

Welcome back!


CS 744: POWERGRAPH

Shivaram Venkataraman

Fall 2022

ADMINISTRIVIA

- Midterm grading in progress → *end of this week*
- Course Project: GCP credits



GRAPH DATA


- Web graph hyperlinks

Datasets

- Social network graph

- Maps . Locations → vertices
Roads → edges

- Knowledge graph



- Chemicals / Molecules

Application

Search (important or relevant pages)

Recommendation

Classify chemical structures

GRAPH ANALYTICS

Perform computations on graph-structured data

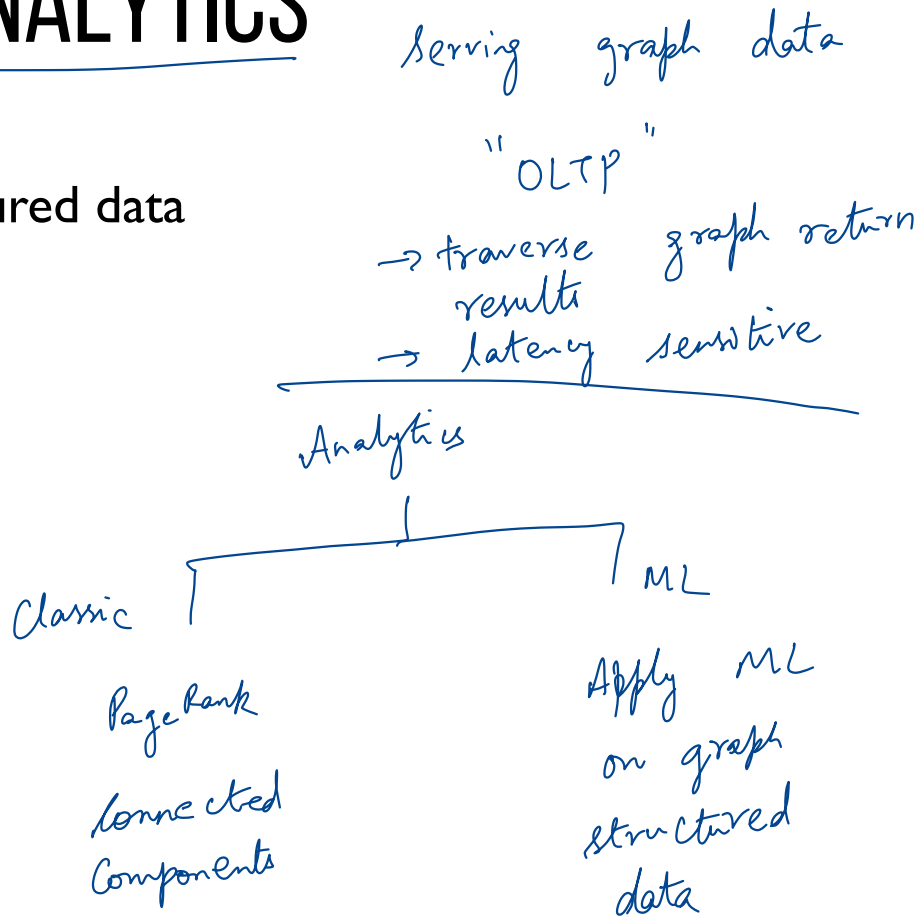
Examples

PageRank

Shortest path

Connected components


...



~ 2008-2009

PREGEL: PROGRAMMING MODEL

"Think like a vertex"



```
Message combiner(Message m1, Message m2):  
    return Message(m1.value() + m2.value());
```

```
void PregelPageRank(Message msg):  
    float total = msg.value();
```

```
    vertex.val = 0.15 + 0.85*total;
```


```
    foreach(nbr in out_neighbors):  
        SendMsg(nbr, vertex.val/num_out_nbrs);
```

① Express graph algorithms.

