

Using Realistic Simulation to Study the Hadoop Ecosystem

Guanying Wang, Ali R. Butt
 {wanggy, butta}@cs.vt.edu




Motivation

- How to build an optimal Hadoop cluster?
 - Lots of configurations to choose
 - Impractical to optimize by experiments
 - Simulation is promising (though not as accurate)
- MRPerf realistically simulates Hadoop behavior
 - Guide designing of a Hadoop cluster
 - Help evaluate innovative designs
 - Allow study on impact of failures

Background: Hadoop/MapReduce

- A framework to run large-scale data-intensive applications
- Key enabler for "Cloud Computing"
- Map side
 - Map phase
 - Sort phase
 - Spill phase
- Reduce side
 - Shuffle/sort phase
 - Reduce phase

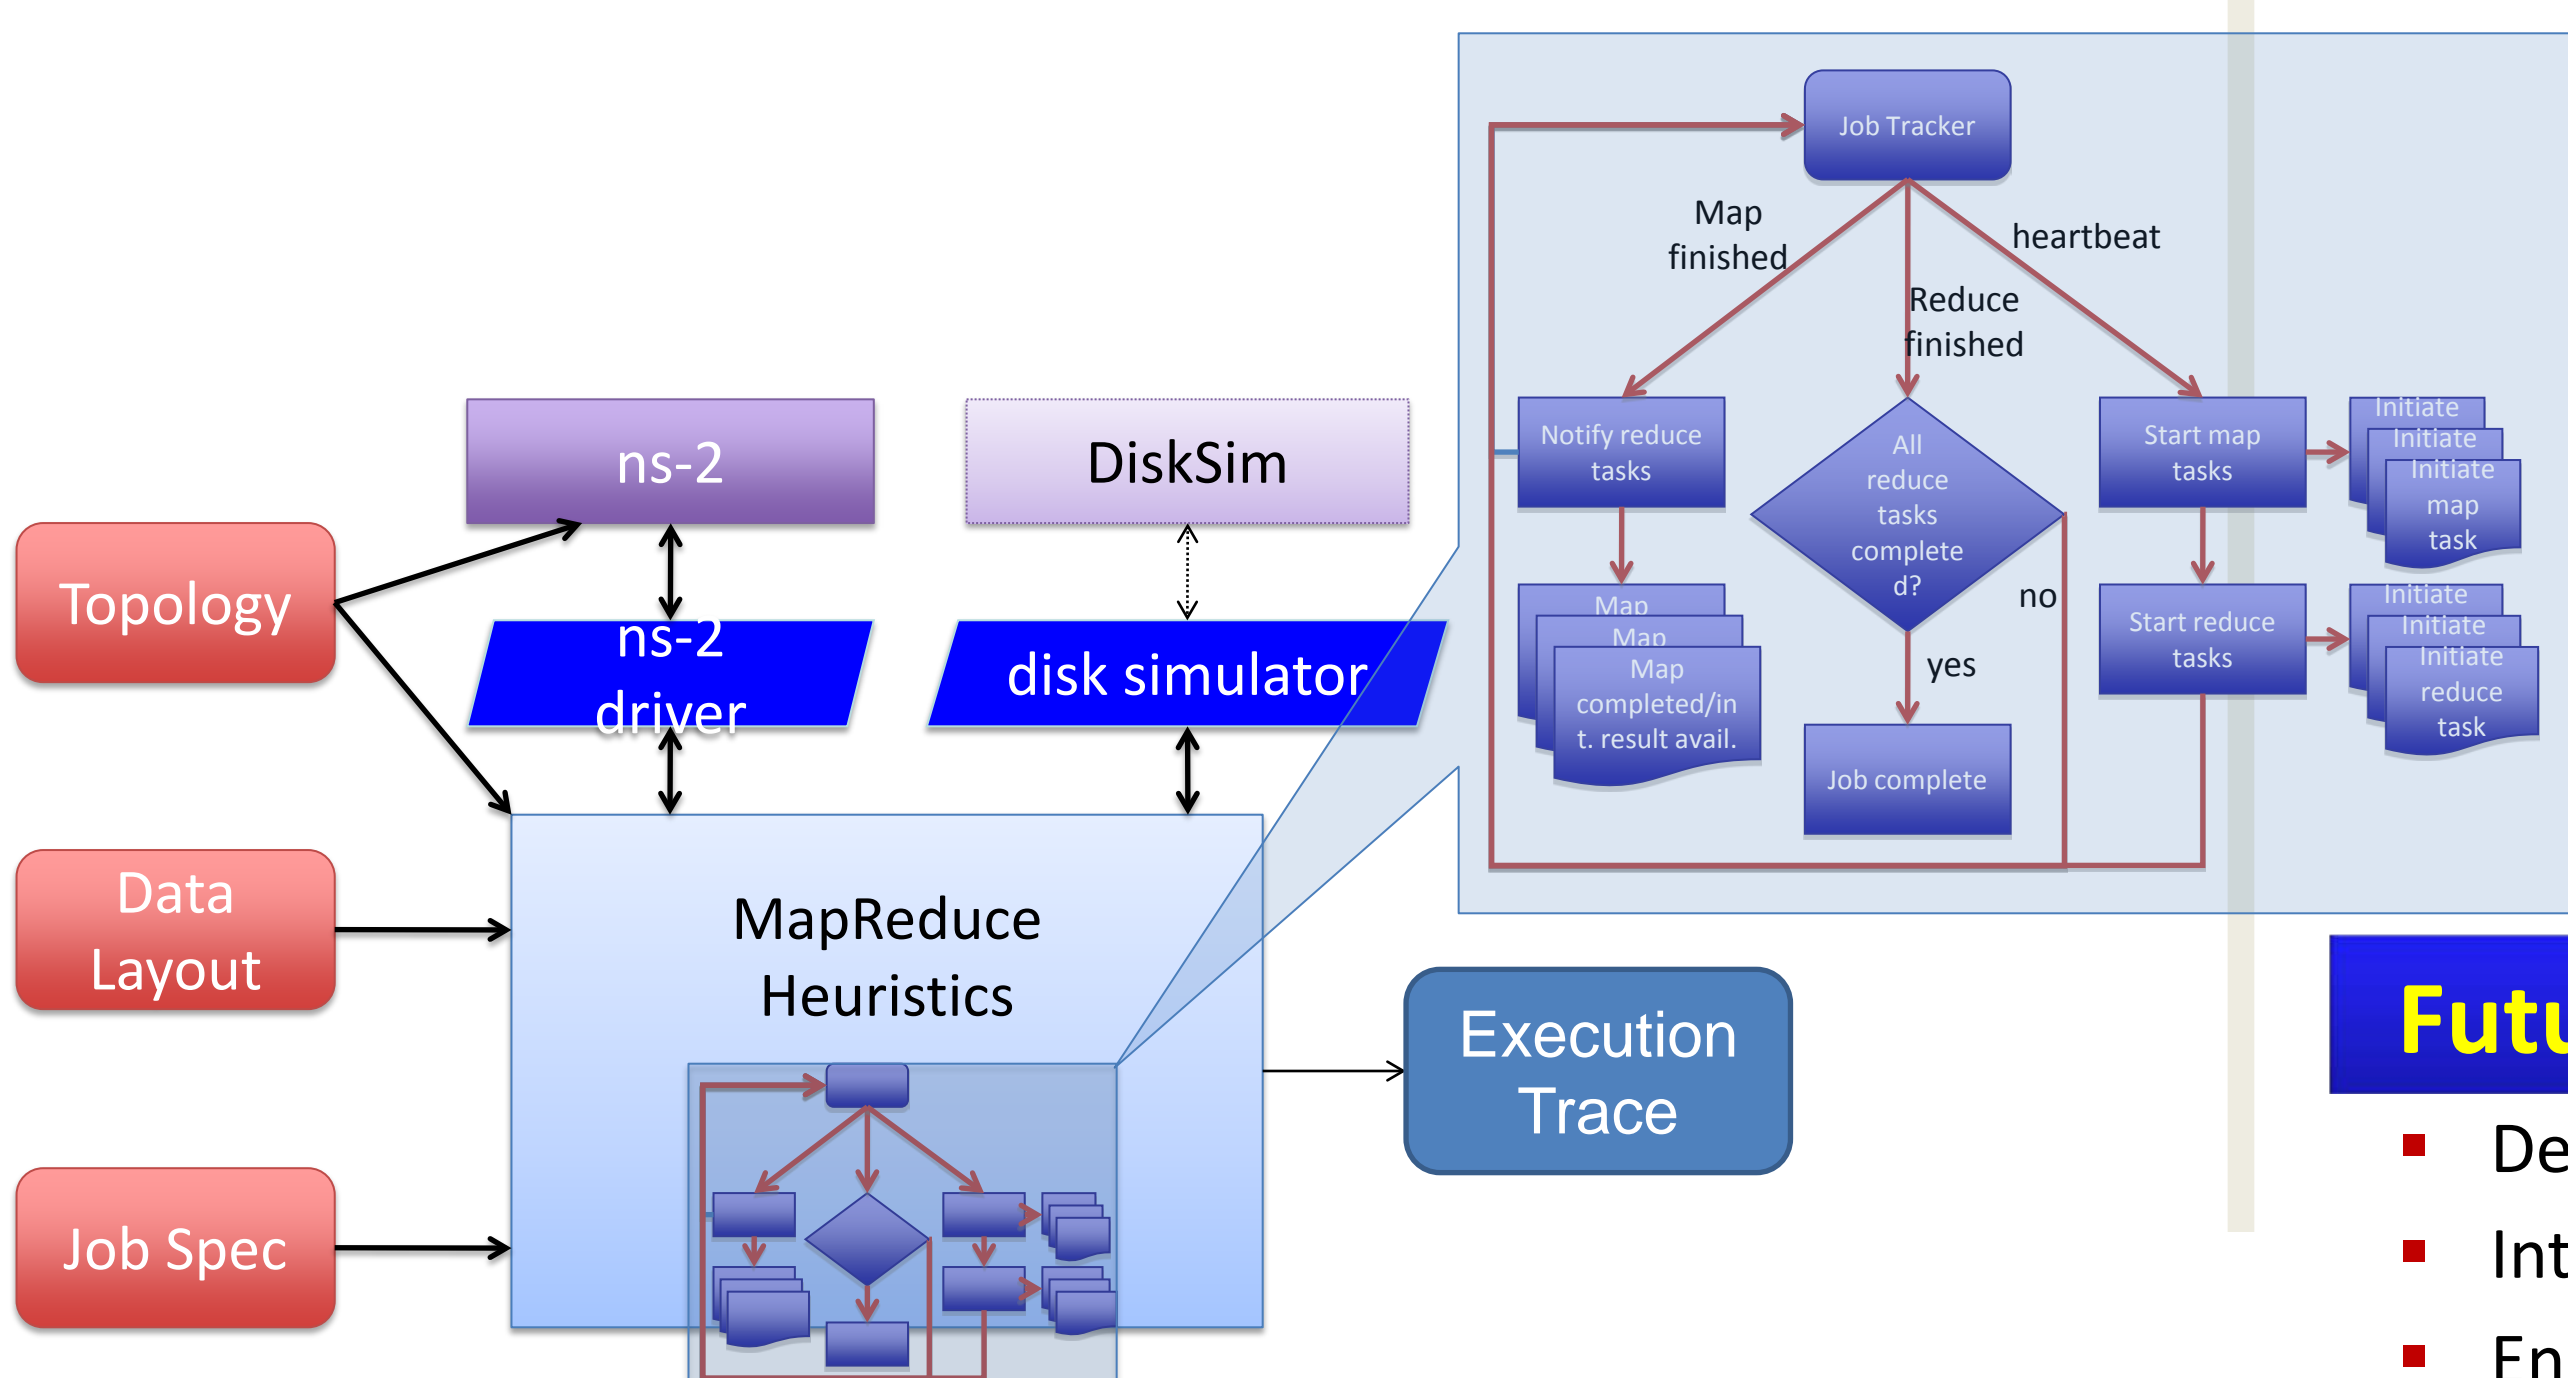


Challenges in Simulator Design

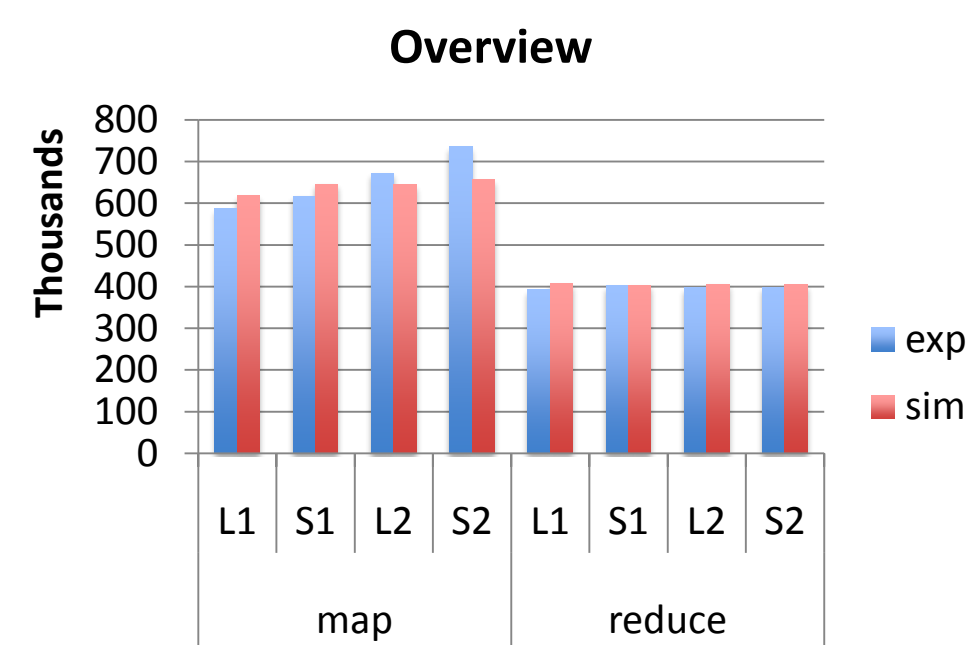
- Selecting right level of abstraction
- Managing application data layout
- Handling skewness of input data
- Simulating failures and recovery mechanism
- Validating the simulator using real test-beds
- Maintaining accuracy of simulator as Hadoop evolves

Design

- All network traffic handled by ns-2
- Simulator implements Hadoop heuristics
- No actual data is transferred, only metadata such as size of data and control protocol are utilized
- Output/computation determined by size of input data

