

Reviewing the Ligra single-node graph processing framework

Thomas Parks

October 24, 2017

U. of Cam

Introduction

Ligra: A Lightweight Graph Processing Framework for Shared Memory


Authors: Julian Shun and Guy E. Blelloch, CMU

A reaction to the availability of large single nodes.

The shift of computing



Interest in processing graph data has been relatively constant over time, whereas cluster computing fluctuates in the published literature.



The RAM capacity for a single server has grown exponentially, with a knee approximately where the use of clusters drops off

Ligra single node graph computations

API inspired by Hybrid BFS¹.

Aims for every high efficiency by using CAS²

Outperforms Pregel on a per core and a absolute basis³.

Also claims superior performance per dollar and Joule⁴.

¹Beamer, Asanovic, et al., *Searching for a parent instead of fighting over children: A fast breadth-first search implementation for graph500*.

²Schweizer, Besta, and Hoefler, "Evaluating the Cost of Atomic Operations on Modern Architectures".

³This was not thoroughly explored in the paper.

⁴This was not mentioned again after claiming improvements in the abstract.

API

Ligra API and motivating example

```
parents = [-1, ..., -1];    // The parent of every node
UPDATE s, d
| return CAS(parents[d], -1, s);
COND i
| return parents[i] == -1;
BFS G, r
| parents[r] := r;
| frontier = r;
| while size(frontier) != 0 do
|   // For every vertex in the frontier,
|     UPDATE all neighbouring j if COND. Add
|     to returned set if UPDATE(i, j).
|   frontier := EDGEMAP(G, frontier, UPDATE, COND);
end
```

EDGEMAP working outwards


Semantics allow for multiple implementations with different performance.

EDGEMAP_SPARSE G, U, F, C

```
result = {};  
/* both loops fully parallel */  
foreach v in U do  
  foreach v2 in out_neighbours(v) do  
    if C(v2) and F(v, v2) // not in the BFS tree  
    then  
      add v2 to result;  
    end  
  end  
end  
return result;
```


EDGEMAP working outwards

Semantics allow for multiple implementations with different performance.



5

⁵Beamer, Asanović, and Patterson, “Direction-optimizing Breadth-first Search”.

EDGEMAP working over all elements

EDGEMAP_DENSE G, U, F, C

```
result = {};  
/* first loop parallel */  
foreach i in [0, ..., |V(G)|] do  
  if C(i) // not in the BFS tree  
  then  
    foreach v in in_neighbours(i) do  
      if v ∈ U and F(v, i) then add i to result;  
      if not c(i) then break;  
    end  
  end  
end  
return result;
```

VERTEXMAP

VERTEXMAP U, F

```
| result = {};
```

```
| /* parallel loop
```

```
*/
```

```
| foreach  $u \in U$  do
```

```
| | if  $F(u)$  then add  $u$  to result;
```

```
| end
```

```
| return result;
```

Performance measurements

Ligra scaling

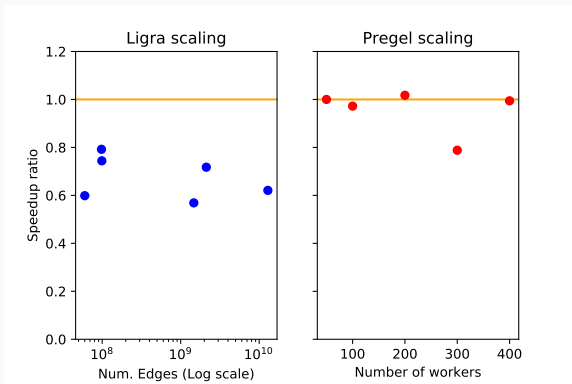


Table 1: 1B vertex binary tree shortest path

Pregel	20 seconds	300 Nodes
Ligra	2 seconds	1 Node

Navigating the maze of Graphs

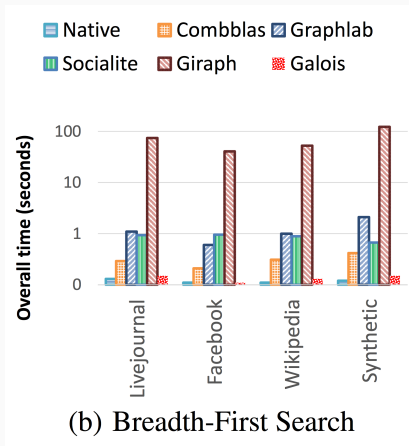


Figure 1: The real performance of algorithms can be hard to find.⁶

⁶Satish et al., “Navigating the Maze of Graph Analytics Frameworks Using Massive Graph Datasets”.

Navigating the maze of Graphs




Figure 2: Galois can implement Ligra simply.⁷

⁷Nguyen, Lenharth, and Pingali, “A Lightweight Infrastructure for Graph Analytics”.

Questions?

Nice algo bits.

Graph diameter estimation.

Algorithm 7 Radii Estimation

```
1: Visited = {0, ..., 0}           ▷ initialized to all 0
2: NextVisited = {0, ..., 0}      ▷ initialized to all 0
3: Radii = {∞, ..., ∞}           ▷ initialized to all ∞
4: round = 0
5:
6: procedure RADIIUPDATE(s, d)
7:   if (Visited[d] ≠ Visited[s]) then
8:     ATOMICOR(&NextVisited[d], Visited[d] | Visited[s])
9:     oldRadii = Radii[d]
10:    if (Radii[d] ≠ round) then
11:      return CAS(&Radii[d], oldRadii, round)
12:    return 0
13:
14: procedure ORCOPY(i)
15:   NextVisited[i] = NextVisited[i] | Visited[i]
16:   return 1
17:
18: procedure RADII(G)
19:   Sample K vertices and for each one set a unique bit in Visited to 1
20:   Initialize Frontier to contain the K sampled vertices
21:   Set the Radii entries of the sampled vertices to 0
22:   while (SIZE(Frontier) ≠ 0) do
23:     round = round + 1
24:     Frontier = EDGEMAP(G, Frontier, RADIIUPDATE, Ctrue)
25:     Frontier = VERTEXMAP(Frontier, ORCOPY)
26:     SWAP(Visited, NextVisited)   ▷ switch roles of bit-vectors
27:   return Radii
```

Associate a bit vector with each vertex for all BFS searches, and bitwise OR the current vertex vector with neighbours.

Vertices that change are on the new multiBFS frontier. Store the iteration number of the last time a vertex changed its vector. This is a lower bound on centrality of that vertex, and $\max(\text{centrality})$ is the diameter.