



A New Parallel Algorithm for Connected Components in Dynamic Graphs

Authored by: Robert McColl, Oded Green,
David A. Bader
Presented by: Omar Obeya


Connected Components Problem




Dynamic Connected Components Problem

- Operations
 - Query node
 - Insert Edge
 - Delete Edge

Connected Components Problem



Connected Components Problem



Goals



- Parallel
- Exact
- Dynamic

Reformatting the problem


Given a deletion classify
whether it is safe

- 100% True positive.
- Minimum false negative.


Approach I: Adjacency List Intersection




Approach II: Spanning Tree




Approach II: Spanning Tree



Approach II: Spanning Tree



Approach II: Spanning Tree



Solution

- Maintain Neighbor-Parent list of *ThreshPN* “neighpars” only.
- At Deletion:
 - if it has a parent, or have a neighbour that has a parent.
 - Otherwise: recalculate.
- Use STINGER

Definitions

Table I
THE DATA STRUCTURES MAINTAINED WHILE TRACKING DYNAMIC CONNECTED COMPONENTS.

Name	Description	Type	Size (Elements)
C	Component labels	array	$O(V)$
$Size$	Component sizes	array	$O(V)$
$Level$	Approximate distance from the root	array	$O(V)$
PN	Parents and neighbors of each vertex	array of arrays	$O(V \cdot thresh_{PN}) = O(V)$
$Count$	Counts of parents and neighbors	array	$O(V)$
$thresh_{PN}$	Maximum count of parents and neighbors for a given vertex	value	$O(1)$
\bar{E}_I	Batch of edges to be inserted into graph	array	$O(batch\ size)$
\bar{E}_R	Batch of edges to be deleted from graph	array	$O(batch\ size)$

Initialization



Algorithm 1 A parallel breadth-first traversal that extracts the parent-neighbor subgraph.

Input: $G(V, E)$

Output: $C_{id}, Size, Level, PN, Count$

```
1: for  $v \in V$  do
2:    $Level[v] \leftarrow \infty, Count[v] \leftarrow 0$ 
3: for  $v \in V$  do
4:   if  $Level[v] = \infty$  then
5:      $Q[0] \leftarrow v, Q_{start} \leftarrow 0, Q_{end} \leftarrow 1$ 
6:      $Level[v] \leftarrow 0, C_{id}[v] \leftarrow v$ 
7:     while  $Q_{start} \neq Q_{end}$  do
8:        $Q_{stop} \leftarrow Q_{end}$ 
9:       for  $i \leftarrow Q_{start}$  to  $Q_{stop}$  in parallel do
10:        for each neighbor  $d$  of  $Q[i]$  do
11:          if  $Level[d] = \infty$  then
12:             $Q[Q_{end}] \leftarrow d$ 
13:             $Q_{end} \leftarrow Q_{end} + 1$ 
14:             $Level[d] \leftarrow Level[Q[i]] + 1$ 
15:             $C_{id}[d] \leftarrow C_{id}[Q[i]]$ 
16:          if  $Count[d] < thresh_{PN}$  then
17:            if  $Level[Q[i]] < Level[d]$  then
18:               $PN_d[Count[d]] \leftarrow Q[i]$ 
19:               $Count[d] \leftarrow Count[d] + 1$ 
20:            else if  $Level[Q[i]] = Level[d]$  then
21:               $PN_d[Count[d]] \leftarrow -Q[i]$ 
22:               $Count[d] \leftarrow Count[d] + 1$ 
23:        $Q_{start} \leftarrow Q_{stop}$ 
24:      $Size[v] \leftarrow Q_{end}$ 
```

- Parallel BFS
- Put neighpars in your list if there is room.
- Use -ve values for neighbours and +ve values for parents.

