

CS 425 / ECE 428

Distributed Systems

Fall 2025

Aishwarya Ganesan

W/ Indranil Gupta (Indy)

Lecture 22 B: Stream Processing

All slides © IG

Stream Processing: What We'll Cover

- Why Stream Processing
- Storm

Stream Processing Challenge

- Large amounts of data => Need for real-time views of data
 - Social network trends, e.g., Twitter real-time search
 - Website statistics, e.g., Google Analytics
 - Intrusion detection systems, e.g., in most datacenters
- Process large amounts of data
 - With latencies of few seconds
 - With high throughput

MapReduce?

- Batch Processing => Need to wait for entire computation on large dataset to complete
- Not intended for long-running stream-processing

Which one of these is NOT a stream processing job?

A) Uber

Calculating surge prices [<https://www.youtube.com/watch?v=YUBPimFvcN4>]

B) LinkedIn

Aggregating updates into one email [<http://www.vldb.org/pvldb/vol10/p1634-noghabi.pdf>]

C) Netflix

Understanding user behavior to improve personalization
[https://www.youtube.com/watch?v=p8qSWE_nAAE]

D) TripAdvisor

Calculating earnings per day & fraud detection [<https://www.youtube.com/watch?v=KQ5OnL2hMBY>]

E) None of them are stream processing

F) → ALL of them are stream processing jobs!



Enter Storm

- Apache Project
- <http://storm.apache.org/>
- Highly active JVM project
- Multiple languages supported via API
 - Python, Ruby, etc.
- Used by over 30 companies including
 - Twitter: For personalization, search
 - Flipboard: For generating custom feeds
 - Weather Channel, WebMD, etc.



Storm Components

- Tuples
- Streams
- Spouts
- Bolts
- Topologies




Tuple

- An ordered list of elements
- E.g., <tweeter, tweet>
 - E.g., <“Miley Cyrus”, “Hey! Here’s my new song!”>
 - E.g., <“Justin Bieber”, “Hey! Here’s MY new song!”>
- E.g., <URL, clicker-IP, date, time>
 - E.g., <coursera.org, 101.102.103.104, 4/4, 10:35:40>
 - E.g., <coursera.org, 101.102.103.105, 4/4, 10:35:42>

Tuple


Stream

- Sequence of tuples
 - Potentially unbounded in number of tuples
- Social network example:
 - <“Miley Cyrus”, “Hey! Here’s my new song!”>,
<“Justin Bieber”, “Hey! Here’s MY new song!”>,
<“Rolling Stones”, “Hey! Here’s my old song that’s still a super-hit!”>, ...
- Website example:
 - <coursera.org, 101.102.103.104, 4/4, 10:35:40>, <coursera.org,
101.102.103.105, 4/4, 10:35:42>, ...



Spout


- A Storm entity (process) that is a source of streams
- Often reads from a crawler or DB



Bolt




- A Storm entity (process) that
 - Processes input streams
 - Outputs more streams for other bolts




Topology

- A directed graph of spouts and bolts (and output bolts)
- Corresponds to a Storm “application”



Topology

- Can have cycles if the application requires it



Bolts come in many Flavors

- Operations that can be performed
 - **Filter:** forward only tuples which satisfy a condition
 - **Joins:** When receiving two streams A and B, output all pairs (A,B) which satisfy a condition
 - **Apply/transform:** Modify each tuple according to a function
 - And many others
- But bolts need to process a lot of data
 - Need to make them fast


Parallelizing Bolts

- Have multiple processes (“tasks”) constitute a bolt
- Incoming streams split among the tasks
- Typically each incoming tuple goes to one task in the bolt
 - Decided by “Grouping strategy”
- Three types of grouping are popular

Grouping

- **Shuffle Grouping**


- Streams are distributed evenly among the bolt's tasks
- Round-robin fashion



Grouping


- **Fields Grouping**

- Group a stream by a subset of its fields
- E.g., All tweets where twitter username starts with [A-H,a-h,0-3] go to task 1, tweets starting with [I-Q,i-q,4-6] go to task 2, tweets starting with [R-Z,r-z,7-9] go to task 3




Grouping

- **All Grouping**
 - All tasks of bolt receive all input tuples



Storm Cluster

- Master (Coordinator or Leader) node
 - Runs a daemon called *Nimbus*
 - Responsible for
 - Distributing code around cluster
 - Assigning tasks to machines
 - Monitoring for failures of machines
- Worker node
 - Runs on a machine (server)
 - Runs a daemon called *Supervisor*
 - Listens for work assigned to its machine
 - Runs “Executors”(which contain groups of tasks)
- Zookeeper
 - Coordinates Nimbus and Supervisors communication
 - All state of Supervisor and Nimbus is kept here



Failures

- A tuple is considered failed when its topology (graph) of resulting tuples fails to be fully processed within a specified timeout
- **Anchoring:** Anchor an output to one or more input tuples
 - Failure of one tuple causes one or more tuples to be replayed

API For Fault-Tolerance (OutputCollector)

- **Emit**(tuple, output)
 - Emits an output tuple, perhaps anchored on an input tuple (first argument)
- **Ack**(tuple)
 - Acknowledge that you (bolt) finished processing a tuple
- **Fail**(tuple)
 - Immediately fail the spout tuple at the root of tuple topology if there is an exception from the database, etc.
- Must remember to ack/fail each tuple
 - Each tuple consumes memory. Failure to do so results in memory leaks.

Twitter's Heron System (Optional Additional Slide)

- Fixes the inefficiencies of Storm's acking mechanism (among other things)
- Uses **backpressure**: a congested downstream tuple will ask upstream tuples to slow or stop sending tuples
- 1. TCP Backpressure: uses TCP windowing mechanism to propagate backpressure
- 2. Spout Backpressure: node stops reading from its upstream spouts
- 3. Stage by Stage Backpressure: think of the topology as stage-based, and propagate back via stages
- Use:
 - Spout+TCP, or
 - Stage by Stage + TCP
- Beats Storm throughput handily (see Heron paper)

Summary: Stream Processing

- Processing data in real-time a big requirement today
- Storm
 - And other sister systems, e.g., Spark Streaming, Heron, (LinkedIn's Samza, "Kafka", etc.)
- Parallelism
- Application topologies
- Fault-tolerance