CMSC 714
Lecture 14
Cloud Computing –
Spark and Mesos

Alan Sussman

Notes

- OpenMP project grades and reports out today or tomorrow

- Feedback on research project proposals in next day or 2

- Exam coming up on November 13
  - sample exam questions posted soon

Spark

- Single engine for distributed data processing
  - SQL
  - stream processing
  - machine learning
  - graph processing
- Basic idea is to enable composing different types of processing into a single application
  - without copying data, so reuse of data and doing operations in memory is fundamental
- Key abstraction is Resilient Distributed Dataset (RDD)
  - a fault tolerant collection of objects (data items) partitioned across a cluster that can be operated on in parallel
- Functional programming API in Scala, Java, Python, and R

Spark (cont.)

- Users/developers write local functions that operate on RDDs
- RDDs evaluated by Spark runtime lazily
  - that means when they are needed, so only when one needs to be instantiated
  - enables creating an execution plan for a whole set of data transformations (like in an RDBMS)
- User can enable sharing an RDD by making it persistent in memory (spilled to disk if too big)
  - this is a big difference from MapReduce implementations
- Fault tolerance – RDDs can be recomputed if lost by keeping track of lineage (how they were computed)
- Can use different external systems for persistent storage
  - e.g., HDFS, S3, Cassandra

Spark (cont.)

- Additional functionality comes from building libraries on top of basic abstractions
  - SparkSQL for relational queries – but no transactions
  - DataFrames – RDDs of records with a known schema, used for tables in R and Python
  - Spark Streaming for incremental stream processing on discretized streams – split input data into small batches (e.g., data that arrives over 200ms) that is combined with state stored in RDDs to produce new results
  - GraphX – graph computation interface – vertex based computations for graphs, and graphs partitioned across nodes
  - Mllib – machine learning library
- Claim is that performance is comparable to specialized systems for each kind of processing
- Last note is that they do admit that synchronization in Spark means it does not work well for latency sensitive computations

Mesos

- A *meta-scheduler* – to enable multiple cluster computing frameworks (e.g., Hadoop, OpenMPI) to share cluster resources
  - an alternative to a centralized scheduler
- Basic idea is that the resources register with Mesos, Mesos offers resources to frameworks, frameworks decide whether to accept or reject the resource offers
  - so frameworks do their own scheduling, once they obtain resources from Mesos
- One catch is that someone has to tell Mesos how to decide which resources to offer to which frameworks
  - this is a policy decision (e.g., fair sharing), and there is a Mesos plugin interface for the policy module
  - similar to how HPC cluster schedulers work – Torque, SLURM

Mesos (cont.)

- Basic architecture is one Mesos *master* that frameworks communicate with, and a Mesos slave daemon on each cluster node
  - each slave process offers resources through its daemon
  - master offers resources to frameworks, which they can accept or reject
  - frameworks decide which offered resources to use – through a scheduler they register with the master
  - framework can then launch tasks on acquired resources through their *executor* process

- Use Zookeeper for fault tolerance
  - a distributed coordination service, to deal with faults in the Mesos master – enables having hot spare copies of the master – *leader election*
  - use *soft state* so new master can reconstruct internal state from slave daemons and framework schedulers

Mesos (cont.)

- Efficiency and robustness
  - Framework can set *filters*, to tell master which offers it will always reject – so master won't even try such offers
  - To give incentive for frameworks to respond quickly to offers, Mesos counts outstanding resource offers toward a framework's allocation of a cluster – so they don't hang onto resources they may not use
  - If a framework does not respond for a while, Mesos rescinds a resource offer

- Performance
  - simulation study shows Mesos provides both good latency to schedulers that need resources, and good cluster utilization, compared to a centralized scheduler
  - Performance best for frameworks that have short tasks to run, and jobs that can scale elastically – so probably not so good for HPC workloads