

# Massive Graph Triangulation

by X. Hu, Y. Tao, and C. Chung, SIGMOD'13

Ilias Giechaskiel

Cambridge University, R212  
ig305@cam.ac.uk

February 21, 2014

## Takeaway Messages

- ▶ Triangle listing important input for graph properties
- ▶ I/O becomes bottleneck for massive graphs
  - ▶ Obvious approach doesn't work
- ▶ MGT algorithm
  - ▶ Total order of vertices guarantees unique triangle orientation
  - ▶ Near optimal asymptotic I/O + CPU performance
  - ▶ Much faster than alternatives in practice

## Definition

Given a graph  $G = (V, E)$ , list exactly once all

$$\Delta_{v_1 v_2 v_3} = \{v_1, v_2, v_3\} \text{ such that } v_i \in V \text{ and } (v_i, v_j) \in E$$

## Motivation

- ▶ Triangle = shortest non-trivial cycle and clique
- ▶ Various metrics
  - ▶ Dense neighborhood discovery
  - ▶ Triangular connectivity
  - ▶  $k$ -truss
  - ▶ Clustering coefficient

## The Algorithm

```
procedure LIST( $G$ )  
   $\Delta(G) \leftarrow \emptyset$   
  loop  $u \in V$   
    loop  $v \in \text{adj}_G(u) \ \& \ v > u$   
      loop  $w \in \text{adj}_G(u) \cap \text{adj}_G(v) \ \& \ w > v$   
         $\Delta(G) \leftarrow \Delta(G) \cup \{\Delta_{uvw}\}$   
  return  $\Delta(G)$ 
```

## The Problem

- ▶ Random access to  $\text{adj}_G(v)$  for  $v \in \text{adj}_G(u)$
- ▶  $\mathcal{O}(|E| \cdot \text{scan}(d_{\max}))$  I/Os in the worst case
  - ▶ When it doesn't fit in the memory of size  $M$
  - ▶ Recall:  $\text{scan}(N) = \Theta(N/B)$  where  $B$  is the disk block size

## Previous Approaches


- ▶ External Memory Compact Forward (EM-CF)
  - ▶  $\mathcal{O}(|E| + |E|^{1.5}/B)$  I/Os
  - ▶  $|E|$  I/O reads
  - ▶ Output insensitive
- ▶ External Memory Node Iterator (EM-NI)
  - ▶  $\mathcal{O}\left(|E|^{1.5}/B \cdot \log_{M/B}(|E|/B)\right)$  I/Os
  - ▶ Almost insensitive to  $M$
  - ▶ Output insensitive
- ▶ Graph Partition [CC12]
  - ▶  $\mathcal{O}(|E|^2/(MB) + K/B)$  I/Os where  $K$  triangles
  - ▶ In practice,  $M > \sqrt{|E|}$
  - ▶ If  $M = c|E|$ , asymptotically optimal
  - ▶ But under a set of assumptions...

## This Approach

- ▶  $\mathcal{O}(|E|^2/(MB) + K/B)$  I/Os in all settings
- ▶  $\mathcal{O}(|E| \log |E| + |E|^2/M + \alpha|E|)$  CPU time
  - ▶  $\alpha$  is the arboricity of the graph
- ▶ Both optimal up to constants
- ▶ Key idea: total order for unique triangle orientation
- ▶ Side note: also improves analysis of previous work

## Defining $G^*$

- ▶ Define  $\prec$  on  $V$  by  $u \prec v$  iff
  - ▶  $d(u) < d(v)$  or  $d(u) = d(v)$  and  $id(u) < id(v)$
  - ▶ Is a total order
- ▶  $G^*$  is  $G$  with edges oriented by  $\prec$ 
  - ▶ Takes  $\mathcal{O}(\text{sort}(|E|))$  I/Os
  - ▶ Recall:  $\text{sort}(N) = \Theta(N/B \log_{M/B} N/B)$
- ▶ Every triangle  $\{u, v, w\}$  has unique orientation  $u \prec v \prec w$



## Initial Idea

1. Load next  $cM$  edges of  $G^*$  into memory ( $E_{mem}$ )
  - ▶ All-or-nothing requirement (small-degree assumption)
2. Find all triangle with pivot edges in  $E_{mem}$




