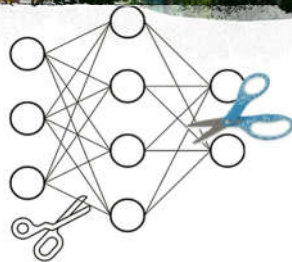


CONFERENCE ON  
PARALLEL  
PROCESSING



# Prune the Unnecessary: Parallel Pull-Push Louvain Algorithms with Automatic Edge Pruning

Jesmin Jahan Tithi ♥

Andrzej Stasiak \*

Sriram Ananthakrishnan\*

Fabrizio Petrini ♥

♥Parallel Computing Labs, Intel,

\*Data Center Group, Intel.

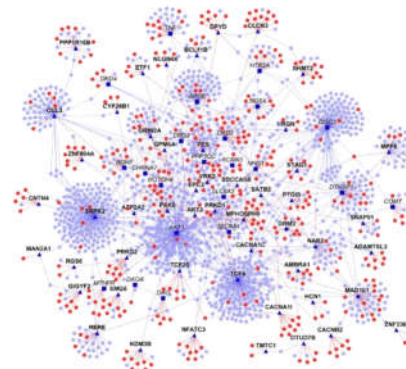
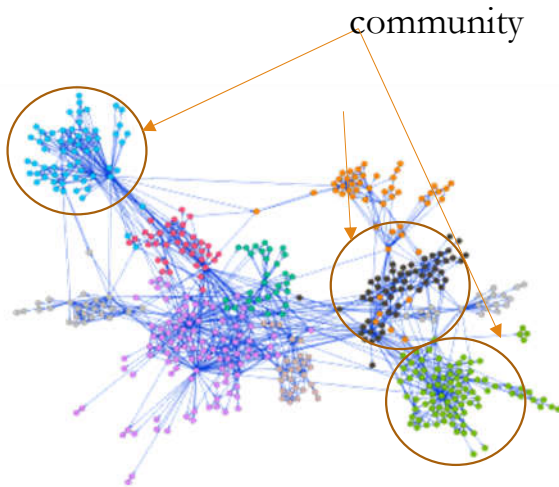


What is community?

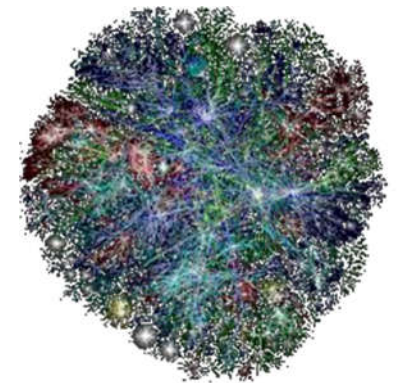
---

# What is Community?

- Sets of vertices that have dense intra-connections, but sparse inter-connections
- Uncover hidden structures inside a graph in a form of coherent modules of vertices
- Strongly correlated to functional and structural properties



Protein-Protein Interaction Network



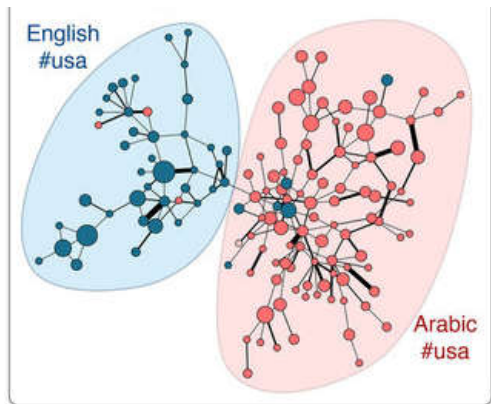
World Wide Web

Image source: Google Image

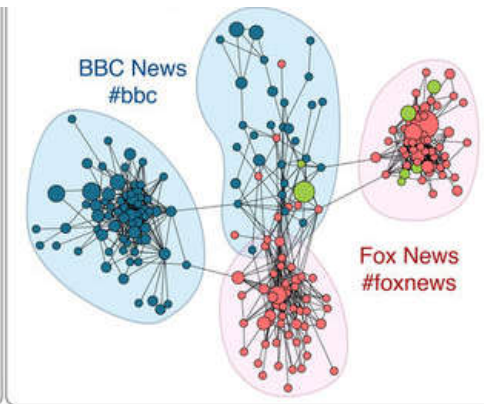
What is community detection?

# What is Community Detection?

- Algorithms to identify communities in a network
- Applications:** network analysis to retrieve information or patterns of the network



Virality Prediction and Community Structure in Social Networks



Nodus Labs Against Putin Facebook protest group visualization, December 2011

**NODUS LABS**  
www.noduslabs.com  
social and sentiment network analysis

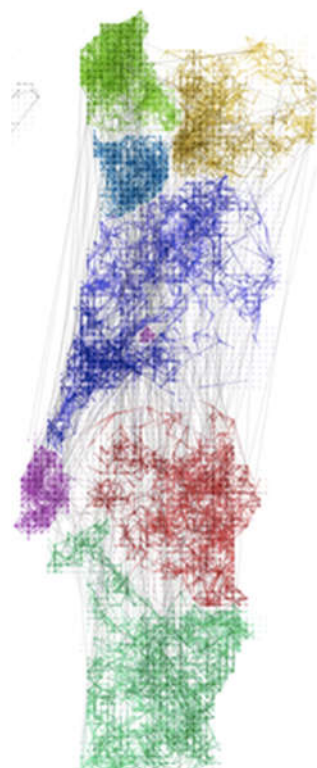
Visualization of Facebook group  
Putin Must Leave  
http://www.facebook.com/groups/putinmustleave/

as of 24 December 2011

1809 members  
5558 connections  
Average path: 3.7  
Average degree: 5.7  
Diameter: 10  
Clustering: 0.184  
Modularity: 0.449  
Graph density: 0.003  
Orphan members: 57%

4 main clusters:

- 15% nodes (magenta)  
political activists close to Kuznetsov
- 11% nodes (blue)  
journalists, people from the media
- 10% nodes (green)  
activists from Georgia
- 8% nodes (white)  
activists from Georgia



[http://senseable.mit.edu/community\\_detection/](http://senseable.mit.edu/community_detection/)

How to measure the quality of  
the detected communities?

---

# A Measure of Solution Quality

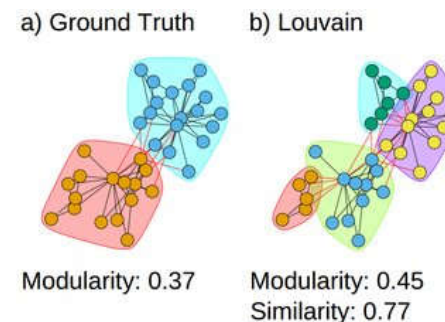
- **Modularity:** A measure of interconnectedness of the communities

$$\text{Modularity, } Q = \sum_{c \in C} \left[ \frac{\sum c_{in}}{2m} - \frac{\sum c_{tot}^2}{4m^2} \right] \text{Max Value of } Q = 1$$

$$\sum c_{in} = \sum W_{u,v}, \text{ for all } u, v \in c$$

$$\sum c_{tot} = \sum W_{u,v}, \text{ for all } u \in c \text{ or } v \in c$$

$$m = \sum_{e(u,v)} W_{u,v}$$



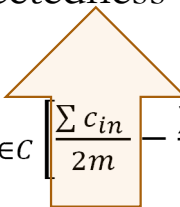
- $|Q| \in (0, 1]$ , and the higher the better
- Community detection algorithm identifies communities in a way that maximizes modularity

How do we maximize  
modularity?

---

# A Recipe of Modularity Optimization

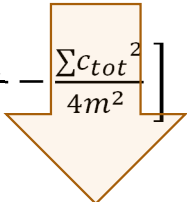
- **Modularity:** A measure of interconnectedness of the communities

$$\text{Modularity, } Q = \sum_{c \in C} \left[ \frac{\sum c_{in}}{2m} - \frac{\sum c_{tot}^2}{4m^2} \right] \text{ Max Value of } Q = 1$$


- Large values of  $Q$  correlate with **high  $\sum c_{in}$**  and low  $\sum c_{tot}$ 
  - Communities that are dense within their structure and weakly coupled among each other
- To **get high  $\sum c_{in}$** , the highest possible number of edges should fall in each community

# A Recipe of Modularity Optimization

- **Modularity:** A measure of interconnectedness of the communities

$$\text{Modularity, } Q = \sum_{c \in C} \left[ \frac{\sum c_{in}}{2m} - \frac{\sum c_{tot}^2}{4m^2} \right] \text{ Max Value of } Q = 1$$


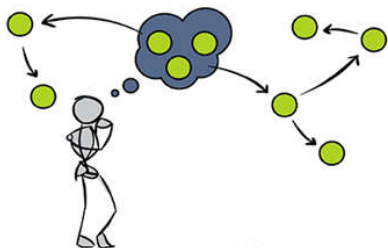
- Large values of  $Q$  correlate with high  $\sum c_{in}$  and **low**  $\sum c_{tot}$ 
  - Communities that are dense within their structure and weakly coupled among each other
- To **decrease**  $\sum c_{tot}$ , divide the network into several communities with small total degrees

# NP-hardness of Modularity Optimization

- **Modularity:** A measure of interconnectedness of the communities

$$\text{Modularity, } Q = \sum_{c \in C} \left[ \frac{\sum c_{in}}{2m} - \frac{\sum c_{tot}^2}{4m^2} \right] \quad \text{Max Value of } Q = 1$$

**Challenge:** Finding communities with optimal modularity is “NP-hard”



**NP-Hard problems**  
(at least as difficult as the  
hardest NP problems)

**NP-Complete  
problems**

**NP problems**  
(require at least  
NP time to solve)

## An abstract geometric composition featuring a variety of shapes and colors. In the top left, a large, light green, stylized arrow points towards the right. Below it, a small, light green circle is visible. To the right of the circle, a large, solid orange semi-circle is positioned. Further right, a small, solid orange semi-circle is located at the top edge. In the bottom left, a large, solid orange circle is partially visible. To its right, a series of four small, light green, curved lines are arranged in a diagonal pattern. In the bottom right, a large, light green square frame is shown. The background is a solid, light gray color.

Maximizes modularity following a greedy algorithm

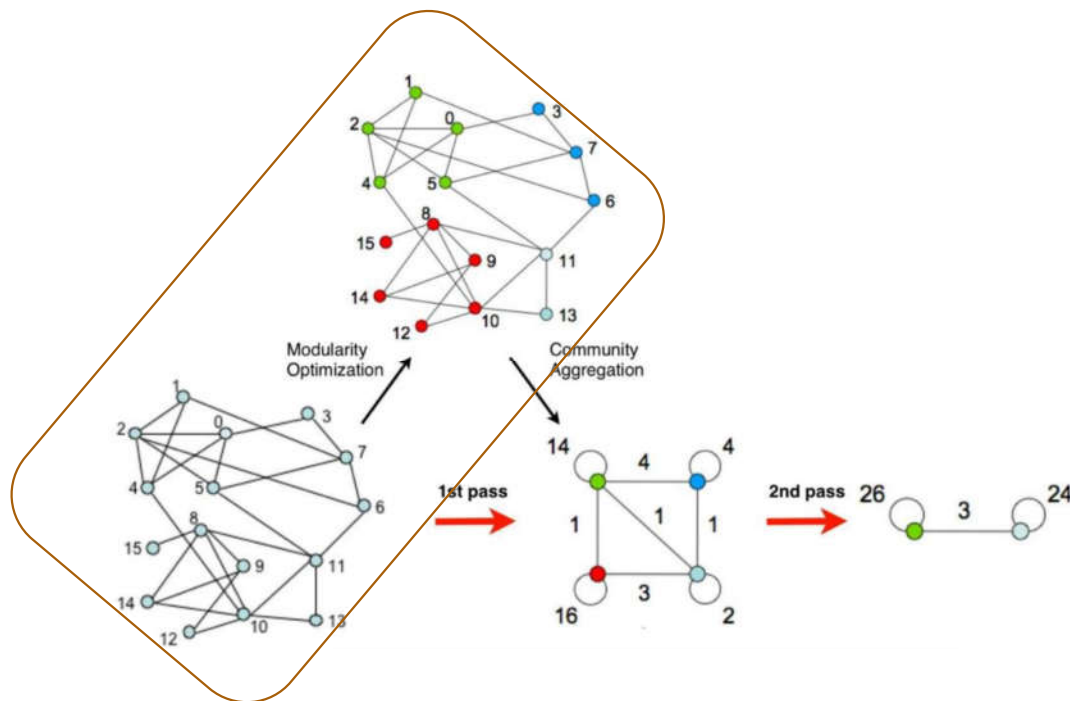
V. D. Blondel, J.-L. Guillaume, R. Lambiotte and E. Lefebvre, "Fast unfolding of communities in large networks," *J. Stat. Mech.* (2008) P10008, p. 12, 2008

## Louvain: Algorithm Steps

- **Outer Loop:** Traverse the graph in several passes to incrementally build communities

# Louvain: Algorithm Steps

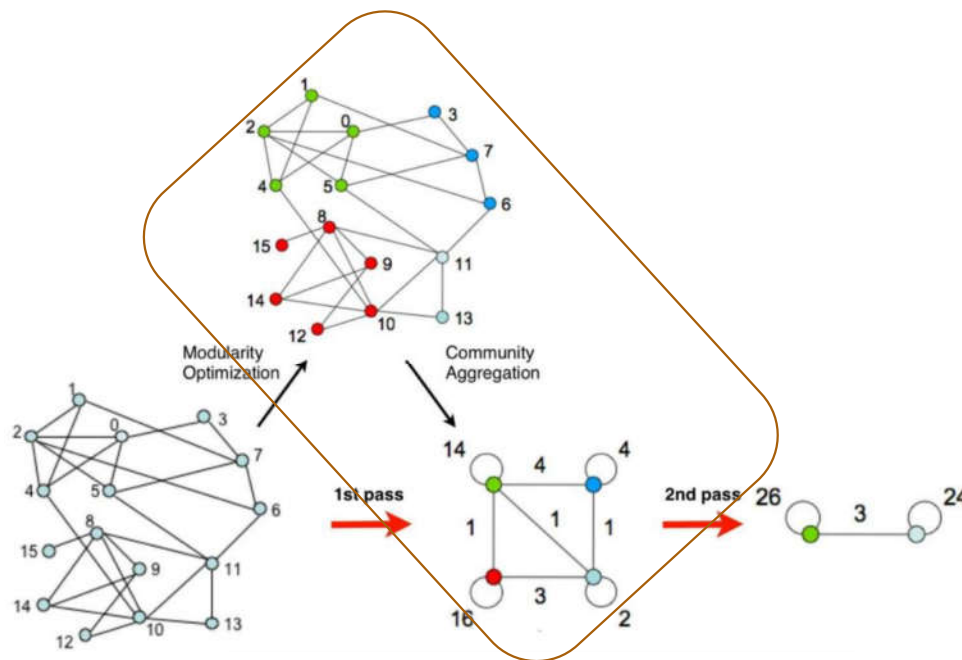
- **Outer Loop:** Traverse the graph in several passes to incrementally build communities
  - **Phase 1:** Modularity Optimization/Inner loop



