

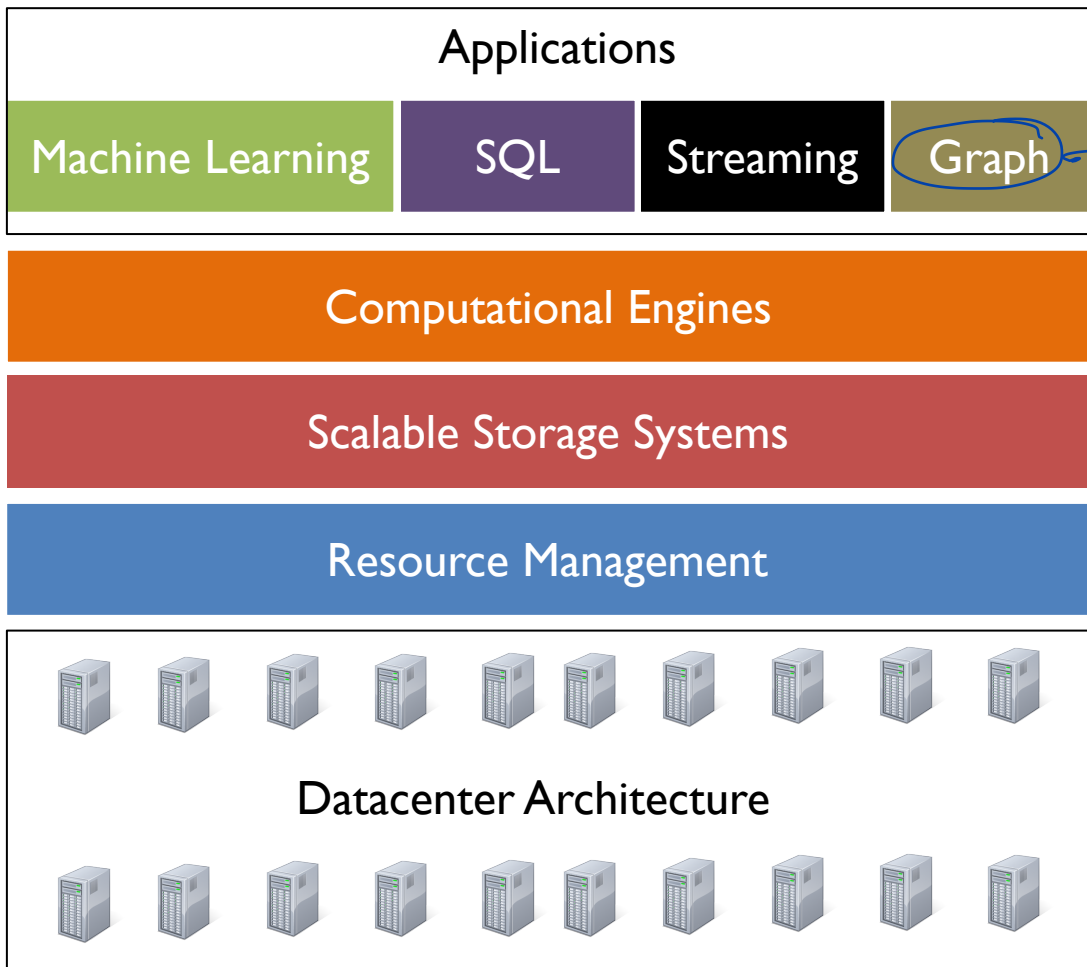
CS 744: GRAPHX

Shivaram Venkataraman

Fall 2019

ADMINISTRIVIA

- Midterm grades are up!
- Course Project: Check in meetings Thu, Mon



Analyze
large
graph data

natural
graphs

POWERGRAPH

Programming Model:
Gather-Apply-Scatter

Better Graph Partitioning
with vertex cuts

Distributed execution
(Sync, Async)

What is different from dataflow system e.g., Spark?

- Fine-grained parallelism
- Less comm using vertex cuts
- Not all are activated

What are some shortcomings?