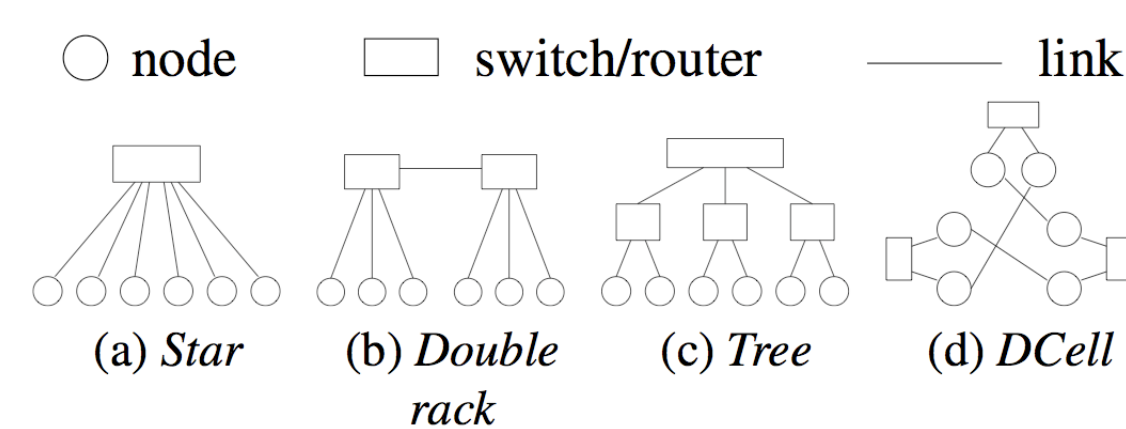


Validation




- Overall, results match (within error range).
- Rack-local map tasks in single-rack experiments show large variation due to disk contention.

Applying MRPerf



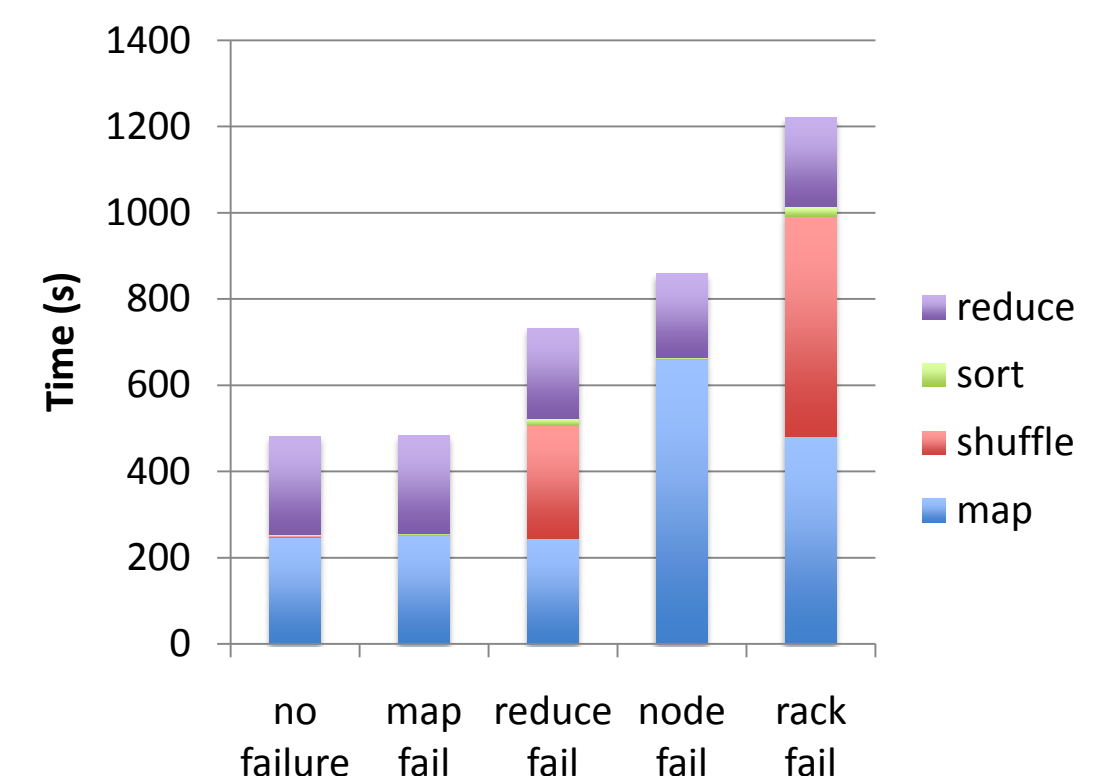
Different topologies evaluated in this work. Double-rack and DCell are evaluated below.

Impact of Data Locality in Different Topologies:




Impact of Failures:

- Four types of failures
 - Map task failure
 - Reduce task failure
 - Node failure
 - Rack failure



Utilization under Different Hadoop Job Schedulers:



Future Work

- Design efficient cloud systems
- Integrating cloud and HPC
- Enabling Cloud to use Accelerators
- Studying energy-cost of providing data reliability in the cloud
- Comparing and evaluating schedulers
- Evaluating innovative design of Hadoop cluster with storage