good morning!!

CS 744: FLINK

Shivaram Venkataraman

Spring 2024

ADMINISTRIVIA

Grading

- Assignment 2 grading → released
- Course Project Proposal feedback → today / tomorrow
- Midterm → next week

Resources for Course Projects

- Cloudlab (Reservations?) → Survey ?
- GCP credits (Email Tzu-Tao and me)
  → link

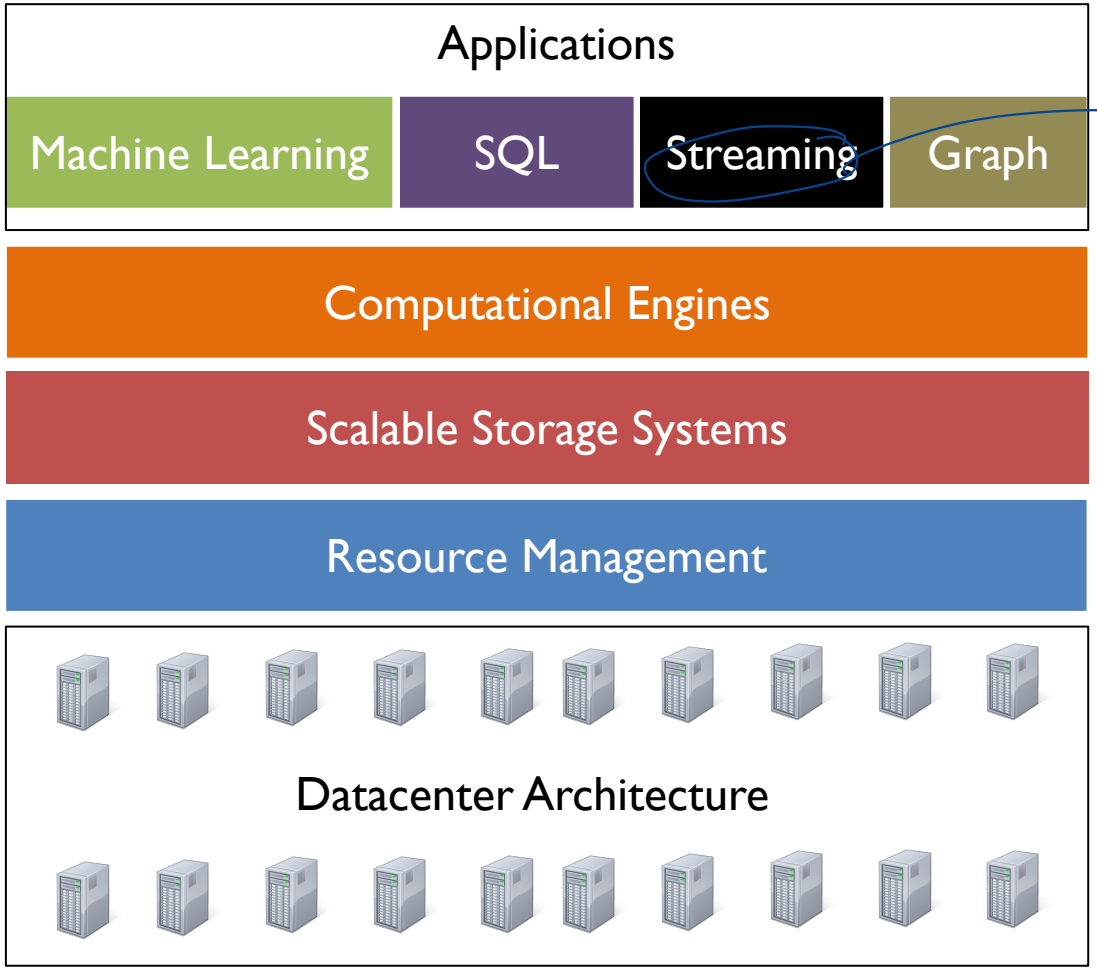

API

features

Windowing

Triggers

- Watermark

DASHBOARDS

update regularly

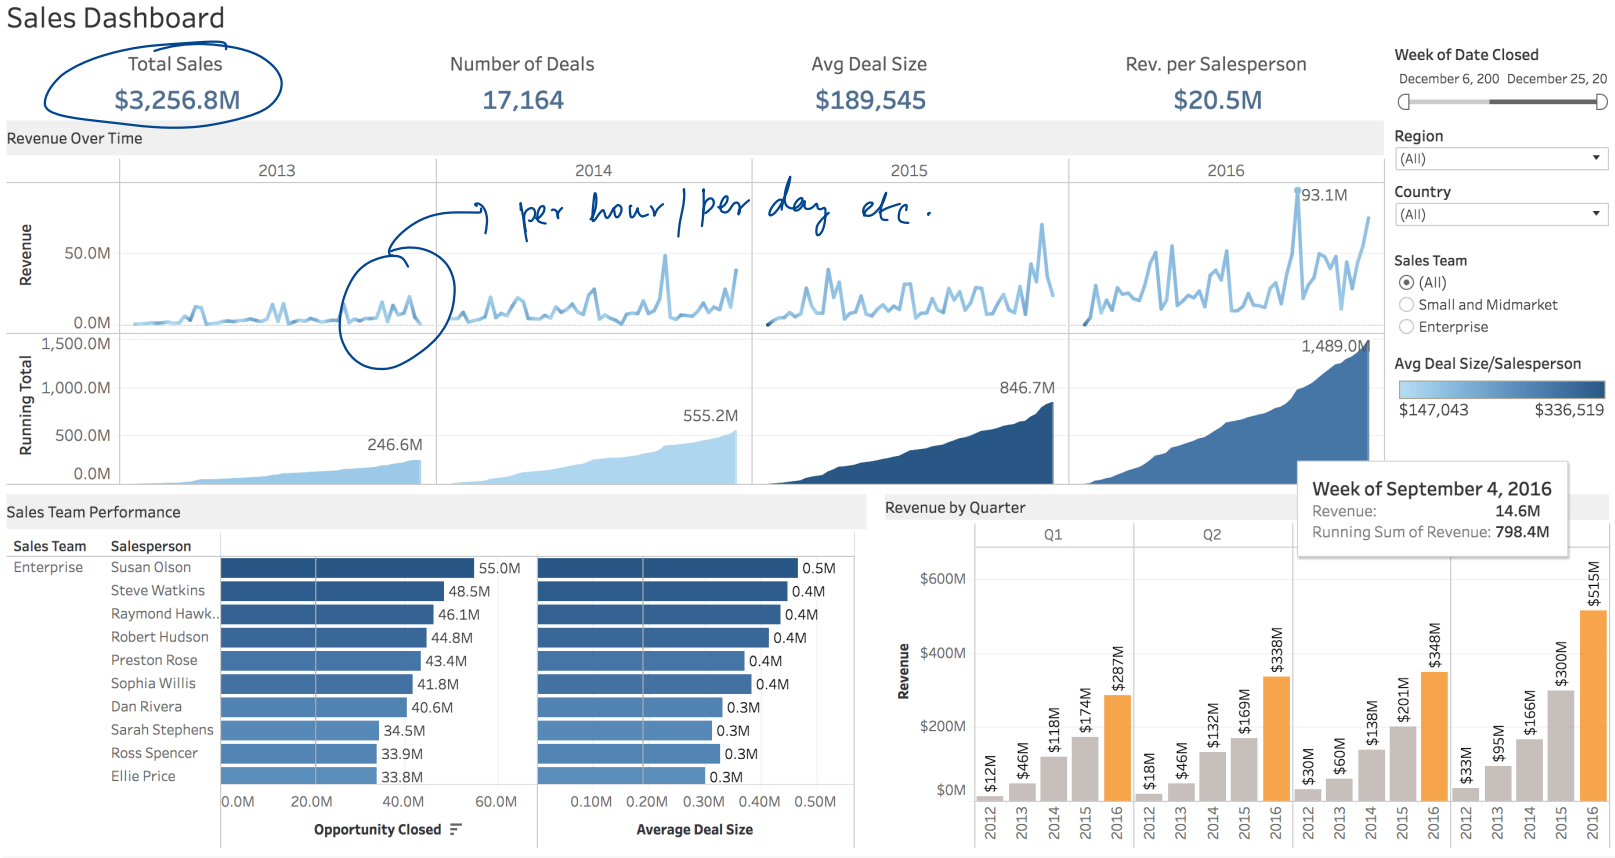

STREAMING COMPUTATION

