SCHOOL of ELECTRONIC ENGINEERING and COMPUTER SCIENCE

Keddah: Capture Hadoop Networking Behaviour

Presenter: Jie Deng

Supervisor: Steve Uhlig, Felix Cuadrado, Gareth Tyson

Queen Mary
University of London

BigData


BIG
DATA
INFORMATION
PETABYTES
STORAGE
LARGE
SEARCH
ANALYSIS
MANAGEMENT
TECHNOLOGIES
SOFTWARE
SETS
MILLION
VOLUME
GOVERNMENT
PROCESS
DEVELOPMENT
TECHNOLOGY
RESEARCH
INTERNET
SCIENCE
FUTURE
NEW
BILLION
SAN
NAS
SOCIAL
YEARS
RESULTS
DRIVE
REPLICATION
HANDLE
SOURCES
HAND
CAPACITY
CHALLENGES
LIMITS
COMPLEX
ANNOUNCED
PAST
INCLUDE
NETWORKS
HYPOTHESIS
END
MOST
NEARLY
MAPREDUCE
CRITIQUES
PARALLEL
STATE
CRITIQUE
PEOPLE
TRAFFIC
USE
CURRENTLY
MANY
APPROACH
OUT
ONE
FACTOR
SYSTEMS
MANAGE
ONLINE
DAY
WORLDWIDE
MAKING
INTEGRATION
MASSIVE
DEFINITION
USING
FRAMEWORK
BETWEEN
STATISTICS
BUSINESS
EFFECTIVE
VISUALIZATION
DATABASES
SIZE
ANALYTICS
PROCESSING
SIMULATIONS
EXABYTES
DETERMINE
WELL
SET
DEPARTMENT
HIGHER
ALGORITHMS
COMPANIES
QUERIES
HIGH
USERS
AMOUNT
REAL
RATE
STRUCTURE
WORK
APPLICATIONS
INFRASTRUCTURE
TERMS
YEAR
PRIVATE
CAPABILITIES
SUCCESSFUL
NEEDED
USES
LEARNING
LESS
HUGE
EVERY
WORLD
VARIOUS
COMPARED
BASED
TIMES
SHARED
ARCHITECTURE
PRIVACY
INTERNATIONAL
IMPLEMENTATION
BECOMES
SECOND
NATIONAL
INTELLIGENCE
APPROACHES
PROCESSES
FLOW
INITIATIVE
ECONOMIC
PROJECT
DISTRIBUTED
PER
RECORDS
COST
PARADIGM
FUNDING
INSIGHT
TOOLS
CENTER
TERABYTES
MARKET
VARIETY
RELEVANT
ANNUAL
RELATED
TIME
SENSOR
WOULD
NOW
USED

Distributed systems

- Network is essential
- But not well studied yet

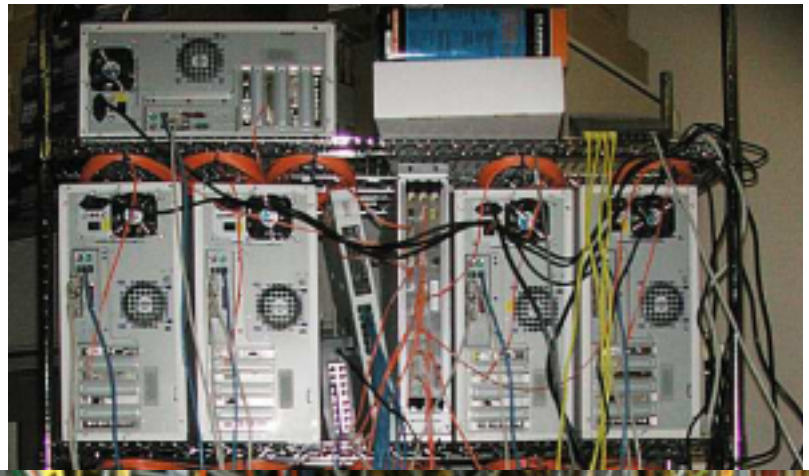

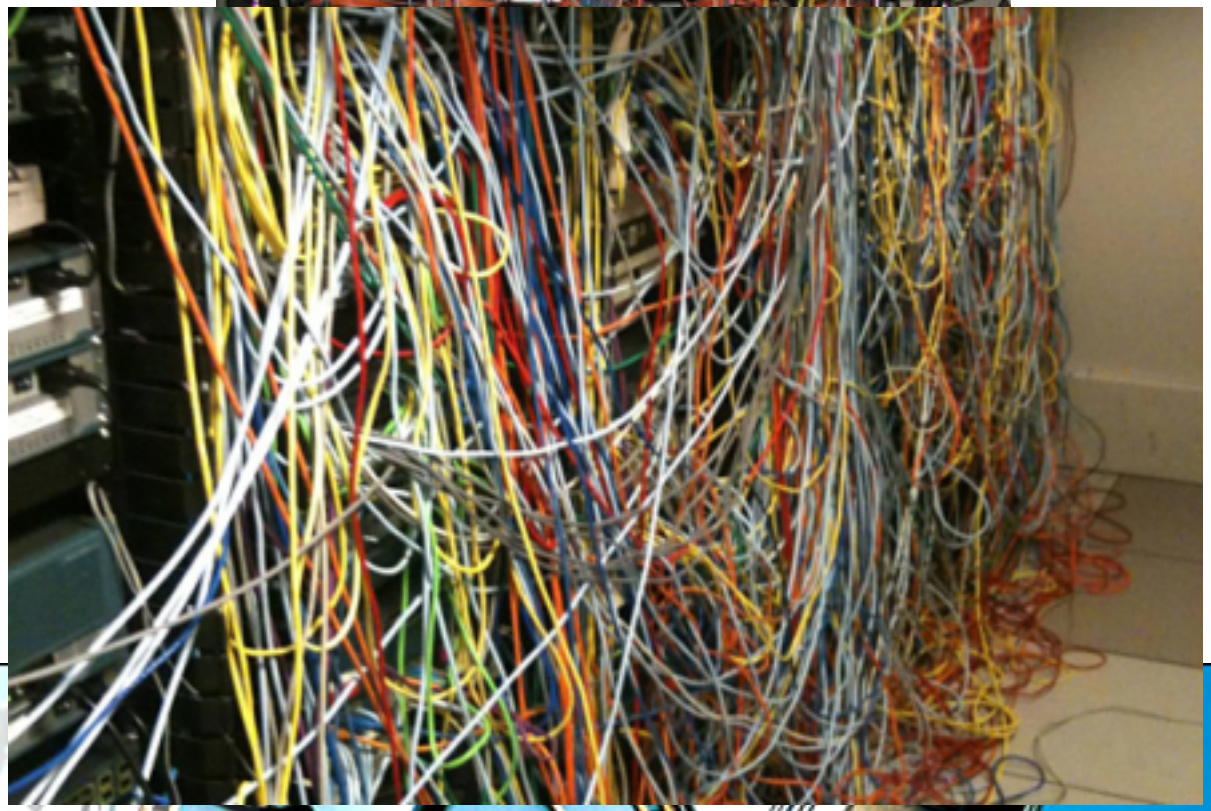

Networks are complicated and costly...

