Reconfigurable Network Testbed for Evaluation of Datacenter Topologies

W. Clay Moody
Clemson University

Co authors: Jason Anderson, K.-C. Wang, Amy Apon

Our Paper (briefly)

- **Word Cloud:**

- **MD5 Hash:** d5d744982be2f62da9972a19f2c0895f

Agenda

- Problem Statement
- Motivation
- Architecture Description
- Use Case: Hadoop
- Experiments and Results
- Conclusion

Problem Statement

Large Datacenters

hadoop

OpenMP

LexisNexis

HPCC Systems

Many Applications

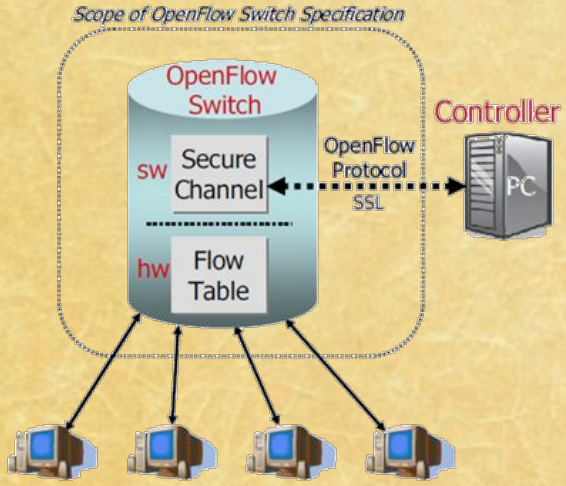

Potential Enhancements

Limited Budgets

Motivation

- Cheap (or free as in beer)
- Reprogrammable
- Representative
- Physically accessible
- Administrative privileges
- Ample Network Connections

Yoda* Cluster

*The placement of these four letters in this specific order does not constitute an endorsement of any little green characters owned by LucasFilm and Disney. This is merely an term chosen to describe our little powerful cluster used by our lab whose informal name is DAGOBA (DAta GOing Beyond Analytics).

Architecture (hardware)

12 Client Workstations (with virtual switch)

Additional LAN ports

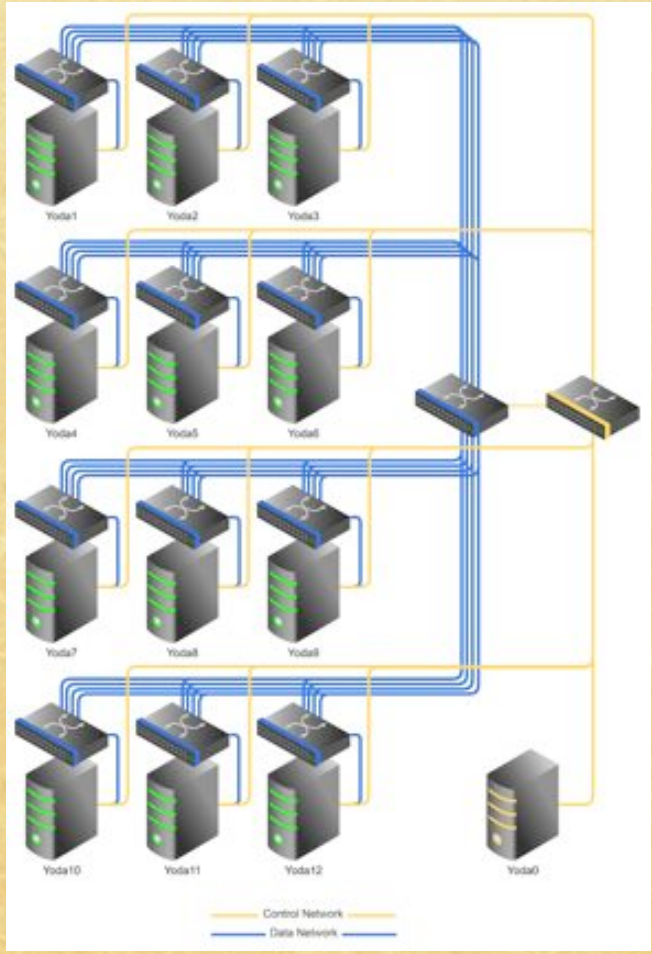

<u>Networks:</u>
control
access
data (reprogramable)

2 SDN switches

Primary Server

Architecture (software)