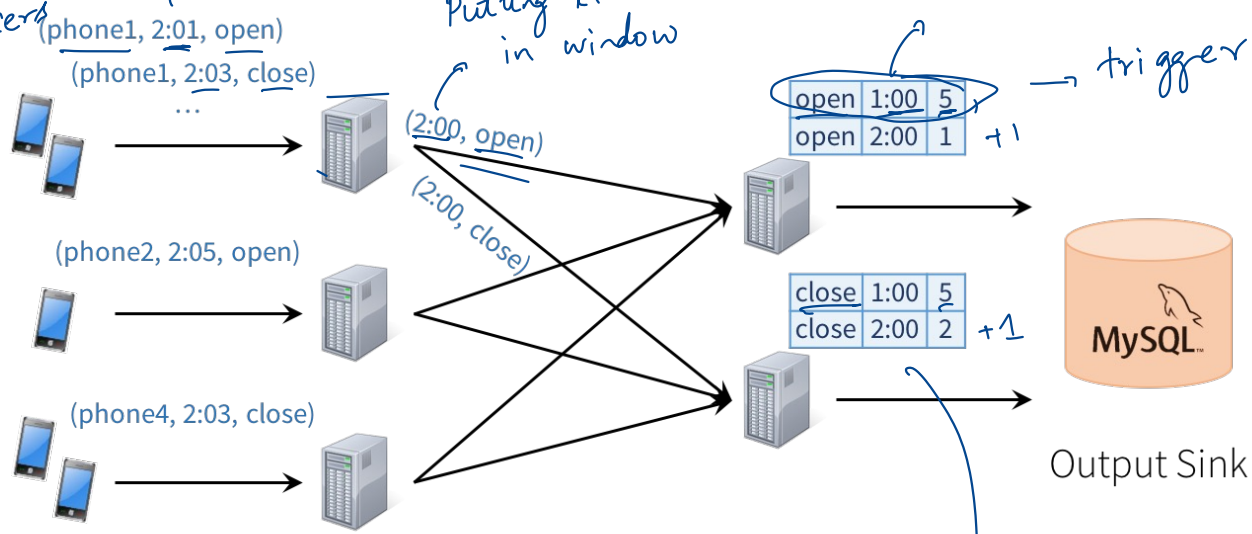

① MR → wait for all mappers before reducers

Intermediate data disk. Streaming them here!

event TS

Putting it in window

windows on event ts

→ trigger

Here: we are doing them concurrently

Database / dashboard etc.

