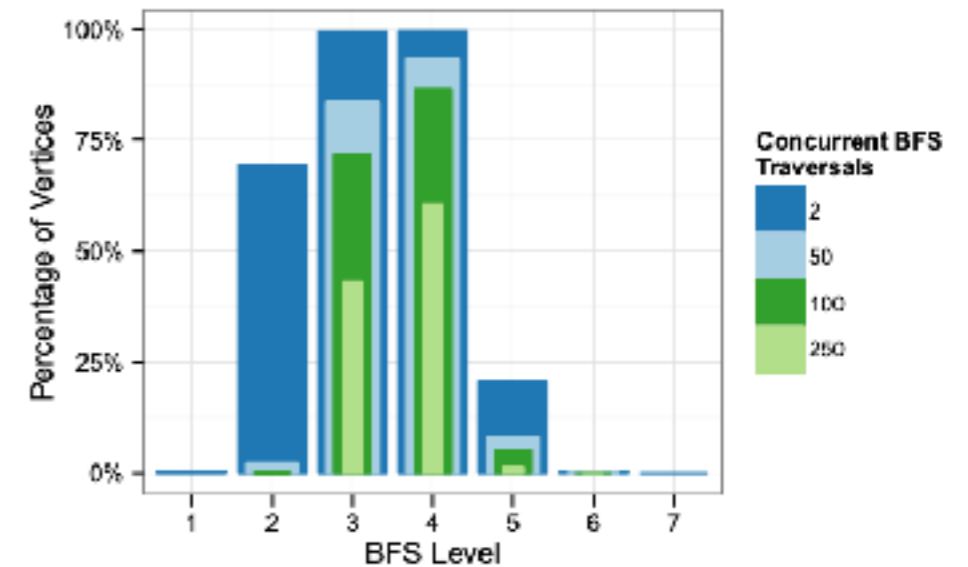



Figure 1: Percentage of vertex explorations that can be shared per level across 512 concurrent BFSs.

## SMALL WORLD GRAPHS

The distance between any two vertices is very small compared to the size of the graph, and the number of vertices discovered in each iteration of the BFS algorithm grows rapidly.

# MS-BFS: EXAMPLE



- ▶ Increase the dimensionality of textbook BFS to allow for multiple BFS at once
- ▶ Shared exploration of nodes
- ▶ Finish all BFSs executions in parallel

**Listing 2: The MS-BFS algorithm.**

```

1 Input:  $G, \mathbb{B}, S$ 
2  $seen_{s_i} \leftarrow \{b_i\}$  for all  $b_i \in \mathbb{B}$ 
3  $visit \leftarrow \bigcup_{b_i \in \mathbb{B}} \{(s_i, \{b_i\})\}$ 
4  $visitNext \leftarrow \emptyset$ 
5
6 while  $visit \neq \emptyset$ 
7   for each  $v$  in  $visit$ 
8      $\mathbb{B}'_v \leftarrow \emptyset$ 
9     for each  $(v', \mathbb{B}') \in visit$  where  $v' = v$ 
10       $\mathbb{B}'_v \leftarrow \mathbb{B}'_v \cup \mathbb{B}'$ 
11     for each  $n \in neighbors_v$ 
12       $\mathbb{D} \leftarrow \mathbb{B}'_v \setminus seen_n$ 
13      if  $\mathbb{D} \neq \emptyset$ 
14         $visitNext \leftarrow visitNext \cup \{(n, \mathbb{D})\}$ 
15         $seen_n \leftarrow seen_n \cup \mathbb{D}$ 
16        do BFS computation on  $n$ 
17    $visit \leftarrow visitNext$ 
18    $visitNext \leftarrow \emptyset$ 

```

# OPTIMIZATIONS FOR MS-BFS

---

- ▶ Bit operations
- ▶ Aggregated neighbor processing
- ▶ Direction optimized
- ▶ Neighbor prefetching
- ▶ Sharing heuristic

# OPTIMIZATIONS FOR MS-BFS

---

Listing 2: The MS-BFS algorithm.

```
1 Input:  $G, \mathbb{B}, S$ 
2  $seen_{s_i} \leftarrow \{b_i\}$  for all  $b_i \in \mathbb{B}$ 
3  $visit \leftarrow \bigcup_{b_i \in \mathbb{B}} \{(s_i, \{b_i\})\}$ 
4  $visitNext \leftarrow \emptyset$ 
5
6 while  $visit \neq \emptyset$ 
7   for each  $v$  in  $visit$ 
8      $\mathbb{B}'_v \leftarrow \emptyset$ 
9     for each  $(v', \mathbb{B}') \in visit$  where  $v' = v$ 
10       $\mathbb{B}'_v \leftarrow \mathbb{B}'_v \cup \mathbb{B}'$ 
11     for each  $n \in neighbors_v$ 
12       $\mathbb{D} \leftarrow \mathbb{B}'_v \setminus seen_n$ 
13      if  $\mathbb{D} \neq \emptyset$ 
14         $visitNext \leftarrow visitNext \cup \{(n, \mathbb{D})\}$ 
15         $seen_n \leftarrow seen_n \cup \mathbb{D}$ 
16        do BFS computation on  $n$ 
17  $visit \leftarrow visitNext$ 
18  $visitNext \leftarrow \emptyset$ 
```