V. D. Blondel, J.-L. Guillaume, R. Lambiotte and E. Lefebvre, "Fast unfolding of communities in large networks," *J. Stat. Mech.* (2008) P10008, p. 12, 2008

# Louvain: Algorithm Steps

- **Outer Loop:** Traverse the graph in several passes to incrementally build communities
  - **Phase 2: Community Aggregation and Graph Reconstruction**



V. D. Blondel, J.-L. Guillaume, R. Lambiotte and E. Lefebvre, "Fast unfolding of communities in large networks," *J. Stat. Mech.* (2008) P10008, p. 12, 2008

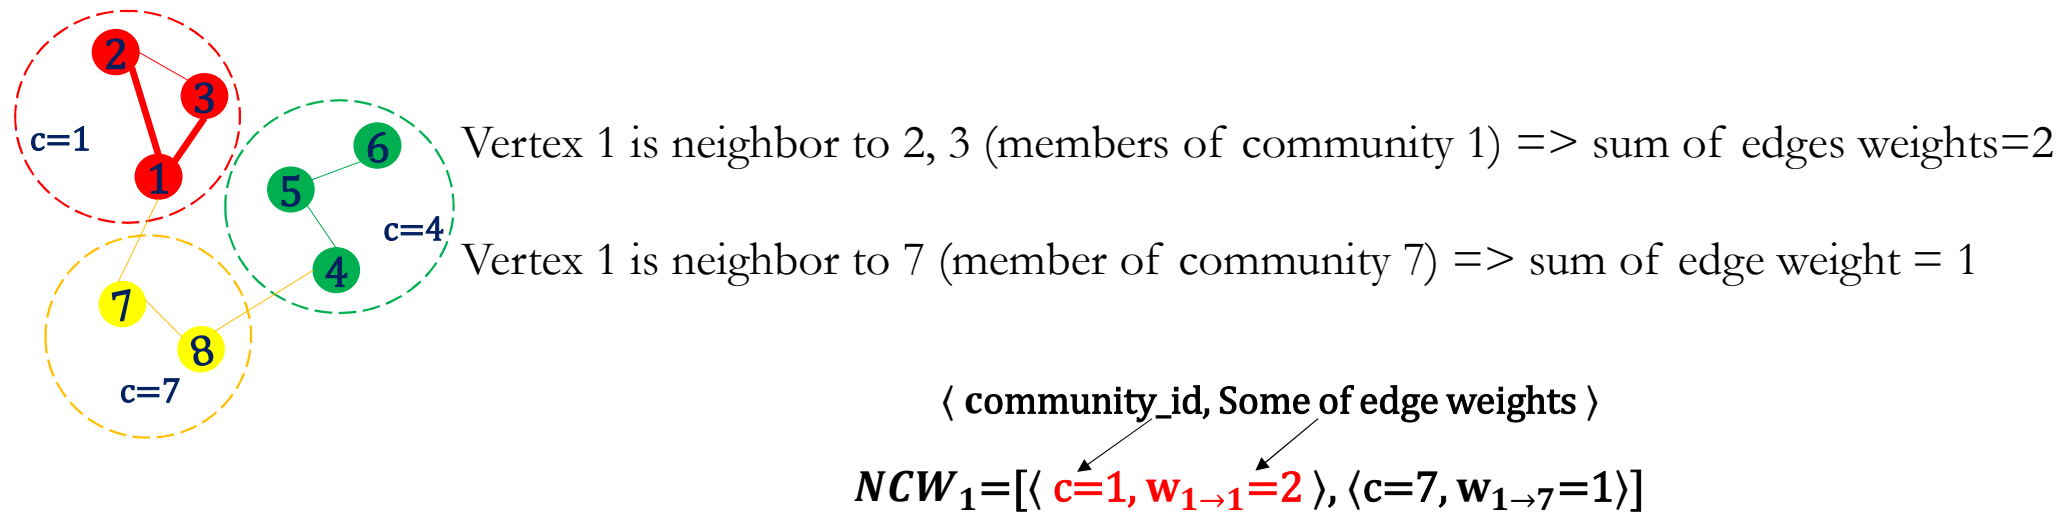
# Louvain: Algorithm Steps

- **Outer Loop:** Traverse the graph in several passes to incrementally build communities
  - **Phase 1:** Modularity Optimization/Inner loop -  $O(k(|V| + |E|))$
  - **Phase 2:** Community Aggregation and Graph Reconstruction -  $O(|V| + |E|)$

A key data structure to decide  
pull or push

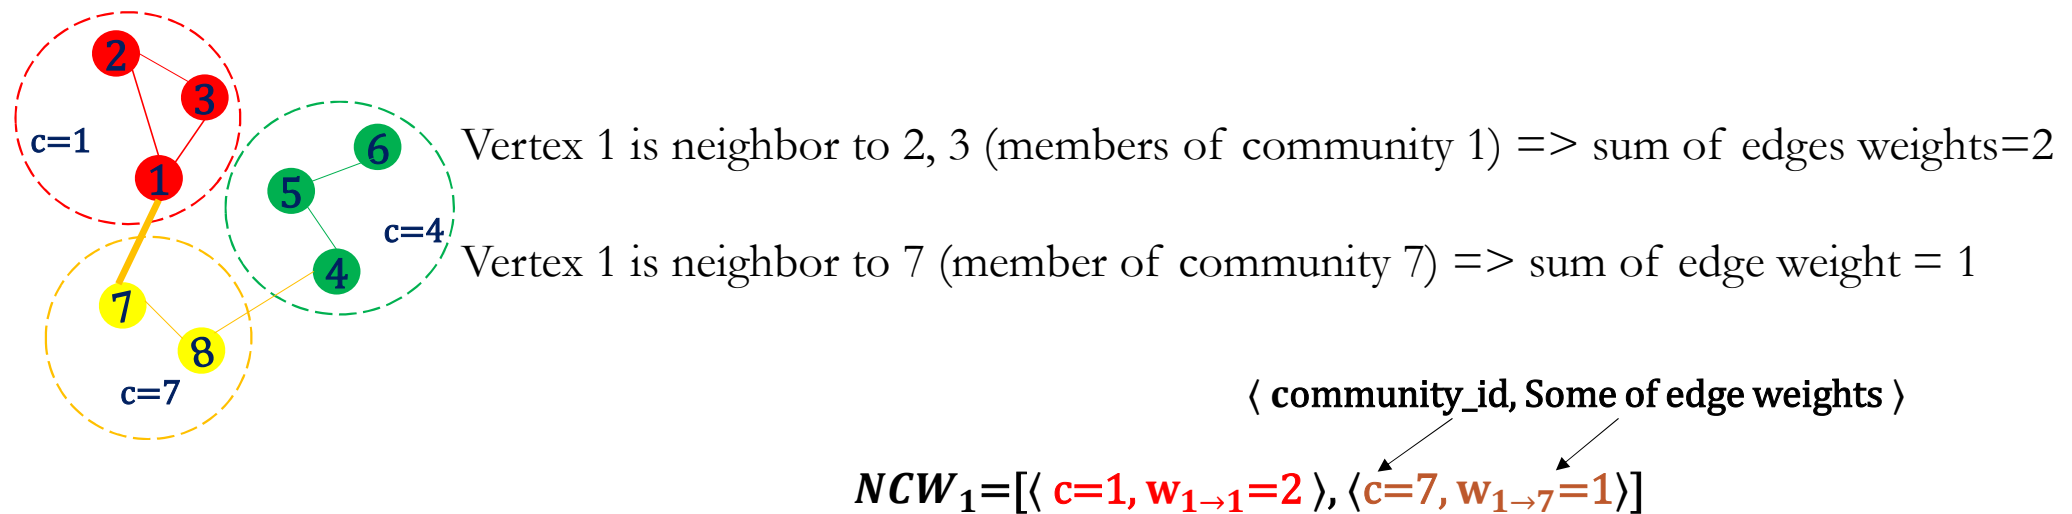
Hash map NCW–  $\langle \text{community\_id}, \text{Some of edge weights} \rangle$

A hash map with  $\langle \text{key} = \text{neighboring community}, \text{val} = \text{sum of edge weights to that community} \rangle$



Hash map NCW–  $\langle \text{community\_id}, \text{Some of edge weights} \rangle$

A hash map with  $\langle \text{key} = \text{neighboring community}, \text{val} = \text{sum of edge weights to that community} \rangle$



# Louvain Pseudocode

Repeat if there is a change in community membership

---

**Algorithm 1:** Sequential Louvain Algorithm.

---

**Data:** Graph  $G = (V, E)$ ,  $\tau = \text{Threshold}$ .

**Result:** Final Community Assignment  $C$ , and Modularity  $Q$

```
1 //Outer Loop
2 while true do
3   changes=0
4   C = initialize each vertex in its own community
5   Q = computeModularity(V, E, C)
6   //Phase 1 (Inner loop): Modularity Optimization
7   while true do
8     Qprev = Q
9     for all  $u \in V$  do
10      //build NCW
11       $NCW_u = \langle c, w_{u \rightarrow c} \rangle \forall c \in NC_u$ 
12      (bestcomm, bestGain) =
13         $\arg \max_{c \in NCW_u} \Delta Q_{u \rightarrow c}$ 
14      if bestGain>0 and bestcomm  $\neq$  oldcomm of u
15        then
16          | Move u to bestcomm and update C
17          // new modularity value
18          Q = computeModularity(V, E, C)
19          if  $Q - Q_{prev} \leq \tau$  then
20            | break
21          changes=1
22      //Phase 2: Aggregation and Graph Reconstruction
23      if changes = 0 then
24        | break
25    reconstructGraph(V, E, C)
```

---

# Louvain Pseudocode

Initialize each vertex in its own community  
Compute initial modularity

---

**Algorithm 1:** Sequential Louvain Algorithm.

---

**Data:** Graph  $G = (V, E)$ ,  $\tau = \text{Threshold}$ .

**Result:** Final Community Assignment  $C$ , and Modularity  $Q$

```
1 //Outer Loop
2 while true do
3   changes=0
4   C = initialize each vertex in its own community
5   Q = computeModularity(V, E, C)
6   //Phase 1 (Inner loop): Modularity Optimization
7   while true do
8     Qprev = Q
9     for all  $u \in V$  do
10      //build NCW
11       $NCW_u = \langle c, w_{u \rightarrow c} \rangle \forall c \in NC_u$ 
12      (bestcomm, bestGain) =
13         $\arg \max_{c \in NCW_u} \Delta Q_{u \rightarrow c}$ 
14      if bestGain > 0 and bestcomm  $\neq$  oldcomm of  $u$ 
15        then
16          Move  $u$  to bestcomm and update  $C$ 
17      // new modularity value
18      Q = computeModularity(V, E, C)
19      if  $Q - Q_{prev} \leq \tau$  then
20        break
21      changes=1
22   //Phase 2: Aggregation and Graph Reconstruction
23   if changes = 0 then
24     break
25   reconstructGraph(V, E, C)
```

---

# Louvain Pseudocode

Phase 1/ inner loop starts

---

**Algorithm 1:** Sequential Louvain Algorithm.

---

**Data:** Graph  $G = (V, E)$ ,  $\tau = \text{Threshold}$ .

**Result:** Final Community Assignment  $C$ , and Modularity  $Q$

```
1 //Outer Loop
2 while true do
3   changes=0
4   C = initialize each vertex in its own community
5   Q = computeModularity(V, E, C)
6   //Phase 1 (Inner loop): Modularity Optimization
7   while true do
8     Qprev = Q
9     for all  $u \in V$  do
10      //build NCW
11       $NCW_u = \langle c, w_{u \rightarrow c} \rangle \forall c \in NC_u$ 
12      (bestcomm, bestGain) =
13         $\arg \max_{c \in NCW_u} \Delta Q_{u \rightarrow c}$ 
14      if bestGain>0 and bestcomm  $\neq$  oldcomm of u
15        then
16          Move u to bestcomm and update C
17      // new modularity value
18      Q = computeModularity(V, E, C)
19      if  $Q - Q_{prev} \leq \tau$  then
20        break
21    changes=1
22  //Phase 2: Aggregation and Graph Reconstruction
23  if changes = 0 then
24    break
25  reconstructGraph(V, E, C)
```

---

# Louvain Pseudocode

For each vertex, build NCW by pulling community info from neighbors

---

**Algorithm 1:** Sequential Louvain Algorithm.

---

**Data:** Graph  $G = (V, E)$ ,  $\tau = \text{Threshold}$ .

**Result:** Final Community Assignment  $C$ , and Modularity  $Q$

```
1 //Outer Loop
2 while true do
3   changes=0
4   C = initialize each vertex in its own community
5   Q = computeModularity(V, E, C)
6   //Phase 1 (Inner loop): Modularity Optimization
7   while true do
8     Qprev = Q
9     for all  $u \in V$  do
10      //build NCW
11       $NCW_u = \langle c, w_{u \rightarrow c} \rangle \forall c \in NC_u$ 
12      (bestcomm, bestGain) =
13         $\arg \max_{c \in NCW_u} \Delta Q_{u \rightarrow c}$ 
14      if bestGain>0 and bestcomm  $\neq$  oldcomm of u
15      then
16        Move u to bestcomm and update C
17      // new modularity value
18      Q = computeModularity(V, E, C)
19      if  $Q - Q_{prev} \leq \tau$  then
20        break
21      changes=1
22  //Phase 2: Aggregation and Graph Reconstruction
23  if changes = 0 then
24    break
25  reconstructGraph(V, E, C)
```