② FT is different

③ Bottlenecks

④ Input source

mutable state that is present in operators

FLINK: COMPUTATION MODEL

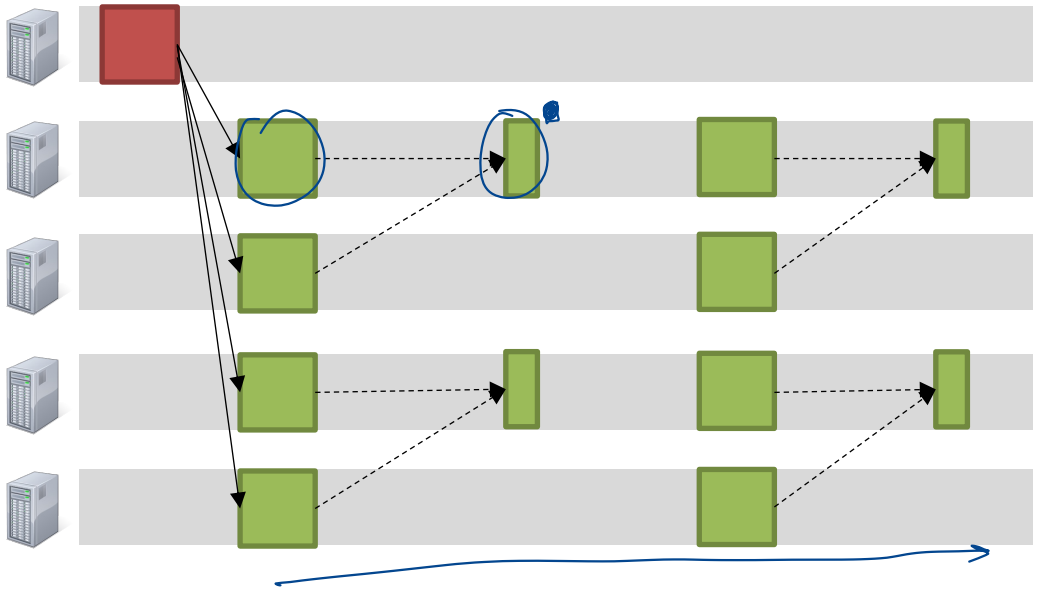

Started at beginning

Long-lived operators

Mutable State

Google MillWheel

Flink

Streaming DBs: Borealis, Flux etc

Naiad

Manager/Driver

Task/Computation

Control Message

Network Transfer

INTERMEDIATE DATA STREAMS

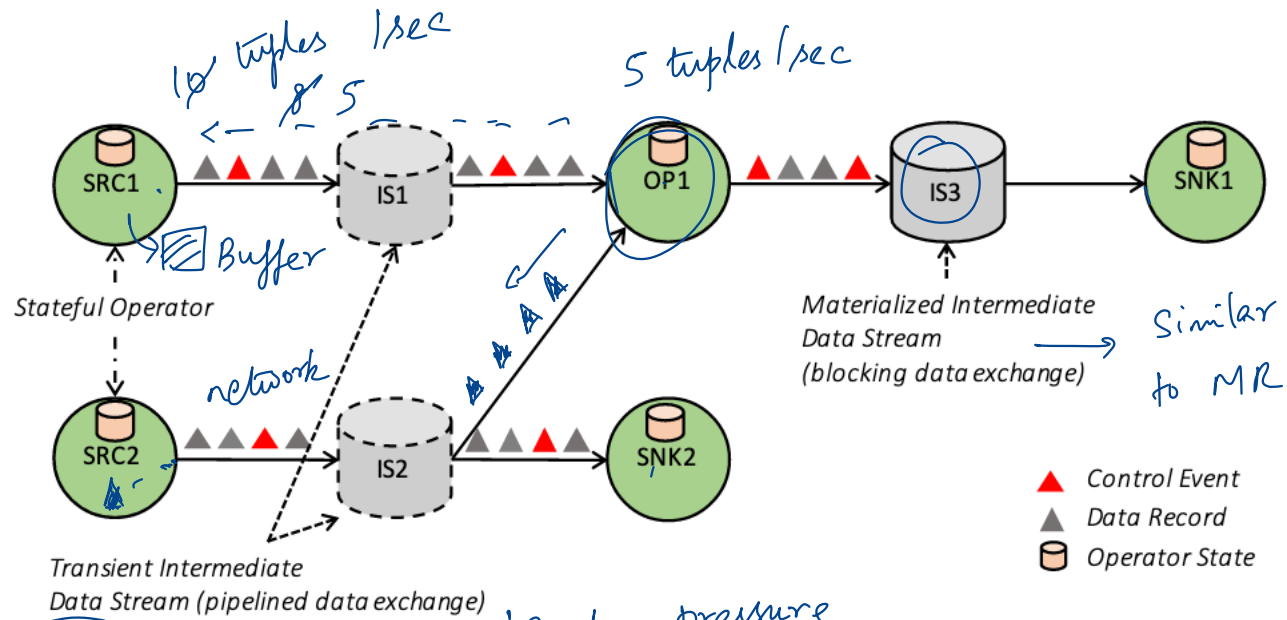

Data Stream

Transient

↳ Pipeline

Comm. with computation

↳ 1-1 or 1-many or many-1 dependency

Trigger

- Buffer full
- Timeout [20 ms]

Back pressure

- signal that goes from dest to src to say slow down

STATEFUL OPERATORS

Examples?