Insertion



Algorithm 2 The algorithm for updating the parent-neighbor subgraph for inserted edges.

```

Input:  $G(V, E)$ ,  $\tilde{E}_I$ ,  $C_{id}$ ,  $Size$ ,  $Level$ ,  $PN$ ,  $Count$ 
Output:  $C_{id}$ ,  $Size$ ,  $Level$ ,  $PN$ ,  $Count$ 
1: for all  $\langle s, d \rangle \in \tilde{E}_I$  in parallel do  $E \leftarrow E \cup \langle s, d \rangle$ 
2:   insert( $E$ ,  $\langle s, d \rangle$ )
3:   if  $C_{id}[s] = C_{id}[d]$  then
4:     if  $Level[s] > 0$  then
5:       if  $Level[d] < 0$  then
6:         // d is not "safe"
7:         if  $Level[s] < -Level[d]$  then
8:           if  $Count[d] < thresh_{PN}$  then
9:              $PN_d[Count[d]] \leftarrow s$ 
10:             $Count[d] \leftarrow Count[d] + 1$ 
11:           else
12:              $PN_d[0] \leftarrow s$ 
13:              $Level[d] \leftarrow -Level[d]$ 
14:           else
15:             if  $Count[d] < thresh_{PN}$  then
16:               if  $Level[s] < Level[d]$  then
17:                  $PN_d[Count[d]] \leftarrow s$ 
18:                  $Count[d] \leftarrow Count[d] + 1$ 
19:               else if  $Level[s] = Level[d]$  then
20:                  $PN_d[Count[d]] \leftarrow -s$ 
21:                  $Count[d] \leftarrow Count[d] + 1$ 
22:               else if  $Level[s] < Level[d]$  then
23:                 for  $i \leftarrow 0$  to  $thresh_{PN}$  do
24:                   if  $PN_d[i] < 0$  then
25:                      $PNV_d[i] \leftarrow s$ .
26:                   Break for-loop
27:              $\tilde{E}_I \leftarrow \tilde{E}_I \setminus \langle s, d \rangle$ 
28:   for all  $\langle s, d \rangle \in \tilde{E}_I$  do
29:     if  $C_{id}[s] \neq C_{id}[d]$  then
30:       if  $Size[s] = 1$  then
31:          $Size[s] \leftarrow 0$ 
32:          $Size[d] \leftarrow Size[d] + 1$ 
33:          $C_{id}[s] \leftarrow C_{id}[d]$ ,  $PN_s[0] \leftarrow d$ 
34:          $Level[s] \leftarrow \mathbf{abs}(Level[d]) + 1$ ,  $Count[s] \leftarrow 1$ 
35:       else
36:         connectComponent(Input,  $s$ ,  $d$ )

```

- Intra-connecting edges
 - In Parallel
 - If don't have a parent, add s if it was a parent.
 - If have a parent.
 - Try add s if it was neighbour
 - Must add s if it was parent (unless list is full of parents).

Algorithm 2 The algorithm for updating the parent-neighbor subgraph for inserted edges.