- Cost consuming infrastructures
- Time consuming cabling

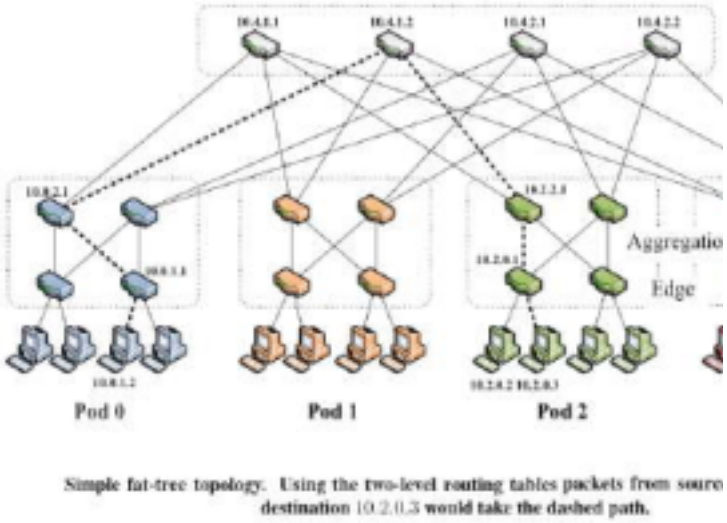

Simple fat-tree topology. Using the two-level routing tables packets from source destination 10.2.0.3 would take the dashed path.

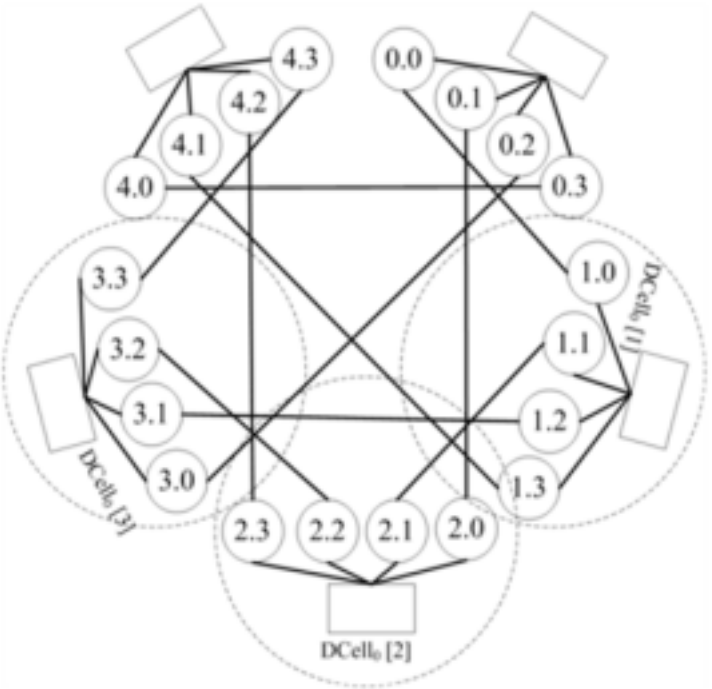

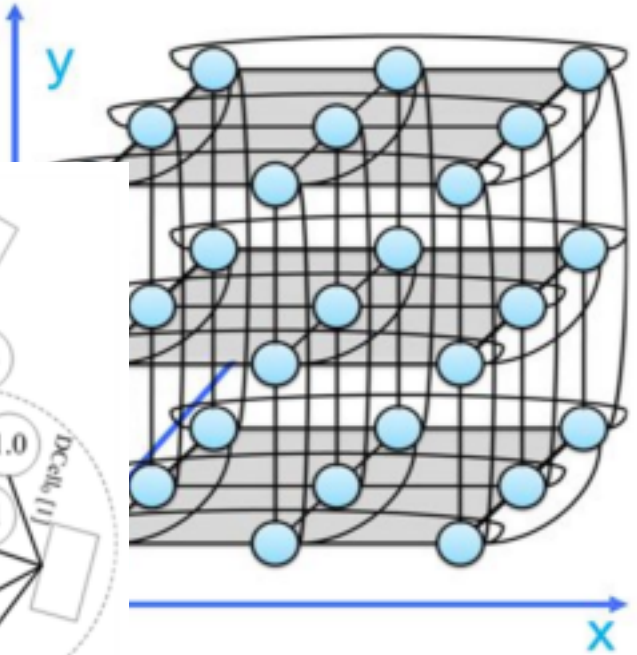

What we want... Traffic generator!

