
Theia: Visual Signatures for Problem Diagnosis in Large Hadoop Clusters

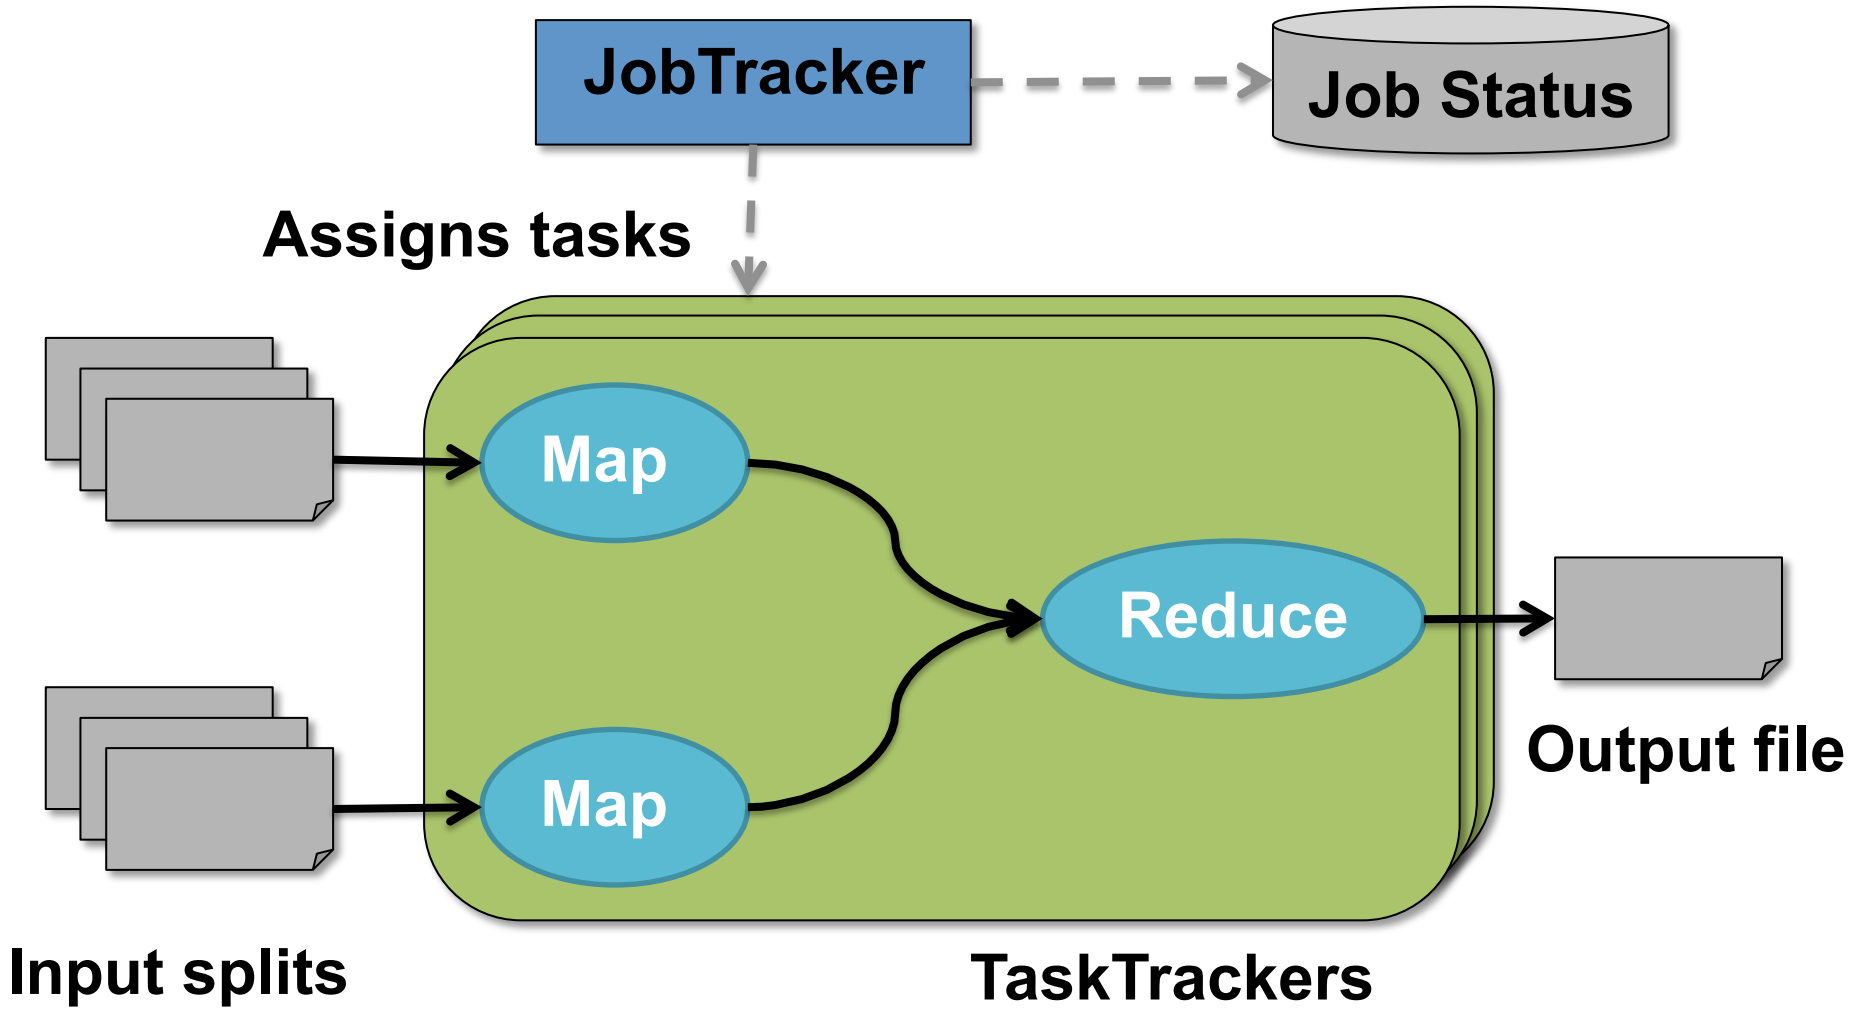
Elmer Garduno, **Soila P. Kavulya**, Jiaqi Tan,
Rajeev Gandhi, Priya Narasimhan