Input: $G(V, E)$, \tilde{E}_I , C_{id} , $Size$, $Level$, PN , $Count$

Output: C_{id} , $Size$, $Level$, PN , $Count$

```

1: for all  $\langle s, d \rangle \in \tilde{E}_I$  in parallel do  $E \leftarrow E \cup \langle s, d \rangle$ 
2:   insert( $E$ ,  $\langle s, d \rangle$ )
3:   if  $C_{id}[s] = C_{id}[d]$  then
4:     if  $Level[s] > 0$  then
5:       if  $Level[d] < 0$  then
6:         //  $d$  is not "safe"
7:         if  $Level[s] < -Level[d]$  then
8:           if  $Count[d] < thresh_{PN}$  then
9:              $PN_d[Count[d]] \leftarrow s$ 
10:             $Count[d] \leftarrow Count[d] + 1$ 
11:           else
12:              $PN_d[0] \leftarrow s$ 
13:             $Level[d] \leftarrow -Level[d]$ 
14:         else
15:           if  $Count[d] < thresh_{PN}$  then
16:             if  $Level[s] < Level[d]$  then
17:                $PN_d[Count[d]] \leftarrow s$ 
18:                $Count[d] \leftarrow Count[d] + 1$ 
19:             else if  $Level[s] = Level[d]$  then
20:                $PN_d[Count[d]] \leftarrow -s$ 
21:                $Count[d] \leftarrow Count[d] + 1$ 
22:             else if  $Level[s] < Level[d]$  then
23:               for  $i \leftarrow 0$  to  $thresh_{PN}$  do
24:                 if  $PN_d[i] < 0$  then
25:                    $PNV_d[i] \leftarrow s$ 
26:                 Break for-loop
27:            $\tilde{E}_I \leftarrow \tilde{E}_I \setminus \langle s, d \rangle$ 
28: for all  $\langle s, d \rangle \in \tilde{E}_I$  do
29:   if  $C_{id}[s] \neq C_{id}[d]$  then
30:     if  $Size[s] = 1$  then
31:        $Size[s] \leftarrow 0$ 
32:        $Size[d] \leftarrow Size[d] + 1$ 
33:        $C_{id}[s] \leftarrow C_{id}[d]$ ,  $PN_s[0] \leftarrow d$ 
34:        $Level[s] \leftarrow \text{abs}(Level[d]) + 1$ ,  $Count[s] \leftarrow 1$ 
35:     else
36:       connectComponent(Input,  $s, d$ )

```

- Inter-connecting edges
 - In Series
 - Special handle singletons.
 - Merge the labeling of the smaller component with the bigger.

Deletion



Algorithm 3 The algorithm for updating the parent-neighbor subgraph for deleted edges.

Input: $G(V, E)$, \tilde{E}_R , C_{id} , $Size$, $Level$, PN , $Count$

Output: C_{id} , $Size$, $Level$, PN , $Count$

```
1: for all  $\langle s, d \rangle \in \tilde{E}_R$  in parallel do
2:    $E \leftarrow E \setminus \langle s, d \rangle$ 
3:    $hasParents \leftarrow false$ 
4:   for  $p \leftarrow 0$  to  $Count[d]$  do
5:     if  $PN_d[p] = s$  or  $PN_d[p] = -s$  then
6:        $Count[d] \leftarrow Count[d] - 1$ 
7:        $PN_d[p] \leftarrow PN_d[Count[d]]$ 
8:     if  $PN_d[p] > 0$  then
9:        $hasParents \leftarrow true$ 
10:  if (not  $hasParents$ ) and  $Level[d] > 0$  then
11:     $Level[d] \leftarrow -Level[d]$ 
12: for all  $\langle s, d \rangle \in \tilde{E}_R$  in parallel do
13:   for all  $p \in PN_d$  do
14:     if  $p \geq 0$  or  $Level[abs(p)] > 0$  then
15:        $\tilde{E}_R \leftarrow \tilde{E}_R \setminus \langle s, d \rangle$ 
16:  $PREV \leftarrow C_{id}$ 
17: for all  $\langle s, d \rangle \in \tilde{E}_R$  do
18:    $unsafe \leftarrow (C_{id}[s] = C_{id}[d] = PREV_s)$ 
19:   for all  $p \in PN_d$  do
20:     if  $p \geq 0$  or  $Level[abs(p)] > 0$  then
21:        $unsafe \leftarrow false$ 
22:   if  $unsafe$  then
23:     if  $\{\langle u, v \rangle \in G(E, V) : u = s\} = \emptyset$  then
24:        $Level[s] \leftarrow 0, C_{id}[s] \leftarrow s$ 
25:        $Size[s] \leftarrow 1, Count[s] \leftarrow 0$ 
26:     else
27:       Algorithm 4
28:        $repairComponent(Input, s, d)$ 
```

- Maintain Data structure
 - In Parallel
 - Remove S from list
 - Recalc Has Parent?
 - Mark safe if possible
 - In Series
 - Mark safe if possible
 - Otherwise repair

Repair



Algorithm 4 The algorithm for repairing the parent-neighbor subgraph when an unsafe deletion is reported.

