

STREAMER: a Distributed Framework for Incremental Closeness Centrality Computation

A. Erdem Sarıyüce ^{1,2}, Erik Saule ⁴, Kamer Kaya ¹, Ümit V. Çatalyürek ^{1,3}

 ¹ Department of Biomedical Informatics
² Department of Computer Science & Engineering
³ Department of Electrical & Computer Engineering The Ohio State University
⁴ Department of Computer Science University of North Carolina Charlotte

IEEE Cluster 2013, Indianapolis, IN

HPC Lab bmi.osu.edu/hpc

Massive Graphs are everywhere



Wexner Medical Center

- Facebook has a billion users and a trillion connections
- Twitter has more than 200 million users

bmi.osu.edu/hpc



Large(r) Networks and Centrality

- Who is more important in a network? Who controls the flow between nodes?
 - Centrality metrics answer these questions
 - Closeness Centrality (CC) is an intriguing metric
- How to handle changes?
 - Incremental algorithms are good but not enough in practice
 - Parallelism is essential

bmi.osu.edu/hpc

IEEE

Cluster'13







Closeness Centrality

- Let G=(V, E) be a graph with vertex set V and edge set E
 - Farness (far) of a vertex is the sum of shortest distances to each vertex

$$\operatorname{far}[u] = \sum_{\substack{v \in V \\ d_G(u,v) \neq \infty}} d_G(u,v)$$

• Closeness centrality (cc) of a vertex :

$$\operatorname{cc}[u] = \frac{1}{\operatorname{\mathtt{far}}[u]}$$

- Best algorithm: All-pairs shortest paths
 - O(|V|.|E|) complexity for unweighted networks
- For large and dynamic networks
 - From scratch computation is infeasible
 - Faster solutions are essential

bmi.osu.edu/hpc

IEEE

CC Algorithm



Wexner Medical Center

Algorithm 1: CC: Basic centrality computation



IEEE

Incremental Closeness Centrality



- Computing cc values from scratch after each edge change is very costly
 - Incremental algorithms are used to handle changes
 - Main idea is to <u>reduce number of SSSPs</u> to be executed
- Three filtering techniques are proposed
 - Filtering with level differences
 - Filtering with biconnected components
 - Filtering with identical vertices
- Details can be found at

bmi.osu.edu/hpc

IFFF

Cluster'13

"A. E. Sarıyuce, K. Kaya, E. Saule, and Umit V. Catalyurek. Incremental algorithms for Closeness Centrality. IEEE BigData Conference, 2013"

Filtering with level differences



Wexner Medical Center

 Upon edge insertion, breadth-first search tree of each vertex will change. Three possibilities:



- Case 1 and 2 will not change cc of s!
 - No need to apply SSSP from them
- Just Case 3

bmi.osu.edu/hpc

IEEE

Cluster'13

• BFSs are executed from u and v and level diff is checked

Filtering with biconnected components



Wexner Medical Center

What if the graph have articulation points?



- Change in A can change cc of any vertex in A and B
- Computing the change for u is enough for finding changes for any vertex v in B (constant factor is added)

Filtering with identical vertices



- Two types of identical vertices:
 - Type I: u and v are identical vertices if N(u) = N(v), i.e., their neighbor lists are same

 Type II: u and v are identical vertices if {u} U N(u) = {v} U N(v), i.e., they are also connected

- If u and v are identical vertices, their cc are the same
 - Same breadth-first search trees!

bmi.osu.edu/hpc

IFFF



IEEE

Cluster'13

bmi.osu.edu/hpc



name|V||E|Time (in sec.)web-NotreDame325K1,090K53.0amazon0601403K2,443K298.1web-Google875K4,322K824.4

- Too slow for real-time processing
- The problem is mostly parallel and graphs are relatively small.
 - Source-level parallelism can be used to fill up a cluster

DataCutter

- Component-based middleware tool
 - Supports filter-stream pro§
 - Implements the computati of components (filters) tha data through logical strean (unidirectional data flows)
- Layout is a filter ontology
 - Describes the set of tasks, streams and the connectio
 - All replicable

bmi.osu.edu/hpc

IEEE





STREAMER Framework



Wexner Medical Center



12

Multicore architecture and NUMA effects





- Preparator makes the actual graph
- Pointers are shared between Executors



• Dataset

IEEE

Cluster'13

Name	V	E	% of computation saved
web-NotreDame	325,729	1,090,008	97.5
amazon0601	403,394	2,443,308	92.3
web-Google	916,428	4,321,958	94.4
soc-pokec	1,632,804	30,622,464	93.9

• 64 node cluster

bmi.osu.edu/hpc

- Each with dual Intel Xeon E5520 Quad-Core processor
- 8MB L3 cache per processor
- 48GB main memory, 20Gbps Infiniband Connection
- Compiled with GCC 4.5.2 with –O3 flag



NUMA awareness

IEEE

Cluster'13



Wexner Medical Center

The performance of Streamer with 31 worker nodes and different node-level configurations normalized to 1 thread case (performance on soc-pokec is normalized to 8 threads, 1 graph/thread). The last column is the advantage of Shared Memory awareness (ratio of columns 5 and 3).

Name	4 threads	8 threads, 1 graph per			Shared Mem.
		thread	node	NUMA	awareness
web-NotreDame	3.69	6.46	7.13	6.99	1.08
amazon0601	3.26	6.75	6.81	7.45	1.10
web-Google	3.69	7.77	7.55	8.06	1.03
soc-pokec	-	1.00	0.92	1.01	1.01

Exploiting multiple cores and properly taking the shared-memory aspect brings significant improvement

Sublinear scaling case, 3 nodes



Wexner Medical Center



bmi.osu.edu/hpc

STREAMER: a Distributed Framework for Incremental Closeness Centrality Computation

Sublinear scaling case, 15 nodes



Wexner Medical Center



bmi.osu.edu/hpc

STREAMER: a Distributed Framework for Incremental Closeness Centrality Computation

18

Sublinear scaling case, 63 nodes







Conclusion

- STREAMER, a distributed-memory framework, proves to be an effective solution for fast and exact incremental closeness centrality computation
 - Exploits replicated and pipelined parallelism
 - Scales well
 - Reaches speedup of 497 with 64 nodes and 8 cores/node
- Future Work

bmi.osu.edu/hpc

IFFF

- StreamingMaster and Aggregator can be replicated and work can be partitioned
- Biconnected Decomposition, main part of Aggregator, can be parallelized as well

Thanks



- For more information
 - Email <u>umit@bmi.osu.edu</u>
 - Visit <u>http://bmi.osu.edu/~umit</u> or <u>http://bmi.osu.edu/hpc</u>
- Acknowledgement of Support

