

CS 744: NAIAD

Shivaram Venkataraman

Fall 2020

ADMINISTRIVIA

- Course Project Proposal feedback today!
- Midterm grading in progress
- Assignment regrades?

Applications

Machine Learning

SQL

Streaming

Graph

Computational Engines

Scalable Storage Systems

Resource Management




Datacenter Architecture




DASHBOARDS


Sales Dashboard



STREAMING + ITERATIVE COMPUTATION




TIMELY DATAFLOW



TIMELY DATAFLOW

Timestamp : $(\overbrace{e \in \mathbb{N}}^{\text{epoch}}, \overbrace{\langle c_1, \dots, c_k \rangle \in \mathbb{N}^k}^{\text{loop counters}})$



Vertex	Input timestamp	Output timestamp
Ingress	$(e, \langle c_1, \dots, c_k \rangle)$	$(e, \langle c_1, \dots, c_k, 0 \rangle)$
Egress	$(e, \langle c_1, \dots, c_k, c_{k+1} \rangle)$	
Feedback	$(e, \langle c_1, \dots, c_k \rangle)$	

VERTEX API

Receiving Messages

`v.OnRecv(e : Edge, m : Msg, t : Time)`

`v.OnNotify(t : Timestamp)`

Sending Messages

`this.SendBy(e : Edge, m : Msg, t : Time)`

`this.NotifyAt(t : Timestamp)`

IMPLEMENTING TIMELY DATAFLOW

Need to track when it is safe to notify

Path Summary

Check if (t_1, l_1) could-result-in (t_2, l_2)

Scheduler

Occurrence and Precursor count


Precursor count = 0 $\rightarrow$ Frontier

ARCHITECHTURE

Workers communicate using
Shared Queue

Batch messages delivered
Account for cycles

Vertex single threaded



DISTRIBUTED PROGRESS TRACKING

Broadcast-based approach

- Maintain local precursor count, occurrence count

- Send progress update ($p \in \text{Pointstamp}, \delta \in \mathbb{Z}$)

- Local frontier tracks global frontier

Optimizations

- Batch updates and broadcast

- Use projected timestamps from logical graph

FAULT TOLERANCE

Checkpoint

Log data as computation goes on

Write a full checkpoint on demand

Pause worker threads

Flush message queues OnRecv

Restore

Reset all workers to checkpoint

Reconstruct state

Resume execution

MICRO STRAGGLERS

What is different from stragglers in MapReduce?

Sources of stragglers

Network

Concurrency

Garbage Collection

HIGH LEVEL API

```
// 1a. Define input stages for the dataflow.  
var input = controller.NewInput<string>();  
// 1b. Define the timely dataflow graph.  
// Here, we use LINQ to implement MapReduce.  
var result = input.SelectMany(y => map(y))  
                  .GroupBy(y => key(y),  
                           (k, vs) => reduce(k, vs));  
// 1c. Define output callbacks for each epoch  
result.Subscribe(result => { ... });  
// 2. Supply input data to the query.  
input.OnNext(/* 1st epoch data */);  
input.OnCompleted();
```

SUMMARY

Stream processing → Increasingly important workload trend


Timely dataflow: Principled approach to model batch, streaming together

Vertex message model

- Compute frontier
- Distributed progress tracking

DISCUSSION

<https://forms.gle/qn8AzxHMT9VtyEc37>



Consider you are implementing a micro-batch streaming API on top of Apache Spark. What are some of the bottlenecks/challenges you might have in building such a system?

SUMMARY

Next class: Spark Streaming

Course project peer feedback due tonight!