

The BFS Kernel: Applications and Implementations

Peter M. Kogge

Graph Exploration

- Common graph problem: “explore” region around some vertex
- Exploration: follow edges to see what’s reachable
- Possible outputs:
 - Identification of reachable vertices
 - “labelling” of vertices
 - Properties of reachable sub-graph
- Options:
 - Constraints on “how far”
 - Constraints on “which edges”

Major Variants of Exploration

- Depth-First: Keep jumping from vertex to vertex until stopping
 - And then back up to last vertex and see if any untried edges
- Breadth-First: Explore in waves
 - Explore all edges from current “Frontier”
 - Mark as all new vertices as “New Frontier”
 - Start over with new frontier when all current one is searched
- This kernel: Breadth-First

Example: Airline Routes

- Consider graph with
 - Vertices: airports (~17,000 in world):
 - Properties: Country, International designation, Control tower, etc
 - Edges as flights between airports (100,000/day):
 - Properties: Airline (5,000 different airlines in world), Flight number, equipment, ...
 - Edges are directional
 - Note graph changes dynamically
- Possible explorations
 - What airports are reachable from some specific one
 - What if we constraint # of stops or airlines,
 - ...

Example: Six Degrees of Kevin Bacon


- IMDb Data base
 - Vertices: Multiple “classes”
 - 8.7M+ people
 - 4.8M+ titles of 10 types
 - Edges: (u,v) between people and titles
 - Person u has had one of 34 roles in title v
 - Again directional
- Possible exploration:
 - Can a chain of $(u_1,t_1), (u_2,t_1), (u_2,t_2), (u_3,t_2), \dots$ connect any one person u_1 to all other people in database?
 - Kevin Bacon: 6 titles away from everyone else
- See <https://oracleofbacon.org/>

Other Interesting Applications

- From: <https://www.geeksforgeeks.org/applications-of-breadth-first-traversal/>
 - Search for neighbors in peer-peer networks
 - Search engine web crawlers
 - Social networks – distance k friends
 - GPS navigation to find “neighboring” locations
 - Patterns for “broadcasting” in networks
- From Wikipedia: https://en.wikipedia.org/wiki/Breadth-first_search
 - Community Detection
 - Maze running
 - Routing of wires in circuits
 - Finding Connected components
 - Copying garbage collection, Cheney's algorithm
 - Shortest path between two nodes u and v
 - Cuthill–McKee mesh numbering
 - Maximum flow in a flow network
 - Serialization/Deserialization of a binary tree
 - Construction of the failure function of the Aho-Corasick pattern matcher.
 - Testing bipartiteness of a graph

Key Kernel: BFS - Breadth First Search

- Given a huge graph
- Start with a root, find all reachable vertices
- Performance metric: **TEPS**: Traversed Edges/sec



Starting at 1: 1, 0, 3, 2, 9, 5

No Flops – just Memory & Networking

Graph500: www.graph500.org

- Several years of reports on performance of BFS implementations on
 - Different size graphs
 - Different hardware configurations
- Standardized graphs for testing
- Standard approach for measuring
 - Generate a graph of certain size
 - Repeat 64 times
 - Select a root
 - Find “level” of each reachable vertex
 - Record execution time
 - TEPS = graph edges / execution time

Graph500 Graphs

- Kronecker graph generator algorithm
 - D. Chakrabarti, Y. Zhan, and C. Faloutsos, R-MAT: A recursive model for graph mining, SIAM Data Mining 2004
- Recursively sub-divides adjacency matrix into 4 partitions A, B, C, D
- Add edges one at a time, choosing partitions probabilistically
 - A = 57%, B = 19%, C = 19%, D = 5%
- # of generated edges = $16 * \#$ vertices
 - Average Vertex Degree is 2X this

Graph Sizes

Level	Scale	Size	Vertices (Billion)	TB	Bytes /Vertex
10	26	Toy	0.1	0.02	281.8048
11	29	Mini	0.5	0.14	281.3952
12	32	Small	4.3	1.1	281.472
13	36	Medium	68.7	17.6	281.4752
14	39	Large	549.8	141	281.475
15	42	Huge	4398.0	1,126	281.475
				Average	281.5162

Scale = $\log_2(\# \text{ vertices})$

Notional Sequential Algorithm

- Top-Down search: Keep a “frontier” of new vertices that have been “touched” but not “explored”
 - Explore them and repeat
- Bottom-up search: look at all “untouched vertices” and see if any of their edges lead to a touched vertex
 - If so, mark as touched, and repeat
- Special considerations
 - Vertices that have huge degrees

Top-Down

Algorithm 1 Top Down BFS:

