

Making caches work for graph analytics

Yunming Zhang, Vladimir Kiriansky, Charith Mendis,
Saman Amarasinghe, Matei Zaharia

Historical Context

Paper at 2017 Paper @ BigData

Yunming Zhang (Google TPU compiler stuff; MIT PhD 2020, adv. Julian, Saman)

Vladimir Kiriansky (VMWare; MIT [B.Sc.](#), M.Eng. 2003, PhD 2019, adv. Saman)

Charith Mendis (Prof. UIUC, Tensor compiler stuff, MIT PhD 2020, adv. Saman)

Saman Amarasinghe (Prof @ MIT for 25+y; Compiler, system stuff)

Matei Zaharia (Prof @ MIT->Stanford->UCB 10y, Billionaire, Cofounder & CTO of Databricks, PhD @ Stanford 2013, Creator of Apache Spark)

William Hasenplaugh (D.E.Shaw Research, MIT PhD 2016, adv. Charles Eric Leiserson)

Julian Shun (Hi!)

Where we're at in class

- Graph
- Non-Graph
- Distributed
- (Single machine) Parallelized
- External (storage) algorithm
- **Cache level optimization**

Motivation

Graph algorithms are irregular

Lots of random memory access

Hard to utilize cache well


PageRank on $V=41\text{M}$, $E=1.5\text{B}$ ($\sim 656\text{MB}$)

LLC (L3, $\sim 30\text{-}55\text{MB}$) miss rate **>45%**

60-80% cycles spent stalled for memory access

Algorithm 1 PageRank

```
1 procedure PAGERANK(Graph  $G$ )
2   parallel for  $v : G.\text{vertexArray}$  do
3     for  $u : G.\text{edgeArray}[v]$  do
4        $G.\text{newRank}[v] +=$ 
5          $G.\text{rank}[u] / G.\text{degree}[u]$ 
6     end for
7   end parallel for
8 end procedure
```



Objective & Overview

Build 'Cagra' framework

EdgeMap (G, ActiveFrontier, EdgeUpdate, Merge)

VertexMap (G, VertexSubset, VertexUpdate)

How?

CSR Segmenting - Split the graph into subgraphs, to fit in cache

Frequency Based Clustering - Reorder CSR to maximize 'local' edges

CSR Segmenting – Split Subgraph

Given (inEdge) CSR, split edges by **source**

Keep all edges (and destinations) in subgraph

Algorithm 2 Preprocessing

Input: Number of vertices per segment N , Graph G

for $v : G.vertices$ **do**

for $inEdge : G.inEdges(v)$ **do**

$segmentID \leftarrow inEdge.src / N$

$subgraphs[segmentID].addInEdge(v, inEdge.src)$

end for

end for

for $subgraph : subgraphs$ **do**




$subgraph.sortByDestination()$

$subgraph.constructIdxMap()$

$subgraph.constructBlockIndices()$

$subgraph.constructIntermBuf()$

end for



CSR Segmenting – Per-Segment Processing

For each subgraph, parallelize on destination

All sources in segment fits in cache

inEdges are in order of destination

Keep results for each dest in each subgraph

Processing segment 1:

source vertices

segment 1

destination vertices

v1

Processing segment 2:

source vertices

segment 2

destination vertices

v2

. . .

Algorithm 3 Parallel Segment Processing

```
for subgraph : subgraphs do  
  parallel for v : subgraph.Vertices do  
    for inEdge : subgraph.inEdges(v) do  
      Process inEdge  
    end for  
  end parallel for  
end for
```

CSR Segmenting – Merging

Each dest have multiple partial results

Merge results from subgraphs

Destination blocks fit in L1 Cache

Algorithm 4 Cache-Aware Merge

```
parallel for block : blocks do  
  for subgraph : G.subgraphs do  
    blockStart  $\leftarrow$  subgraph.blockStarts[block]  
    blockEnd  $\leftarrow$  subgraph.blockEnds[block]  
    intermBuf  $\leftarrow$  subgraph.intermBuf  
    for localIdx from blockStart to blockEnd do  
      globalIdx  $\leftarrow$  subgraph.idxMap[localIdx]  
      localUpdate = intermBuf[localIdx]  
      merge(output[globalIdx], localUpdate)  
    end for  
  end for  
end parallel for  
return output
```




Fig. 4: Cache-aware merge

Parameters & Analysis

Source segments to fit in LLC (L3 Cache)

For PageRank, 30MB LLC can fit 4M vertices

How many merge happens for each vertex?


