

# Ligra: A Lightweight Graph Processing Framework for Shared Memory

Julian Shun

Carnegie Mellon University  
jshun@cs.cmu.edu

Guy E. Blelloch

Carnegie Mellon University  
guyb@cs.cmu.edu

*Categories and Subject Descriptors* D.1.3 [*Programming Techniques*]: Concurrent Programming—Parallel Programming

*General Terms* Algorithms, Performance

*Keywords* Shared Memory, Graph Algorithms, Parallel Programming

# What is Ligra?

- Ligra is a lightweight and efficient framework for graph processing in shared memory. It was designed to make it easier for developers to implement graph algorithms that can run efficiently on multi-core machines
- Well suited for graph traversal problems

# Implementations

---

**Algorithm 1** EDGEMAP

---

```
1: procedure EDGEMAP( $G, U, F, C$ )
2:   if ( $|U| + \text{sum of out-degrees of } U > \text{threshold}$ ) then
3:     return EDGEMAPDENSE( $G, U, F, C$ )
4:   else return EDGEMAPSPARSE( $G, U, F, C$ )
```

---

---

**Algorithm 2** EDGEMAPSPARSE

---

```
1: procedure EDGEMAPSPARSE( $G, U, F, C$ )
2:   Out = {}
3:   parfor each  $v \in U$  do
4:     parfor ngh  $\in N^+(v)$  do
5:       if ( $C(\text{ngh}) == 1$  and  $F(v, \text{ngh}) == 1$ ) then
6:         Add ngh to Out
7:   Remove duplicates from Out
8:   return Out
```

---

---

**Algorithm 3** EDGEMAPDENSE

---

```
1: procedure EDGEMAPDENSE( $G, U, F, C$ )
2:   Out = {}
3:   parfor  $i \in \{0, \dots, |V| - 1\}$  do
4:     if ( $C(i) == 1$ ) then
5:       for ngh  $\in N^-(i)$  do
6:         if ( $\text{ngh} \in U$  and  $F(\text{ngh}, i) == 1$ ) then
7:           Add  $i$  to Out
8:     if ( $C(i) == 0$ ) then break
9:   return Out
```

---

---

**Algorithm 4** VERTEXMAP

---

```
1: procedure VERTEXMAP( $U, F$ )
2:   Out = {}
3:   parfor  $u \in U$  do
4:     if ( $F(u) == 1$ ) then Add  $u$  to Out
5:   return Out
```

---

# Application (Breadth-First Search)

```
1: Parents = { -1, ..., -1 }           ▷ initialized to all -1's
2:
3: procedure UPDATE(s, d)
4:   return (CAS(&Parents[d], -1, s))
5:
6: procedure COND(i)
7:   return (Parents[i] == -1)
8:
9: procedure BFS(G, r)                 ▷ r is the root
10:  Parents[r] = r
11:  Frontier = {r}                     ▷ vertexSubset initialized to contain only r
12:  while (SIZE(Frontier) ≠ 0) do
13:    Frontier = EDGEMAP(G, Frontier, UPDATE, COND)
```

**Figure 1.** Pseudocode for Breadth-First Search in our framework. The compare-and-swap function  $CAS(loc, oldV, newV)$  atomically checks if the value at location  $loc$  is equal to  $oldV$  and if so it updates  $loc$  with  $newV$  and returns *true*. Otherwise it leaves  $loc$  unmodified and returns *false*.