Input: $G(V, E)$, \tilde{E}_R , C_{id} , $Size$, $Level$, PN , $Count$, s , d
Output: C_{id} , $Size$, $Level$, PN , $Count$

```
1:  $Q[0] \leftarrow d$ ,  $Q_{start} \leftarrow 0$ ,  $Q_{end} \leftarrow 1$ 
2:  $SLQ \leftarrow \emptyset$ ,  $SLQ_{start} \leftarrow 0$ ,  $SLQ_{end} \leftarrow 0$ 
3:  $Level[d] \leftarrow 0$ ,  $C_{id}[d] \leftarrow d$ 
4:  $disconnected \leftarrow true$ 
5: while  $Q_{start} \neq Q_{end}$  do
6:    $Q_{stop} \leftarrow Q_{end}$ 
7:   for  $i \leftarrow Q_{start}$  to  $Q_{stop}$  in parallel do
8:      $u \leftarrow Q[i]$ 
9:     for each neighbor  $v$  of  $u$  do
10:      if  $C_{id}[v] = C_{id}[s]$  then
11:        if  $Level[v] \leq \text{abs}(Level[d])$  then
12:           $C_{id}[v] \leftarrow C_{id}[d]$ 
13:           $disconnected \leftarrow false$ 
14:           $SLQ[SLQ_{end}] \leftarrow v$ 
15:           $SLQ_{end} \leftarrow SLQ_{end} + 1$ 
16:        else
17:           $C_{id}[v] \leftarrow C_{id}[d]$ 
18:           $Count[v] \leftarrow 0$ 
19:           $Level[v] \leftarrow Level[u] + 1$ 
20:           $Q[Q_{end}] \leftarrow v$ 
21:           $Q_{end} \leftarrow Q_{end} + 1$ 
22:        if  $Count[v] < \text{thresh}_{PN}$  then
23:          if  $Level[u] < Level[v]$  then
24:             $PN_v[Count[v]] \leftarrow u$ 
25:             $Count[v] \leftarrow Count[v] + 1$ 
26:          else if  $Level[v] = Level[u]$  then
27:             $PN_v[Count[v]] \leftarrow -u$ 
28:             $Count[v] \leftarrow Count[v] + 1$ 
29:       $Q_{start} \leftarrow Q_{stop}$ 
30: if  $disconnected$  then
31:    $Size[d] \leftarrow Q_{end}$ 
32: else
33:   for  $i \leftarrow SLQ_{start}$  to  $SLQ_{end}$  in parallel do
34:      $C_{id}[i] \leftarrow C_{id}[s]$ 
35:   while  $SLQ_{start} \neq SLQ_{end}$  do
36:      $SLQ_{stop} \leftarrow SLQ_{end}$ 
37:     for  $i \leftarrow SLQ_{start}$  to  $SLQ_{stop}$  in parallel do
38:        $u \leftarrow SLQ[i]$ 
39:       for each neighbor  $v$  of  $u$  do
40:         if  $C_{id}[v] = C_{id}[d]$  then
41:            $C_{id}[v] \leftarrow C_{id}[u]$ 
42:            $Count[v] \leftarrow 0$ 
43:            $Level[v] \leftarrow Level[u] + 1$ 
44:            $SLQ[SLQ_{end}] \leftarrow v$ 
45:            $SLQ_{end} \leftarrow SLQ_{end} + 1$ 
46:         if  $Count[v] < \text{thresh}_{PN}$  then
47:           if  $Level[u] < Level[v]$  then
48:              $PN_v[Count[v]] \leftarrow u$ 
49:              $Count[v] \leftarrow Count[v] + 1$ 
50:           else if  $Level[v] = Level[u]$  then
51:              $PN_v[Count[v]] \leftarrow -u$ 
52:              $Count[v] \leftarrow Count[v] + 1$ 
53:        $Q_{start} \leftarrow Q_{stop}$ 
```

