Combining Duplication & Reordering to Accelerate Parallel Graph Processing

Vignesh Balaji

Brandon Lucia





Graph Processing Has Many Important Applications



Large Graphs Can Be Processed in a Single Node





Cores	8 - 64 🏠
Memory	(100s GB - TBs) 合

Large Graphs Can Be Processed in a Single Node



Large Graphs Can Be Processed in a Single Node



For Graphs that fit in main memory, single node graph processing is more efficient

[1] "Scalability! But at what COST?" Usenix HOTOS'15

Outline

Single-node Graph Processing is Sub-optimal

- Existing Optimizations have Overheads
- RADAR: Combining the Benefits of Duplication and Reordering
- Advantages of RADAR over Push-Pull Direction Switching



Outline

- Single-node Graph Processing is Sub-optimal
 - ▶ **Expensive Atomic Updates** \Rightarrow Coherence & serialization overheads
 - > **Poor LLC locality** \Rightarrow DRAM accesses dominate runtime
- Existing Optimizations have Overheads
- RADAR: Combining the Benefits of Duplication and Reordering
- Advantages of RADAR over Push-Pull Direction Switching



Shared-memory Graph Processing Overview



Out-neighbors

Shared-memory Graph Processing Overview



Out-neighbors



9

Shared-memory Graph Processing Overview



Out-neighbors

Vertex Centric Graph Processing (Push)

for src in Frontier:
 for dst in out_neigh(src):
 vtxData[dst] += auxData[src]



The Need for Atomic Updates



Vertex Centric Graph Processing (Push)



parallel_for src in Frontier:
 for dst in out_neigh(src):
 atomic {vtxData[dst] += auxData[src]}



The Need for Atomic Updates



Bottleneck #1: Atomic Updates Hurt Performance



Bottleneck #1: Atomic Updates Hurt Performance





Out-neighbors

Vertex Centric Graph Processing (Push)

for src in Frontier:
 for dst in out_neigh(src):
 vtxData[dst] += auxData[src]





Out-neighbors

Vertex Centric Graph Processing (Push)

```
for src in Frontier:
  for dst in out_neigh(src):
      vtxData[dst] += auxData[src]
```







Bottleneck #2: Graph Applications are DRAM-latency bound

LLC Miss Rate (%)



Cycles stalled on DRAM / Total Cycles



Problem: Poor LLC locality ⇒ Many long-latency DRAM accesses



Figure from "Optimizing Cache Performance for Graph Analytics" ArXiv v1;

Outline

- Single-node Graph Processing is Sub-optimal V
 - ▶ **Expensive Atomic Updates** \Rightarrow Coherence & serialization overheads
 - > **Poor LLC locality** \Rightarrow DRAM accesses dominate runtime
- Existing Optimizations incur Overheads
- RADAR: Combining the Benefits of Duplication and Reordering
- Advantages of RADAR over Push-Pull Direction Switching



Outline

- Single-node Graph Processing is Sub-optimal
- **Existing Optimizations incur Overheads**
 - > Duplication incurs Hub Identification Overhead
 - **Reordering incurs False Sharing Overhead**
- RADAR: Combining the Benefits of Duplication and Reordering
- Advantages of RADAR over Push-Pull Direction Switching



Optimizations for Eliminating Performance Bottlenecks

BUT

- Data Duplication : Improves Scalability
- Graph Reordering : Reduces DRAM accesses



- Duplication incurs Hub Identification Overhead
- Reordering incurs False Sharing Overhead



Carnegie Mellon

Optimization #1: Data Duplication

parallel_for src in Frontier:
for dst in neigh(src):
 atomic {vtxData[dst] +=
 auxData[src]}



vtxData



Carnegie Mellon

Optimization #1: Data Duplication



parallel_for src in Frontier:
for dst in neigh(src):
 atomic {vtxData[dst] +=
 auxData[src]}



parallel_for src in Frontier:

vtxDataDup[tid][dst] +=

auxData[src]

for dst in neigh(src):



vtxData



Naive Duplication Imposes High Memory Overhead



For a graph with 64M vertices, 4B / vtx, 16 threads

Memory footprint after Duplication = **4GB!**



Power-Law Graphs Allow Memory-Efficient Duplication

Hubs Air Traffic Network



Carnegie Mellon

Power-Law Graphs Allow Memory-Efficient Duplication





Majority of Atomic Updates are to Hub Vertices

Duplicate *only* the Hub Vertex Data



HUBDUP: Duplication for Power-Law Graphs





HUBDUP: Duplication for Power-Law Graphs



Carnegie Mellon

HUBDUP Incurs Overheads





HUBDUP Incurs Overheads





Summary Of Overheads In HUBDUP





Summary Of Overheads In HUBDUP



Problem: Hubs can have *arbitrary* vertex IDs



Optimizations for eliminating performance bottlenecks

- Data Duplication : Improves Scalability
- Graph Reordering : Reduces DRAM accesses