---

**Algorithm 6** Betweenness Centrality

---

```
1: NumPaths = {0, ..., 0}           ▷ initialized to all 0
2: Visited = {0, ..., 0}           ▷ initialized to all 0
3: NumPaths[r] = 1
4: Visited[r] = 1
5: currLevel = 0
6: Levels = []
7: Dependencies = {0.0, ..., 0.0}   ▷ initialized to all 0.0
8:
9: procedure VISIT(i)
10:   Visited[i] = 1
11:   return 1
12:
13: procedure PATHSUPDATE(s, d)
14:   repeat
15:     oldV = NumPaths[d]
16:     newV = oldV + NumPaths[s]
17:   until (CAS(&NumPaths[d], oldV, newV) == 1)
18:   return (oldV == 0)
19:
20: procedure DEPUPDATE(s, d)
21:   repeat
22:     oldV = Dependencies[d]
23:     newV = oldV +  $\frac{\text{NumPaths}[s]}{\text{NumPaths}[d]} \times (1 + \text{Dependencies}[s])$ 
24:   until (CAS(&Dependencies[d], oldV, newV) == 1)
25:   return (oldV == 0.0)
26:
27: procedure COND(i)
28:   return (Visited[i] == 0)
29:
30: procedure BC(G, r)
31:   Frontier = {r}                 ▷ vertexSubset initialized to contain only r
32:   while (SIZE(Frontier) ≠ 0) do     ▷ Phase 1
33:     Frontier = EDGEMAP(G, Frontier, PATHSUPDATE, COND)
34:     Levels[currLevel] = Frontier
35:     Frontier = VERTEXMAP(Frontier, VISIT)
36:     currLevel = currLevel + 1
37:
38:   Visited = {0, ..., 0}           ▷ reinitialize to all 0
39:   currLevel = currLevel - 1
40:   TRANSPOSE(G)                   ▷ transpose graph
41:
42:   while (currLevel ≥ 0) do         ▷ Phase 2
43:     Frontier = Levels[currLevel]
44:     VERTEXMAP(Frontier, VISIT)
45:     EDGEMAP(G, Frontier, DEPUPDATE, COND)
46:     currLevel = currLevel - 1
47:   return Dependencies
```

---

# Application (Betweenness Centrality)

# Application (Radii 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
```

---

# Application (Connected Components)

---

## Algorithm 8 Connected Components

---

```
1:  $IDs = \{0, \dots, |V| - 1\}$   $\triangleright$  initialized such that  $IDs[i] = i$ 
2:  $prevIDs = \{0, \dots, |V| - 1\}$   $\triangleright$  initialized such that  $prevIDs[i] = i$ 
3:
4: procedure CCUPDATE( $s, d$ )
5:    $origID = IDs[d]$ 
6:   if (WRITEMIN(&IDs[ $d$ ], IDs[ $s$ ])) then
7:     return ( $origID == prevIDs[d]$ )
8:   return 0
9:
10: procedure COPY( $i$ )
11:    $prevIDs[i] = IDs[i]$ 
12:   return 1
13:
14: procedure CC( $G$ )
15:    $Frontier = \{0, \dots, |V| - 1\}$   $\triangleright$  vertexSubset initialized to  $V$ 
16:   while (SIZE(Frontier)  $\neq$  0) do
17:      $Frontier = VERTEXMAP(Frontier, COPY)$ 
18:      $Frontier = EDGEMAP(G, Frontier, CCUPDATE, C_{true})$ 
19:   return IDs
```

---

# Application (PageRank)

---

## Algorithm 9 PageRank

---

```
1:  $p_{curr} = \{\frac{1}{|V|}, \dots, \frac{1}{|V|}\}$  ▷ initialized to all  $\frac{1}{|V|}$ 
2:  $p_{next} = \{0.0, \dots, 0.0\}$  ▷ initialized to all 0.0
3:  $diff = \{\}$  ▷ array to store differences
4:
5: procedure PRUPDATE( $s, d$ )
6:   ATOMICINCREMENT( $\&p_{next}[d], \frac{p_{curr}[s]}{deg^+(s)}$ )
7:   return 1
8:
9: procedure PRLocalCOMPUTE( $i$ )
10:   $p_{next}[i] = (\gamma \times p_{next}[i]) + \frac{1-\gamma}{|V|}$ 
11:   $diff[i] = |p_{next}[i] - p_{curr}[i]|$ 
12:   $p_{curr}[i] = 0.0$ 
13:  return 1
14:
15: procedure PAGERANK( $G, \gamma, \epsilon$ )
16:  Frontier =  $\{0, \dots, |V| - 1\}$  ▷ vertexSubset initialized to  $V$ 
17:  error =  $\infty$ 
18:  while (error >  $\epsilon$ ) do
19:    Frontier = EDGEMAP( $G, Frontier, PRUPDATE, C_{true}$ )
20:    Frontier = VERTEXMAP(Frontier, PRLocalCOMPUTE)
21:    error = sum of diff entries
22:    SWAP( $p_{curr}, p_{next}$ )
23:  return  $p_{curr}$ 
```

---



# Application (Bellman-Ford)

---

## Algorithm 10 Bellman-Ford

---