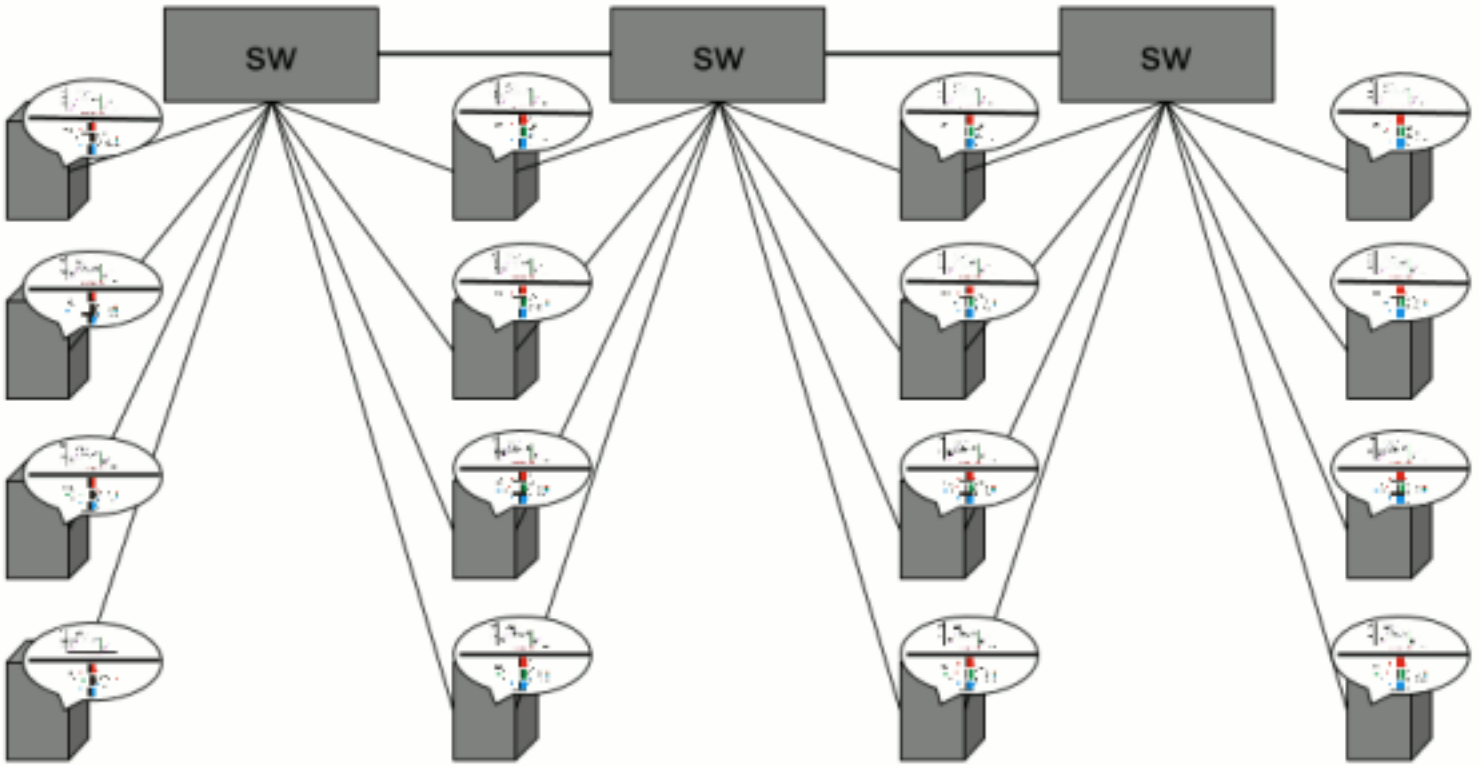

Methodology

Testbed:

- 16 nodes w/ Hadoop & sflow installed

Reproduce Hadoop traffic:

- General over jobs
- Extendable over parameters
- Replicable over clusters

Hadoop

- Storage: HDFS
- Processing: MapReduce jobs

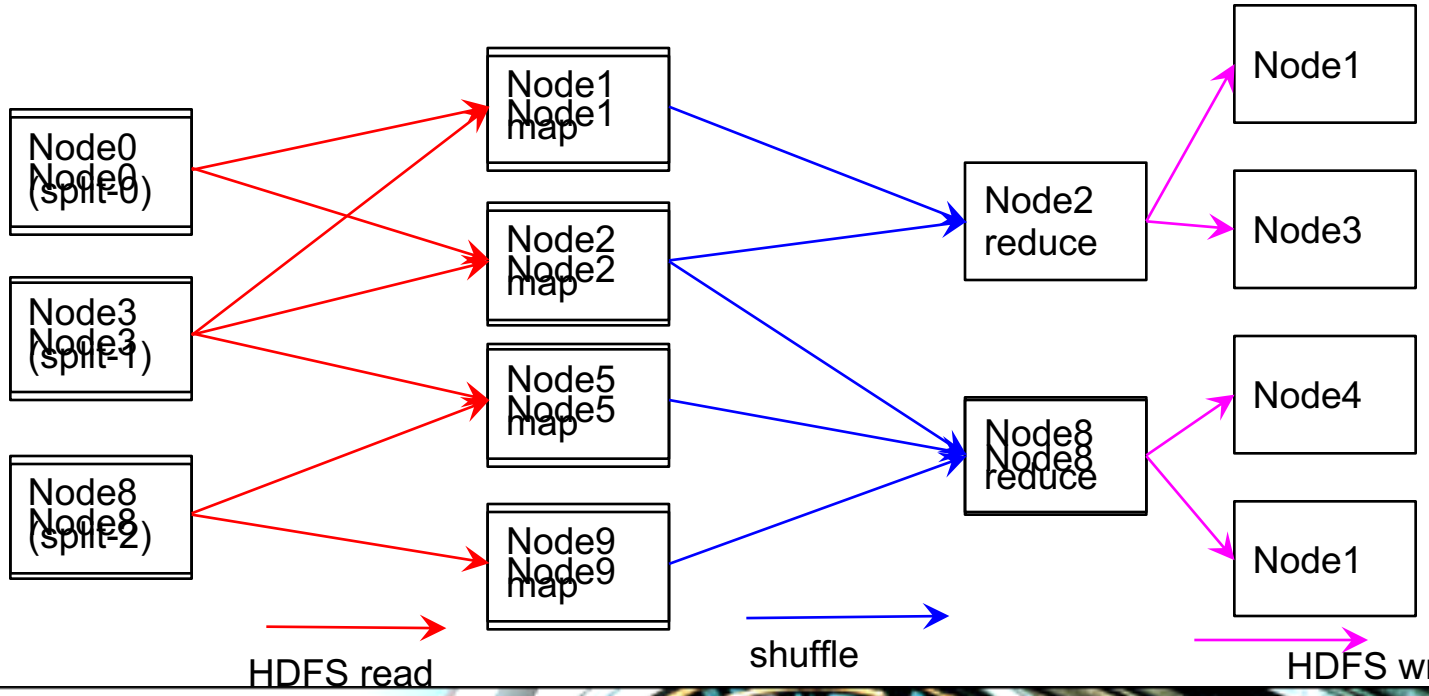

Capture the randomness...

- Nodes are randomly selected in Hadoop, thus the distribution.