---

# Louvain Pseudocode

Find the best community to move into by iterating through all entries of NCW

---

**Algorithm 1:** Sequential Louvain Algorithm.

---

**Data:** Graph  $G = (V, E)$ ,  $\tau = \text{Threshold}$ .  
**Result:** Final Community Assignment  $C$ , and Modularity  $Q$

```
1 //Outer Loop
2 while true do
3   changes=0
4   C = initialize each vertex in its own community
5   Q = computeModularity(V, E, C)
6   //Phase 1 (Inner loop): Modularity Optimization
7   while true do
8     Qprev = Q
9     for all  $u \in V$  do
10      //build NCW
11       $NCW_u = \langle c, w_{u \rightarrow c} \rangle \forall c \in NC_u$ 
12      (bestcomm, bestGain) =
13         $\arg \max_{c \in NCW_u} \Delta Q_{u \rightarrow c}$ 
14      if bestGain>0 and bestcomm  $\neq$  oldcomm of u
15        then
16          Move u to bestcomm and update C
17          // new modularity value
18          Q = computeModularity(V, E, C)
19          if  $Q - Q_{prev} \leq \tau$  then
20            break
21          changes=1
22      //Phase 2: Aggregation and Graph Reconstruction
23      if changes = 0 then
24        break
25    reconstructGraph(V, E, C)
```

---

# Louvain Pseudocode

---

**Algorithm 1:** Sequential Louvain Algorithm.

---

**Data:** Graph  $G = (V, E)$ ,  $\tau = Threshold$ .

**Result:** Final Community Assignment  $C$ , and Modularity  $Q$

```
1 //Outer Loop
2 while true do
3   changes=0
4   C = initialize each vertex in its own community
5   Q = computeModularity(V, E, C)
6   //Phase 1 (Inner loop): Modularity Optimization
7   while true do
8     Qprev = Q
9     for all  $u \in V$  do
10      //build NCW
11       $NCW_u = \langle c, w_{u \rightarrow c} \rangle \forall c \in NC_u$ 
12      (bestcomm, bestGain) =
13         $\arg \max_{c \in NCW_u} \Delta Q_{u \rightarrow c}$ 
14      if bestGain>0 and bestcomm  $\neq$  oldcomm of u
15        then
16          Move u to bestcomm and update C
17          // new modularity value
18          Q = computeModularity(V, E, C)
19          if  $Q - Q_{prev} \leq \tau$  then
20            break
21          changes=1
22      //Phase 2: Aggregation and Graph Reconstruction
23      if changes = 0 then
24        break
25  reconstructGraph(V, E, C)
```

---

Move to the best community and update community info

# Louvain Pseudocode

---

**Algorithm 1:** Sequential Louvain Algorithm.

---

**Data:** Graph  $G = (V, E)$ ,  $\tau = \text{Threshold}$ .

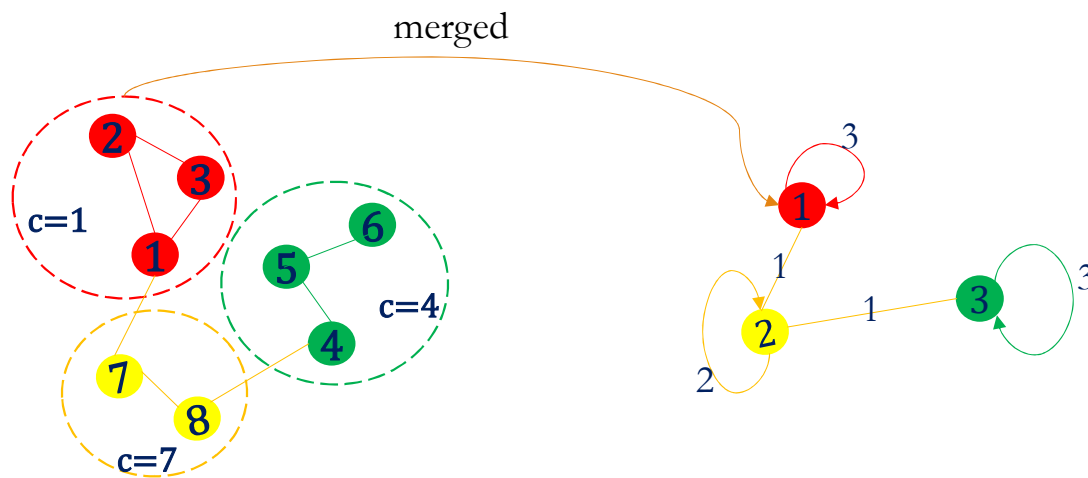
**Result:** Final Community Assignment  $C$ , and Modularity  $Q$

```
1 //Outer Loop
2 while true do
3   changes=0
4   C = initialize each vertex in its own community
5   Q = computeModularity(V, E, C)
6   //Phase 1 (Inner loop): Modularity Optimization
7   while true do
8     Qprev = Q
9     for all  $u \in V$  do
10      //build NCW
11       $NCW_u = \langle c, w_{u \rightarrow c} \rangle \forall c \in NC_u$ 
12      (bestcomm, bestGain) =
13         $\arg \max_{c \in NCW_u} \Delta Q_{u \rightarrow c}$ 
14      if bestGain > 0 and bestcomm  $\neq$  oldcomm of  $u$ 
15        then
16          Move  $u$  to bestcomm and update  $C$ 
17          // new modularity value
18          Q = computeModularity(V, E, C)
19          if  $Q - Q_{prev} \leq \tau$  then
20            break
21      changes=1
22   //Phase 2: Aggregation and Graph Reconstruction
23   if changes = 0 then
24     break
25   reconstructGraph(V, E, C)
```

---

Once done for all vertices, compute new modularity and repeat if modularity increased by a threshold

# Louvain Pseudocode



When modularity stabilizes, create a new graph by merging all vertices in same community into one

## Algorithm 1: Sequential Louvain Algorithm.

**Data:** Graph  $G = (V, E)$ ,  $\tau = \text{Threshold}$ .

**Result:** Final Community Assignment  $C$ , and Modularity  $Q$

```

1 //Outer Loop
2 while true do
3   changes=0
4   C = initialize each vertex in its own community
5   Q = computeModularity(V, E, C)
6   //Phase 1 (Inner loop): Modularity Optimization
7   while true do
8     Qprev = Q
9     for all  $u \in V$  do
10      //build NCW
11       $NCW_u = \langle c, w_{u \rightarrow c} \rangle \forall c \in NC_u$ 
12      (bestcomm, bestGain) =
13         $\arg \max_{c \in NCW_u} \Delta Q_{u \rightarrow c}$ 
14      if bestGain>0 and bestcomm  $\neq$  oldcomm of u
15      then
16        Move u to bestcomm and update C
17      // new modularity value
18      Q = computeModularity(V, E, C)
19      if  $Q - Q_{prev} \leq \tau$  then
20        break
21      changes=1
22 //Phase 2: Aggregation and Graph Reconstruction
23 if changes = 0 then
24   break
25 reconstructGraph(V, E, C)
  
```



# We call the standard Louvain Algorithm a Pull-based Louvain Algorithm

To build  $NCW$  at each iteration, it pulls latest info from neighbors

---

**Algorithm 1:** Sequential Louvain Algorithm.

---

**Data:** Graph  $G = (V, E)$ ,  $\tau = \text{Threshold}$ .

**Result:** Final Community Assignment  $C$ , and Modularity  $Q$

---

```
1 //Outer Loop
2 while true do
3   changes=0
4   C = initialize each vertex in its own community
5   Q = computeModularity(V, E, C)
6   //Phase 1 (Inner loop): Modularity Optimization
7   while true do
8     Qprev = Q
9     for all  $u \in V$  do
10      //build NCW
11       $NCW_u = \langle c, w_{u \rightarrow c} \rangle \forall c \in NC_u$ 
12      (bestcomm, bestGain) =
13         $\arg \max_{c \in NCW_u} \Delta Q_{u \rightarrow c}$ 
14      if bestGain>0 and bestcomm  $\neq$  oldcomm of u
15      then
16        Move u to bestcomm and update C
17      // new modularity value
18      Q = computeModularity(V, E, C)
19      if  $Q - Q_{prev} \leq \tau$  then
20        break
21      changes=1
22 //Phase 2: Aggregation and Graph Reconstruction
23 if changes = 0 then
24   break
25 reconstructGraph(V, E, C)
```

---

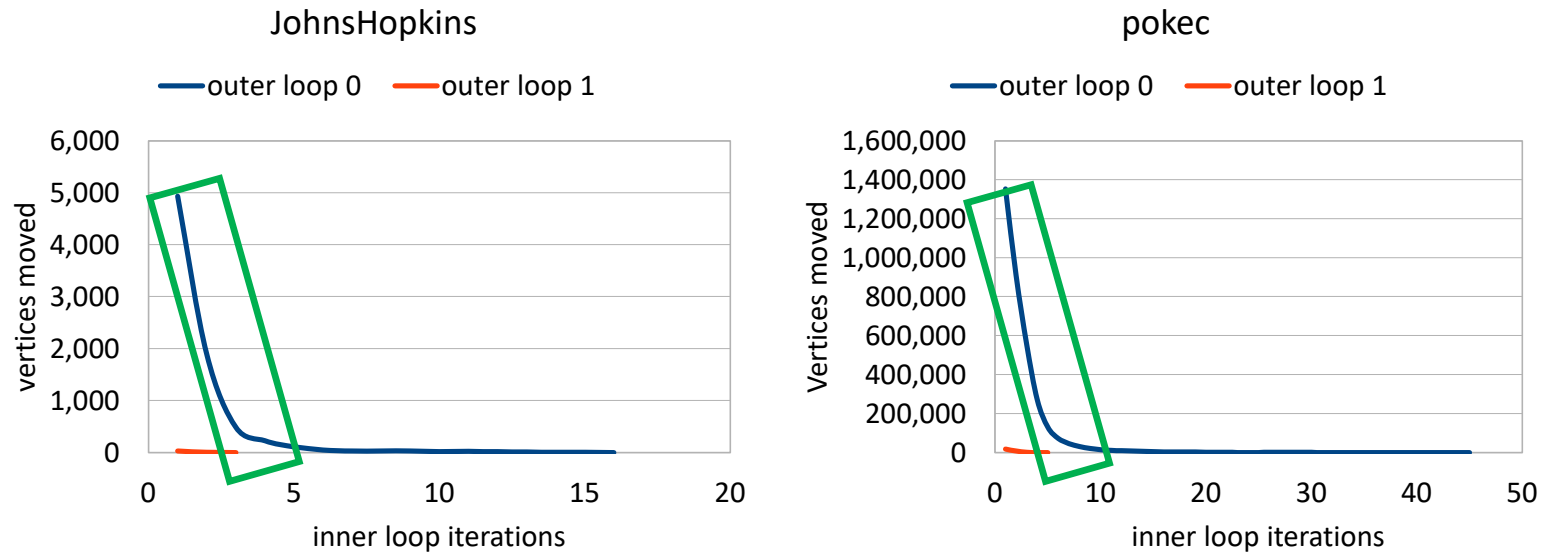
# Unnecessary work in Louvain

---

# Observations

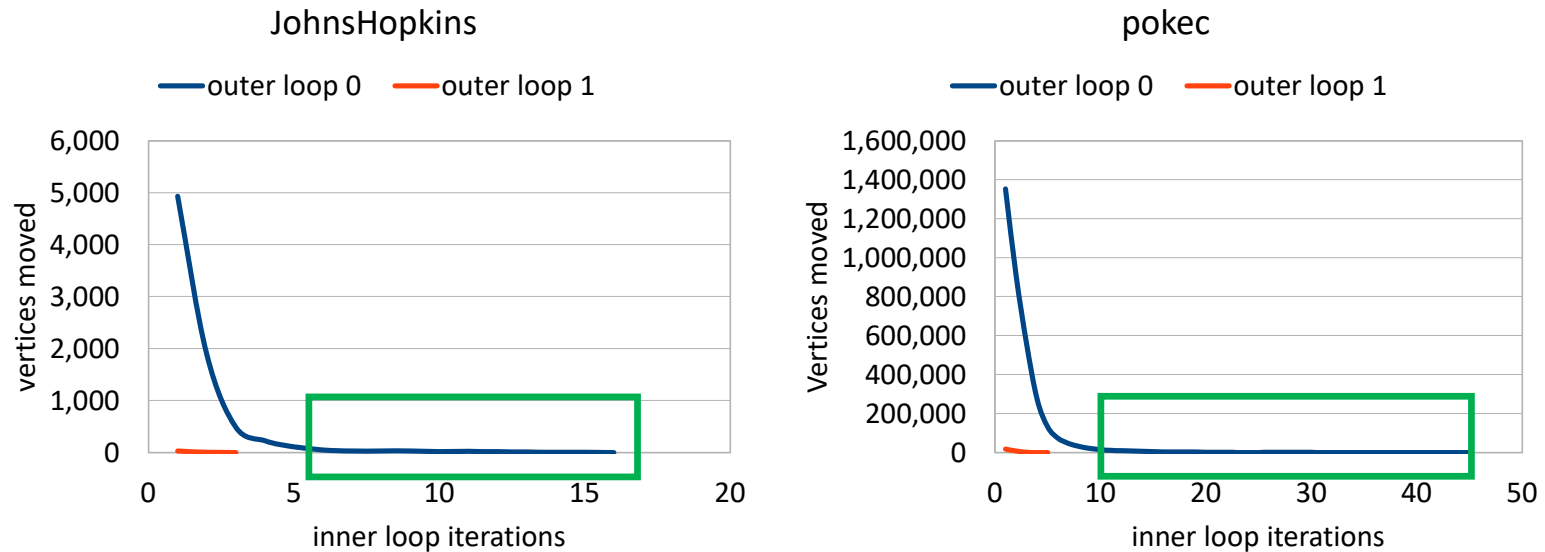


# Number of vertex moves drops significantly after the first few iterations of phase1



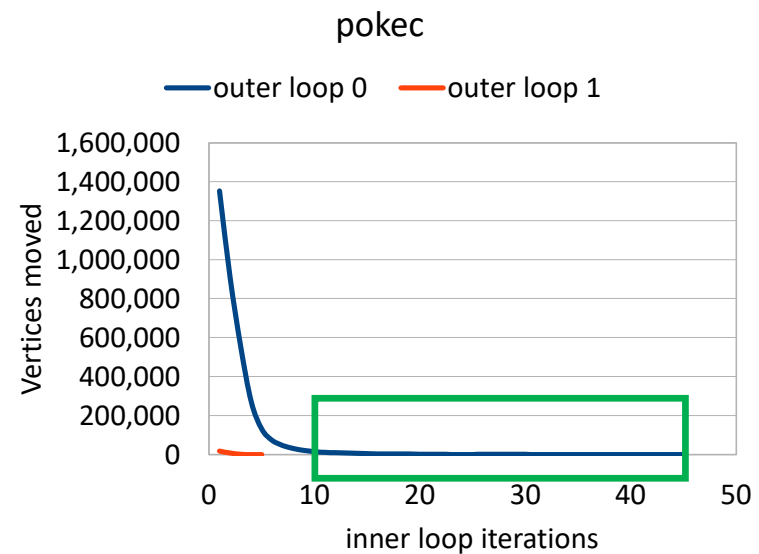
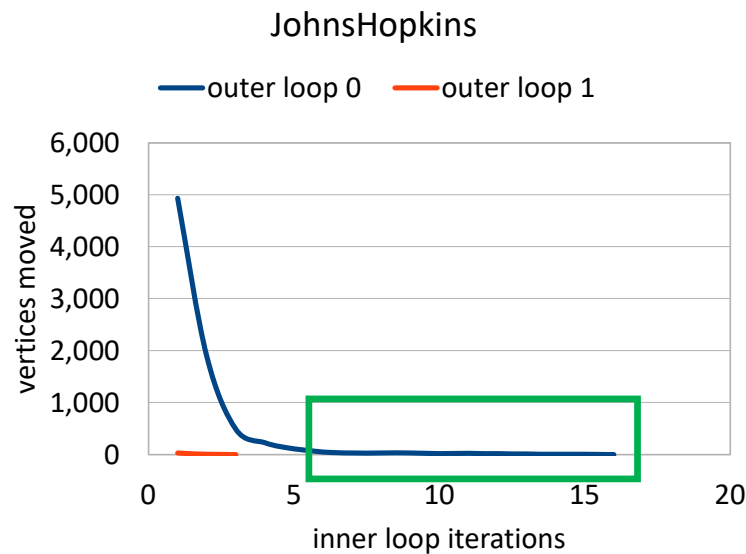
- For a particular outer loop, the number of vertices that change communities drops drastically after the first few inner loop iterations (e.g., 5).