- Static graphs
 - Fault Tolerance
 - Cross-graph analytics
- ↑ Summary

THIS CLASS

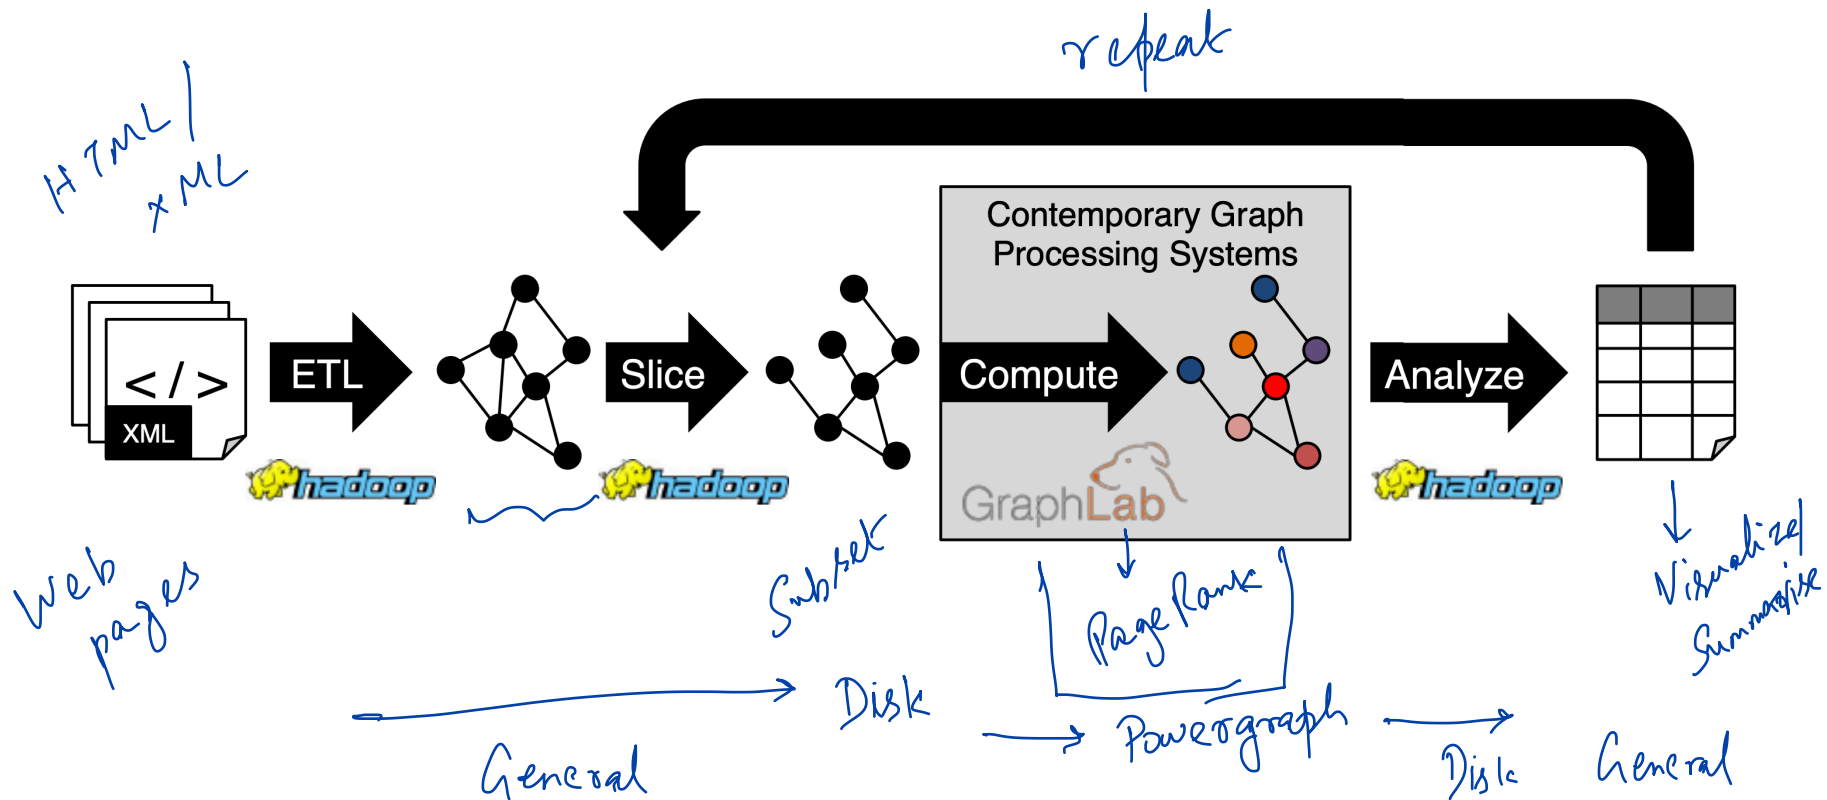
GraphX

Can we efficiently map graph abstractions to dataflow engines?

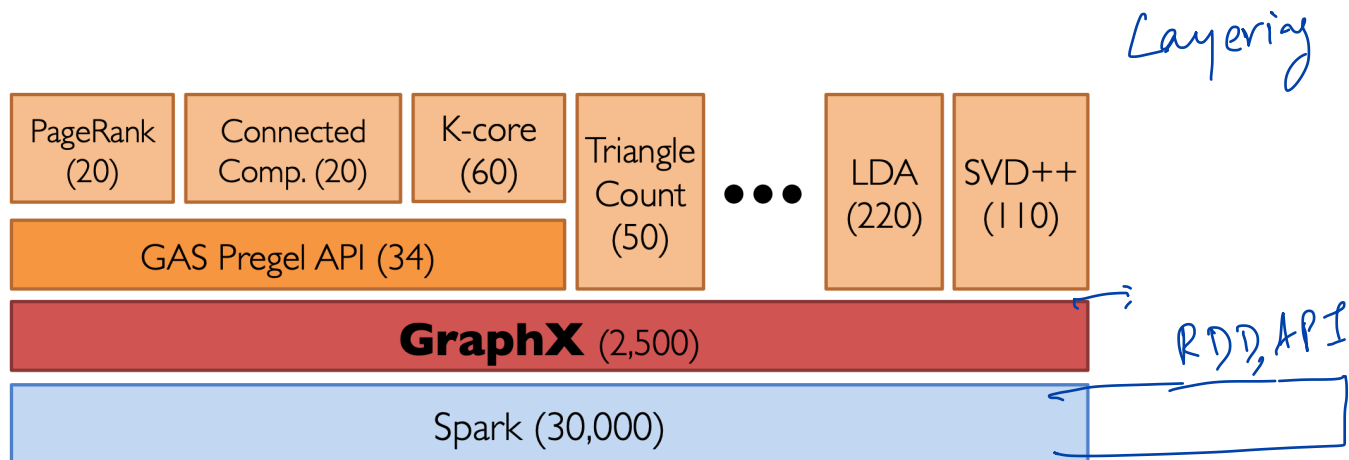
Scalability! But at what COST?

When should we distribute graph processing?

MOTIVATION



SYSTEM OVERVIEW



Advantages?

↳ leverage lower-level

↳ Cluster reuse

Codebase sharing

↳ View data graph / table

PROGRAMMING MODEL

```
class Graph[V, E] {  
  // Constructor  
  def Graph(v: Collection[(Id, V)],  
            e: Collection[(Id, Id, E)])  
  // Collection views  
  def vertices: Collection[(Id, V)]  
  def edges: Collection[(Id, Id, E)]  
  def triplets: Collection[Triplet]  
  // Graph-parallel computation  
  def mrTriplets(f: (Triplet) => M,  
                 sum: (M, M) => M): Collection[(Id, M)]  
  // Convenience functions  
  def mapV(f: (Id, V) => V): Graph[V, E]  
  def mapE(f: (Id, Id, E) => E): Graph[V, E]  
  def leftJoinV(v: Collection[(Id, V)],  
                f: (Id, V, V) => V): Graph[V, E]  
  def leftJoinE(e: Collection[(Id, Id, E)],  
                f: (Id, Id, E, E) => E): Graph[V, E]  
  def subgraph(vPred: (Id, V) => Boolean,  
               ePred: (Triplet) => Boolean)  
    : Graph[V, E]  
  def reverse: Graph[V, E]  
}
```