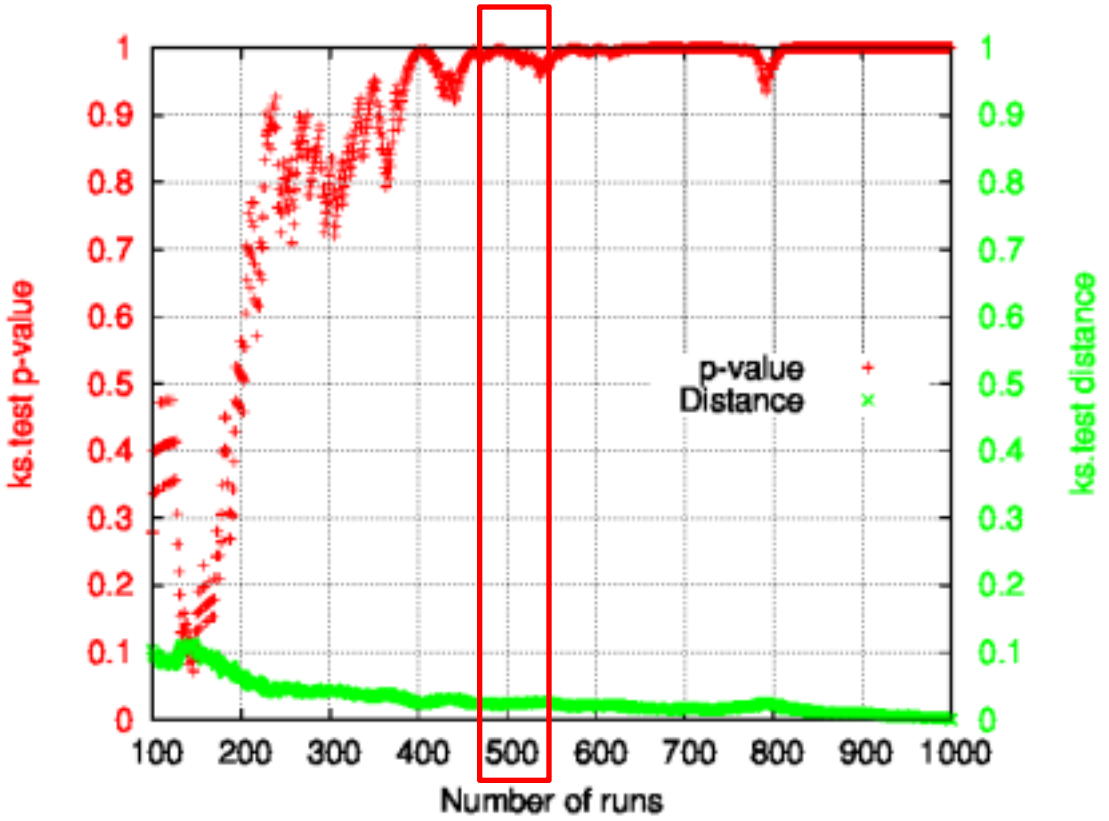

Traffic Model

HDFS import
->Shuffle
->HDFS export

Job	HDFS import	Shuffle	HDFS export
TeraSort	0.1	0.9	0
PageRank	0.25	0.66	0.09
kmeans	0.31	0	0.67