Motivation

- Hadoop: Open-source MapReduce framework
 - MapReduce is popular for data-intensive computing
 - 100+ companies powered by Hadoop
 - Yahoo, Facebook, Hulu, ...
- Problem diagnosis in Hadoop is challenging
 - Overwhelming volume of monitoring data
 - Complex component interactions obscure root-cause

Goal: Generate visualizations that aid diagnosis

Background: Hadoop MapReduce




OpenCloud Cluster

- OpenCloud: Cluster for data-intensive research
 - 64-nodes worker nodes
 - Each node has 8 cores, 16 GB DRAM, 4 1TB disks and 10GbE connectivity
- Users are researchers familiar with Hadoop
 - Run diverse kinds of real workloads
 - Graph mining, natural language processing, ...
- Tools available for diagnosis
 - Hadoop web console and Ganglia
 - Sysadmins also use Nagios to receive alerts

Shortcomings of Current Tools

Ganglia



Large number of monitoring panels can overwhelm user

Hadoop Web Console

Kind	% Complete	Failed/Killed Task Attempts
map	65.00% <div style="width: 65%;"><div style="background-color: blue; height: 10px;"></div></div>	0 / 8
reduce	0.00% <div style="width: 0%;"><div style="background-color: blue; height: 10px;"></div></div>	0 / 0

Large number of tasks/jobs --- need to click through multiple screens to troubleshoot problem

What Do Users Want?