# Number of vertex moves drops significantly after the first few iterations of phase1

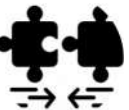


- The number of vertices that change communities in the later inner loop iterations is minimal

# Implications



- Wasteful to scan all neighbors to compute  $NCW$ , if no change in neighborhood
- Wasteful to iterate over all vertices for each iteration of phase 1, vertices do not move



# Pruning Unnecessary Work in Louvain

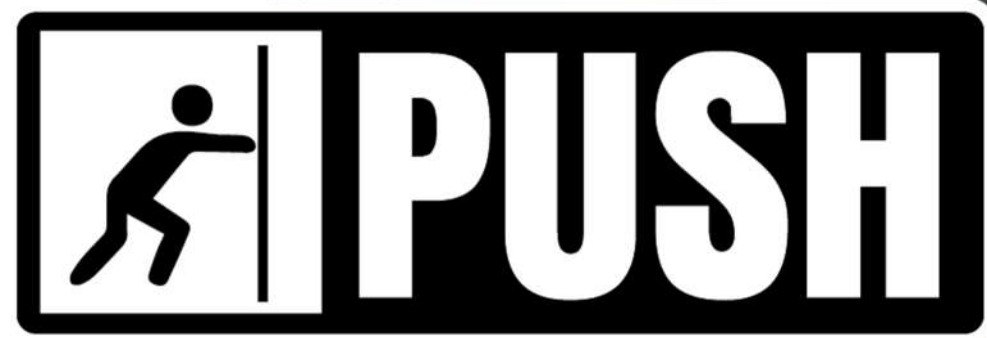
---

Prune vertices that are unlikely to move

Prune unnecessary neighborhood exploration

# Push-based Louvain Algorithm

Vertex does not pull, rather neighbors actively push any changes



# Push-based Louvain

The Push-based algorithm starts with an initialized  $NCW$ , assuming each vertex is in its own community

---

**Algorithm 2:** A push-based sequential Louvain algorithm.

---

**Data:** Graph  $G = (V, E)$ ,  $\tau = Threshold$ .

**Result:** Final Community Assignment  $C$ , and Modularity  $Q$

```
1 //Outer Loop
2 while true do
3   changes=0 C = initialize each vertex in its own
   community
4   Q = computeModularity(V, E, C)
5   //build initial  $NCW_u$  for all vertices
6    $NCW_u = \langle c, w_{u \rightarrow c} \rangle \forall c \in NC_u \forall u \in V$ 
7   //Phase 1 (Inner loop): Modularity Optimization;
8   while true do
9     Qprev = Q
10    for all  $u \in V$  do
11      (bestcomm, bestGain) =
         $\arg \max_{c \in NCW_u} \Delta Q_{u \rightarrow c}$ 
12      if bestGain > 0 and bestcomm  $\neq$  oldcomm of  $u$ 
        then
13        Move  $u$  to bestcomm and update  $C$ 
14        Update  $NCW_u$  for bestcomm and oldcomm
15        Update  $NCW_x \forall x \in N_u$ , for bestcomm and
        oldcomm
16      // new modularity value
17      Q = computeModularity(V, E, C)
18      if  $Q - Q_{prev} \leq \tau$  then
19        break
20    changes=1
21  //Phase 2: Aggregation and Graph Reconstruction
22  if changes = 0 then
23    break
24  reconstructGraph(V, E, C)
```

---

# Push-based Louvain

During Phase 1, it never recreates  $NCW$

---

**Algorithm 2:** A push-based sequential Louvain algorithm.

---

**Data:** Graph  $G = (V, E)$ ,  $\tau = Threshold$ .

**Result:** Final Community Assignment  $C$ , and Modularity  $Q$

```

1 //Outer Loop
2 while true do
3   changes=0 C = initialize each vertex in its own
   community
4   Q = computeModularity(V, E, C)
5   //build initial  $NCW_u$  for all vertices
6    $NCW_u = \langle c, w_{u \rightarrow c} \rangle \forall c \in NC_u \forall u \in V$ 
7   //Phase 1 (Inner loop): Modularity Optimization;
8   while true do
9     Qprev = Q
10    for all  $u \in V$  do
11      (bestcomm, bestGain) =
         $\arg \max_{c \in NCW_u} \Delta Q_{u \rightarrow c}$ 
12      if bestGain>0 and bestcomm  $\neq$  oldcomm of u
        then
13        Move u to bestcomm and update C
14        Update  $NCW_u$  for bestcomm and oldcomm
15        Update  $NCW_x \forall x \in N_u$ , for bestcomm and
        oldcomm
16      // new modularity value
17      Q = computeModularity(V, E, C)
18      if  $Q - Q_{prev} \leq \tau$  then
19        break
20      changes=1
21  //Phase 2: Aggregation and Graph Reconstruction
22  if changes = 0 then
23    break
24  reconstructGraph(V, E, C)
```

---

# Push-based Louvain

If there is a change in community membership

---

**Algorithm 2:** A push-based sequential Louvain algorithm.

---

**Data:** Graph  $G = (V, E)$ ,  $\tau = \text{Threshold}$ .

**Result:** Final Community Assignment  $C$ , and Modularity  $Q$

```

1 //Outer Loop
2 while true do
3   changes=0 C = initialize each vertex in its own
   community
4   Q = computeModularity(V, E, C)
5   //build initial  $NCW_u$  for all vertices
6    $NCW_u = \langle c, w_{u \rightarrow c} \rangle \forall c \in NC_u \forall c \in V$ 
7   //Phase 1 (Inner loop): Modularity Optimization;
8   while true do
9     Qprev = Q
10    for all  $u \in V$  do
11      (bestcomm, bestGain) =
         $\arg \max_{c \in NCW_u} \Delta Q_{u \rightarrow c}$ 
12      if bestGain>0 and bestcomm  $\neq$  oldcomm of  $u$ 
        then
13        Move  $u$  to bestcomm and update C
14        Update  $NCW_u$  for bestcomm and oldcomm
15        Update  $NCW_x \forall x \in N_u$ , for bestcomm and
        oldcomm
16        // new modularity value
17        Q = computeModularity(V, E, C)
18        if  $Q - Q_{prev} \leq \tau$  then
19          break
20        changes=1
21  //Phase 2: Aggregation and Graph Reconstruction
22  if changes = 0 then
23    break
24  reconstructGraph(V, E, C)
  
```

---

# Push-based Louvain

Update  $NCW$  for the vertex itself, and push updates to all its neighbors

---

**Algorithm 2:** A push-based sequential Louvain algorithm.

---

**Data:** Graph  $G = (V, E)$ ,  $\tau = Threshold$ .

**Result:** Final Community Assignment  $C$ , and Modularity  $Q$

---

```

1 //Outer Loop
2 while true do
3   changes=0 C = initialize each vertex in its own
      community
4   Q = computeModularity(V, E, C)
5   //build initial  $NCW_u$  for all vertices
6    $NCW_u = \langle c, w_{u \rightarrow c} \rangle \forall c \in NC_u \forall u \in V$ 
7   //Phase 1 (Inner loop): Modularity Optimization;
8   while true do
9     Qprev = Q
10    for all  $u \in V$  do
11      (bestcomm, bestGain) =
12         $\arg \max_{c \in NCW_u} \Delta Q_{u \rightarrow c}$ 
13      if bestGain > 0 and bestcomm  $\neq$  oldcomm of  $u$ 
14        then
15          Move  $u$  to bestcomm and update  $C$ 
16          Update  $NCW_u$  for bestcomm and oldcomm
17          Update  $NCW_x \forall x \in N_u$ , for bestcomm and
18            oldcomm
19          // new modularity value
20          Q = computeModularity(V, E, C)
21          if  $Q - Q_{prev} \leq \tau$  then
22            break
23          changes=1
24    //Phase 2: Aggregation and Graph Reconstruction
25    if changes = 0 then
26      break
27    reconstructGraph(V, E, C)
  
```

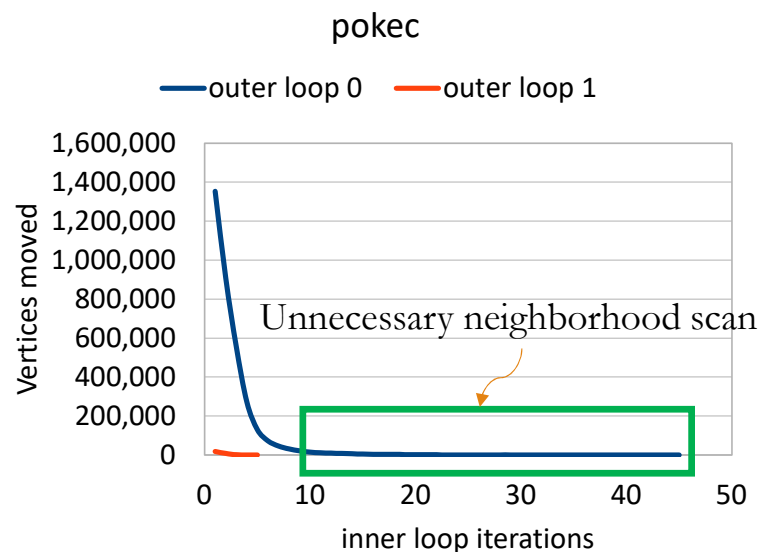
---

# Pros and Cons of Pull and Push

---

# Pull – Cons

Does redundant memory read  
by scanning all vertices and their neighbors to rebuild  
 $NCW$  for each inner loop,  
even when the vertex's neighborhood has not changed



## Algorithm 1: Sequential Louvain Algorithm.

**Data:** Graph  $G = (V, E)$ ,  $\tau = Threshold$ .

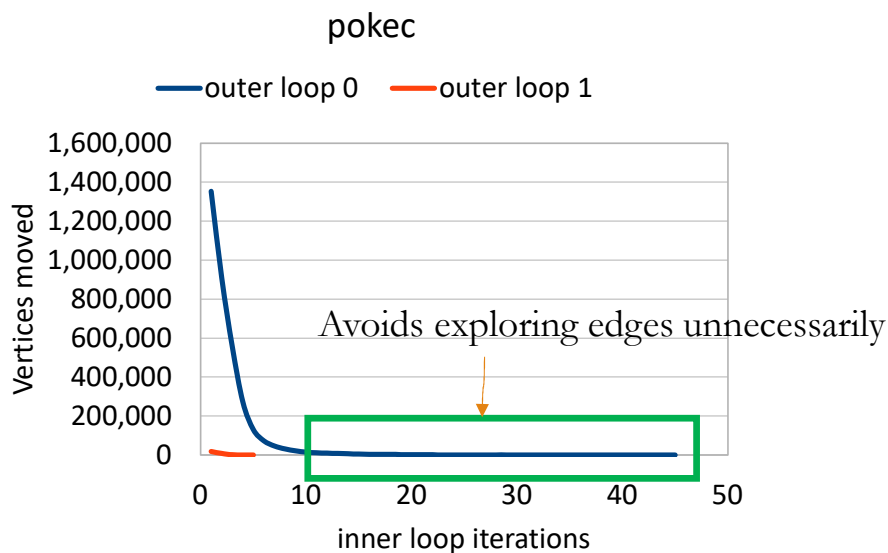
**Result:** Final Community Assignment  $C$ , and Modularity  $Q$

```

1 //Outer Loop
2 while true do
3   changes=0
4   C = initialize each vertex in its own community
5   Q = computeModularity(V, E, C)
6   //Phase 1 (Inner loop): Modularity Optimization
7   while true do
8     Qprev = Q
9     for all u ∈ V do
10      //build NCW
11       $NCW_u = \langle c, w_{u \rightarrow c} \rangle \forall c \in NC_u$ 
12      (bestcomm, bestGain) =
13         $\arg \max_{c \in NCW_u} \Delta Q_{u \rightarrow c}$ 
14      if bestGain>0 and bestcomm ≠ oldcomm of u
15        then
16          Move u to bestcomm and update C
17          // new modularity value
18          Q = computeModularity(V, E, C)
19          if Q - Qprev ≤ τ then
20            break
21      changes=1
22 //Phase 2: Aggregation and Graph Reconstruction
23 if changes = 0 then
24   break
25 reconstructGraph(V, E, C)
  
```

# Push – Pros

Scans through all neighbors of a vertex only when a vertex changes its community to update  $NCW$