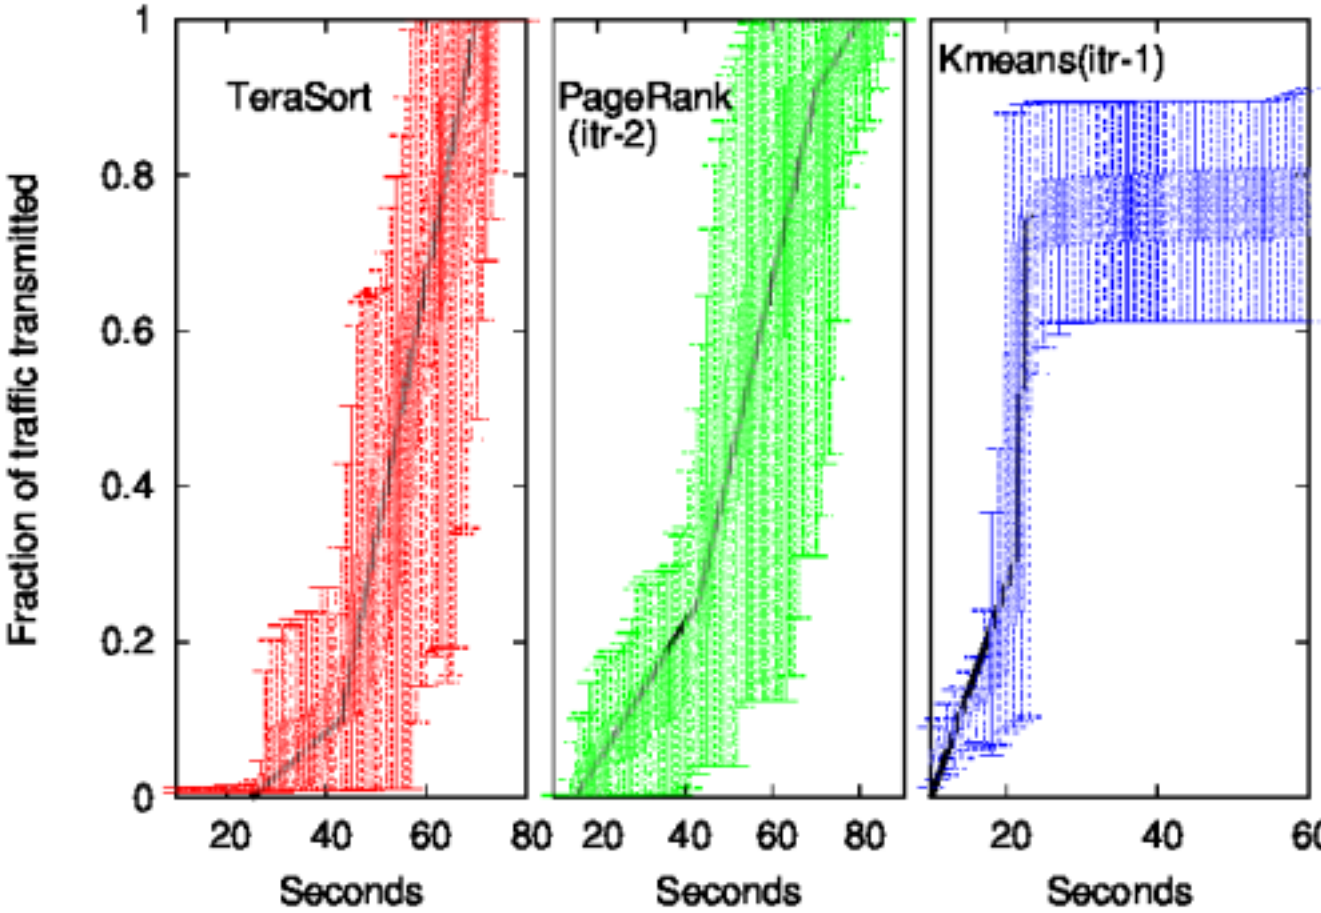

Modelling

- Number of nodes
- Number of flows
- Flow size
- Flow start time

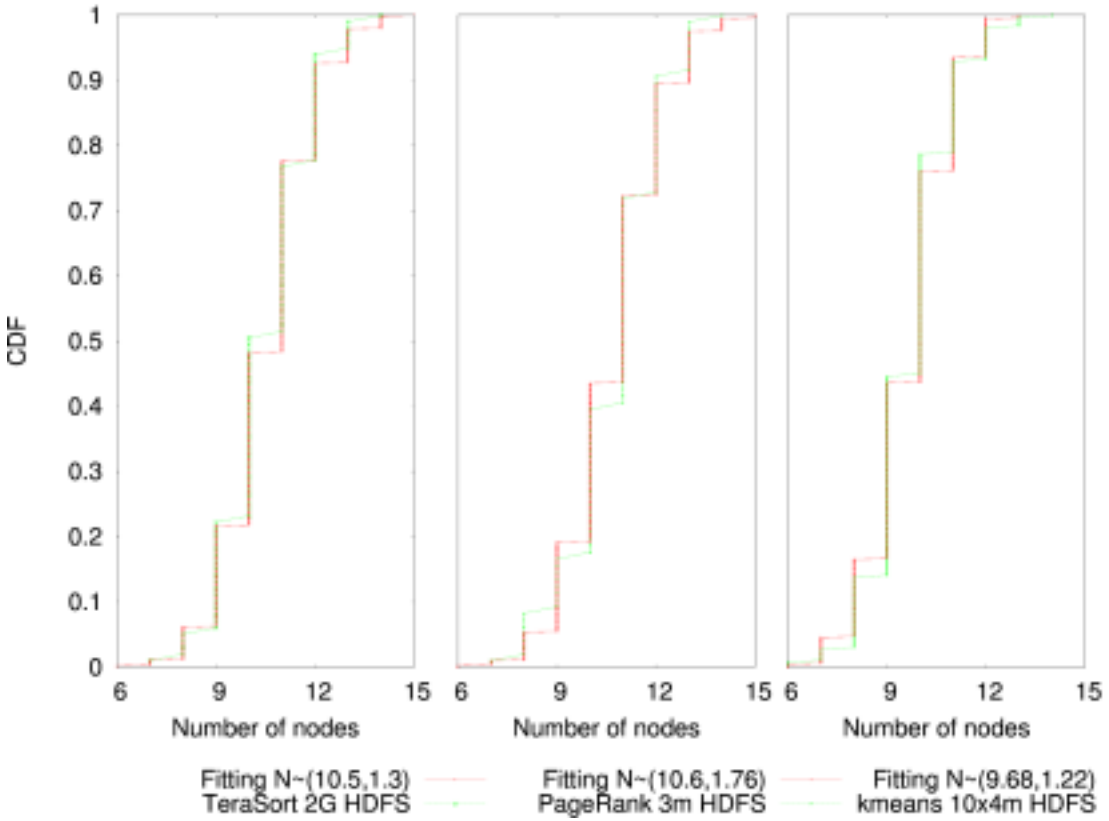