→ Windowing operations

→ Aggregation

Stateless operators

→ map() → input tuple → output tuple

→ filter()

Challenge

How to ensure fault tolerance?

→  → What do we do?

Explicitly register local variables

StateBackends that are automatically saved/recovered

→ annotate some variables as state variables in operator → Periodically checkpoint state variables

FAULT TOLERANCE: CHECKPOINTING

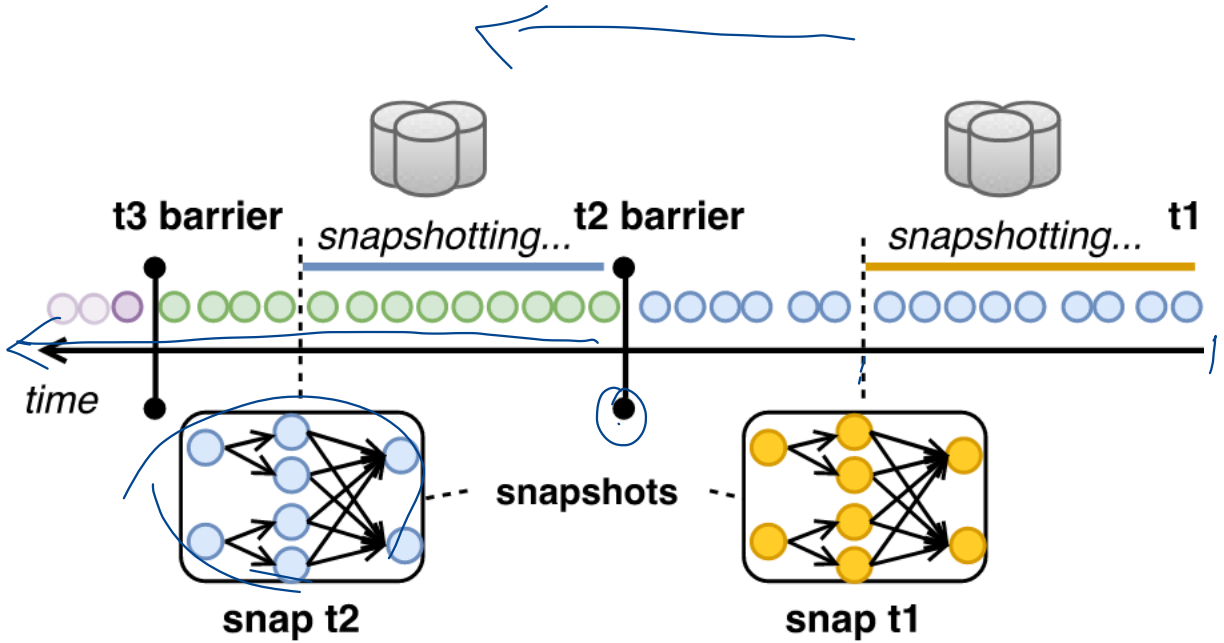

Guarantee that we want:

→ Exactly once semantics

↳ useful

each input's effect is seen exactly once in the output!

Periodically take checkpoints → Restart the operator and replay tuples that came after checkpoint. → Input is replayable!