---

**Algorithm 2:** A push-based sequential Louvain algorithm.

---

**Data:** Graph  $G = (V, E)$ ,  $\tau = Threshold$ .

**Result:** Final Community Assignment  $C$ , and Modularity  $Q$

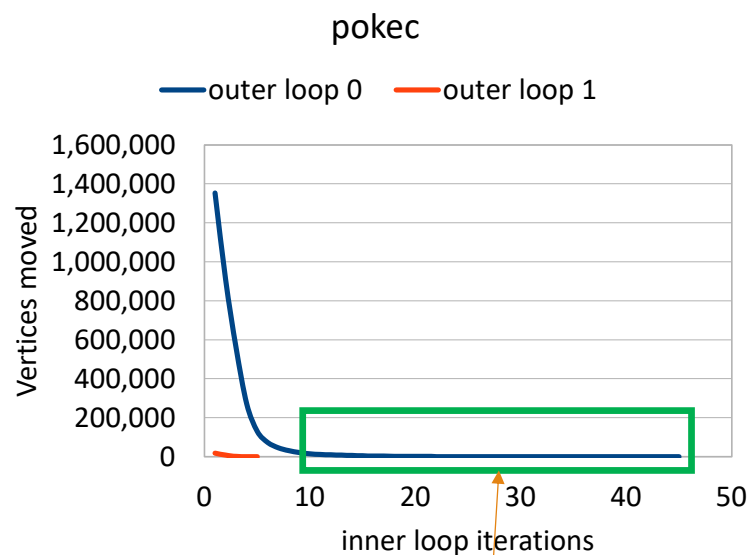
```

1 //Outer Loop
2 while true do
3   changes=0 C = initialize each vertex in its own
     community
4   Q = computeModularity(V, E, C)
5   //build initial  $NCW_u$  for all vertices
6    $NCW_u = \langle c, w_{u \rightarrow c} \rangle \forall c \in NC_u \forall u \in V$ 
7   //Phase 1 (Inner loop): Modularity Optimization;
8   while true do
9     Qprev = Q
10    for all  $u \in V$  do
11      (bestcomm, bestGain) =
         $\arg \max_{c \in NCW_u} \Delta Q_{u \rightarrow c}$ 
12      if bestGain > 0 and bestcomm  $\neq$  oldcomm of  $u$ 
        then
13        Move  $u$  to bestcomm and update  $C$ 
14        Update  $NCW_u$  for bestcomm and oldcomm
15        Update  $NCW_x \forall x \in N_u$ , for bestcomm and
          oldcomm
16      // new modularity value
17      Q = computeModularity(V, E, C)
18      if  $Q - Q_{prev} \leq \tau$  then
19        break
20    changes=1
21  //Phase 2: Aggregation and Graph Reconstruction
22  if changes = 0 then
23    break
24  reconstructGraph(V, E, C)
  
```

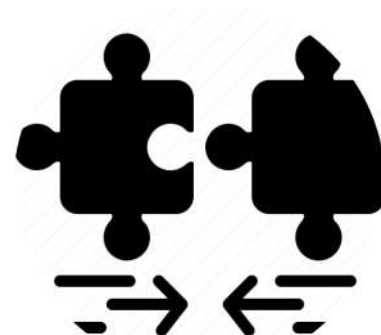
---

# Implications

A push-based Louvain algorithm is likely to do fewer edge explorations compared to a pull-based

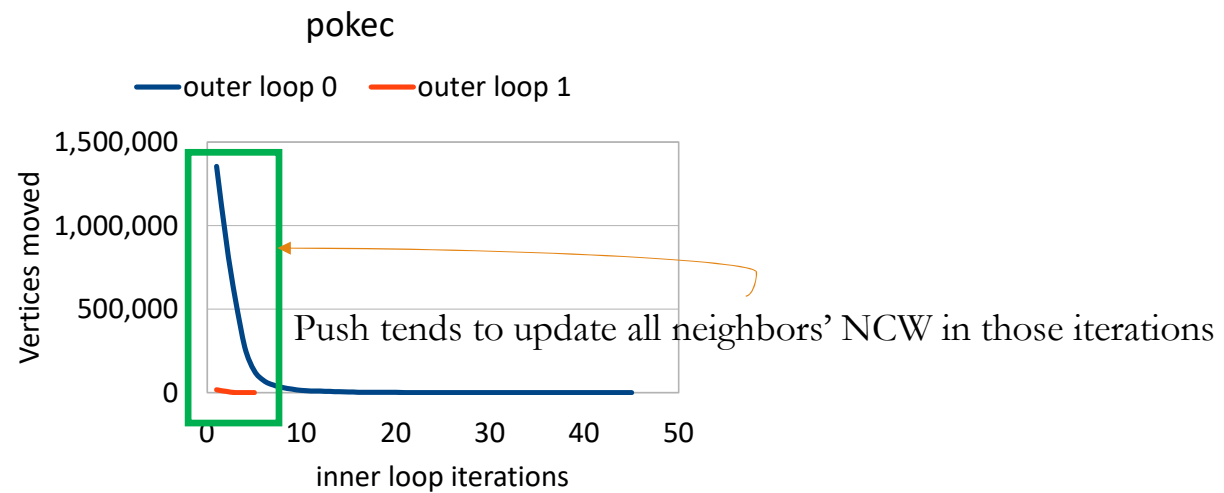


in the later inner loop iterations



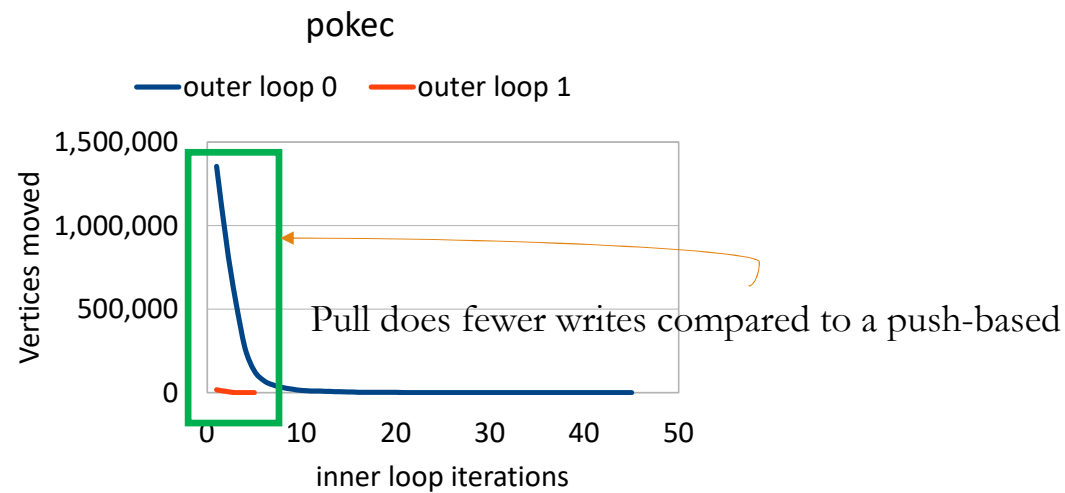
# Push – Cons

Push does more writes to memory compared to a pull-based when there is a lot of moves



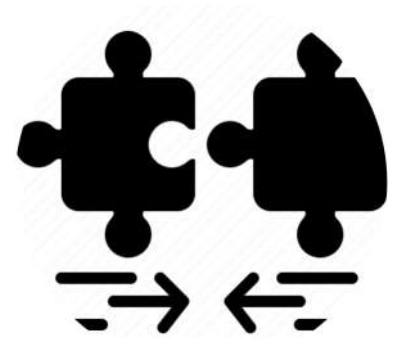
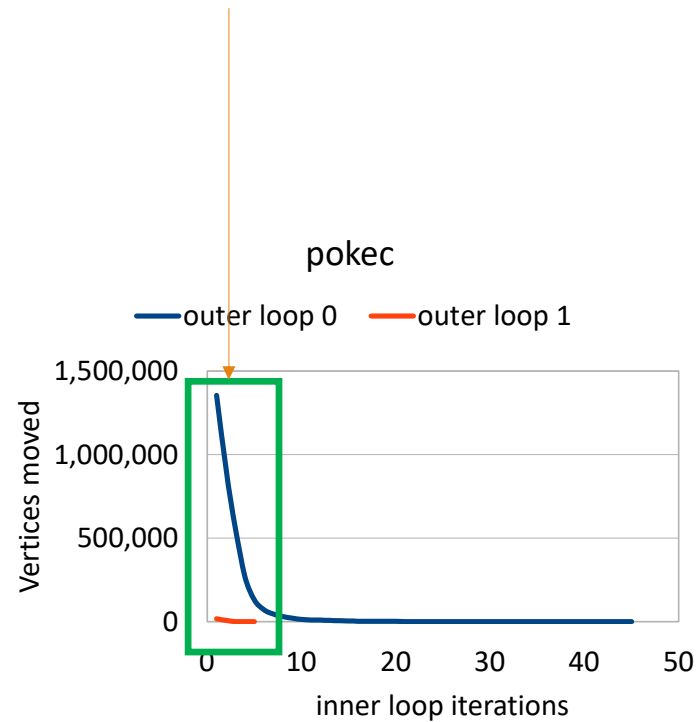
# Pull – Pros

Pull **does fewer writes** compared to a push-based algorithm when there is a lot of moves



# Implications

Using a push-based algorithm in the first few inner loop iterations **might not be beneficial**



# Take-home Message

---

Neither pull nor push performs best across all iteration space



# Pull-Push/Hybrid Louvain

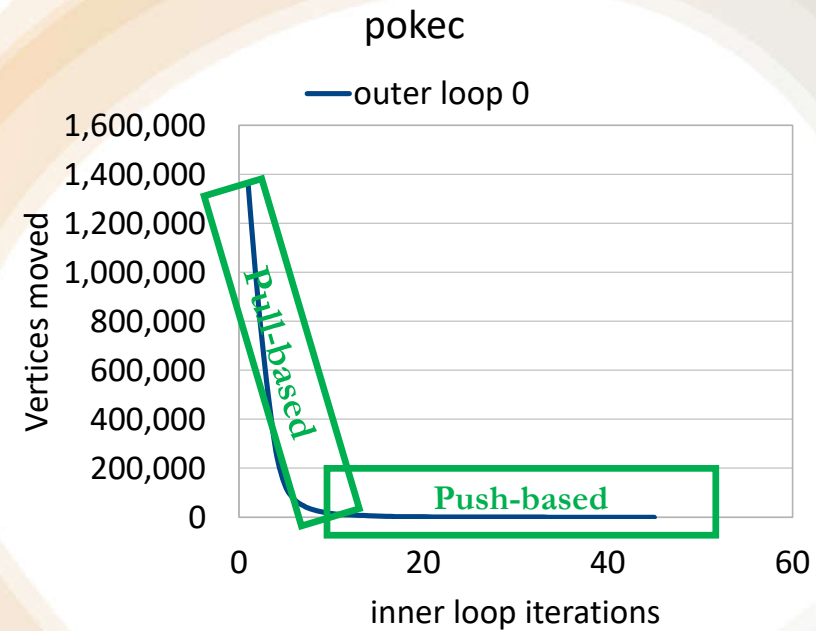
Best of both worlds

# Pull-Push Louvain Algorithm

## How it works

For a given outer loop

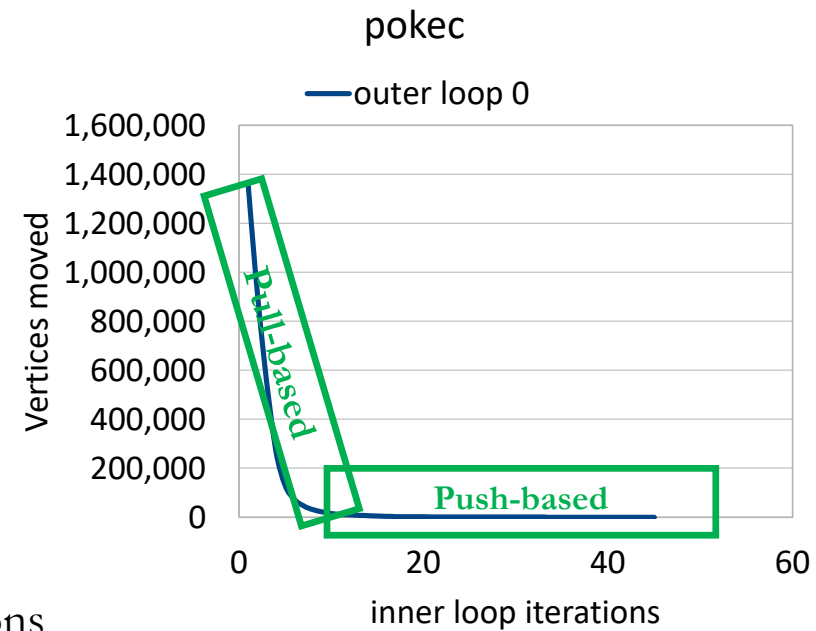
- Start with a pull-based
- Switch to a push-based after a given #of iterations



# Pull-Push Louvain Algorithm

## Benefits

- Explores a vertex's neighborhood when there a change
- Automatically prunes a significant number of edge-explorations



# Automatic Edge Pruning of Pull-Push Algorithm

Graph: POKEC

Algorithmic Improvement		pull	hybrid
Vertices	Visited	833M	833M
	<b>Reduction</b>	<b>1.00x</b>	<b>1.00x</b>
Edges	Visited	2.34G	0.18G
	<b>Reduction</b>	<b>1.00x</b>	<b>12.82x</b>

Graph: Hollywood

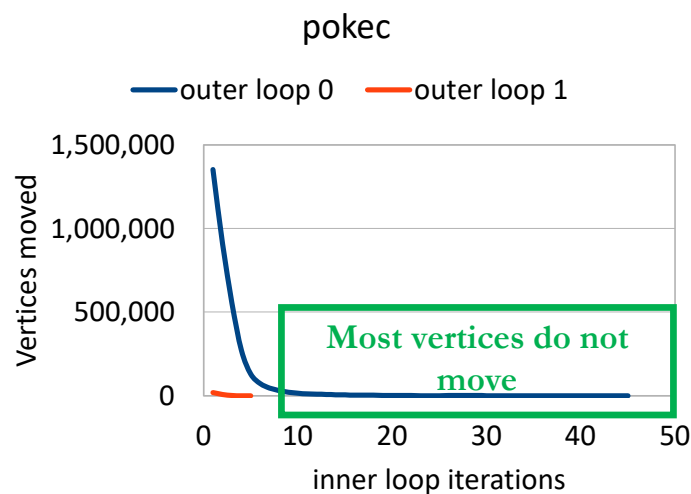
Algorithmic Improvement		pull	hybrid
Vertices	Visited	27.7M	27.7M
	<b>Reduction</b>	<b>1.00x</b>	<b>1.00x</b>
Edges	Visited	2.82G	0.45G
	<b>Reduction</b>	<b>1.00x</b>	<b>6.20x</b>

Prunes **6-13 × edges** compared to a standard (pull-based) Louvain

# Vertex Pruning

# Vertex Pruning

- It is unnecessary to iterate over all vertices - number of vertices changing community drops significantly after the first few inner loop iterations



- We show analytical and intuitive derivation of the vertices that can be pruned with minimal sacrifice

# Vertex Pruning: Analytical Derivation

Modularity gain by moving a vertex  $u$  to a community  $c$ :

$$Q_{u \rightarrow c} = \left[ \frac{w_{u \rightarrow c}}{2m} - \frac{\sum_{tot}^c w(u)}{2m^2} \right] = \frac{1}{2m} \left[ w_{u \rightarrow c} - \frac{\sum_{tot}^c w(u)}{m} \right]$$