② Recv inputs from nbrs → Comp on the vertex → send msg to nbrs

NATURAL GRAPHS

1. Exp. distribution of
nbr hood sizes
 - few vertices very large degree
 - large number with very small
2. Lack of symmetry
 - Imbalance in work done per
vertex
 - Imbalance lead to \rightarrow low utilization
 \hookrightarrow stragglers



(a) Twitter In-Degree

POWERGRAPH

Programming Model:
Gather-Apply-Scatter

→ builds on think like a
vertex

Sync / Async execution

→ single machine

Better Graph Partitioning
with vertex cuts

→ Distributed

GATHER-APPLY-SCATTER

Gather: Accumulate info from nbrs

```
// gather_nbrs: IN_NBRS
gather(Du, D(u,v), Dv):
    return Dv.rank / #outNbrs(v)
```

state for source → `D(u,v)`
dest → `Dv`
edge state → `D(u,v)`
Accumulator ← `return Dv.rank / #outNbrs(v)`

Apply: Accumulated value to vertex

```
sum(a, b): return a+b
```

```
apply(Du, acc):
```

Vertex → `Du`
accumulated results of gather → `acc`

```
    rnew = 0.15 + 0.85 * acc
```

```
    Du.delta = (rnew - Du.rank) /  
                #outNbrs(u)
```

```
    Du.rank = rnew
```

Scatter: Update adjacent edges

Gather → input vertex, edge, states
returns accumulator

Apply → return / update vertex state

Scatter → takes update vertex state
and can update nbrs.

```
// scatter_nbrs: OUT_NBRS
```

```
scatter(Du, D(u,v), Dv):
```

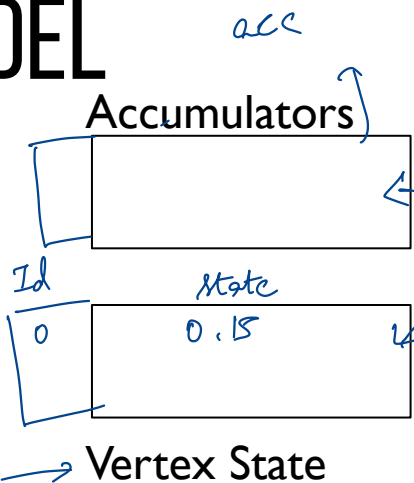
```
    if(|Du.delta| > ε) Activate(v)
```

```
    return delta
```

EXECUTION MODEL

At beginning

- Activate all vertices



Gather

gather(v_0)
[list all nbrs (v_0)]
get their state
accumulate

Apply


apply(v_0)