V is the set of vertices; E a set of edges

```
1: procedure TOPDOWN-BFS( $G, \text{ROOT}$ )
2:    $Touched \leftarrow \{\text{root}\}$ 
3:    $Frontier \leftarrow \{\text{root}\}$ 
4:    $Labels \leftarrow$  N-vector of a large integer
5:    $Label[\text{root}] \leftarrow 0$ 
6:    $Level \leftarrow 0$ 
7:   while  $Frontier$  not empty do
8:      $Level += 1$ 
9:      $TopDown - Pass(Frontier, Touched, Level)X$ 
10:  return
11: procedure TOPDOWNPASS( $\text{TOUCHED}, \text{LEVEL}$ )
12:   $Next \leftarrow \{\}$ 
13:  for  $u$  in  $Frontier$  do
14:    for all edges  $(u, v)$  in  $E$  do
15:      if  $v$  not in  $Touched$  then
16:         $Touched \leftarrow Touched \cup \{v\}$ 
17:         $Next \leftarrow Next \cup \{v\}$ 
18:         $Label[v] \leftarrow Level$ 
19:   $Frontier \leftarrow Next$ 
20:  return
```

Notional Complexity:
 $O(M)$

Bottom Up

Algorithm 2 BottomUp BFS:

V is the set of vertices; E a set of edges

```
1: procedure BOTTOMUP-BFS( $G, \text{root}$ )
2:    $Touched \leftarrow \{\text{root}\}$ 
3:    $Labels \leftarrow$  N-vector of a large integer
4:    $Label[\text{root}] \leftarrow 0$ 
5:    $Level \leftarrow 0$ 
6:    $TouchedFlag \leftarrow \text{True}$ 
7:   while  $TouchedFlag$  do
8:      $Level += 1$ 
9:      $TouchedFlag \leftarrow \text{BackwardPass}(Touched, Level)$ 
10:  return
11: procedure BOTTOMUPPASS(INOUT  $Touched, Level$ )
12:    $TouchedFlag \leftarrow \text{False}$ 
13:   for  $v$  not in  $Touched$  do
14:     for all edges  $(u, v)$  in  $E$  do
15:       if  $u$  in  $Touched$  then
16:          $TouchedFlag \leftarrow \text{True}$ 
17:          $Touched \leftarrow Touched \cup v$ 
18:          $Label[v] \leftarrow Level$ 
19:   return  $TouchedFlag$ 
```

Notional Complexity:
 $O(NM)$

Key Observation

- Forward direction requires investigation of *every edge* leaving a frontier vertex
 - Each edge can be done in parallel
- Backwards direction can stop investigating edges *as soon as 1 vertex* in current frontier is found
 - If search edges sequentially, potentially significant work avoidance
- In any case, can still parallelize over vertices in frontier