# Vertex Pruning: Analytical Derivation

Modularity gain by moving a vertex  $u$  to a community  $c$ :

$$Q_{u \rightarrow c} = \left[ \frac{w_{u \rightarrow c}}{2m} - \frac{\sum_{tot}^c w(u)}{2m^2} \right] = \frac{1}{2m} \left[ w_{u \rightarrow c} - \frac{\sum_{tot}^c w(u)}{m} \right]$$

$w_{u \rightarrow c}$  = sum of edge weights from  $u$  to community  $c$

$$\sum_{tot}^c = \sum W_{u,v} , \text{ for all } u \in c \text{ or } v \in c, m = \sum_{e(u,v)} W_{u,v}$$

# Vertex Pruning: Analytical Derivation

Modularity gain by moving a vertex  $u$  to a community  $c$ :

$$Q_{u \rightarrow c} = \left[ \frac{w_{u \rightarrow c}}{2m} - \frac{\sum_{tot}^c w(u)}{2m^2} \right] = \frac{1}{2m} \left[ w_{u \rightarrow c} - \frac{\sum_{tot}^c w(u)}{m} \right]$$

$$\text{Let, } Cm = \frac{\sum_{tot}^c}{m}, Cm = (0, 1)$$

Total graph edge

Total community edge

# Vertex Pruning: Analytical Derivation

Modularity gain by moving a vertex  $u$  to a community  $c$ :

$$Q_{u \rightarrow c} = \left[ \frac{w_{u \rightarrow c}}{2m} - \frac{\sum_{tot}^c w(u)}{2m^2} \right] = \frac{1}{2m} \left[ w_{u \rightarrow c} - \frac{\sum_{tot}^c w(u)}{m} \right]$$

$$\text{if, } Cw = Cm * w(u) \Rightarrow Q_{u \rightarrow c} \sim w_{u \rightarrow c} - Cw$$

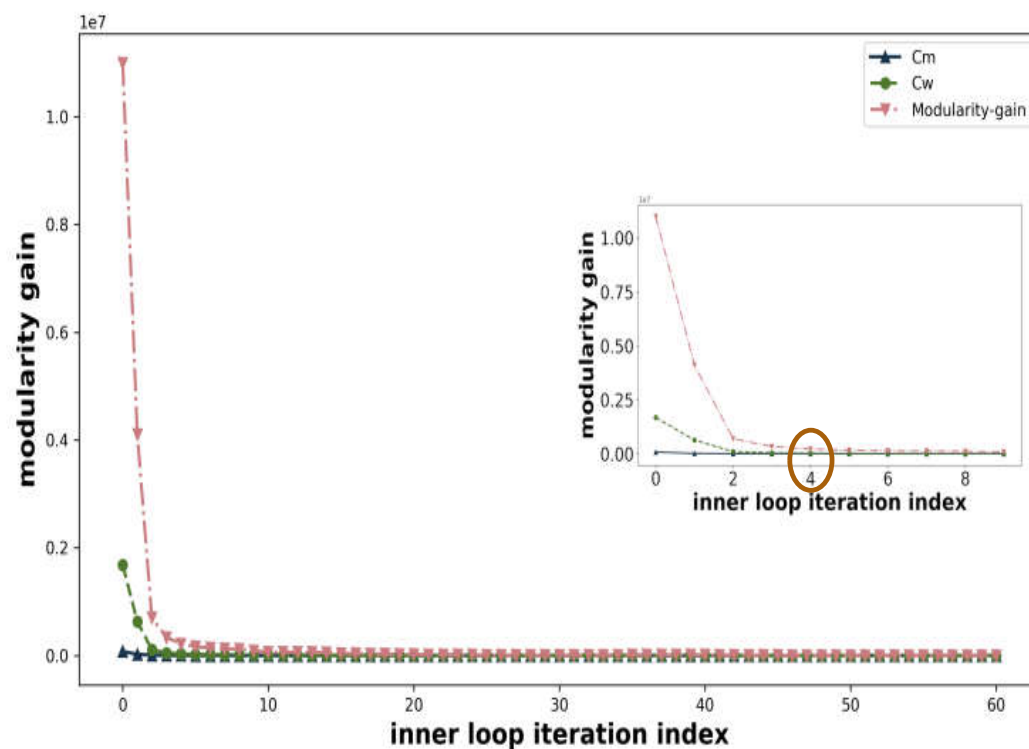
total edges of vertex  $u$

# Vertex Pruning: Analytical Derivation

Modularity gain by moving a vertex  $u$  to a community  $c$ :

$$\text{Let, } Cm = \frac{\Sigma_{tot}^c}{m}, Cm = (0, 1)$$

$$Cw = Cm * w(u) \Rightarrow Q_{u \rightarrow c} \sim w_{u \rightarrow c} - Cw$$



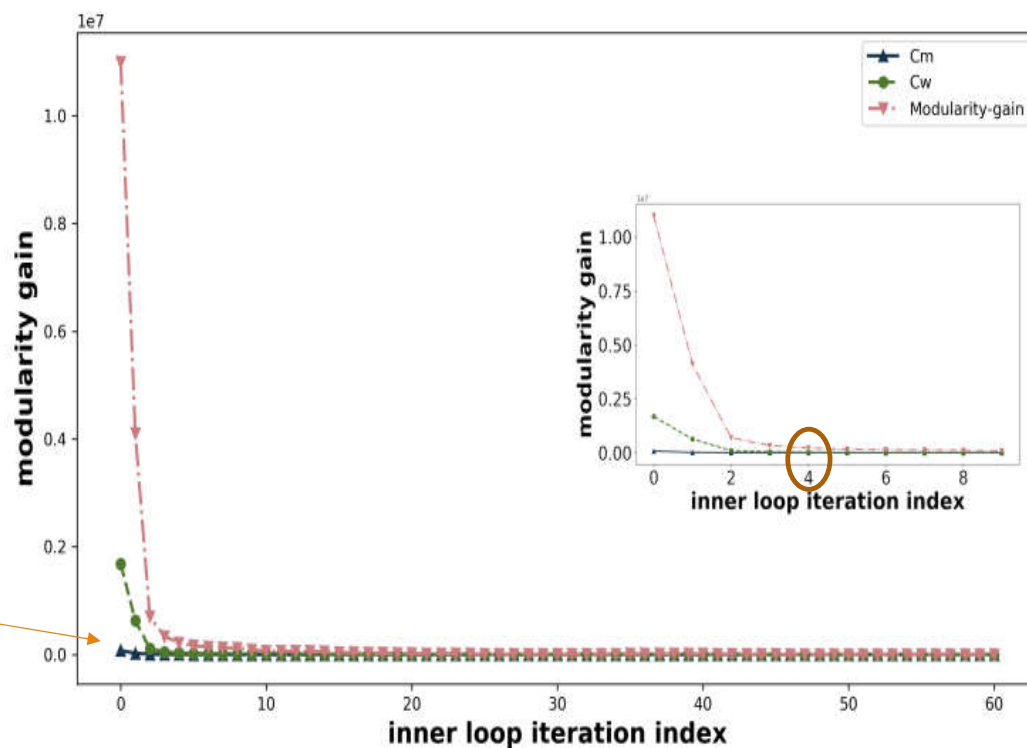
# Vertex Pruning: Analytical Derivation

Modularity gain by moving a vertex  $u$  to a community  $c$ :

$$\text{Let, } Cm = \frac{\sum_{tot}^c}{m}, Cm = (0, 1)$$

$$Cw = Cm * w(u) \Rightarrow Q_{u \rightarrow c} \sim w_{u \rightarrow c} - Cw$$

$Cm$  does not play an important role in the modularity gain



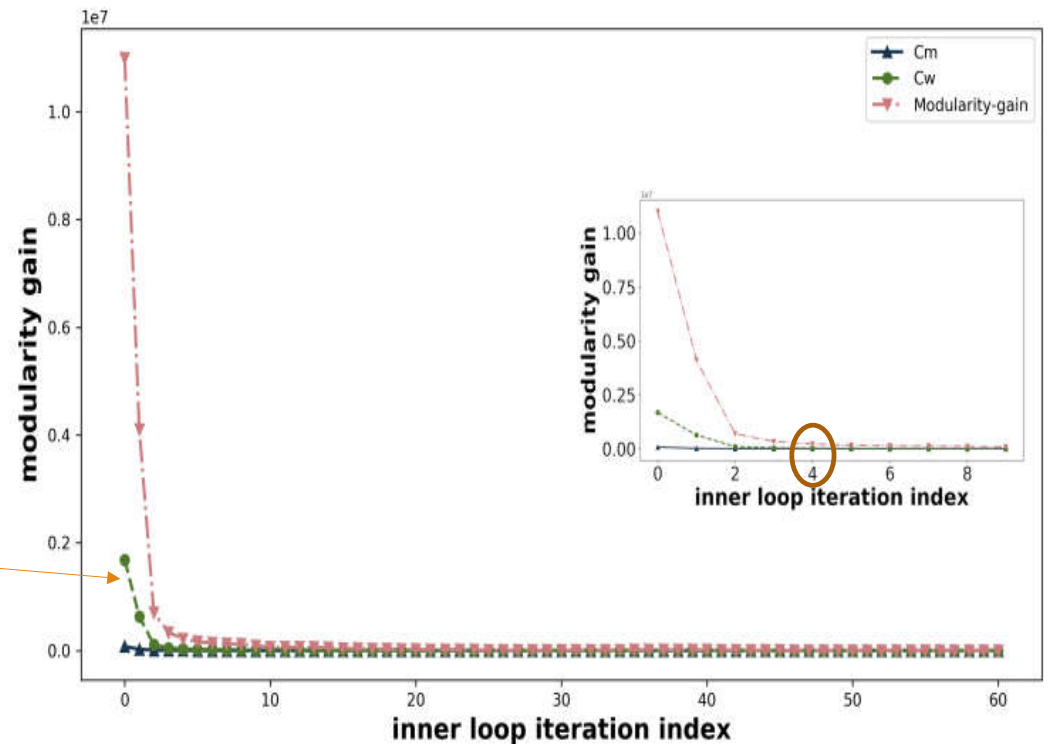
# Vertex Pruning: Analytical Derivation

Modularity gain by moving a vertex  $u$  to a community  $c$ :

$$\text{Let, } Cm = \frac{\Sigma_{tot}^c}{m}, Cm = (0, 1)$$

$$Cw = Cm * w(u) \Rightarrow Q_{u \rightarrow c} \sim w_{u \rightarrow c} - Cw$$

After first few iterations,  $Cw$  does not play an important role in the modularity gain



# Vertex Pruning: Analytical Derivation

Modularity gain by moving a vertex  $u$  to a community  $c$ :

$$\text{Let, } Cm = \frac{\Sigma_{tot}^c}{m}, Cm = (0, 1)$$

$$Cw = Cm * w(u) \Rightarrow Q_{u \rightarrow c} \sim w_{u \rightarrow c} - Cw$$

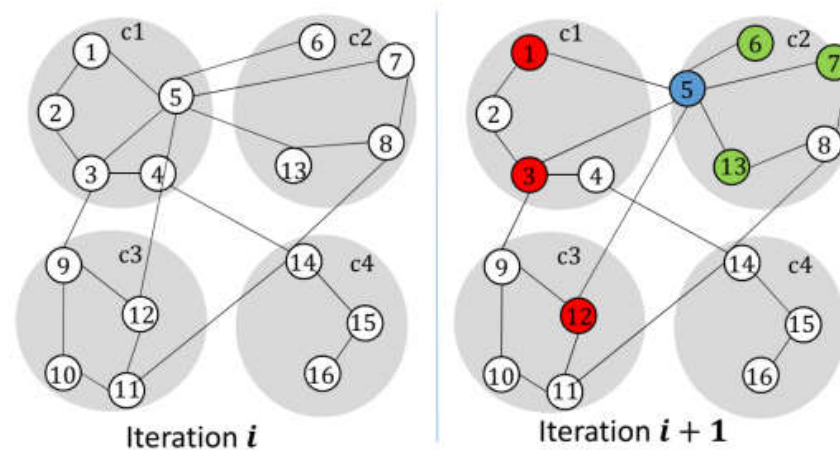


**Intuitive:** Focus on the vertices whose  $w_{u \rightarrow c}$  decreases

**Intuitive:** Skip the vertices whose  $\Sigma_{tot}^c$  is impacted by a move

Impact on modularity is small if applied on the push-phases – later inner loop iterations

# Vertex Pruning: Analytical Derivation



## What to recompute? (Analytical Derivation in Paper)

If a vertex moves, only recompute for its first level neighbors that are **\*not\*** in its new community  
=> recompute **red** neighbors, impacts on green and blue are minimal, no impact on white

# Algorithms and Impact of Vertex & Edge Pruning

Algorithm name	What it does
<b>Pull</b>	Standard pull-based Louvain
<b>Pull-prune</b>	Pull + <b>vertex pruning in all iterations</b>
<b>Hybrid</b>	Switching between pull and push
<b>Hybrid-prune</b>	Hybrid + <b>vertex pruning in push phases only</b>

Graph: POKEC

Algorithmic Improvement		pull	pull prune	hybrid	hybrid prune
Vertices	Visited	833M	6.09M	833M	9.49M
	<b>Reduction</b>	<b>1.00x</b>	<b>13.68x</b>	<b>1.00x</b>	<b>8.78x</b>
Edges	Visited	2.34G	0.3G	0.18G	0.182G
	<b>Reduction</b>	<b>1.00x</b>	<b>7.70x</b>	<b>12.82x</b>	<b>12.82x</b>

Graph: Hollywood

Algorithmic Improvement		pull	pull prune	hybrid	hybrid prune
Vertices	Visited	27.7M	5.44M	27.7M	6.12M
	<b>Reduction</b>	<b>1.00x</b>	<b>5.09x</b>	<b>1.00x</b>	<b>4.52x</b>
Edges	Visited	2.82G	0.95G	0.45G	0.45G
	<b>Reduction</b>	<b>1.00x</b>	<b>2.98x</b>	<b>6.20x</b>	<b>6.20x</b>

- Prune 4 to 12× vertices

# Algorithms and Impact of Vertex & Edge Pruning

Algorithm name	What it does
<b>Pull</b>	Standard pull-based Louvain
<b>Pull-prune</b>	Pull + <b>vertex pruning in all iterations</b>
<b>Hybrid</b>	Switching between pull and push
<b>Hybrid-prune</b>	Hybrid + <b>vertex pruning in push phases only</b>

Graph: POKEC

Algorithmic Improvement		pull	pull prune	hybrid	hybrid prune
Vertices	Visited	833M	6.09M	833M	9.49M
	<b>Reduction</b>	<b>1.00x</b>	<b>13.68x</b>	<b>1.00x</b>	<b>8.78x</b>
Edges	Visited	2.34G	0.3G	0.18G	0.182G
	<b>Reduction</b>	<b>1.00x</b>	<b>7.70x</b>	<b>12.82x</b>	<b>12.82x</b>

Graph: Hollywood

Algorithmic Improvement		pull	pull prune	hybrid	hybrid prune
Vertices	Visited	27.7M	5.44M	27.7M	6.12M
	<b>Reduction</b>	<b>1.00x</b>	<b>5.09x</b>	<b>1.00x</b>	<b>4.52x</b>
Edges	Visited	2.82G	0.95G	0.45G	0.45G
	<b>Reduction</b>	<b>1.00x</b>	<b>2.98x</b>	<b>6.20x</b>	<b>6.20x</b>

- Does not prune additional edges compared to hybrid

# Performance Benefit using Single Thread

Graphs	Pull		Hybrid		Pull-Prune		Hybrid-Prune	
	Q	T	Q	T	Q	T	Q	T
Wikipedia	0.57	98.9	0.57	74.1	0.57	65.3	0.57	61.9
Hollywood	0.73	57.2	0.73	14.8	0.73	31.5	0.73	12.9
POKEC	0.68	19.3	0.68	11.1	0.68	4.82	0.68	5.7

Q= Modularity, T= Time (s)

- Edge pruning :  $1.3\times - 3.9\times$
- Vertex pruning :  $1.5\times - 4\times$
- Vertex pruning on top of edge pruning: upto  $1.9\times$

# Take-home Message

---

Even without any parallelization, edge and vertex pruning gives up to 4x speedup over the standard Louvain algorithm.