ASYNCHRONOUS BARRIER SNAPSHOTTING

all operators have processed upto some tuple → consistent snapshot

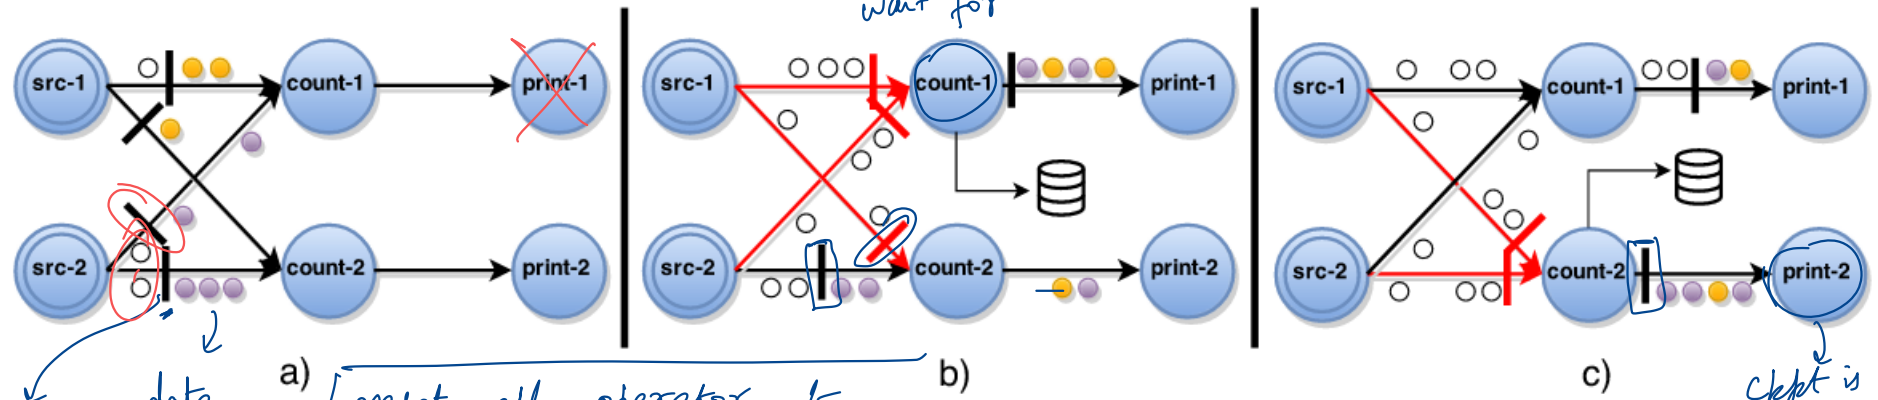

check point

data records

reset all operator to checkpoint state

ckpt is complete

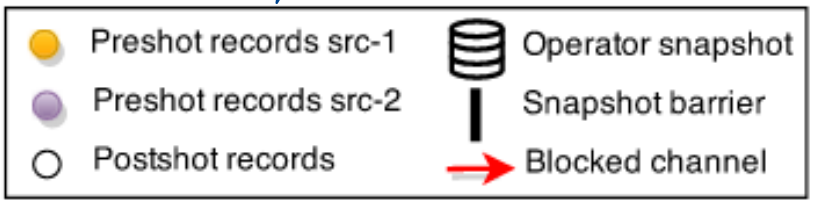

① Control message to ckpt state is saved when message is recvd.

② checkpointing takes time / slowdown during ckpt

WATERMARKS, WINDOWS

Implements similar model as Dataflow

"Watermarks originate at the sources of a topology"

outside of Flink

Propagate through the other operators of dataflow

Windows based on event-time, processing time, ingest time(?)

same as in Dataflow paper.

```
stream
.window(SlidingTimeWindows.of(
        Time.of(6, SECONDS), Time.of(2, SECONDS))
.trigger(EventTimeTrigger.create())
```