Input: $G(V, E), \tilde{E}_R, C_{id}, Size, Level, PN, Count, s, d$

Output: $C_{id}, Size, Level, PN, Count$

```
1:  $Q[0] \leftarrow d, Q_{start} \leftarrow 0, Q_{end} \leftarrow 1$ 
2:  $SLQ \leftarrow \emptyset, SLQ_{start} \leftarrow 0, SLQ_{end} \leftarrow 0$ 
3:  $Level[d] \leftarrow 0, C_{id}[d] \leftarrow d$ 
4:  $disconnected \leftarrow true$ 
5: while  $Q_{start} \neq Q_{end}$  do
6:    $Q_{stop} \leftarrow Q_{end}$ 
7:   for  $i \leftarrow Q_{start}$  to  $Q_{stop}$  in parallel do
8:      $u \leftarrow Q[i]$ 
9:     for each neighbor  $v$  of  $u$  do
10:      if  $C_{id}[v] = C_{id}[s]$  then
11:        if  $Level[v] \leq \text{abs}(Level[d])$  then
12:           $C_{id}[v] \leftarrow C_{id}[d]$ 
13:           $disconnected \leftarrow false$ 
14:           $SLQ[SLQ_{end}] \leftarrow v$ 
15:           $SLQ_{end} \leftarrow SLQ_{end} + 1$ 
16:        else
17:           $C_{id}[v] \leftarrow C_{id}[d]$ 
18:           $Count[v] \leftarrow 0$ 
19:           $Level[v] \leftarrow Level[u] + 1$ 
20:           $Q[Q_{end}] \leftarrow v$ 
21:           $Q_{end} \leftarrow Q_{end} + 1$ 
22:        if  $Count[v] < \text{thresh}_{PN}$  then
23:          if  $Level[u] < Level[v]$  then
24:             $PN_v[Count[v]] \leftarrow u$ 
25:             $Count[v] \leftarrow Count[v] + 1$ 
26:          else if  $Level[u] = Level[v]$  then
27:             $PN_v[Count[v]] \leftarrow -u$ 
28:             $Count[v] \leftarrow Count[v] + 1$ 
29:    $Q_{start} \leftarrow Q_{stop}$ 
```

- Parallel BFS
 - Search for a connection back to the s component


```

30: if disconnected then
31:    $Size[d] \leftarrow Q_{end}$ 
32: else
33:   for  $i \leftarrow SLQ_{start}$  to  $SLQ_{end}$  in parallel do
34:      $C_{id}[i] \leftarrow C_{id}[s]$ 
35:   while  $SLQ_{start} \neq SLQ_{end}$  do
36:      $SLQ_{stop} \leftarrow SLQ_{end}$ 
37:     for  $i \leftarrow SLQ_{start}$  to  $SLQ_{stop}$  in parallel do
38:        $u \leftarrow SLQ[i]$ 
39:       for each neighbor  $v$  of  $u$  do
40:         if  $C_{id}[v] = C_{id}[d]$  then
41:            $C_{id}[v] \leftarrow C_{id}[u]$ 
42:            $Count[v] \leftarrow 0$ 
43:            $Level[v] \leftarrow Level[u] + 1$ 
44:            $SLQ[SLQ_{end}] \leftarrow v$ 
45:            $SLQ_{end} \leftarrow SLQ_{end} + 1$ 
46:         if  $Count[v] < thresh_{PN}$  then
47:           if  $Level[u] < Level[v]$  then
48:              $PN_v[Count[v]] \leftarrow u$ 
49:              $Count[v] \leftarrow Count[v] + 1$ 
50:           else if  $Level[v] = Level[u]$  then
51:              $PN_v[Count[v]] \leftarrow -u$ 
52:              $Count[v] \leftarrow Count[v] + 1$ 
53:    $Q_{start} \leftarrow Q_{stop}$ 

```

- Disconnected -> done.
- Relabel all vertices discovered.
- Second Parallel BFS
 - Relabel more undiscovered vertices.