Keddah in ns3

- HDFS imp

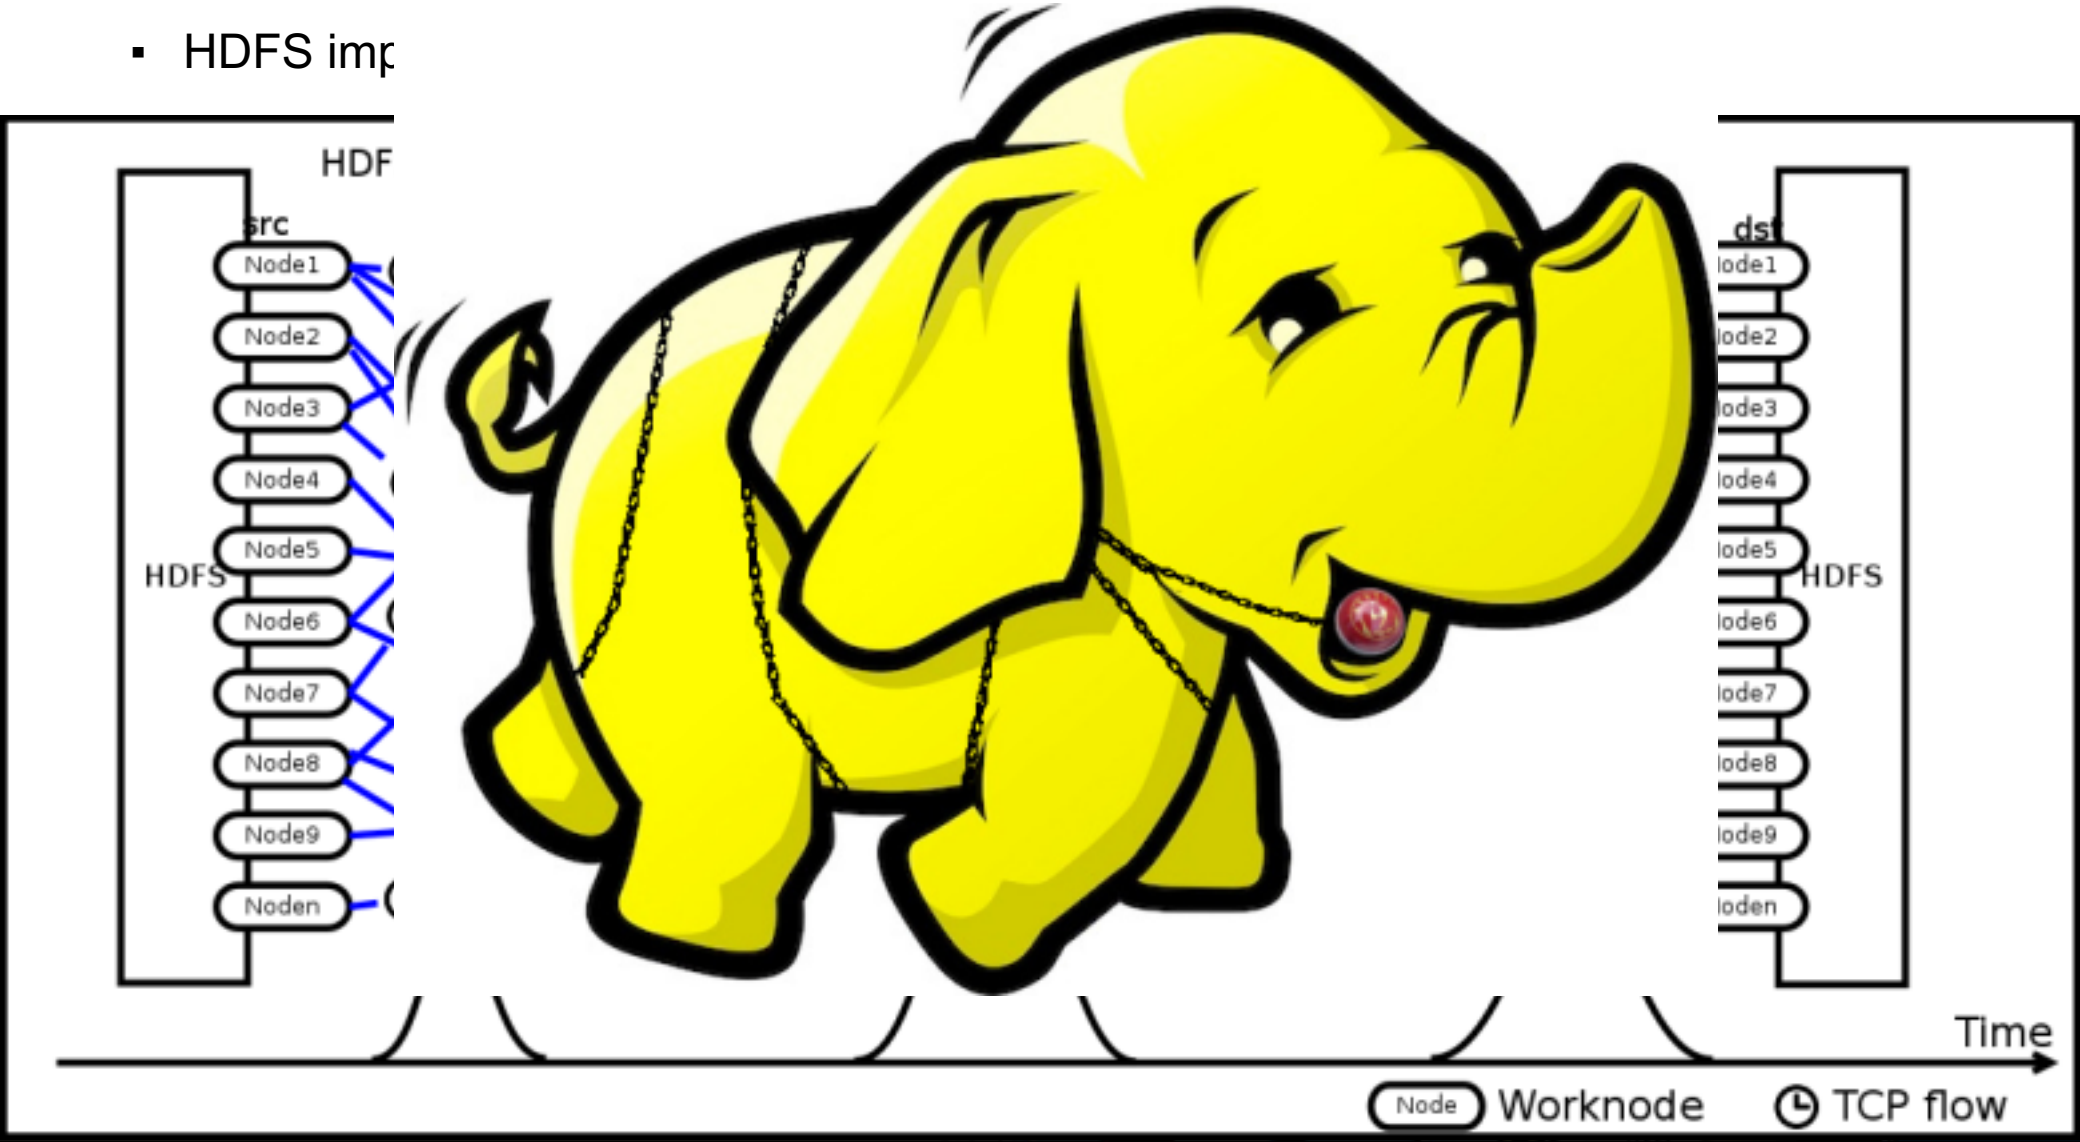

Usefulness validation

- DCN topologies: FatTree(8x8) CamCube(4) Dcell(2,3)