Open Source

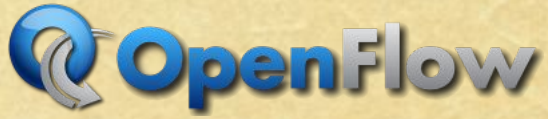

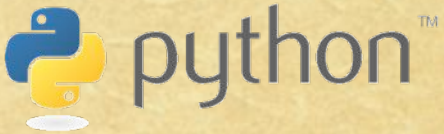

Custom

Powered by NetworkX

Virtual Topology Engine

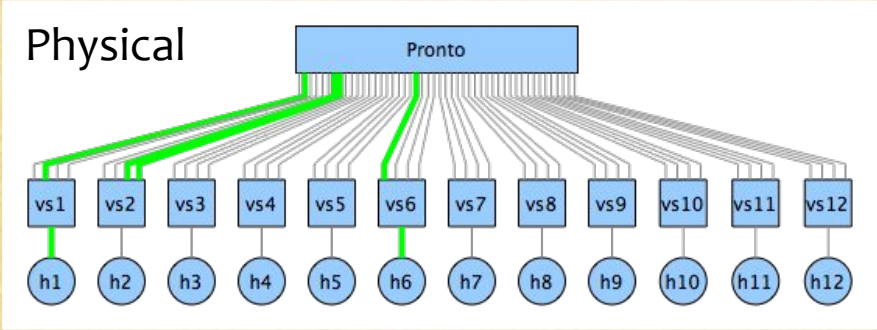

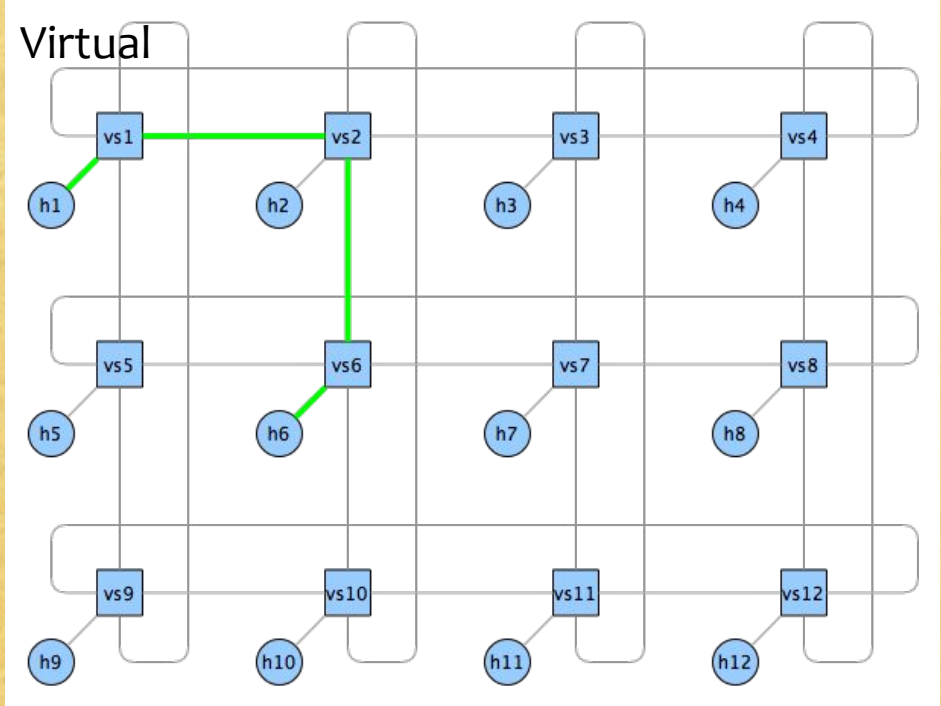

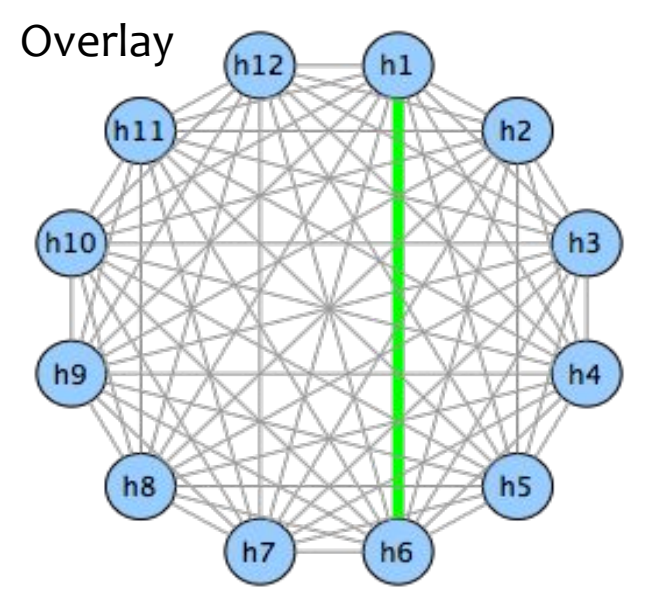

Congestion Matrix

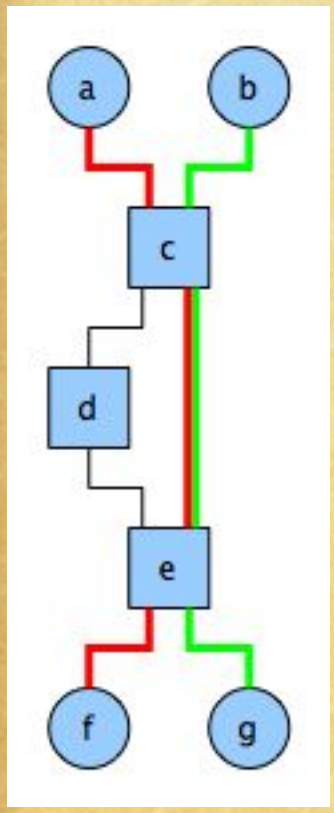

Higher

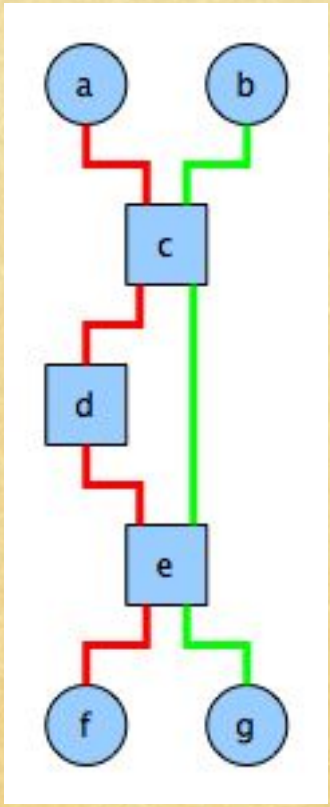

Lower

Algorithm 1 Link Sharing Score

```
1: procedure LSS(pairs, topo)
2:     for all (src, dst) in pairs do
3:         path ← topo.path(src, dst)
4:         path_rate ← min_link(path)
5:         for all edge in path do
6:             edge.usage ← edge.usage + path_rate
7:         end for
8:     end for
9:     for all (src, dst) in pairs do
10:        path ← topo.path(src, dst)
11:        path_rate ← min_link(path)
12:        rate ← path_rate
13:        for all edge in path do
14:            scaled_rate ← path_rate × max(1, edge.cap / edge.usage)
15:            rate ← min(rate, scaled_rate)
16:        end for
17:        total ← total + rate / path_rate
18:    end for
19:    return total / len(pairs)
20: end procedure
```