COMBINING BATCH, STREAMING

Blocked DataStreams → intermediate data to disk

Turn off periodic snapshots

Blocking operators (e.g., sort) → batch specific operators

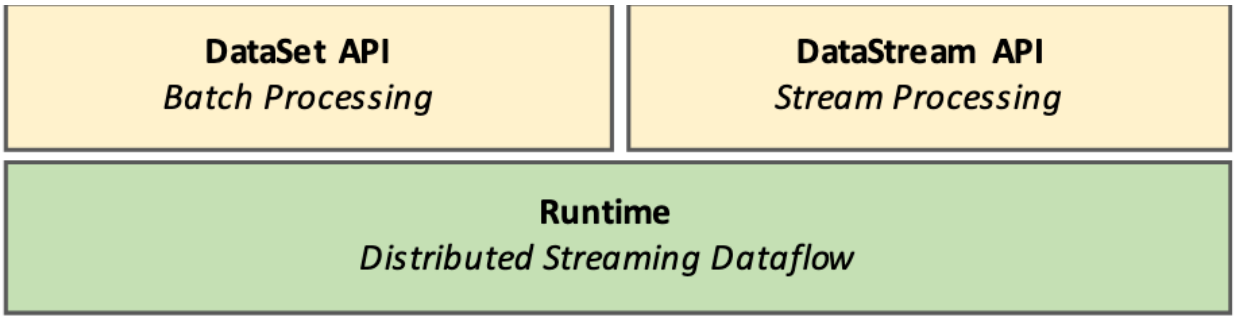

OVERALL ARCHITECTURE

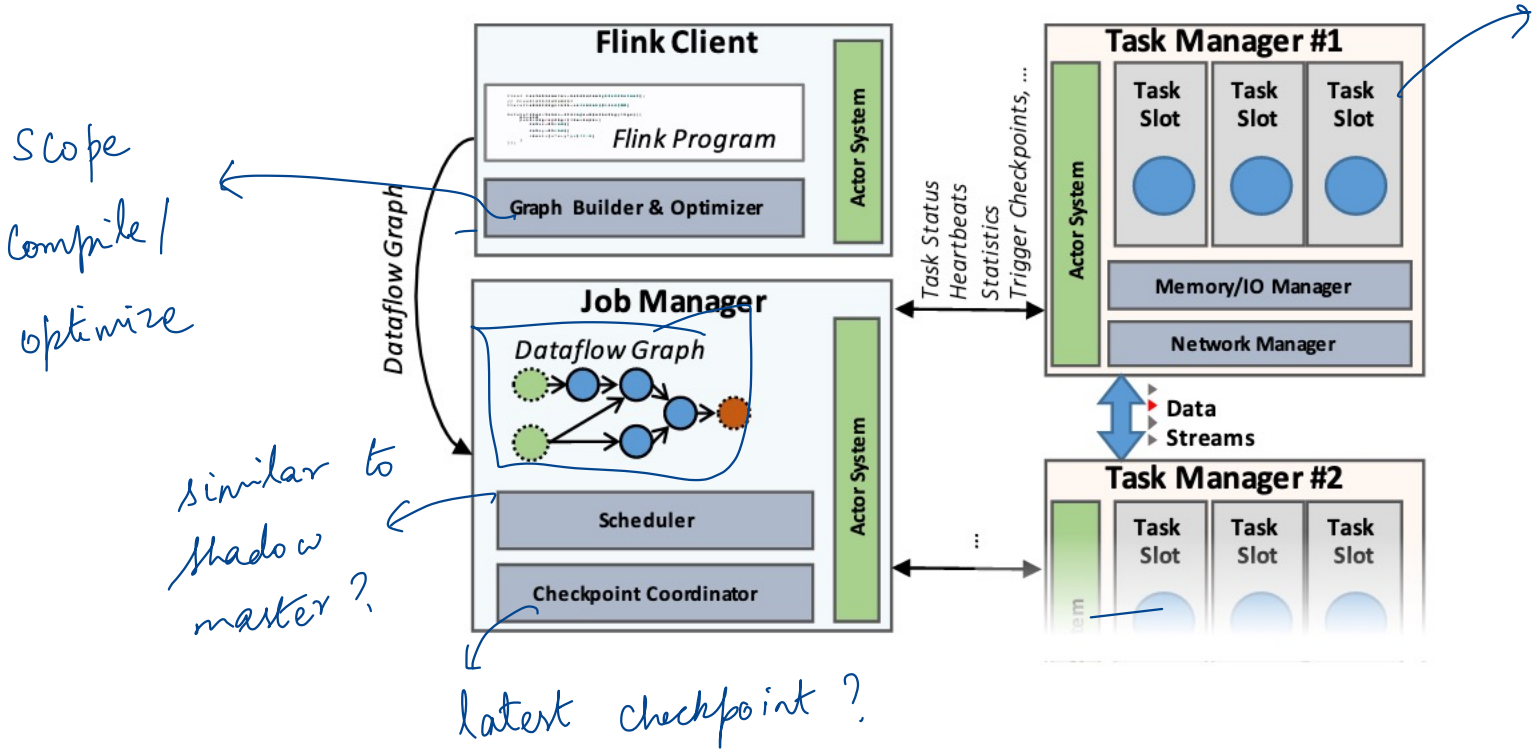

SUMMARY

Stream processing → Increasingly important workload trend

Flink: Distribtuted streaming dataflow to run streaming, batch, iterative

Distribtuted streaming dataflow

- Stateful operators
- Checkpointing based FT

DISCUSSION

https://forms.gle/j9Z7rm4qQpogbz5W8

① Timeout ↑