Constructor

Vertex & Edge
Collection

Vertex

ID	V

Edges

S	D	E

Triplets

Join vertex &

S.ID, D.ID, E, V.S, V.D

select from the edge table

and join vertex E.Source:
V.ID

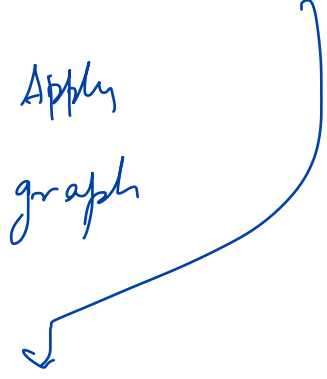
and join vertex E.dest: V.ID

MR TRIPLETS

`mrTriplets(f: (Triplet) => M, sum: (M, M) => M): Collection[(Id, M)]`

`map: Triplet → Message → Apply`


`sum: combine → Sum Powergraph`



Dest vertices → Final value
of
message
associated
with it

Limit Sync execution

PREGEL USING GRAPHX



```
def Pregel(g: Graph[V, E],  
  vprog: (Id, V, M) => V,  
  sendMsg: (Triplet) => M,  
  gather: (M, M) => M): = {
```

`g.mapV((id, v) => (v, halt=false))` → All vertices are active

→ `while (g.vertices.exists(v => !v.halt)) {` → While there are active vertices

`val msgs: Collection[(Id, M)] =`

`g.subgraph(ePred=(s,d,sP,eP,dP)=>!sP.halt)`
`.mrTriplets(sendMsg, gather)`

Filter to get active vertices

`g = g.leftJoinV(msgs).mapV(vprog)`

→ }

`return g.vertices`

}

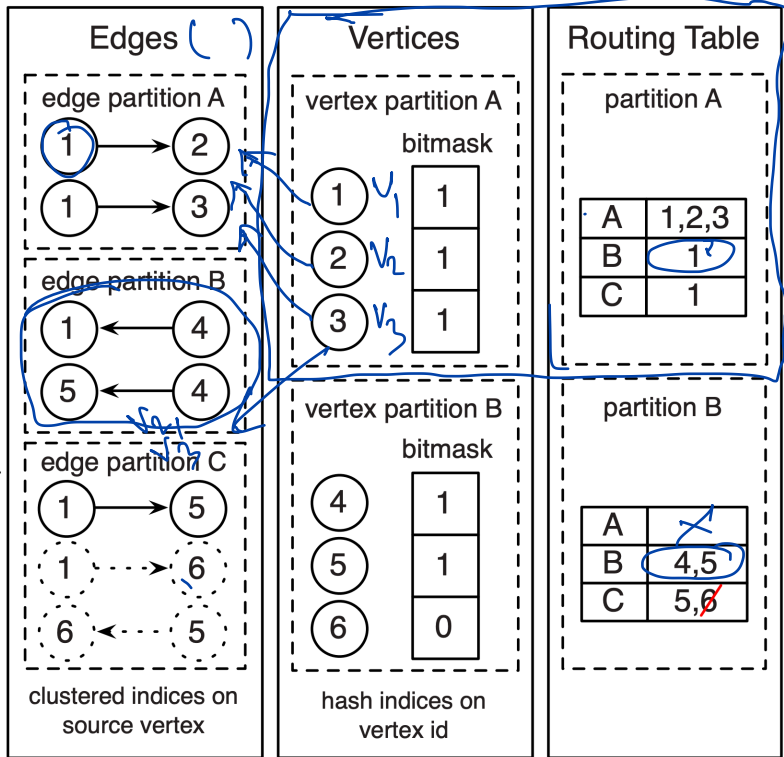
Derive new vertex property given old vertex and msgs