Users want to distinguish between*:

1. Problems due to their job

- Software bugs, job misconfigurations
- Legitimately slow tasks due to data skew
- *Action:* Fix themselves

2. Problems due to infrastructure

- Resource contention
- Hardware failures
- *Action:* Report to sysadmins

**Results from user study published in CHIMIT 2011*

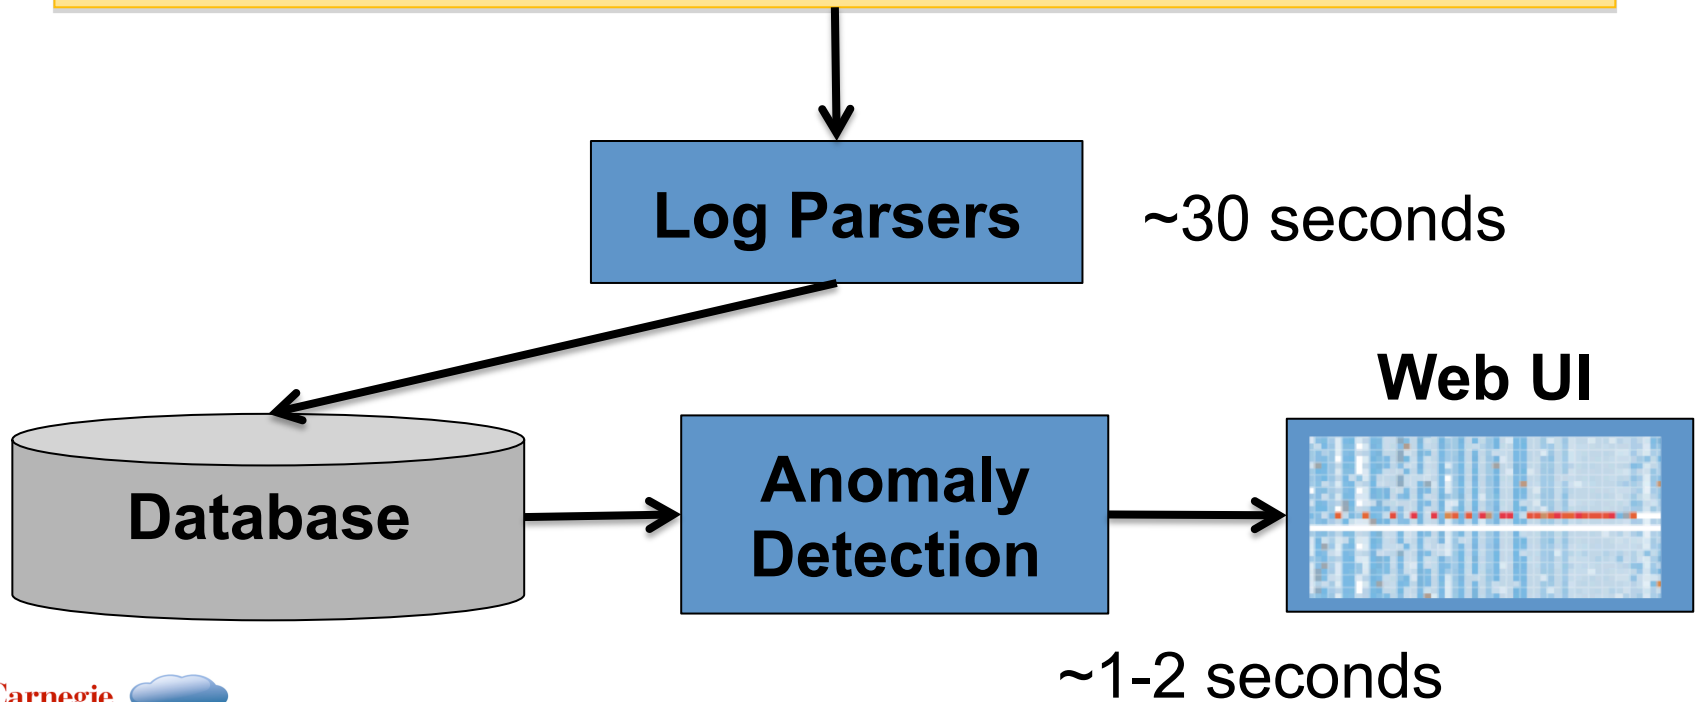
Theia: Visualization Tool

- Visualizes anomalies in Hadoop clusters
 - Maps anomalies observed to broad problem classes
 - i.e., hardware failures, application issue, data skew
- Interactive interface supports data exploration
 - Users drill-down from cluster- to job-level displays
 - Hovering over the visualization gives more context
- Compact representation for scalability
 - Can support clusters with 100s of nodes