# Parallel Pull, Push, Pull-Push Algorithms

# Parallel Pull-based Louvain

Private hashmap for each thread

---

**Algorithm 3:** Parallel Pull-based inner loop of Louvain.

---

**Data:** Graph  $G = (V, E)$ ,  $\tau = \text{Threshold}$ .

**Result:** Final Community Assignment  $C$ , and Modularity  $Q$

```
1 //Phase 1 (Inner loop): Modularity Optimization
2 while true do
3    $Q_{\text{prev}} = Q$ 
4   Parallel for all  $u \in V$  do
5     //build NCW
6      $NCW_u = \langle c, w_{u \rightarrow c} \rangle \forall c \in NC_u$ 
7      $(\text{bestcomm}, \text{bestGain}) = \arg \max_{c \in NCW_u} \Delta Q_{u \rightarrow c}$ 
8     if  $\text{bestGain} > 0$  and  $\text{bestcomm} \neq \text{oldcomm of } u$  then
9       | Move  $u$  to  $\text{bestcomm}$  and update  $C$  atomically
10    //new modularity value
11     $Q = \text{Parallel computeModularity}(V, E, C)$ 
12    if  $Q - Q_{\text{prev}} \leq \tau$  then
13      | break
14     $\text{changes} = 1$ 
```

---

# Parallel Pull-based Louvain

---

**Algorithm 3:** Parallel Pull-based inner loop of Louvain.

---

**Data:** Graph  $G = (V, E)$ ,  $\tau = \text{Threshold}$ .

**Result:** Final Community Assignment  $C$ , and Modularity  $Q$

1 //Phase 1 (Inner loop): Modularity Optimization

2 **while** *true* **do**

3    $Q_{\text{prev}} = Q$

4   **Parallel for** *all*  $u \in V$  **do**

5       //build NCW

6        $NCW_u = \langle c, w_{u \rightarrow c} \rangle \forall c \in NC_u$

7        $(\text{bestcomm}, \text{bestGain}) = \arg \max_{c \in NCW_u} \Delta Q_{u \rightarrow c}$

8       **if**  $\text{bestGain} > 0$  *and*  $\text{bestcomm} \neq \text{oldcomm of } u$  **then**

9           | Move  $u$  to  $\text{bestcomm}$  and update  $C$  **atomically**

10      //new modularity value

11       $Q = \text{Parallel computeModularity}(V, E, C)$

12      **if**  $Q - Q_{\text{prev}} \leq \tau$  **then**

13          | break

14      changes=1

---

For each vertex in parallel

# Parallel Pull-based Louvain

---

**Algorithm 3:** Parallel Pull-based inner loop of Louvain.

---

**Data:** Graph  $G = (V, E)$ ,  $\tau = \text{Threshold}$ .

**Result:** Final Community Assignment  $C$ , and Modularity  $Q$

1 //Phase 1 (Inner loop): Modularity Optimization

2 **while** *true* **do**

3      $Q_{\text{prev}} = Q$

4     **Parallel for** *all*  $u \in V$  **do**

5         //build NCW

6          $NCW_u = \langle c, w_{u \rightarrow c} \rangle \forall c \in NC_u$

7          $(\text{bestcomm}, \text{bestGain}) = \arg \max_{c \in NCW_u} \Delta Q_{u \rightarrow c}$

8         **if** *bestGain* > 0 and *bestcomm*  $\neq$  *oldcomm* of *u* **then**

9             |     |     Move *u* to *bestcomm* and update *C* **atomically**

10        //new modularity value

11      $Q = \text{Parallel computeModularity}(V, E, C)$

12     **if**  $Q - Q_{\text{prev}} \leq \tau$  **then**

13         |     break

14     changes=1

---

Change community membership atomically

# Parallel Pull-based Louvain

---

**Algorithm 3:** Parallel Pull-based inner loop of Louvain.

---

**Data:** Graph  $G = (V, E)$ ,  $\tau = \text{Threshold}$ .

**Result:** Final Community Assignment  $C$ , and Modularity  $Q$

```
1 //Phase 1 (Inner loop): Modularity Optimization
2 while true do
3    $Q_{\text{prev}} = Q$ 
4   Parallel for all  $u \in V$  do
5     //build NCW
6      $NCW_u = \langle c, w_{u \rightarrow c} \rangle \forall c \in NC_u$ 
7      $(\text{bestcomm}, \text{bestGain}) = \arg \max_{c \in NCW_u} \Delta Q_{u \rightarrow c}$ 
8     if  $\text{bestGain} > 0$  and  $\text{bestcomm} \neq \text{oldcomm of } u$  then
9       | Move  $u$  to bestcomm and update  $C$  atomically
10    //new modularity value
11     $Q = \text{Parallel computeModularity}(V, E, C)$ 
12    if  $Q - Q_{\text{prev}} \leq \tau$  then
13      | break
14    changes=1
```

---

Compute modularity using parallel reduction

# Parallel Push-based

Shared hashmap of size  $O(E)$

Update hashmaps using Locks

---

**Algorithm 4:** Parallel Push-based inner loop of Louvain.

---

**Data:** Graph  $G = (V, E)$ ,  $\tau = \text{Threshold}$ .

**Result:** Final Community Assignment  $C$ , and Modularity  $Q$

```
1 //build initial  $NCW_u$  for all vertices
2 Parallel  $\forall u \in V$  do  $NCW_u = \langle c, w_{u \rightarrow c} \rangle \forall c \in NC_u$ 
3 //Phase 1 (Inner loop): Modularity Optimization
4 while true do
5    $Q_{prev} = Q$ 
6   Parallel for all  $u \in V$  do
7      $(bestcomm, bestGain) = \arg \max_{c \in NCW_u} \Delta Q(u)$ 
8     moving to  $c$ 
9     if  $bestGain > 0$  and  $bestcomm \neq oldcomm$  of  $u$  then
10      Move  $u$  to  $bestcomm$  and update  $C$  atomically
11      lock  $NCW_u$ 
12      Update  $NCW_u$  for  $bestcomm$  and  $oldcomm$ 
13      unlock  $NCW_u$ 
14      lock  $NCW_x$ 
15      Update  $NCW_x \forall x \in N_u$ , for  $bestcomm$  and
16      oldcomm
17      unlock  $NCW_x$ 
18   //new modularity value
19    $Q = \text{Parallel computeModularity}(V, E, C)$ 
20   if  $Q - Q_{prev} \leq \tau$  then
21     break
22   changes=1
```

---

# Experimental Results

Graphs	V	E	Graphs	V	E
CA	1.08E+05	1.87E+05	CitationCiteSeer	2.68E+05	2.31E+06
CaidaRouterLevel	1.92E+05	1.22E+06	CoAuthorsDBLP	2.99E+05	1.96E+06
POKEC	5.40E+05	3.05E+07	CoPapersCiteSeer	4.34E+05	3.21E+07
Hollywood	1.14E+06	1.13E+08	Amazon	5.49E+05	1.85E+06
Wikipedia	3.97E+07	9.01E+07	As-Skitter	1.70E+06	2.22E+07
Uk-2005	1.68E+07	3.96E+08	Rgg_n_2_24_s0	1.68E+07	2.65E+08
Friendster	6.65E+07	1.89E+09	Webbase-2001	1.18E+08	1.02E+09

# Input Graphs

# Performance Analysis Platform

## Experimental Platforms

Platform Metric	Platform 1	Platform 2
Processor	Intel <sup>(R)</sup> Xeon <sup>(R)</sup> Platinum 8180	Intel(R) Xeon(R) CPU E7-8880 v3
CPU Clock	2.50GHz	2.30GHz
Sockets	2	4
Cores	56 (each socket has 28)	72 (each socket with 18 cores)
L3 Cache	97 MB	46.1 MB
Memory Speed	2666 MHz	1200 MHz
Memory Size	196.7GB	1 TB
Compiler	Intel ICC 18.0	
Parallel Program	C with OpenMP	

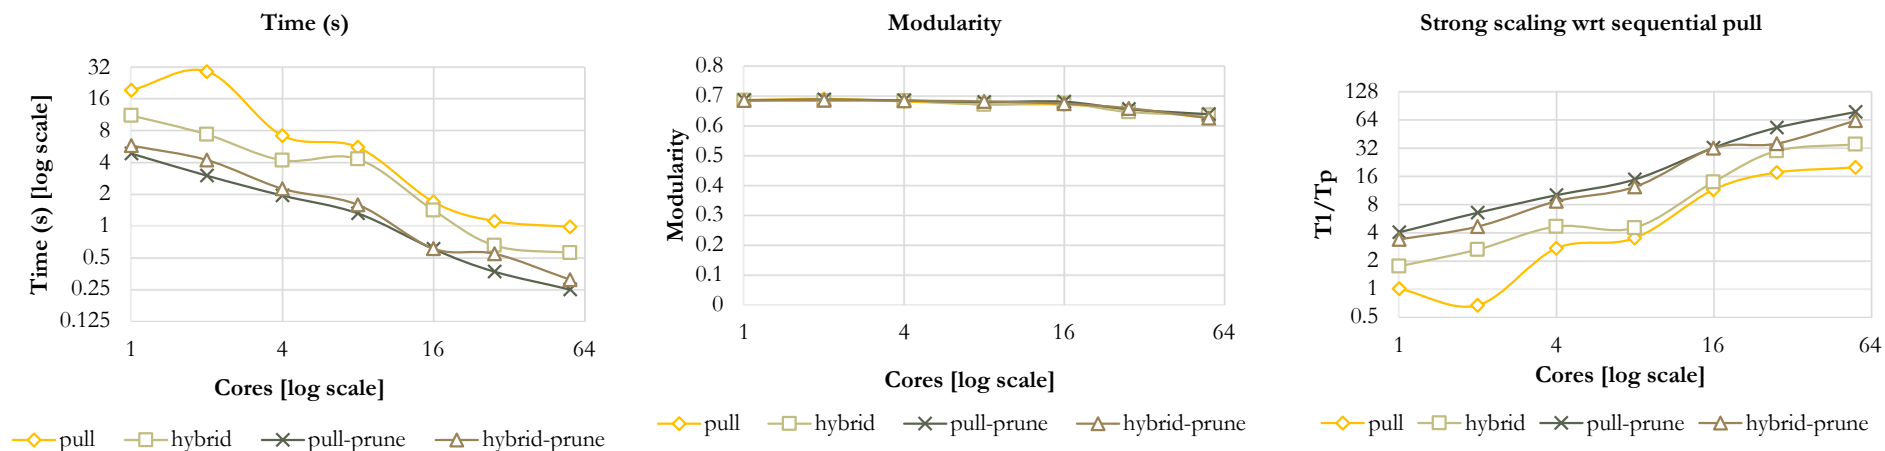


# Algorithms

Algorithm name	What is does
<b>Pull</b>	Standard pull-based Louvain
<b>Pull-prune</b>	Pull + vertex pruning in all iterations
<b>Hybrid</b>	Switching between pull and push
<b>Hybrid-prune</b>	Hybrid + vertex pruning in push phases only

# Hybrid Pull-Push vs Pull Based Louvain

Dataset: POKEC, Outer loop 0



## On the 56 cores of Skylake

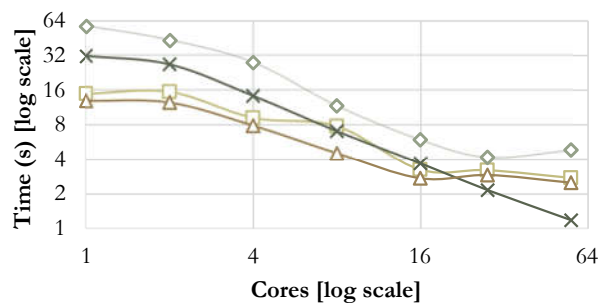
- Pull algorithm gets 19.8×
- Pull-prune gets 78×
- Hybrid gets 35×
- Hybrid-prune gets 63× speedup

Graphs	V	E
POKEC	5.40E+05	3.05E+07

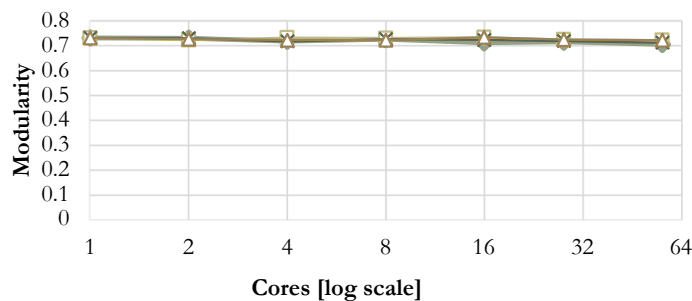
# Hybrid Pull-Push vs Pull Based Louvain