BUT

- HUBDUP incurs Hub Identification Overhead
- Reordering incurs False Sharing Overhead
















Degree Sorting Is Effective For Serial Graph Processing



Degree Sorting Introduces False Sharing





Degree Sorting Introduces False Sharing



41

Degree Sorting Introduces False Sharing



False Sharing Hurts Performance



Optimizations for eliminating performance bottlenecks

- Data Duplication : Improves Scalability
- ✤ Graph Reordering : Reduces DRAM accesses



BUT

- HUBDUP incurs Hub Identification Overhead
- Reordering incurs False Sharing Overhead



Summary of Duplication And Reordering



HUBDUP

No Atomics for Hub Vertices









Degree Sorting





Introduces False Sharing





Outline

- Single-node Graph Processing is Sub-optimal V
- Existing Optimizations have Overheads
 - > Duplication incurs Hub Identification Overhead
 - > Reordering incurs False Sharing Overhead
- RADAR: Combining the Benefits of Duplication and Reordering
- Advantages of RADAR over Push-Pull Direction Switching



Outline

- Single-node Graph Processing is Sub-optimal V
- Existing Optimizations have Overheads
- RADAR: Combining the Benefits of Duplication and Reordering
 - HUBDUP & Degree Sorting are Mutually Enabling
 - RADAR = HUBDUP + Degree Sorting
 - > RADAR outperforms HUBDUP & Degree Sorting
- Advantages of RADAR over Push-Pull Direction Switching



HUBDUP And Degree Sorting Are Mutually Enabling





Degree Sorting Improves HUBDUP





Degree Sorting Improves HUBDUP



HUBDUP Improves Degree Sorting





HUBDUP Improves Degree Sorting



Threads update a private copies of hubs \Rightarrow No False Sharing



<u>Reordering Assisted Duplication/Duplication Assisted Reordering</u>



No Atomics for Hub Vertices Costs incurred for detecting Hubs 👘



Degree Sorting

Improves Cache Locality Introduces False Sharing







RADAR



No Atomics for Hub Vertices (with easy hub detection)





Improves Cache Locality (without false-sharing for hubs)

Outline

- Single-node Graph Processing is Sub-optimal V
- Existing Optimizations have Overheads
- **ADAR: Combining the Benefits of Duplication and Reordering**
 - > HUBDUP & Degree Sorting are Mutually Enabling
 - RADAR = HUBDUP + Degree Sorting
 - > RADAR outperforms HUBDUP & Degree Sorting <
- Advantages of RADAR over Push-Pull Direction Switching



Evaluation Space



Evaluation Platform

- ✤ 28 Cores / 56 Threads across 2 socket ⇒ All versions run with 56 threads
- ✤ 35MB LLC per socket ⇒ Hubs selected to fit LLC
- ✤ 64GB DRAM



RADAR Outperforms Both HUBDUP And Degree Sorting



RADAR Outperforms Both HUBDUP And Degree Sorting



RADAR Outperforms Both HUBDUP And Degree Sorting



Outline

- Single-node Graph Processing is Sub-optimal V
- Existing Optimizations have Overheads
- RADAR: Combining the Benefits of Duplication and Reordering V
 - > HUBDUP & Degree Sorting are Mutually Enabling
 - RADAR = HUBDUP + Degree Sorting
 - RADAR outperforms HUBDUP & Degree Sorting
- Advantages of RADAR over Push-Pull Direction Switching



Outline

- Single-node Graph Processing is Sub-optimal V
- Existing Optimizations have Overheads
- RADAR: Combining the Benefits of Duplication and Reordering
- Advantages of RADAR over Push-Pull Direction Switching
 - RADAR does not compromise work-efficiency
 - **RADAR can support larger input graphs**
 - RADAR incurs lower preprocessing overheads