Theia Implementation

Job History Logs (100-300MB/day)

```
TASK_TYPE="REDUCE" TASKID="X"  
TASK_STATUS="SUCCESS" HOSTNAME=node32  
COUNTERS= [(FILE_BYTES_READ)(229524)] ....
```



Quantifying Anomalies

- Hadoop balances load across nodes in cluster
 - Typically, tasks in a job perform similarly
 - Significant deviations may signal problems
- Indicators of anomalous behavior used by Theia
 - Large task duration
 - Large deviations in data processed (data skew)
 - Large task failure rates
- Quantify anomalies using z-score
 - Deviation of each observation from the mean
 - Normalized by standard deviation

$$z = \frac{x - \mu}{\sigma}$$

Heuristics for Visual Signatures

Mapping anomalies to problem classes

	Application Problem	Infrastructural Problem
Time	<i>Single user/job</i> over time	<i>Multiple users/jobs</i> over time
Space	Span multiple nodes	Affects single node , but correlated failures also occur
Value	Slow/failed tasks, or data skew	Slow/failed tasks

Types of Visualizations in Theia

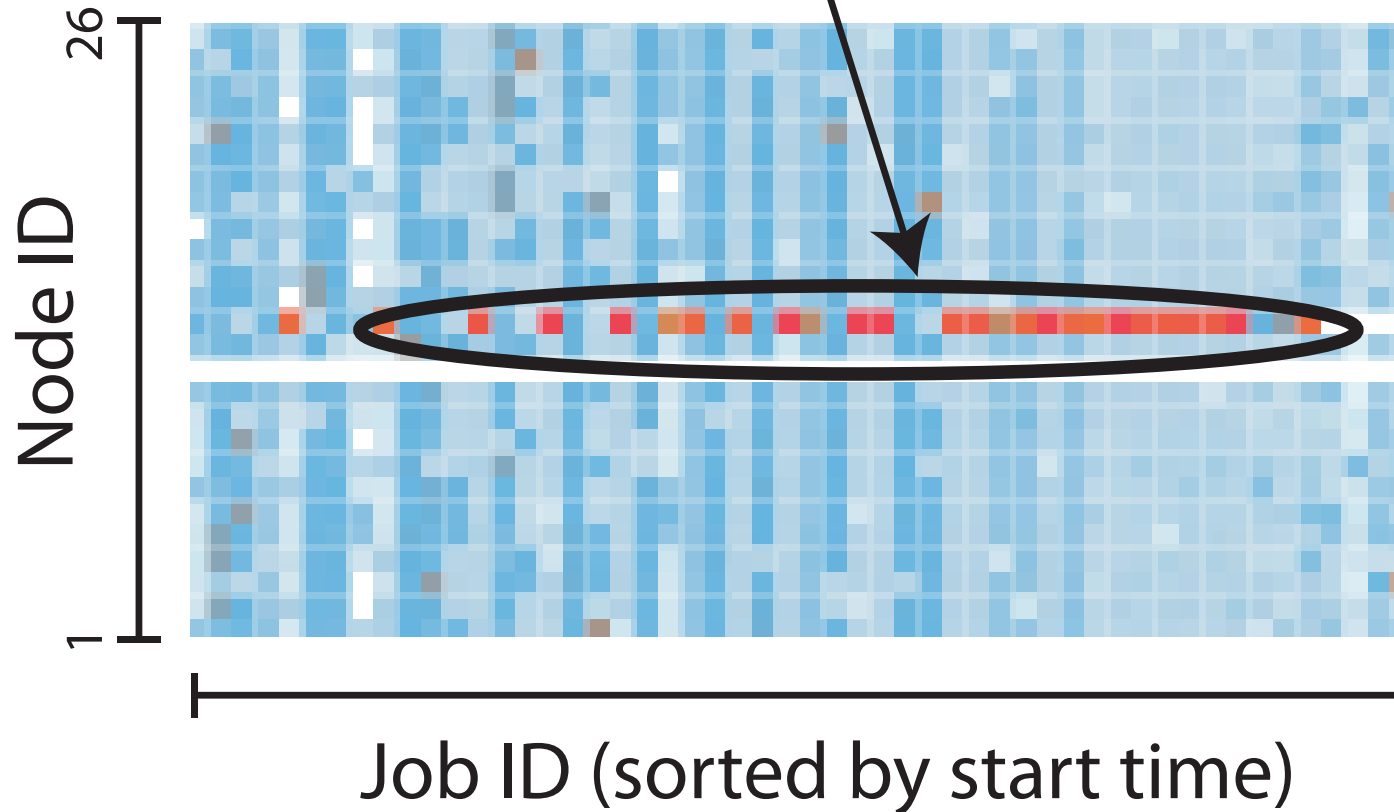
- Anomaly heatmap
 - Visualizes cluster-level anomalies
 - High density overview of nodes in the cluster
- Job execution stream
 - Visualizes job-level anomalies
 - Compact view of aggregate job performance
- Job execution detail
 - Visualizes job-level anomalies
 - Detailed view of job progress over time