Edges Explored per Level

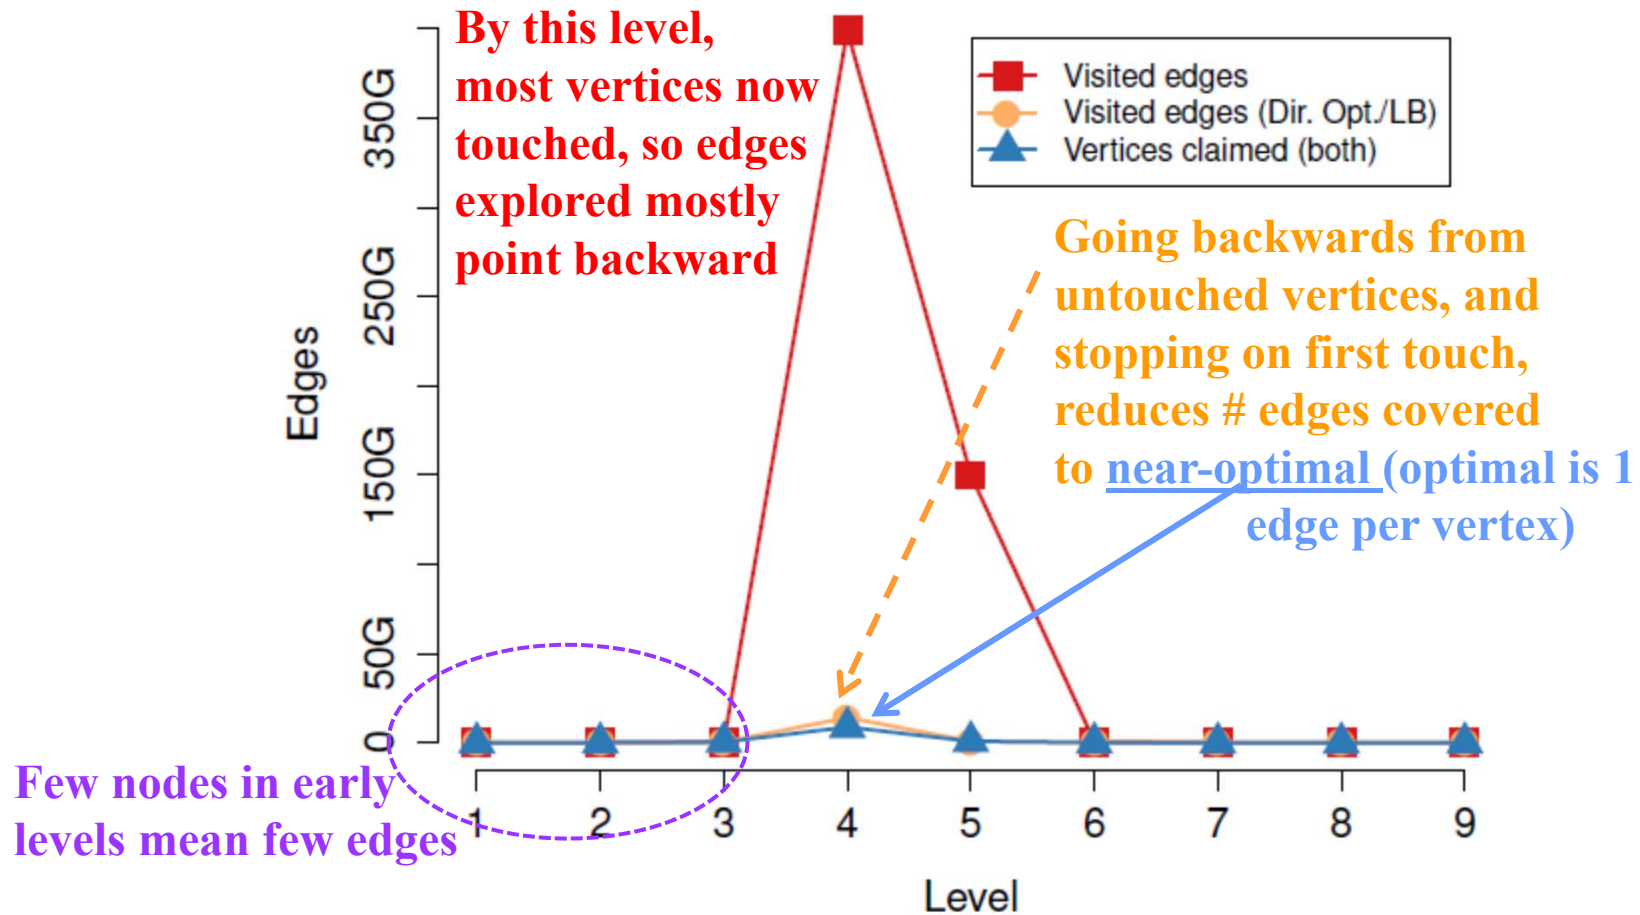


Fig. 5: Graph properties at each exploration level.

Checoni and Petrini, "Traversing Trillions ..."

Beamer's Hybrid Algorithm

- Switch between forward & backward steps
 - Use forward iteration as long as In is small
 - Use backward iteration when Vis is large
- Advantage: when
 - # edges from vertices in $!Vis$
 - are less than # edges from vertices in In
 - then we follow fewer edges overall
- Estimated savings if done optimally: up to 10X reduction in edges
- <http://www.scottbeamer.net/pubs/beamer-sc2012.pdf>

Hybrid Algorithm

Algorithm 3 Hybrid BFS:

V is the set of vertices; E a set of edges

```

1: procedure HYBRID-BFS(G,ROOT)
2:   Touched ← {root}
3:   Frontier ← {root}
4:   Labels ← N-vector of a large integer
5:   Label[root] ← 0
6:   Level ← 0
7:   Nf ← 1
8:   Mf ← outdegree(root)
9:   Mu ← M - Mf
10:  while Frontier not empty do
11:    Level += 1
12:    Next ← {}
13:    Nf ← 0
14:    Mf ← 0
15:    if (Mf ≥ Mu/α) or (Nf < N/β) then
16:      for all u in Frontier do
17:        for all edges (u,v) in E do
18:          if v not in Touched then
19:            Touched ← Touched ∪ {v}
20:            Next ← Next ∪ {v}
21:            Label[v] ← Level
22:            Nf += 1
23:            Mf += outdegree(v)
24:            Mu -= indegree(v)
25:    else
26:      for all v not in Touched do
27:        for all edges (u,v) in E do
28:          if u in Touched then
29:            Next ← Next ∪ v
30:            Label[v] ← Level
31:            Nf += 1
32:            Mf += outdegree(v)
33:            Mu -= indegree(v)
34:      Touched ← Touched ∪ Next
35:      Frontier ← Next
36:  return

```

- N_f = # vertices in current frontier
- M_f = # outgoing edges from current frontier
- M_u = # incoming edges to current untouched
- α = edge reduction factor in bottom-up pass
- β = vertex reduction factor when going from bottom-up to top-down

Switch from **top-down** to **bottom-up** when:

$$M_f < M_u/\alpha$$

Switch back from **bottom-up** to **top-down** when

$$N_f < N/\beta$$

BFS