- higher latency ↑
- higher tput ↑

② Tput wins flatten out

- You might buffer size limit before ?

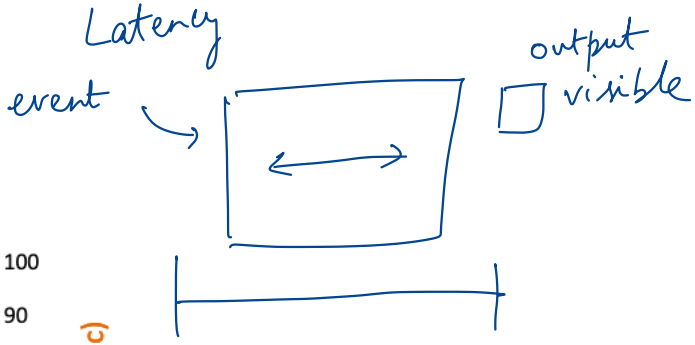

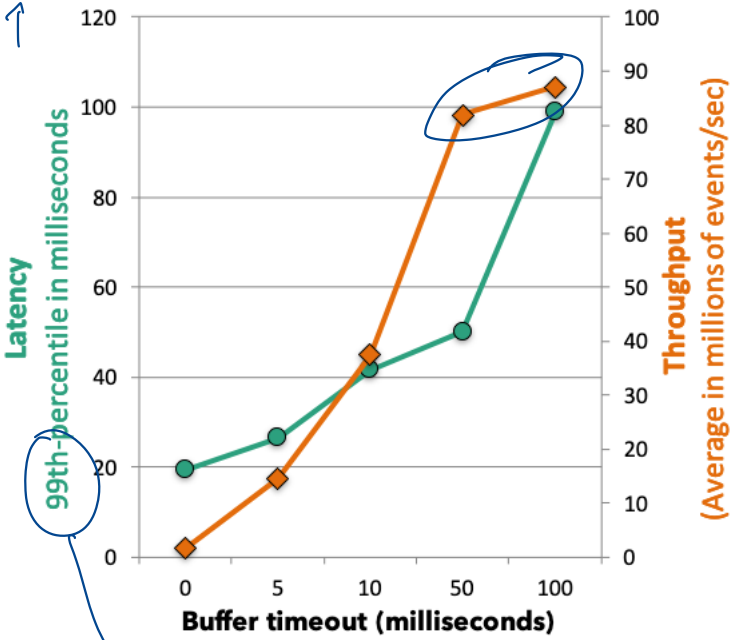

worst case behavior

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?

→ Events to be processed fast
  → Tasks launched when prev. stage finishes ??

→ Windows ?
  →  → output for window → data arrives out of order ??

→ Data shuffling ? → each batch is new data
  → State across batches ? → memory ?

SUMMARY

Next week: Spring break!!

Next class: Spark Streaming