Anomaly Heatmap

- High density overview of cluster performance
 - Summarizes job performance across nodes
 - Uses color variations to visualize anomalies
- Compact representation supports scale
 - Uses 2x2 pixels to represent job/node
 - Fits 1200 jobs x 700 nodes on a 27-inch display
- Best-suited for visualizing
 - Infrastructural problems

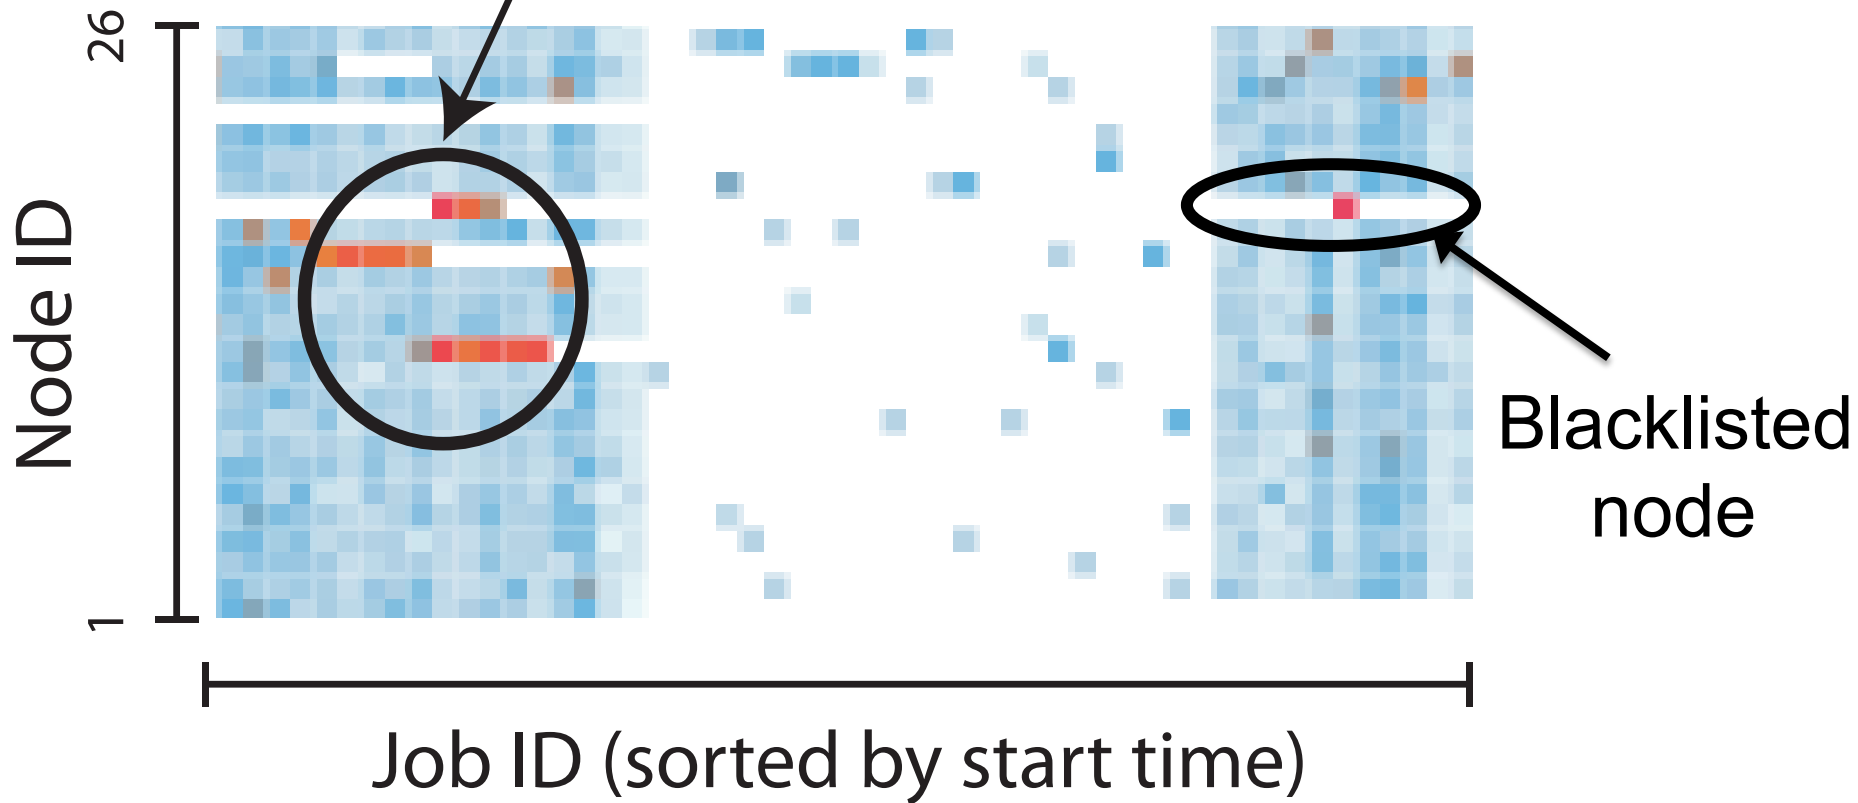
Heatmap of Hardware Problem

Performance degradation due to failing disk



Heatmap of Application Problem

Resource contention due to buggy job




Job Execution Stream

- Visualizes per-job performance across nodes
 - Scrollable stream of jobs sorted by start time
 - Displays performance of Map and Reduce phases
 - Uses color variations to visualize anomalies
- Shows job execution trace in context
 - Job name, duration, and status
 - Task failure rates
 - Task duration anomalies
- Best-suited for detecting application problems
 - Software bugs, data skew

Job Execution Stream

Job_9438/HALFCAMS_kmeans_iteration_0 – **SUCCESS**
21:54:47 / 0h 5m 44s / Oct 10



Darker colored squares indicate anomalous task durations


Red borders indicate failed tasks

Yellow borders indicate speculative tasks were killed

Visualizing Application Problems

job_9035 / depParseFeatExtract -- FAILED

08:57:16 / 0h 6m 41s / Oct-06



Bug in reduce phase indicated by failed tasks

job_9146 / DocPerLineSentenceExtractor -- SUCCESS

20:25:39 / 0h 17m 36s / Oct-07



Performance degradation due to data skew

Job Execution Detail


- Detailed view of task execution
 - Less compact than job execution stream
 - Displays progress of job (using color variation)
 - Displays volume of I/O (using size variation)
- Best-suited for detecting application problems
 - Software bugs, data skew