Use Case: Hadoop

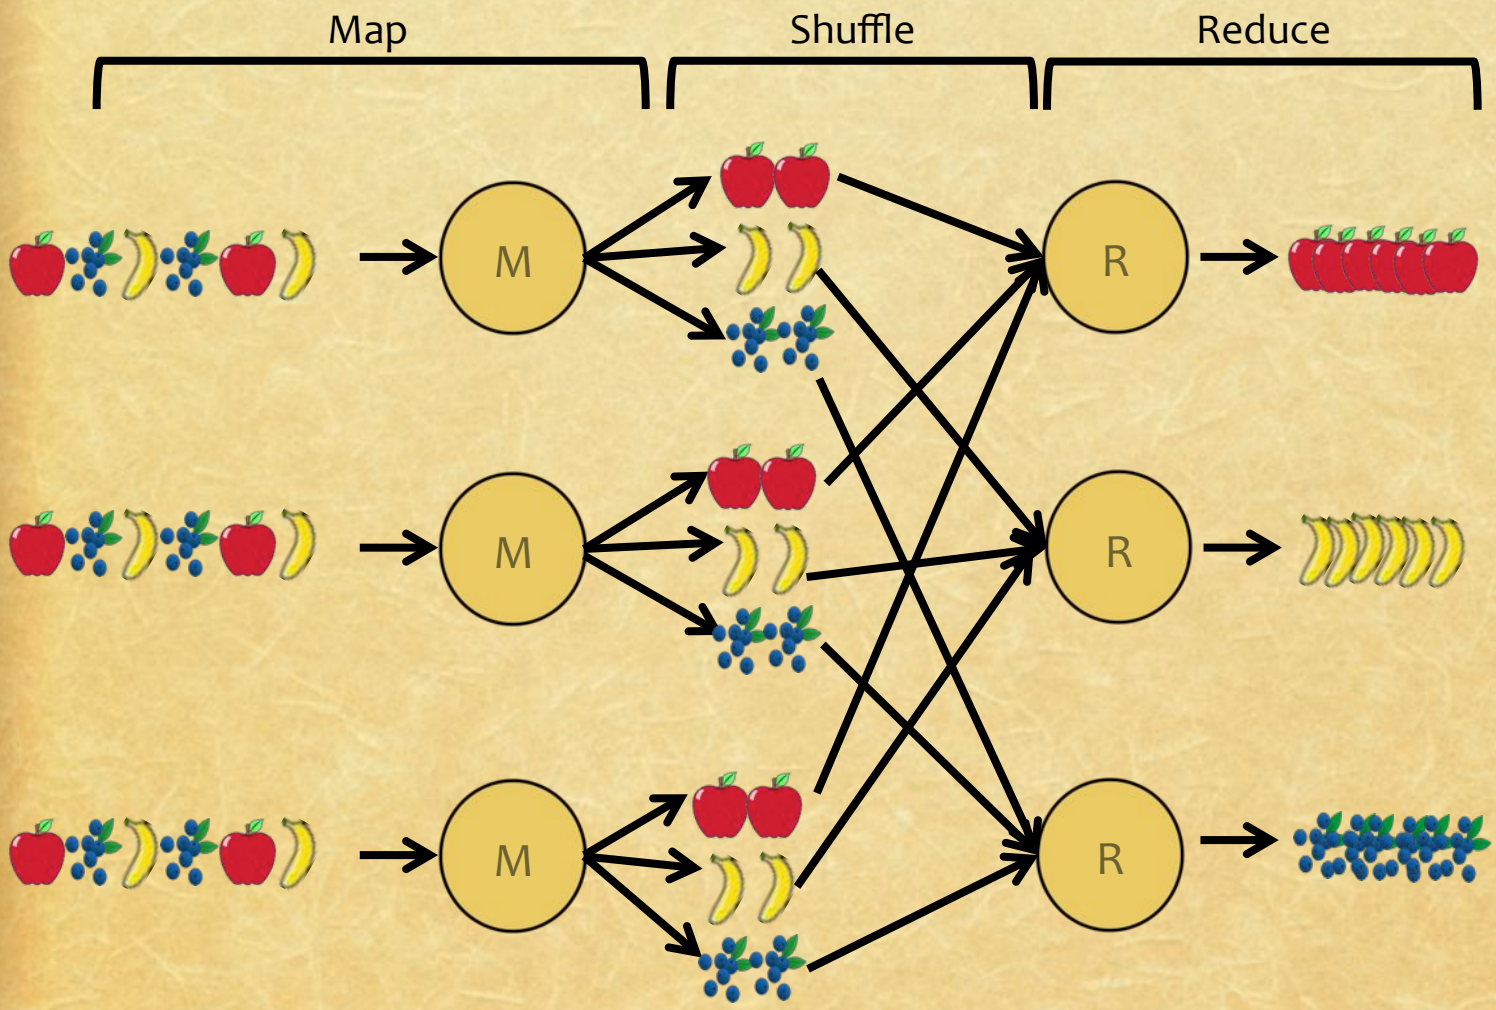

Hadoop Simulator

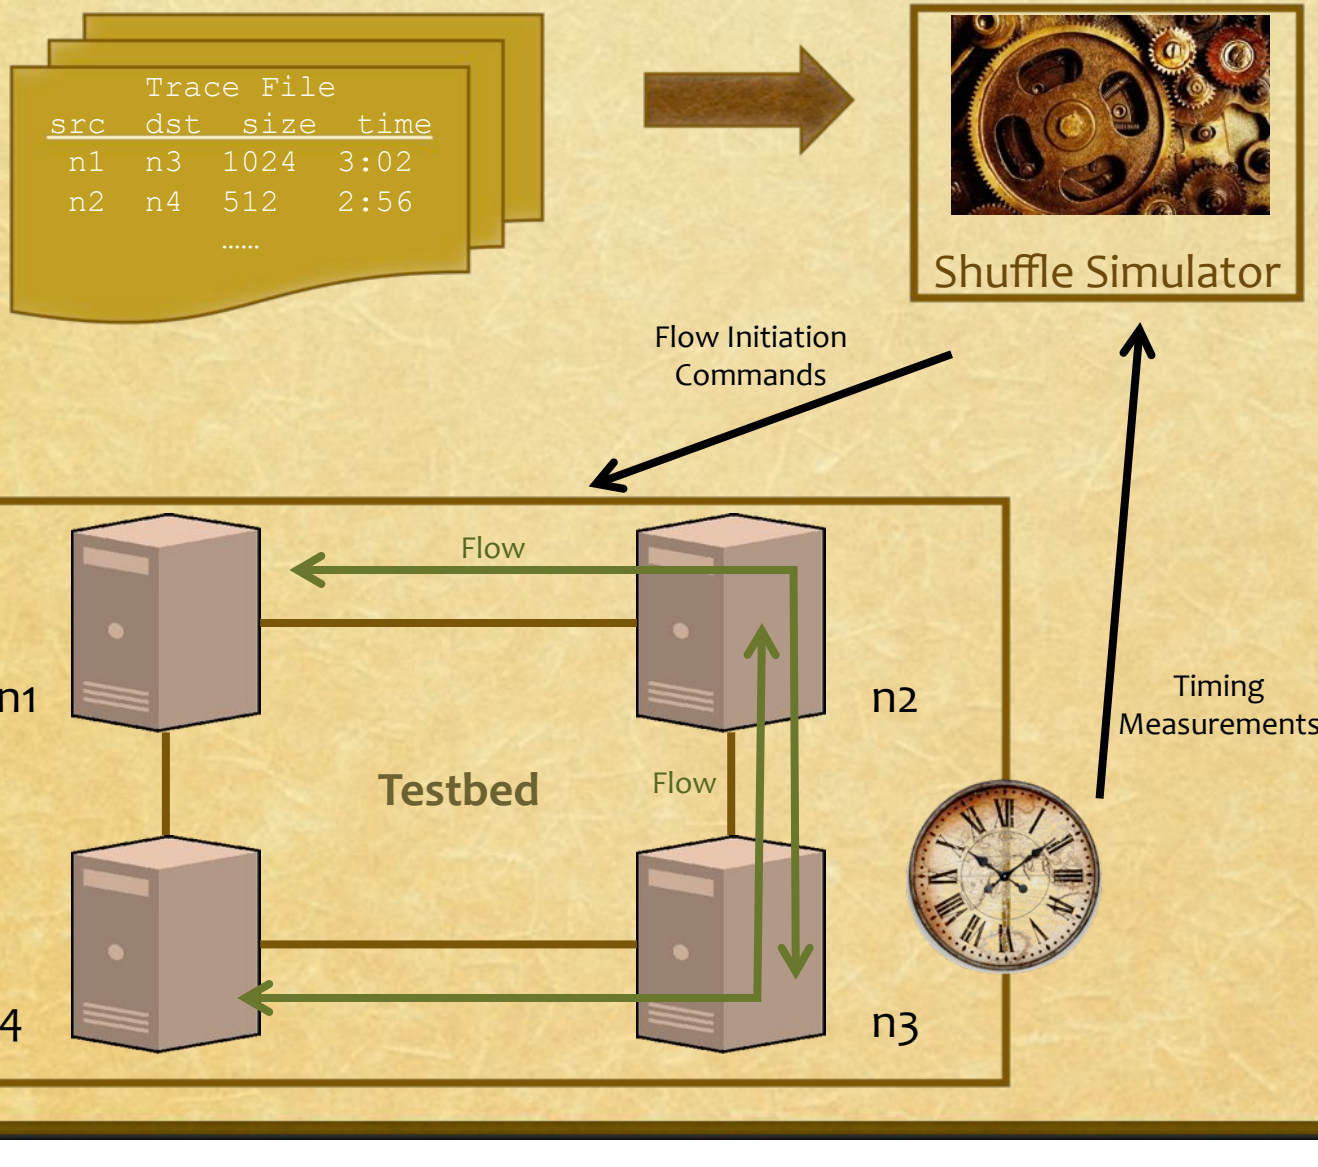

Experiment on Random Torus

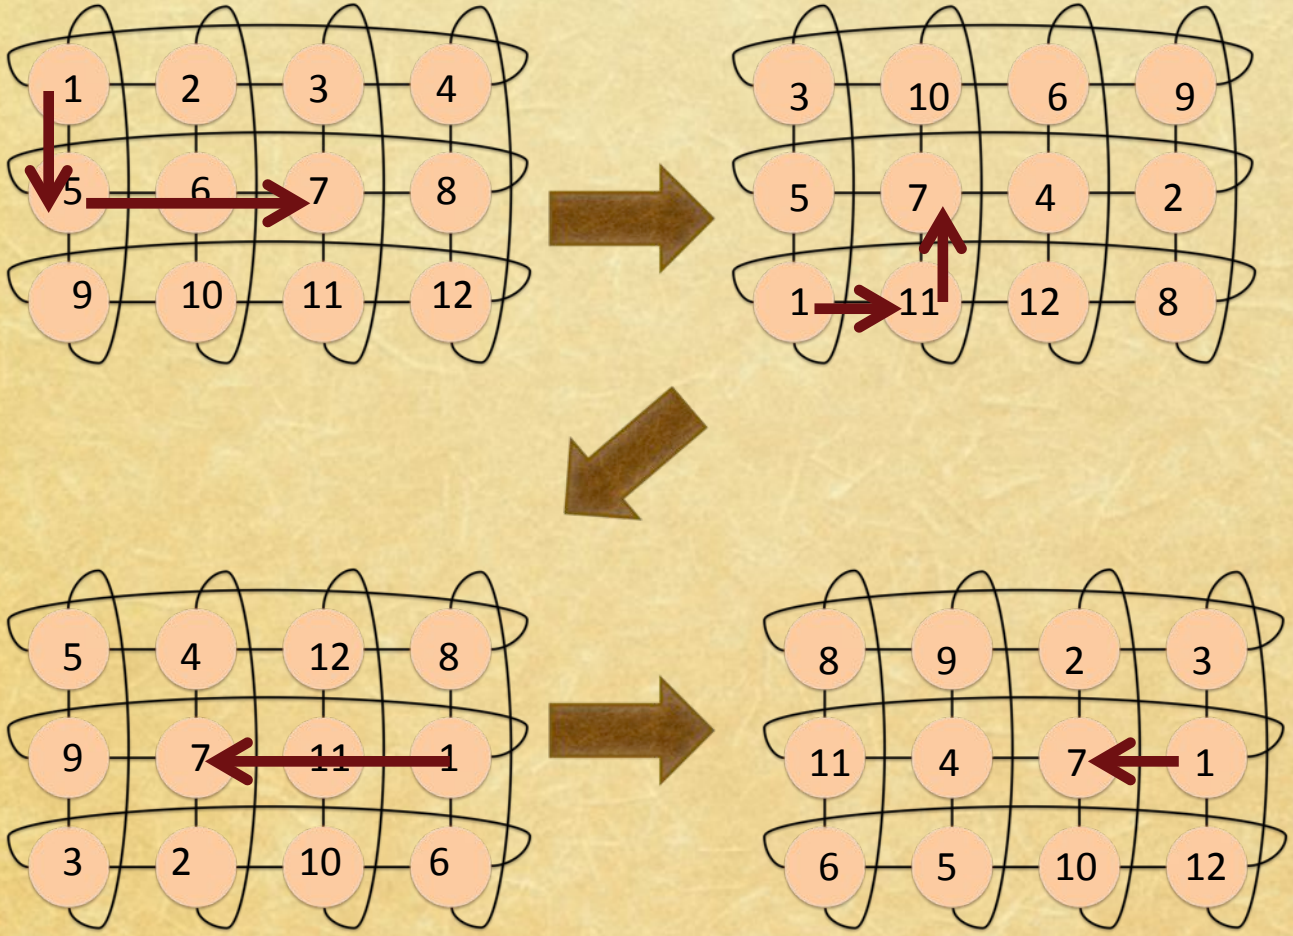