Dataset: Hollywood, Outer loop 0

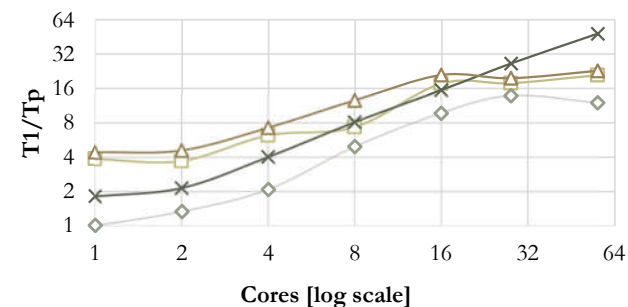
Time (s)



Modularity



Strong scaling wrt sequential pull



## On the 56 cores of Skylake

- Pull algorithm gets 12×
- Pull-prune gets 26×
- Hybrid gets 21×
- Hybrid-prune gets 23×

Graphs

V

E

Hollywood

1.14E+06

1.13E+08



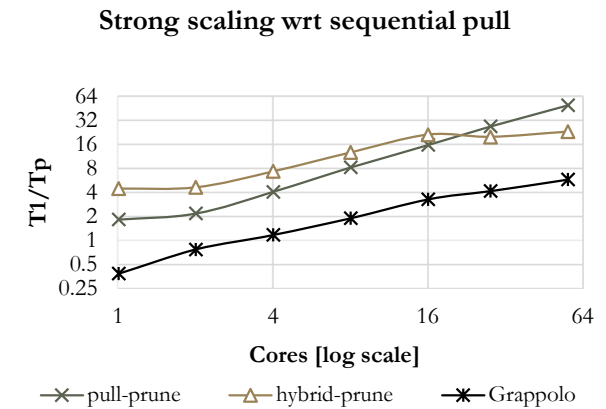
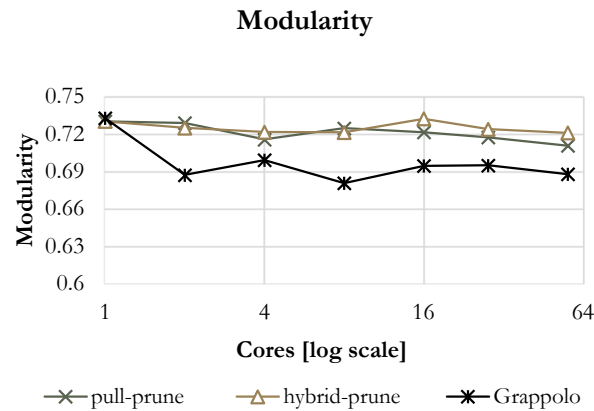
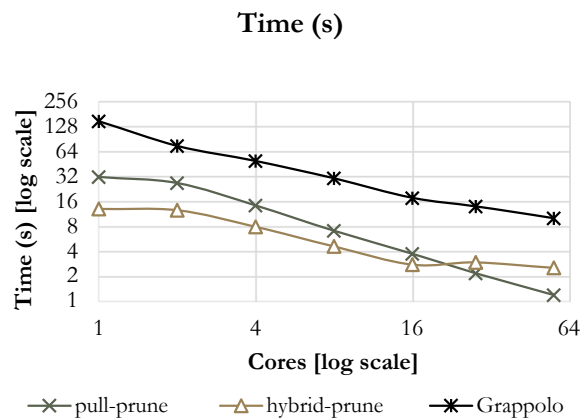
# Comparison with Prior State-of-the-art

## The Louvain in Grappolo

Hao Lu, Mahantesh Halappanavar, and Ananth Kalyanaraman. 2015. Parallel heuristics for scalable community detection. *Parallel Comput.* 47 (2015), 19–37

# 5-11 × faster than the Louvain in Grappolo

Dataset: Hollywood, Outer loop 0



## Compared to Grappolo

- Pull-prune 5-11 × faster
- Hybrid-prune 5-8 × faster, **best modularity**
- Modularity is higher in pull-prune and hybrid-prune

Graphs	V	E
Hollywood	1.14E+06	1.13E+08

# 2-8 × faster than the Louvain in Grappolo

Graph	hybrid-prune		pull-prune		Grappolo		Grappolo vs hybrid-prune	
	Q	T	Q	T	Q	T	Modularity	Speedup
caidaRouterLevel	0.68	0.05	0.65	0.02	0.68	0.11	1.00	2.20
citationCiteeer	0.6	0.06	0.62	0.04	0.59	0.3	1.02	5.00
coPaperDBLP	0.77	0.41	0.77	0.11	0.71	0.85	1.08	2.07
coPaperCiteeer	0.84	0.37	0.84	0.12	0.8	0.85	1.05	2.30
as-Skitter	0.72	0.9	0.71	1.37	0.69	2.12	1.04	2.36
uk-2005	0.95	21.01	0.88	17.71	0.83	136.95	1.14	6.52
rgg_n_2_24_0	0.92	1.71	0.89	1.75	0.74	13.54	1.24	7.92

Q= Modularity, T= Time (s)

The higher the better

## Compared to Grappolo (on 56 cores of Skylake)

- Pull-prune is 2- 8 × faster
- Hybrid-prune is 2 - 8 × faster, provides **best modularity**



skyLake Core	Graph	Grappolo		Pull		Hybrid		Pull-prune		Hybrid-prune		Speedup
		Q	T	Q	T	Q	T	Q	T	Q	T	
1	amazon	0.67	3.76	0.69	0.64	0.69	0.69	0.68	0.23	0.69	0.53	16.05
8		0.67	0.79	0.68	0.12	0.68	0.11	0.68	0.09	0.68	0.08	9.49
1	ca	0.54	0.30	0.56	0.10	0.56	0.09	0.56	0.04	0.56	0.07	4.22
8		0.54	0.08	0.56	0.02	0.56	0.01	0.56	0.01	0.56	0.01	8.21

Q= Modularity, T= Time (s)

The higher the better

### Compared to Grappolo (on 56 cores of Skylake)

- Hybrid-prune is 4-16 × faster
- Modularity is always higher or the same

4-16 × faster than the Louvain in Grappolo

# Quality: Normalized Mutual Information (NMI)

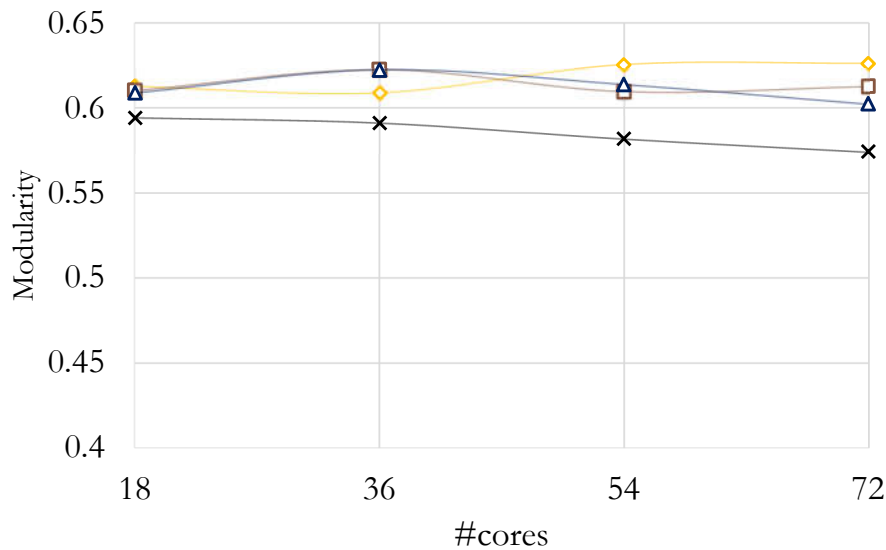
NMI score  $>0.8$  is considered good

Algorithm		Pull		Pull_prune		Hybrid		Hybrid_prune		Grappolo	
Threads		1	56	1	56	1	56	1	56	1	56
NMI Score	ca	1.000	0.995	1.000	0.995	0.999	0.995	1.000	0.995	0.996	0.980
	amazon	1.000	0.991	0.999	0.990	0.998	0.991	0.999	0.990	0.991	0.946

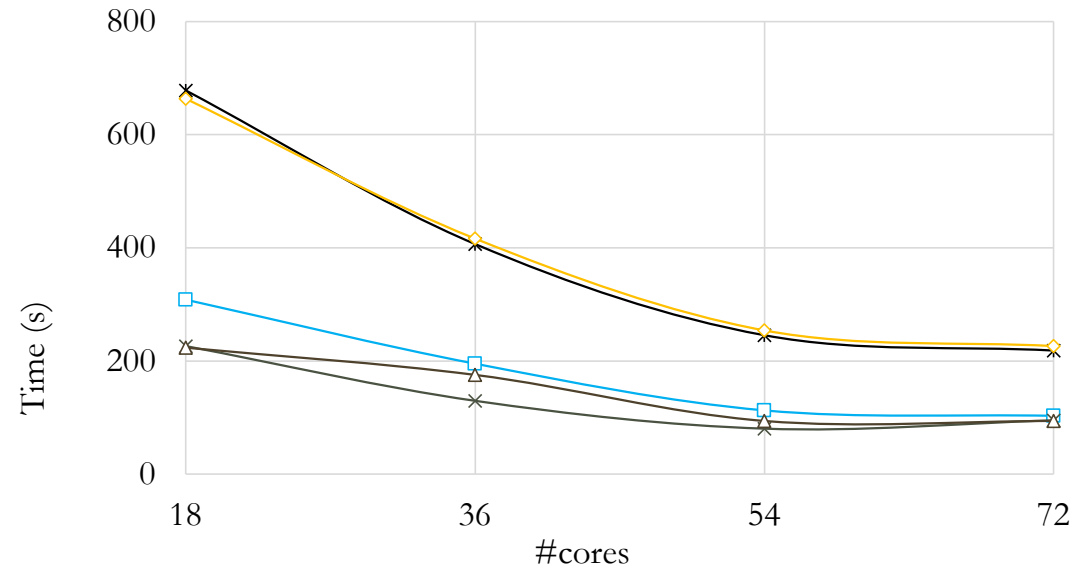
Our algorithms are better in NMI score than Grappolo, baseline is sequential Louvain Algorithm

# Louvain on Large Graphs

FRIENDSTER,  $V = 65,608,366$   $E = 3,612,134,270$



—◇— pull-prune —□— hybrid —△— hybrid-prune —×— Grappolo



—×— Grappolo —◇— pull —□— hybrid —×— pull-prune —△— hybrid-prune

**Compared to Grappolo (on 72 cores of Haswell)**

Platform: 72 core Haswell machine

- Pull-prune and Hybrid-prune is 2-4x faster
- Better Modularity



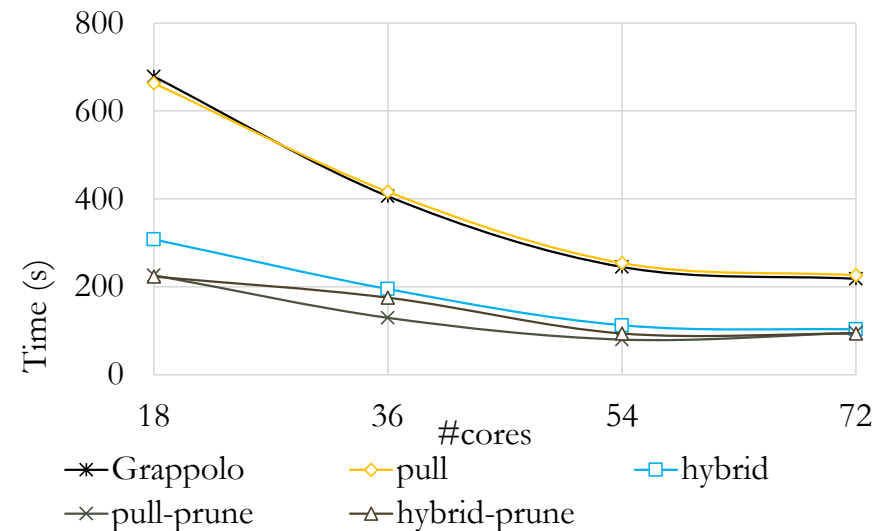
# Comparison with Recent Distributed Memory Algorithm

“Our MPI+OpenMP implementation yields about **7x speedup (on 4K processes)** for **soc-friendster** network (1.8B edges) over **Grappolo on 64 threads** on NERSC CORI system), without compromising output quality”

TABLE III  
DISTRIBUTED MEMORY VS SHARED MEMORY (GRAPPOLO)  
PERFORMANCE (RUNTIME) OF LOUVAIN ALGORITHM ON A SINGLE CORI  
NODE USING 4-64 THREADS. THE INPUT GRAPH IS SOC-FRIENDSTER  
(1.8B EDGES).

#Threads	Distributed memory (sec.)	Shared memory (sec.)
4	6,082.25	1,216.54
8	3,615.52	843.37
16	2,252.09	725.26
32	1,515.24	689.38
64	1,303.98	554.52

2.3x



# Comparison with Recent Distributed Memory Algorithm

“Our MPI+OpenMP implementation yields about **7x speedup (on 4K processes)** for **soc-friendster** network (1.8B edges) over **Grappolo on 64 threads on** NERSC CORI system), without compromising output quality”

TABLE III  
DISTRIBUTED MEMORY VS SHARED MEMORY (GRAPPOLO)  
PERFORMANCE (RUNTIME) OF LOUVAIN ALGORITHM ON A SINGLE CORI  
NODE USING 4-64 THREADS. THE INPUT GRAPH IS SOC-FRIENDSTER  
(1.8B EDGES).

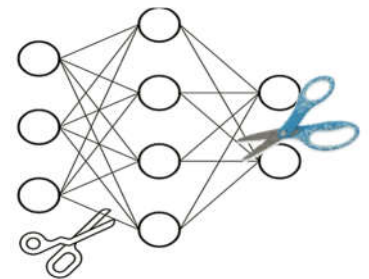
#Threads	Distributed memory (sec.)	Shared memory (sec.)
4	6,082.25	1,216.54
8	3,615.52	843.37
16	2,252.09	725.26
32	1,515.24	689.38
64	1,303.98	554.52

2.3x

Quick math says our approach could be 4 - 8x faster than this algorithm

# Conclusion – a new state-of-art for Louvain

- Prune unnecessary edge and vertex exploration during community detection
- Edges pruned by 6 to 13 $\times$  without sacrificing quality – up to 4x speedup
- Vertex pruned by 4 to 12 $\times$  with minimal sacrifice quality – up to 4x speedup
- Parallel algorithms 2-16x faster than prior state-of-the-art without sacrificing quality



We will be happy to make the code public. Please contact: [jesmin.jahan.tithi@intel.com](mailto:jesmin.jahan.tithi@intel.com)