## Step 2 (Initial)

**procedure** LIST( $G, E_{mem}$ )

**loop**  $u \in V$

$V_{mem}(u) \leftarrow N^+(u) \cap V_{mem}$

    Find triangles with  $u$  cone in  $E_{mem}(u) \cup E_{mem}$



## Step 2 (Details)

**procedure** LIST( $G^*, E_{mem}$ )

Build hash structures

**loop**  $u \in V$

$V_{mem}(u) \leftarrow N^+(u) \cap V_{mem}$

**loop**  $v \in V_{mem}^+(u)$

**loop**  $w \in V_{mem}(u)$

**if**  $v \neq w$  &  $(v, w) \in E_{mem}$  **then**

Output  $\Delta_{uvw}$

## Analysis


- ▶  $\mathcal{O}(|E|^2/(MB) + K/B)$  I/O
  - ▶  $\Theta(|E|/M)$  iterations
  - ▶  $\mathcal{O}(|E|/B)$  I/Os for scanning
  - ▶  $\mathcal{O}(K/B)$  for listing
- ▶  $\mathcal{O}(|E| \log |E| + |E|^2/M + \alpha|E|)$  CPU
  - ▶  $\mathcal{O}(|E| \log |E|)$  for  $G^*$  sorting
  - ▶  $\Theta(|E|/M)$  iterations
  - ▶  $\mathcal{O}(|N^+(u)| + |N^+(u)| \cdot |V_{mem}^+(u)|)$
  - ▶  $\sum |N^+(u)| = |E|$
  - ▶  $\sum_{v \in V} d^+(v)^2 = \mathcal{O}(\alpha|E|)$
- ▶ Optimality comes from considering the complete graph

## Small-Degree Assumption

- ▶ What if  $\exists v$  such that  $d^+(v) > cM/2$ ?
  1. Find one
  2. Load a set  $S$  of  $cM/2$  of its out-edges
  3. Report all triangles involving one of the edges in  $S$
  4. Remove  $S$  from the graph
  5. Repeat

## Small-Degree Assumption

- ▶ How to implement step 3
  - ▶ Create hash table of loaded vertices
  - ▶ Scan all  $|E|$  edges
  - ▶ Also scan  $N(v)$  for each  $v \neq u$  with  $u \in N(v)$
  - ▶ Does not change complexity




## Experimental Setup

- ▶ 8GB memory (but memory conscious)
- ▶ Graphs unoriented
- ▶ Real data
  - ▶ 364MB to 7.5GB
  - ▶ 4.8 to 165 million vertices
  - ▶ 28 to 938 million edges
  - ▶  $|E|/|V|$  from 1.2 to 15.1
  - ▶ Varied  $M$  from 5% to 25% of disk size
- ▶ Synthetic data
  - ▶ Random, Recursive Matrix, Small World
  - ▶  $m = 16n$ ,  $n$  from 16 to 80 million
  - ▶ 2.1GB to 10.6GB


## Real Data

- ▶ MGT always better for CPU
- ▶ MGT almost always better for I/O
- ▶ RGP higher hidden constant in complexity!




# Evaluation


—□— MGT    -◇- DGP    -×- RGP    -△- EM-NI




(a) I/O cost (*RAND*)




(b) I/O cost (*R-MAT*)




(c) I/O cost (*S-WORLD*)



(d) Overall cost (*RAND*)



(e) Overall cost (*R-MAT*)



(f) Overall cost (*S-WORLD*)



## Criticism

- ▶ I/O analysis excludes cost of sorting
- ▶ Algorithm does not exploit parallelism
  - ▶ Is inherently sequential
  - ▶ Not applicable to distributed environment
  - ▶ Or across cores
  - ▶ RGP ideas applied in this case [PC13]
- ▶ Block I/O model for SSDs and parallel environment?
- ▶ Behavior for large-degree vertices
- ▶ Experiments lacking when  $M$  bigger percentage of graph

## Key Insights

- ▶ Total order of vertices guarantees unique triangle orientation
- ▶ Key idea simple, but multiple tricks
- ▶ Near optimal asymptotic I/O + CPU performance
- ▶ Much faster than alternatives in practice

## Key Questions

- ▶ Can you parallelize the algorithms non-trivially on a single PC?
- ▶ How can you extend the I/O model to different environments?
- ▶ How can you minimize data transfers in a distr. environment?
- ▶ Your questions?

-  Shumo Chu and James Cheng, *Triangle listing in massive networks*, ACM Trans. Knowl. Discov. Data **6** (2012), no. 4, 17:1–17:32.
-  Xiaocheng Hu, Yufei Tao, and Chin-Wan Chung, *Massive graph triangulation*, Proceedings of the 2013 ACM SIGMOD International Conference on Management of Data (New York, NY, USA), SIGMOD '13, ACM, 2013, pp. 325–336.
-  Ha-Myung Park and Chin-Wan Chung, *An efficient mapreduce algorithm for counting triangles in a very large graph*, Proceedings of the 22Nd ACM International Conference on Conference on Information & Knowledge Management (New York, NY, USA), CIKM '13, ACM, 2013, pp. 539–548.