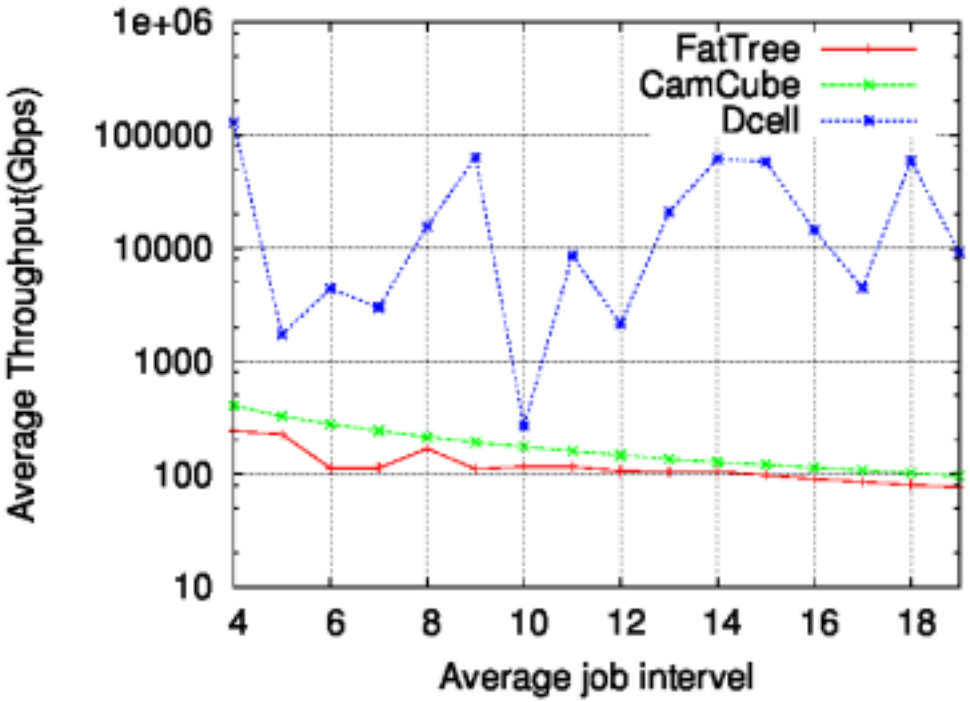

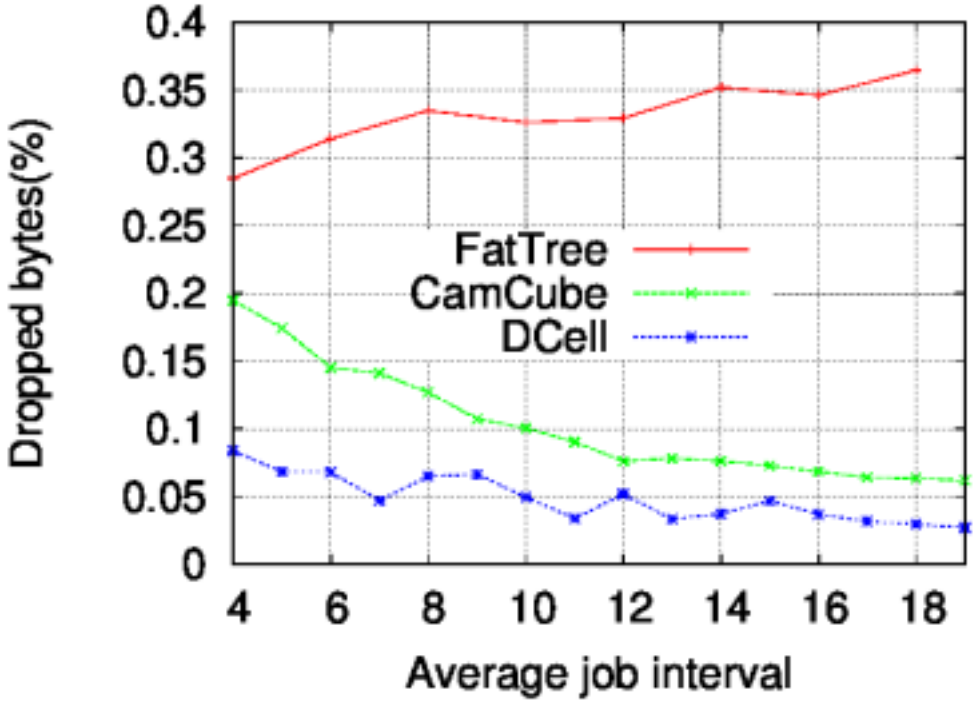

- GitHub: https://git.io/vKeIv

- Current work
  - Produce fine-grained workload traffic
  - Network(bandwidth) as a resource