Quantitative Results



Unsafe Deletions

As we store more neighbors we get lower false negative




Figure 2. Average number of unsafe deletes in PN data structure for batches of $100K$ updates as a function of the average degree (x-axis) and $thresh_{PN}$ (bars).

Performance Results



Scaling

Scales sublinearly.

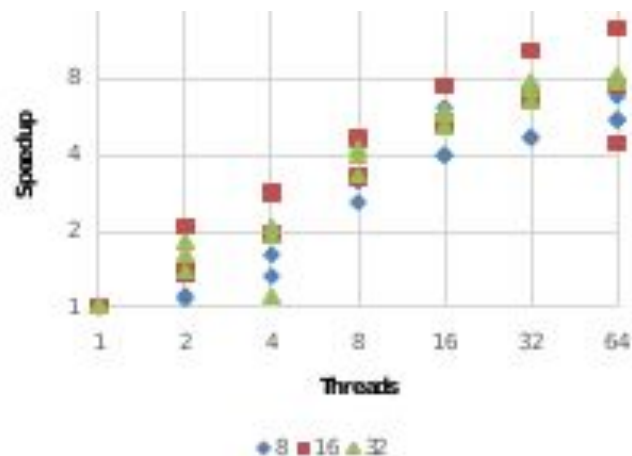


Figure 3. Strong scaling results on RMAT-22 graphs with different average degree as a function of the number of threads. Results include three graphs

Speed up

Speed up (over static) depends on graph density and not no. of threads.

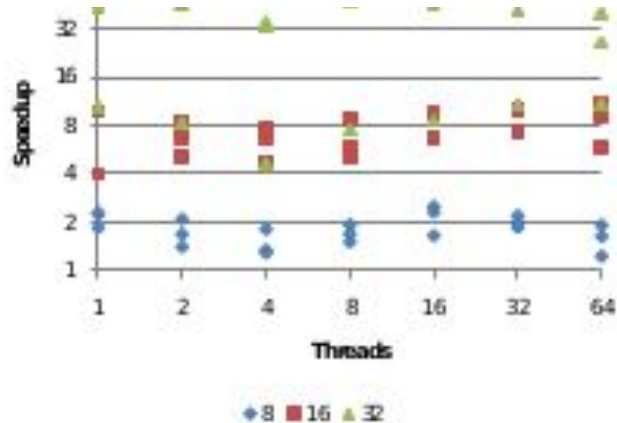


Figure 4. Speed up of the new algorithm over performing parallel static recomputation after each batch on three different RMAT-22 graphs with each average degree as a function of the number of threads.

Graph Dependency

The algorithm
speedup is very graph
dependent

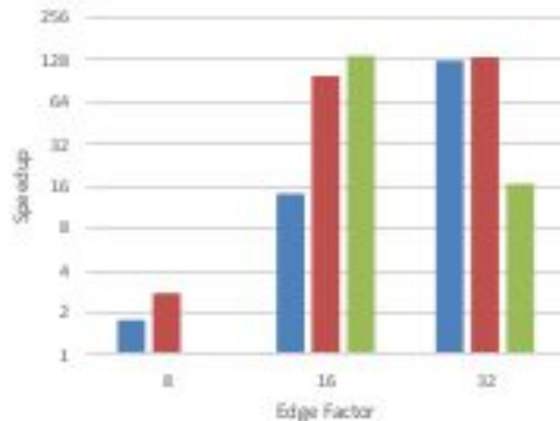


Figure 6. Speed up over performing static recomputation after each batch on scale 24 graphs for three graphs at each edge factor using 64 threads.

Weak Points

- Percentage of deletion is 6.25%
- Graph Dependent
- No probabilistic theoretical bounds

Future Work

- Threshold can be a function of no. of edges.
- What about adjacency matrix intersection using a popular node?

References

- McColl, Robert, Oded Green, and David A. Bader. "A new parallel algorithm for connected components in dynamic graphs." In High Performance Computing (HiPC), 2013 20th International Conference on, pp. 246-255. IEEE, 2013.