Expansion Factor $q = (\text{average inDeg for segment}) / (\text{average } |V| \text{ for segment})$

Rough proxy for number of merges

LLC <> DRAM traffic

k segments, expansion factor q

$E + 2qV + V$



Frequency Based Clustering

Motivation:

- want to maximize edges within subgraph,

- reduce inter-segment merging (=extra work, DRAM access)

Stable sort vertices based on [outdegree/threshold]

- Cluster hot nodes together, while keeping some natural order

Got 2x higher expansion factor, reducing merges

Experimental Evaluation

5x faster than papers, 3x than engineers

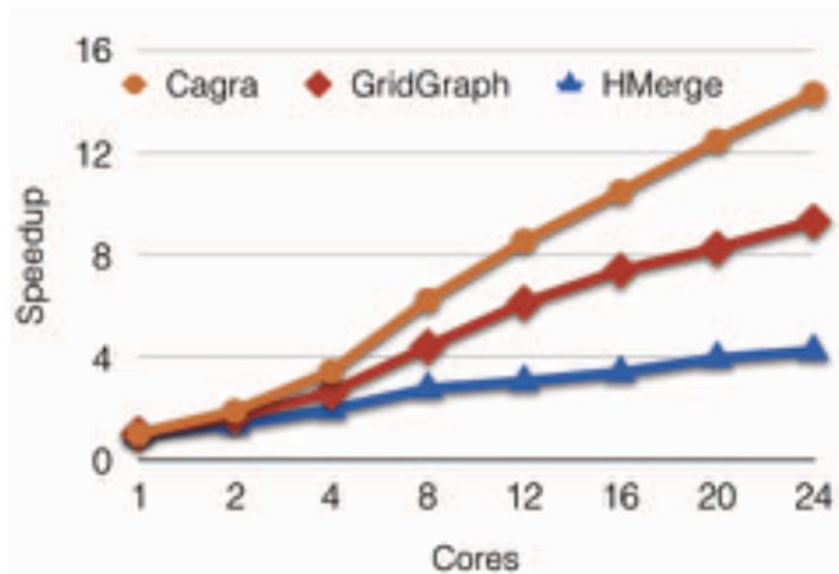


Fig. 10: Scalability for PageRank on Twitter

Algorithm 5 PageRank in Cagra

```
typedef double vertexDataType
contrib  $\leftarrow$  {1/outDegree[v], ...}
newRank  $\leftarrow$  {0.0, ...}

procedure EDGEUPDATE(bufVal, srcVal, dstVal)
    bufVal+ = srcVal
    return true
end procedure

procedure MERGE(newDstVal, bufVal)
    newDstVal+ = bufVal
end procedure

procedure VERTEXUPDATE(v)
    newRank[v]  $\leftarrow$  0.15 + 0.85 * newRank[v]
    newRank[v]  $\leftarrow$  newRank[v]/outDegree[v]
    contrib[v]  $\leftarrow$  0.0
    return true
end procedure

procedure PAGERANK(G, maxIter)
    iter  $\leftarrow$  0
    A  $\leftarrow$  V
    while iter  $\neq$  maxIter do
        A  $\leftarrow$  EdgeMap(G, A, EdgeUpdate, EdgeMerge)
        A  $\leftarrow$  VertexMap(G, A, VertexUpdate)
        Swap(contrib, newRank)
        iter  $\leftarrow$  iter + 1
    end while
end procedure
```





Fig. 7: Speedups of optimizations on PageRank, Label Propagation, Collaborative Filtering



Reviewing the paper

No push based (atomic?)

e.g. For BFS, we can't do dense method

No destination vertex condition check

Very good for PageRank and algos using most edges most times

Can be used in distributed systems as well?

Merge step might be inconvenient