```
1:  $SP = \{\infty, \dots, \infty\}$  ▷ initialized to all  $\infty$ 
2:  $Visited = \{0, \dots, 0\}$  ▷ initialized to all 0
3:
4: procedure BFUPDATE( $s, d, edgeWeight$ )
5:   if (WRITEMIN(&SP[ $d$ ], SP[ $s$ ] + edgeWeight)) then
6:     return CAS(&Visited[ $d$ ], 0, 1)
7:   else return 0
8:
9: procedure BFRRESET( $i$ )
10:  Visited[ $i$ ] = 0
11:  return 1
12:
13: procedure BELLMAN-FORD( $G, r$ )
14:  SP[ $r$ ] = 0
15:  Frontier = { $r$ } ▷ vertexSubset initialized to contain just  $r$ 
16:  round = 0
17:  while (SIZE(Frontier)  $\neq$  0 and round <  $|V|$ ) do
18:    round = round + 1
19:    Frontier = EDGEMAP( $G, Frontier, BF-UPDATE, C_{true}$ )
20:    Frontier = VERTEXMAP(Frontier, BF-RESET)
21:  if (round ==  $|V|$ ) then return "negative-weight cycle"
22:  else return SP
```

---

# Running Times


Application	3D-grid			random-local			rMat24			rMat27			Twitter			Yahoo		
	(1)	(40h)	(SU)	(1)	(40h)	(SU)	(1)	(40h)	(SU)	(1)	(40h)	(SU)	(1)	(40h)	(SU)	(1)	(40h)	(SU)
Breadth-First Search	2.9	0.28	10.4	2.11	0.073	28.9	2.83	0.104	27.2	11.8	0.423	27.9	6.92	0.321	21.6	173	8.58	20.2
Betweenness Centrality	9.15	0.765	12.0	8.53	0.265	32.2	11.3	0.37	30.5	113	4.07	27.8	47.8	2.64	18.1	634	23.1	27.4
Graph Radii	351	10.0	35.1	25.6	0.734	34.9	39.7	1.21	32.8	337	12.0	28.1	171	7.39	23.1	1280	39.6	32.3
Connected Components	51.5	1.71	30.1	14.8	0.399	37.1	14.1	0.527	26.8	204	10.2	20.0	78.7	3.86	20.4	609	29.7	20.5
PageRank (1 iteration)	4.29	0.145	29.6	6.55	0.224	29.2	8.93	0.25	35.7	243	6.13	39.6	72.9	2.91	25.1	465	15.2	30.6
Bellman-Ford	63.4	2.39	26.5	18.8	0.677	27.8	17.8	0.694	25.6	116	4.03	28.8	75.1	2.66	28.2	255	14.2	18.0

**Table 2.** Running times (in seconds) of algorithms over various inputs on a 40-core machine (with hyper-threading). (SU) indicates the speedup of the application (single-thread time divided by 40-core time).

# Setup of Experiments



All experiments were performed on:



10-core E7-8870 Xeon processors, a 1066MHz bus, (256GB) main memory



- 40-core Intel machine (with hyper-threading), 4x2:4GHz Intel



- icpc compiler (version 12.1.0) using CilkPlus with the -O3 flag

# Conclusion

- **Designed for shared-memory machines.**
- **Implementations using Ligra:**
  - Efficient and scalable
  - Often outperform other graph libraries/systems
- **Potential uses for other algorithms:**
  - Maximum flow
  - Biconnected components
  - Belief propagation
  - Markov clustering
- **Current Limitations:**
  - Doesn't support algorithms that modify the input graph
- **Future Directions:**
  - Extend Ligra to support graph modifications
  - Explore adaptability for GPU systems