IMPLEMENTING TRIPLETS VIEW

RDD [Edge]

RDD [Vertex]

Hash partition



Join strategy

- Send vertices to the edge site

Default: Use FS partitions

- Number of vertices $\ll |E|$

Multicast join

Using routing table

A. 1, 2, 3 to machine A
I will send 1, 2, 3

B. 1 Incremental / Partial materialization

Edge out
→ Put all of edges for vertex in same machine

Vertex
→ Greedy "locality"

Broadcast
→ All vertices


OPTIMIZING MR TRIPLETS

Filtered Index Scanning \rightarrow Only process active vertices

Store edges clustered on source vertex id

Filter triplets using user-defined predicate

List of triplets/edges
map only on active




Automatic Join Elimination

Some UDFs don't access source or dest properties

Inspect JVM byte code to avoid joins


$f(b.\text{Triple} \rightarrow M)$:

if ($t.sp == 10$)
return 1;
return 0;



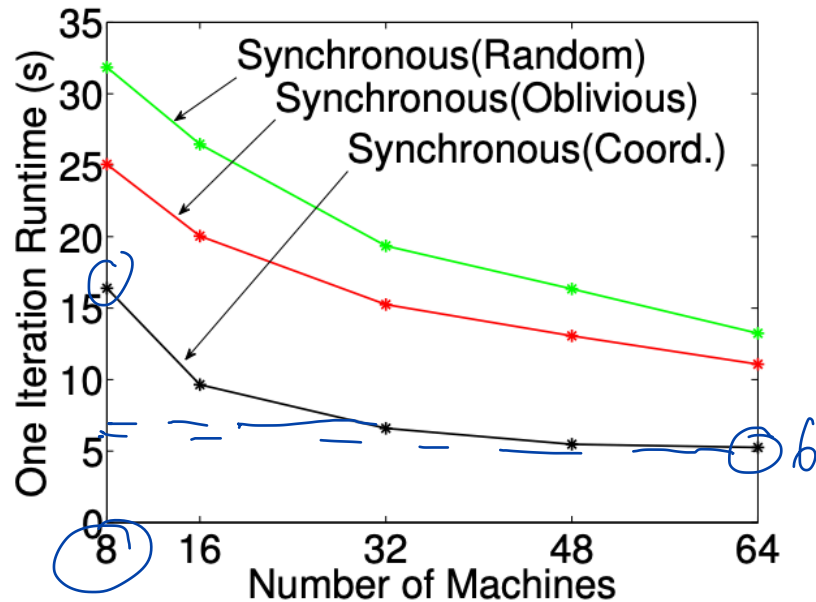
SCALABILITY VS. ABSOLUTE PERFORMANCE

Trend is similar, Twitter, PageRank



GraphX

3x from 8 to 32 machines (4x)



PowerGraph

2.6x from 8 to 32

DISCUSSION

<https://forms.gle/ARaU8Ce9XCpkZznn6>

Consider a single-threaded PageRank implementation as shown and the performance comparison shown in the corresponding table. What could be some reasons for this performance gap?

- Graph X slowest → immutable data lineage tracking

- Single threaded RAM → Avoid SSD lookup

- Overhead of "distribution" → High!
↳ Sync in Pregel

↳

- Will it get better or worse every iteration?

- Graphs are hard to split!

Twitter

→ Graph fits in memory!


100 M vertices	400 MB
----------------	--------

1 B edges	4 GB
-----------	------

≈ 5 GB

PAGERANK has not much compute

Now consider a distributed QR decomposition workload shown in Figure below with corresponding performance breakdown. How would you expect a single-thread implementation to perform here?



What are some workload properties that could explain the difference?

SUMMARY

GraphX: Combine graph processing with relational model

COST

- Configuration that outperforms single-thread
- Measure scalability AND absolute performance
 - Computation model of scalable frameworks might be limited
 - Hardware efficiency matters
 - System/Language overheads

NEXT STEPS

Next class: Weld

Project check-in meetings