Listing 3: MS-BFS using bit operations.

```
1 Input:  $G, \mathbb{B}, S$ 
2 for each  $b_i \in \mathbb{B}$ 
3    $seen[s_i] \leftarrow 1 \ll b_i$ 
4    $visit[s_i] \leftarrow 1 \ll b_i$ 
5 reset  $visitNext$ 
6
7 while  $visit \neq \emptyset$ 
8   for  $i = 1, \dots, N$ 
9     if  $visit[v_i] = \mathbb{B}_\emptyset$ , skip
10    for each  $n \in neighbors[v_i]$ 
11       $\mathbb{D} \leftarrow visit[v_i] \& \sim seen[n]$ 
12      if  $\mathbb{D} \neq \mathbb{B}_\emptyset$ 
13         $visitNext[n] \leftarrow visitNext[n] | \mathbb{D}$ 
14         $seen[n] \leftarrow seen[n] | \mathbb{D}$ 
15        do BFS computation on  $n$ 
16  $visit \leftarrow visitNext$ 
17 reset  $visitNext$ 
```

Listing 4: MS-BFS algorithm using ANP.

```
1 Input:  $G, \mathbb{B}, S$ 
2 for each  $b_i \in \mathbb{B}$ 
3    $seen[s_i] \leftarrow 1 \ll b_i$ 
4    $visit[s_i] \leftarrow 1 \ll b_i$ 
5 reset  $visitNext$ 
6
7 while  $visit \neq \emptyset$ 
8   for  $i = 1, \dots, N$ 
9     if  $visit[v_i] = \mathbb{B}_\emptyset$ , skip
10    for each  $n \in neighbors[v_i]$ 
11       $visitNext[n] \leftarrow visitNext[n] | visit[v_i]$ 
12
13   for  $i = 1, \dots, N$ 
14     if  $visitNext[v_i] = \mathbb{B}_\emptyset$ , skip
15      $visitNext[v_i] \leftarrow visitNext[v_i] \& \sim seen[v_i]$ 
16      $seen[v_i] \leftarrow seen[v_i] | visitNext[v_i]$ 
17     if  $visitNext[v_i] \neq \mathbb{B}_\emptyset$ 
18       do BFS computation on  $v_i$ 
19  $visit \leftarrow visitNext$ 
20 reset  $visitNext$ 
```



# EVALUATION AND RESULTS

- ▶ Running BFS from all nodes as number of vertices increases
- ▶ Traversed edges per second
- ▶ Improvement benefits from various optimizations

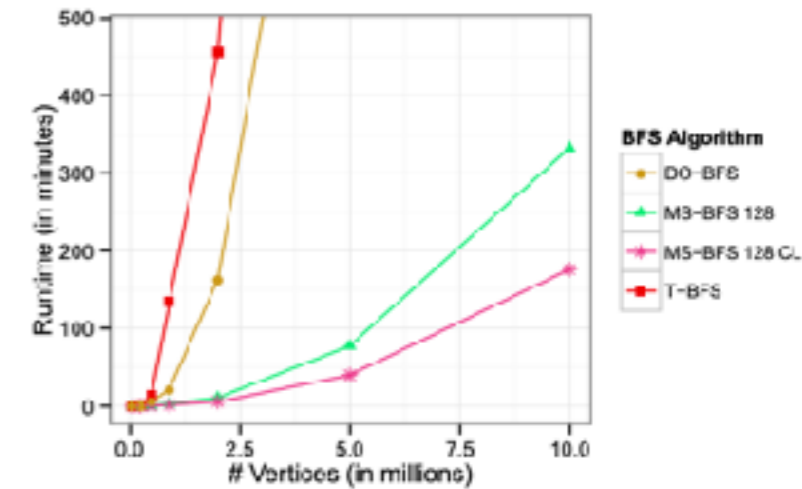


Figure 4: Data size scalability results.




Figure 5: Multi-core scalability results.




Figure 7: Speedup achieved by cumulatively applying different tuning techniques to MS-BFS.

# STRENGTHS

---

- ▶ Comparison with existing approaches
- ▶ Leverage existing optimizations
- ▶ Large scale evaluations

# WEAKNESSES

---

- ▶ Must be overlapping during the same iterations
  - ▶ No “memory” of previously searched nodes
- ▶ Evaluation on non “small-world” graphs
- ▶ Perform optimizations independently
- ▶ Evaluation in a distributed system
- ▶ MS-BFS with parallelization at each level

# DISCUSSION

---

- ▶ What did you guys think were the strengths and weaknesses?
- ▶ On what types of graphs is MS-BFS NOT useful
  - ▶ How could it be improved to be useful on these graphs?
  - ▶ How does MS-BFS perform compared to textbook BFS in these scenarios