Comp. ← read write
acc for v_0
state for v_0
book updated


Scatter

activate nbrs in
next iteration

Active Queue




Next Iter Queue




CACHING

→ Accumulators




*Persist across
iters
(reuse)*

→ Vertex State



Active Queue



Delta caching

Cache accumulator value for vertex

Optionally scatter returns a delta

Accumulate deltas

Avoid for running gather that vertex

Reuse acc computed in previous iteration

If few nbors have changed



SYNC VS ASYNC

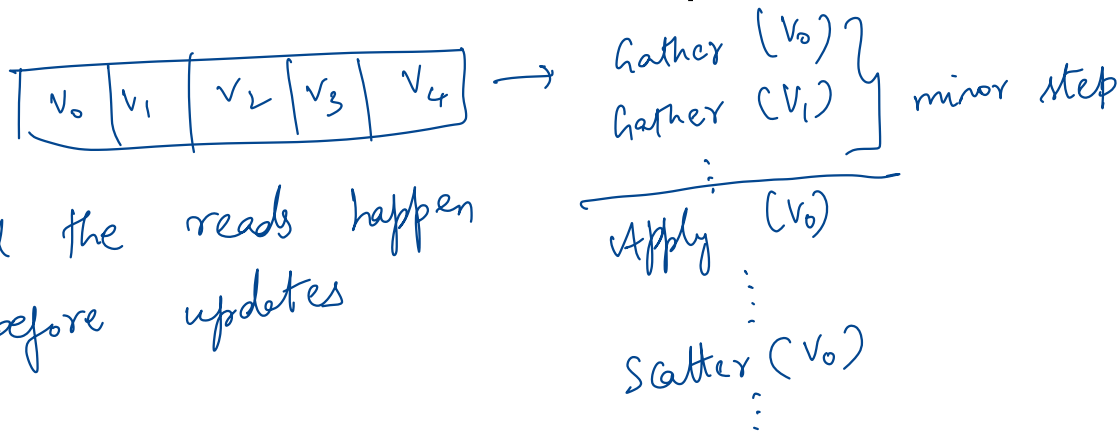
Sync Execution

Gather for all active vertices,
followed by Apply, Scatter

Async Execution

Execute active vertices,
as cores become available

Barrier after each minor-step

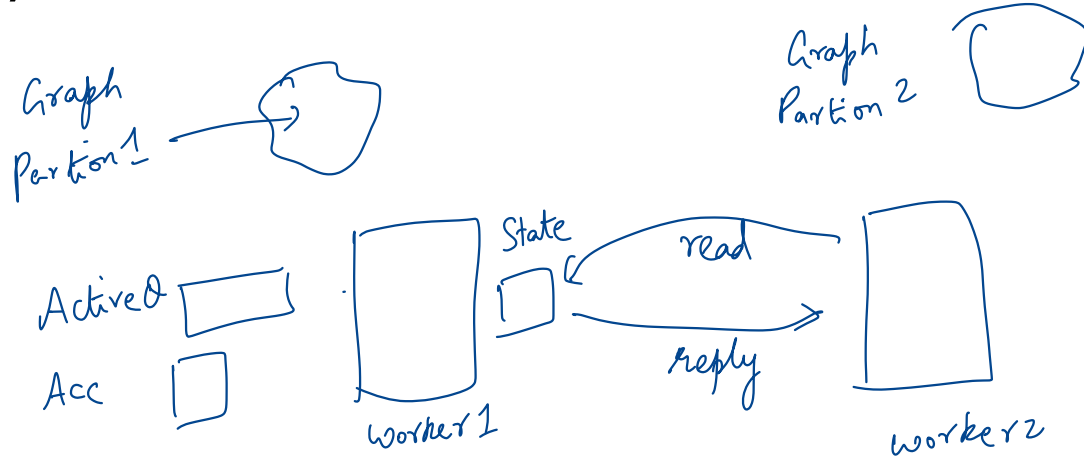


No Barriers! Optionally serializable

Async + Serializable
- Ensure that connected vertices are not processed concurrently

DISTRIBUTED EXECUTION

Symmetric system, no coordinator




Partition graph across machines

Communicate to spread updates, read state

GRAPH PARTITIONING


"read-only" replica
ghost vertices



(a) Edge-Cut

- Assign a vertex to a machine
 - Implies edges are cut across machines - Minimize this
- Induce imbalance for natural graphs

primary
ghost vertices



(b) Vertex-Cut

- Assign edge to particular machine
 - When vertex is on many machines, one primary

RANDOM, GREEDY OBLIVIOUS

Assign edges to machines

Three distributed approaches:

Random Placement


↳ given an edge, choose random machine

Coordinated Greedy Placement

↳ If either vertex is already placed favor those machines

Oblivious Greedy Placement

↳ Avoid coordination, only track vertices present locally



OTHER FEATURES

Async Serializable engine

- Preventing adjacent vertex from running simultaneously

- Acquire locks for all adjacent vertices

Fault Tolerance

- Checkpoint at the end of super-step for sync

SUMMARY

Gather-Apply-Scatter programming model


Vertex cuts to handle power-law graphs

Balance computation, minimize communication


DISCUSSION

<https://forms.gle/K7xk2KybTXf3XX3b6>

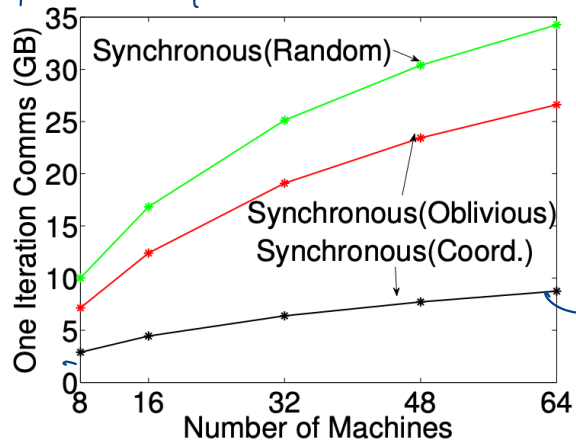
Consider the PageRank implementation in Spark vs synchronous PageRank in PowerGraph. What are some reasons why PowerGraph might be faster?

- 
- ① You can activate a subset of vertices across iterations \Rightarrow work goes down as iters progress
 - ② Partition in Spark \rightarrow random partitions
smarter partitioning here? \rightarrow lower communication
 - ③ Delta caching \rightarrow computation for active vertices can be lowered.

partitioning time?



(a) Twitter PageRank Runtime



(b) Twitter PageRank Comms

diminishing returns get close to same

Comm goes up with machine

Sub-linear with more machines

NEXT STEPS

Next class: Marius