Some topologies place source and destination racks closer in the switching overlay

Shuffle Traffic Scenario

Experiment Design

- 3 simultaneous jobs
- 1 GB data transferred from each map to single reducer
- All flows concurrent
- 1000 runs each under a different placement of nodes in topology; record times, throughput and LSS for each

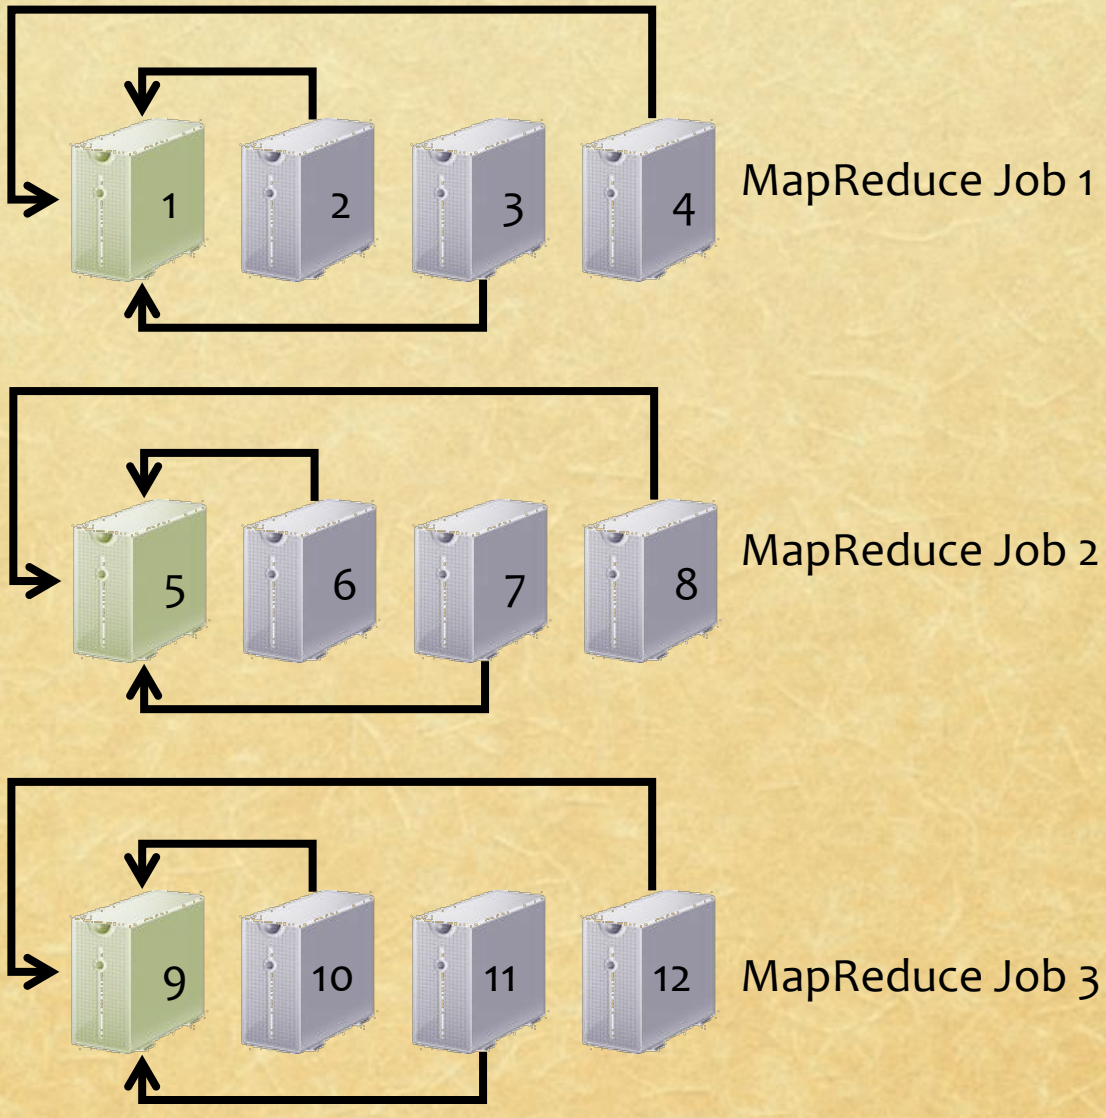

Mapper
Reducer

Shuffle Simulation Results

Node Placement	Trials	$\bar{x}$	s
same	500	620.7 Mb/s	2.27 Mb/s
random	1000	563.6 Mb/s	79.78 Mb/s