Job Execution Detail

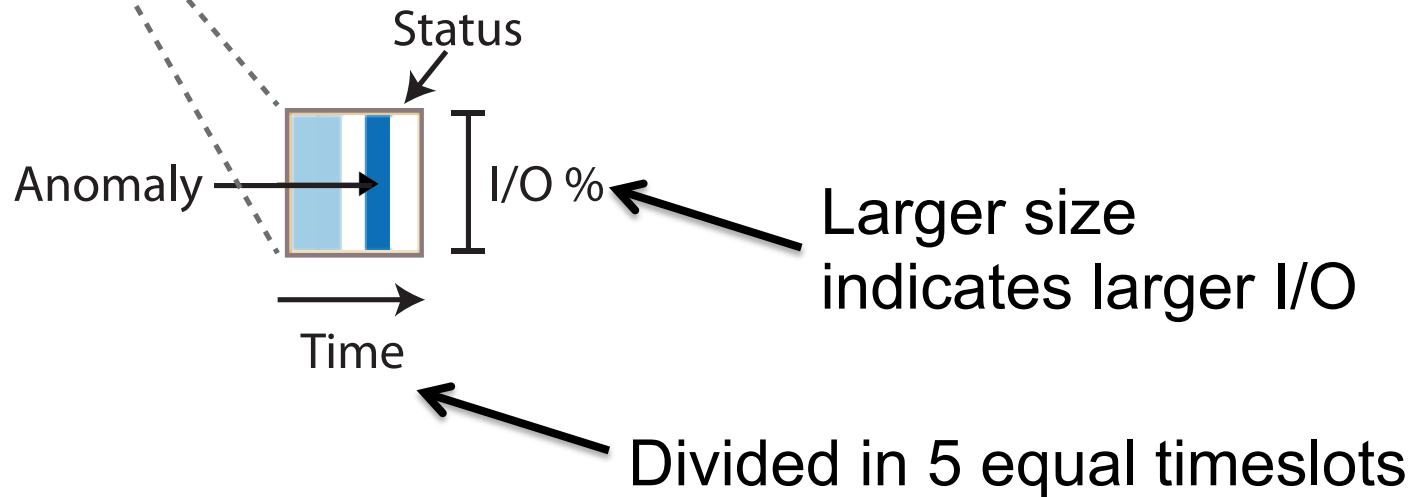

job_9036 / depParseFeatExtract -- SUCCESS

09:06:48 / 0h 1m 32s / Oct -10

Maps



Reduces



Data Skew in Job

job_9146 /

DocPerLineSentenceExtractor -- SUCCESS

20:25:39 / 0h 17m 36s / Oct -07

Maps




Imbalanced workload due to data skew

Bad NIC Degrades Performance

job_9627 /
RandomShuffle -- KILLED
06:28:29 / 0h 20m 57s / Oct -09

Maps



Node: nodeX
Successful Tasks: 6
Failed Tasks: 6
Failure Ratio: 0.5
Total I/O: 490 MB
Total Time: 0h 23m 28s

Evaluation of Tool

- Analyzed one-month of data on OpenCloud
 - 1,373 jobs comprising of ~1.83 million tasks
 - Manually identified 157 application problems and 2 data skews
- 45 hardware failures identified by sysadmins
 - 42 disk controller failures
 - 2 full hard drive failures
 - 1 network interface controller (NIC) failure
- Evaluated performance by manually verifying that visualizations matched our heuristics

Performance of Visualizations

Type	Total	Heatmap	Job Execution Stream
Application problem	157	0	157
Infrastructural problem	45	33	10
Data skew	2	2	2

Conclusion

- Theia: Visualizations for Hadoop
 - Compact, interactive visualizations of job behavior
 - Distinguishes h/w failures, s/w bugs, and data skew
 - Evaluated on real incidents in Hadoop cluster
- Next steps
 - Deploy on production cluster
 - User study
 - Evaluate the effectiveness of our UI for diagnosis

Questions?

Problem Diagnosis Website

<http://www.ece.cmu.edu/~fingerpointing/>



Related Work

- [Mochi2009] J. Tan, X. Pan, S. Kavulya, R. Gandhi, and P. Narasimhan, “**Mochi: visual log-analysis based tools for debugging Hadoop**,” Proceedings of the 2009 conference on Hot topics in cloud computing, ser. HotCloud’09, 2009
- [CHIMIT2011] J. D. Campbell, A. B. Ganesan, B. Gotow, S. P. Kavulya, J. Mulholland, P. Narasimhan, S. Ramasubramanian, M. Shuster, and J. Tan, “**Understanding and improving the diagnostic workflow of MapReduce users**,” ACM Symposium on Computer Human Interaction for Management of Information Technology (CHIMIT), 2011
- [Tufte2001] E. R. Tufte, “**The Visual Display of Quantitative Information**”, 2nd ed. Graphics Pr., May 2001
- [Bostock2011] M. Bostock, V. Ogievetsky, and J. Heer, “**D3: Data-Driven documents**,” IEEE Trans. Visualization & Comp. Graphics (Proc. InfoVis), 2011.