- Thanks!
- Q&A

Resistance from systems

- Time consuming software engineering

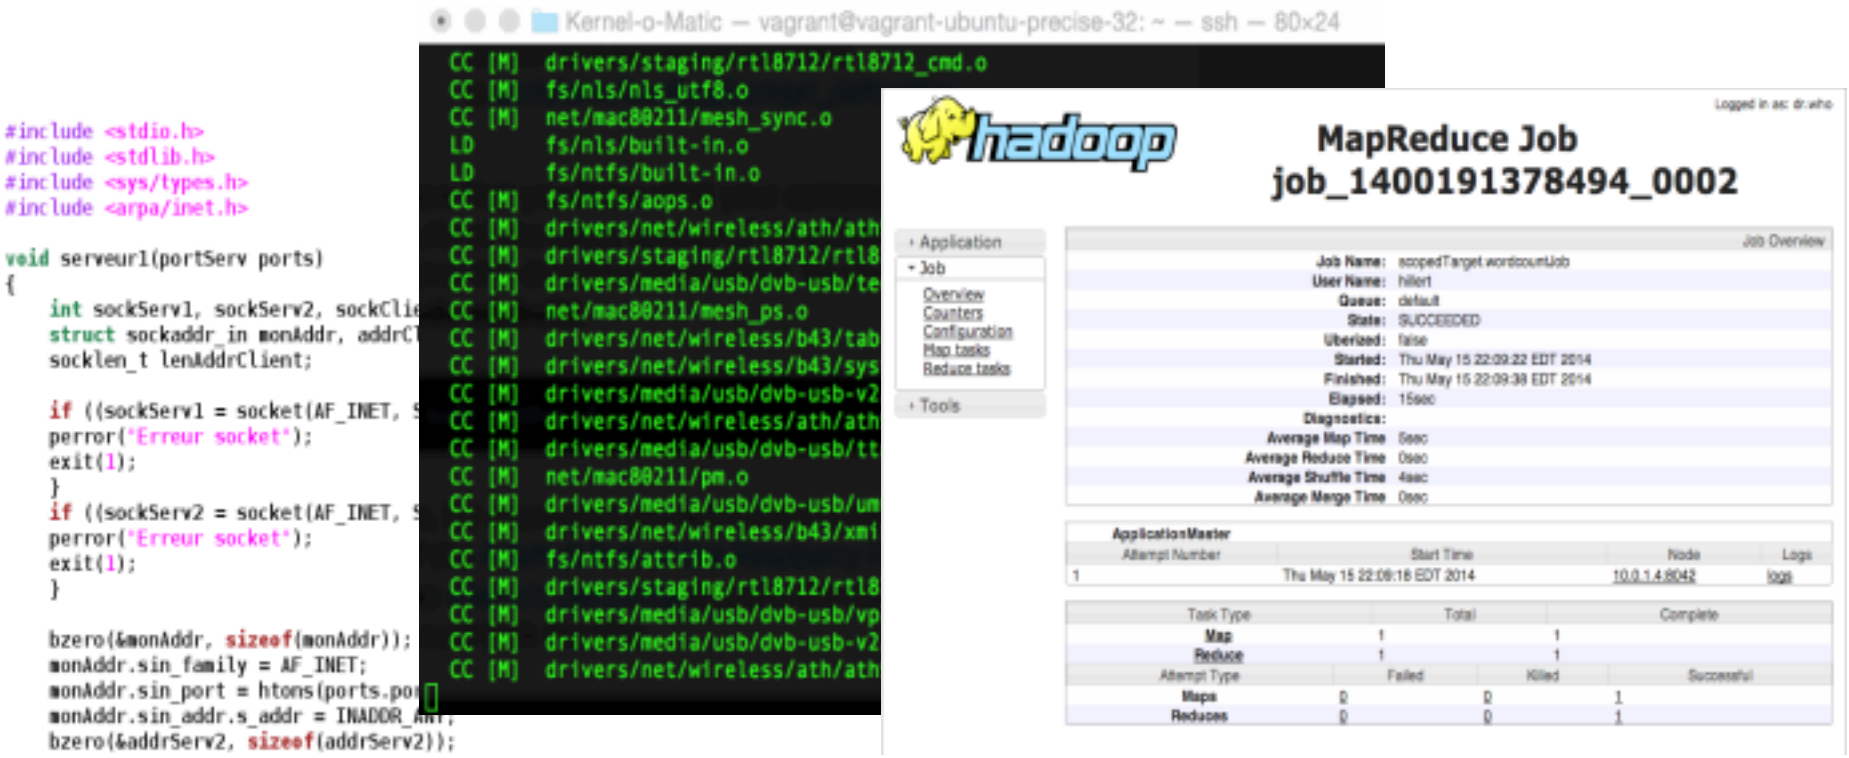

Extend the parameters:

- Job parameters
- Schedular

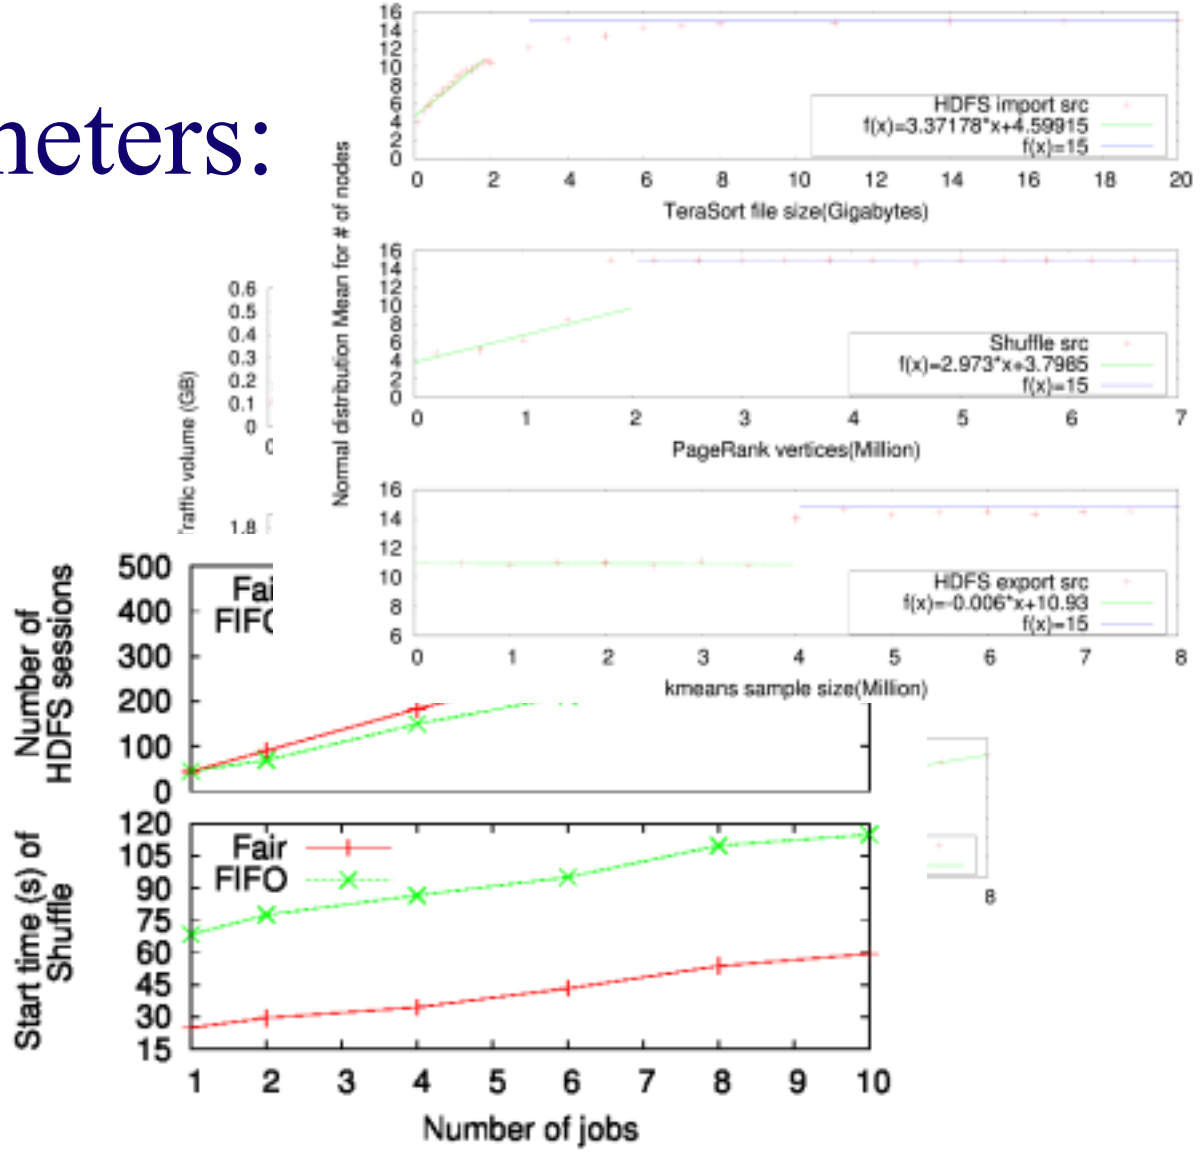

Use cases2

- TCP vs DCTCP