State Of The Art: Push-Pull Direction Switching

Push



parallel_for **src** in Frontier: for dst in out_neigh(src): atomic {vtxData[dst] += auxData[src]}

Pull

```
parallel_for dst in G:
for src in in_neigh(dst):
  if src in Frontier:
    vtxData[dst] += auxData[src]
```



State Of The Art: Push-Pull Direction Switching

Push



parallel_for src in Frontier:
 for dst in out_neigh(src):
 atomic {vtxData[dst] += auxData[src]}

Pull



Pull phase trade-off \Rightarrow work efficiency vs eliminating atomic updates

Advantages of RADAR over Push-Pull

	Push-Pull Direction Switching	RADAR
Work-Efficiency	Work-inefficient	No change
Memory Footprint	2X (outCSR + inCSR)	~1X (reordered outCSR)
Preprocessing Cost	O(E.logE)	O(V.logV + E)











RADAR Can Support Larger Input Graphs

Push-Pull increases memory footprint by 2x



RADAR Can Support Larger Input Graphs

Push-Pull increases memory footprint by 2x



Outline

- Single-node Graph Processing is Sub-optimal V
- Existing Optimizations have Overheads
- RADAR: Combining the Benefits of Duplication and Reordering
- Advantages of RADAR over Push-Pull Direction Switching V
 - RADAR does not compromise work-efficiency
 - RADAR can support larger input graphs
 - > RADAR incurs lower preprocessing overheads









- Single-node Graph Processing Performance is Sub-optimal
- HUBDUP and Degree Sorting optimizations incur overheads
- RADAR offers the best of HUBDUP & Degree Sorting *without their overheads*
- RADAR has many advantages over the state-of-the-art Push-Pull optimization








- Single-node Graph Processing Performance is Sub-optimal
- HUBDUP and Degree Sorting optimizations incur overheads
- RADAR offers the best of HUBDUP & Degree Sorting *without their overheads*
- RADAR has many advantages over the state-of-the-art Push-Pull optimization

https://github.com/CMUAbstract/RADAR-Graph



Thank You!



Combining Duplication & Reordering to Accelerate Parallel Graph Processing

Vignesh Balaji

Brandon Lucia

vigneshb@andrew.cmu.edu

blucia@andrew.cmu.edu

https://github.com/CMUAbstract/RADAR-Graph











- Single-node Graph Processing Performance is Sub-optimal
- HUBDUP and Degree Sorting optimizations incur overheads
- RADAR offers the best of HUBDUP & Degree Sorting *without their overheads*
- RADAR has many advantages over the state-of-the-art Push-Pull optimization

https://github.com/CMUAbstract/RADAR-Graph



Backup Slides



Importance of sizing hubs to LLC size





Carnegie Mellon

All Results for RADAR v/s HUBDUP & DegSort





All Results for RADAR v/s Push-Pull





RADAR is invariant to Graph Order





Push-Pull is better for BC

Algorithm 1 Typical graph processing kernel

- 1: par_for src in Frontier do
- 2: **for** dst in *out_neigh*(src) **do**
- 3: AtomicUpd (vtxData[dst]), auxData[src])

Algorithm 3 Pseudocode for push-phase of BC

- 1: par_for src in Frontier do
- 2: **for** dst in *out_neigh*(src) **do**
- 3: **if** *Visited*[dst] is True **then**
- 4: AtomicUpd (vtxData[dst]), auxData[src])

RADAR Imposes Lower Preprocessing Overhead

All Systems require some form of preprocessing over input file (outCSR):

- HUBDUP: Populate bitvector, map, and inv_map
- Degree-Sorting/RADAR: Reorder input graph by decreasing degrees
- Push-Pull: Construct the inCSR of the graph

	Complexity	DBP	GPL	PLD	TWIT	MPI	KRON	WEB	SD1
HUBDUP	O(V)	0.06s	0.11s	0.14s	0.24s	0.22s	0.23s	0.19s	0.37s
DegSort/RADAR	O(VlogV + E)	0.88s	2.37s	2.29s	8.26s	19.06s	2.94s	3.49s	7.42s
Push-Pull	O(ElogE)	2.96s	7.03s	3.91s	9.68s	47.71s	0s	10.86s	12.51s

HUBDUP	RADAR	Push-Pull		
0.17s	3.93s	9.08		