#### CS 744: POWERGRAPH

Shivaram Venkataraman Fall 2020

#### **ADMINISTRIVIA**

- Midterm update ---- Tonight ||
- Course Project reminders



#### **GRAPH DATA**



#### **GRAPH ANALYTICS**

Perform computations on graph-structured data

-> SQL queries on Tobular data

Examples

• • •

#### NATURAL GRAPHS (1) Distribution of degree is skewed! - most vertices have small degree - some vertices have very high degree Number of Vertices (2) High degree vertices lead to skew in $\alpha = 1.7$ Communication Ly Memory pressure (state) Ly computation $10^{0}$ 10<sup>0</sup> 2 Hard to partition such graphs 10<sup>2</sup> In Degree (a) Twitter In-Degree

#### POWERGRAPH

Programming Model: Gather-Apply-Scatter

- Execution

Better Graph Partitioning with vertex cuts

Distributed execution (Sync, Async)



## A1 A1 A2 A3 GATHER-APPLY-SCATTER state edge

Gather: Accumulate info from nbrs

Apply: Accumulated value to vertex

Scatter: Update adjacent edges, vertices



Activate on a reighboring vertex from scatter(Du,D(u,v Activate on a reighboring vertex from scatter(Du,D(u,v if(|Du.delta process necessary vertices in next return delta

// gather\_nbrs: IN\_NBRS
gather(Du, D(u,v), Dv):
 return Dv.rank / #outNbrs(v)

state

 $\rightarrow$  sum(a, b): return a+b

```
apply(Du, acc):
    rnew = 0.15 + 0.85 * acc
    Du.delta = (rnew - Du.rank)/
        #outNbrs(u) _____ change in value
    Du.rank = rnew
```

// scatter\_nbrs: OUT\_NBRS
scatter(Du,D(u,v),Dv):
 if(|Du.delta|> ε) Activate(v)
next return delta
iteration





#### SYNC VS ASYNC

Async Execution

Sync Execution

Barrier followed by Apply, Scatter

Barrier after each minor-stepread $G(V_1)$ ensuresreighbor $G(V_2)$ ensuresstate $G(V_2)$ edgeiBarrierstateupdate $A(v_1)$ statelocal $A(v_2)$ isstatei $A(v_2)$ stateiBarrierstateiiBarrierirextii</t

Execute active vertices, as cores become available

No Barriers! Optionally serializable  $G(V_i)$   $A(V_i) \longrightarrow updates$  vertex state  $G(V_i)$   $G(V_i)$  $G(V_i)$ 

#### **DISTRIBUTED EXECUTION**

Symmetric system, no coordinator

Load graph into each machine



Communicate across machines to spread updates, read state

#### **GRAPH PARTITIONING**





(b) Vertex-Cut
 Fvery edge is placed on
 a mechine
 Vertices night be across
 Machines
 Betler balance for
 netural graphs

# RANDOM, GREEDY OBLIVIOUS

Three distributed approaches:

Ly stream through edges send ælge to a random machine Random Placement

Is send edge to a machine that already has one of its vertices **Coordinated Greedy Placement** 

**Oblivious Greedy Placement** 

La greedy in parallel so you don't have perfect knowledge of vertex -> machine

#### OTHER FEATURES

Async Serializable engine

Preventing adjacent vertex from running simultaneously  $\rightarrow$  Acquire locks for all adjacent vertices

Fault Tolerance

Checkpoint at the end of super-step for sync  $\begin{bmatrix} G & dher \\ A & phy \end{bmatrix}$  Super step  $\begin{bmatrix} G & dher \\ A & phy \end{bmatrix}$ 

#### SUMMARY

Gather-Apply-Scatter programming model Vertex cuts to handle power-law graphs Balance computation, minimize communication

### DISCUSSION

https://forms.gle/rKB5hcJgT4NQsFgq8

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



#### NEXT STEPS