Throughput Mean and Standard Deviation (bigger is better)

Histogram of Shuffle Times (small is better)

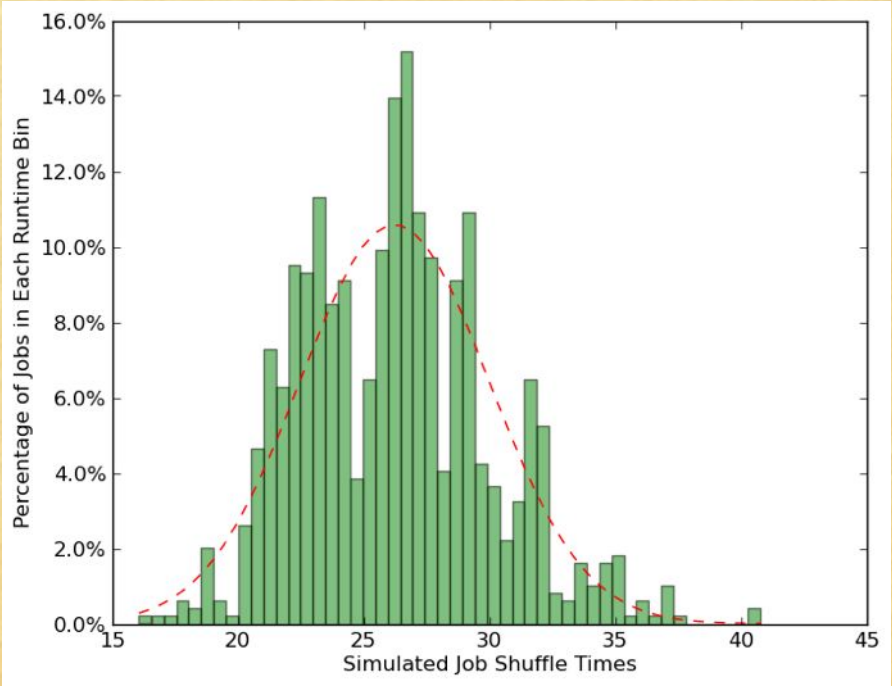

Throughput and LSS

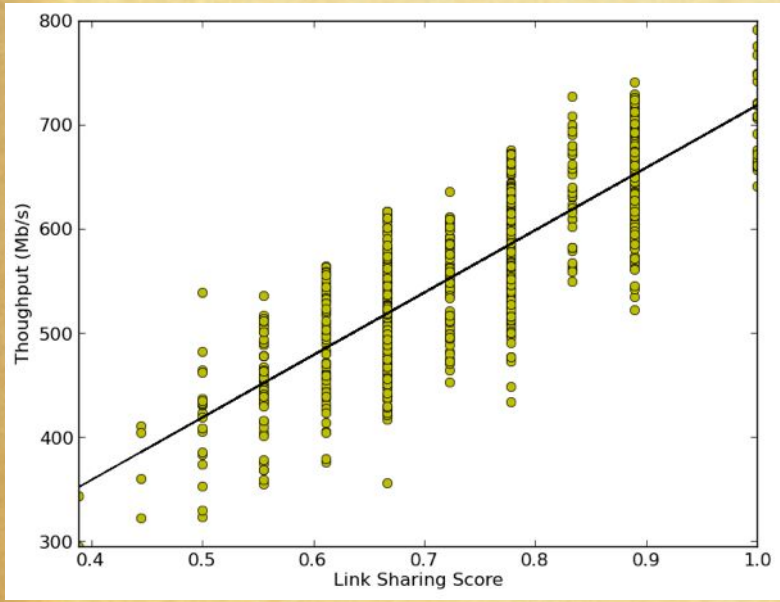

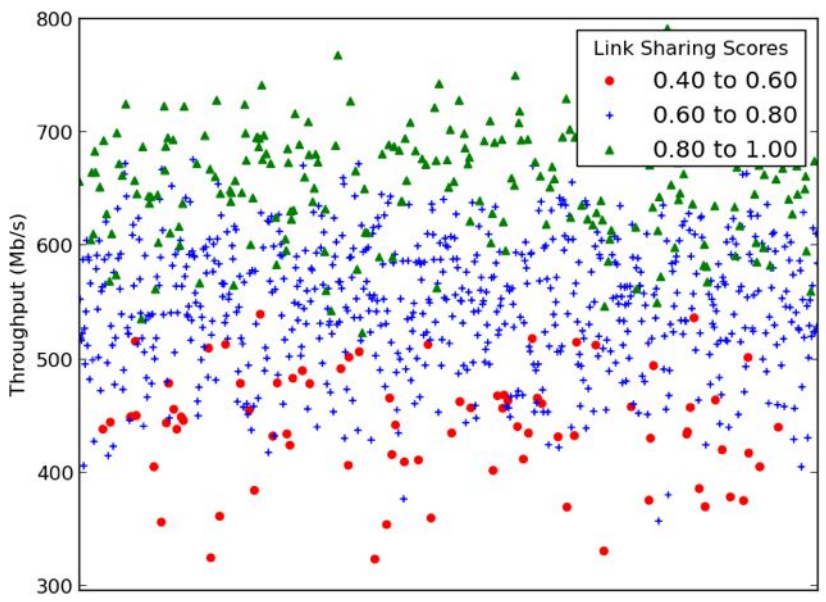

LSS correlates well with Throughput as is expected

Shuffle Time and LSS

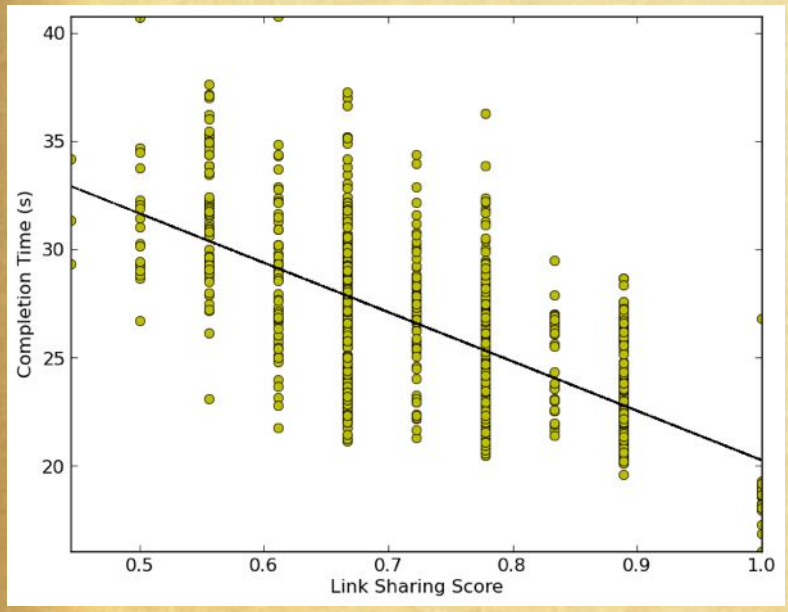

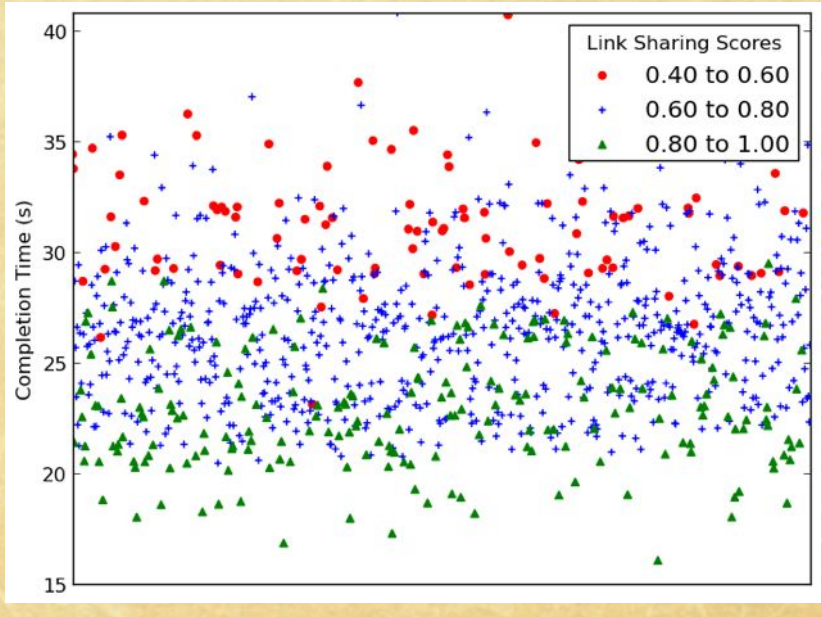

LSS correlates well with Shuffle time as is expected.

The difference between optimal and suboptimal configurations can have significant effect on the overall time taken

Conclusion

- We have presented the Flow Optimized Route Configuration Engine (FORCE), a datacenter testbed emulator with a programmable interconnection controlled by an SDN controller. The FORCE allows researchers to get an early indication of the worthiness of data center topology hypothesis.
- These experimental results come without the cost in time or funding of building production level data centers.
- Additionally, the system features a Virtual Topology Engine, a Flow Network Evaluation System, and a Hadoop shuffle traffic simulator.
- We have presented initial experimental results to suggest that datacenter topology, specifically placement within a 4x3 2-D torus network, can impact the time to shuffle intermediate results from a MapReduce job.
- In the future, we plan to build a complete Hadoop traffic simulator, upgrade the emulated rack workstations, and develop a system that will provide execution time adaptability and maneuverability of datacenter topology to steer away from worst case scenarios. We also plan to deploy and validate our hypotheses in production data centers with SDN capabilities.