



X-Tracing Hadoop

Andy Konwinski, Matei Zaharia,
Randy Katz, Ion Stoica

CS division, UC Berkeley, andyk@cs.berkeley.edu



Motivation & Objectives

- **Motivation:**
 - Hadoop style parallel processing masks failures and performance problems
- **Objectives:**
 - Help Hadoop *developers* **debug** and **profile** Hadoop
 - Help *operators* **monitor** and **optimize** MapReduce jobs



About the RAD Lab



Reliable, Adaptive Distributed Systems



Setup for X-Tracing

- Instrument Hadoop using X-Trace framework
- Trace analysis
 - Visualization via web-based UI
 - Statistical analysis and anomaly detection
- Identify potential problems



Outline

- X-Trace overview
- Applications of trace analysis
- Conclusion



Checkpoint

- X-Trace overview
- Applications of trace analysis
- Conclusion




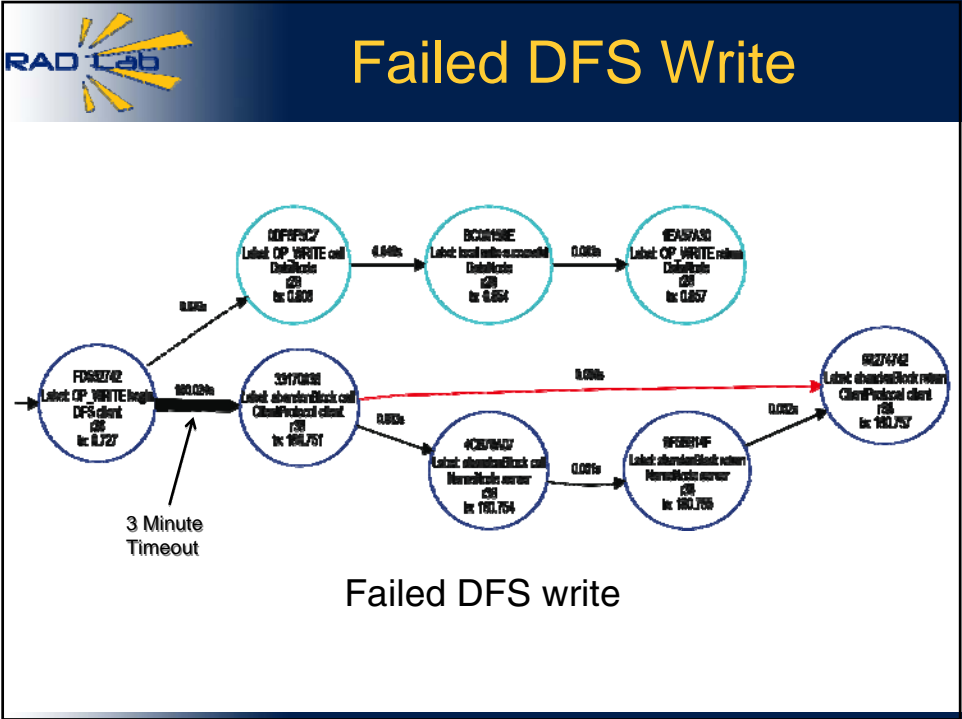
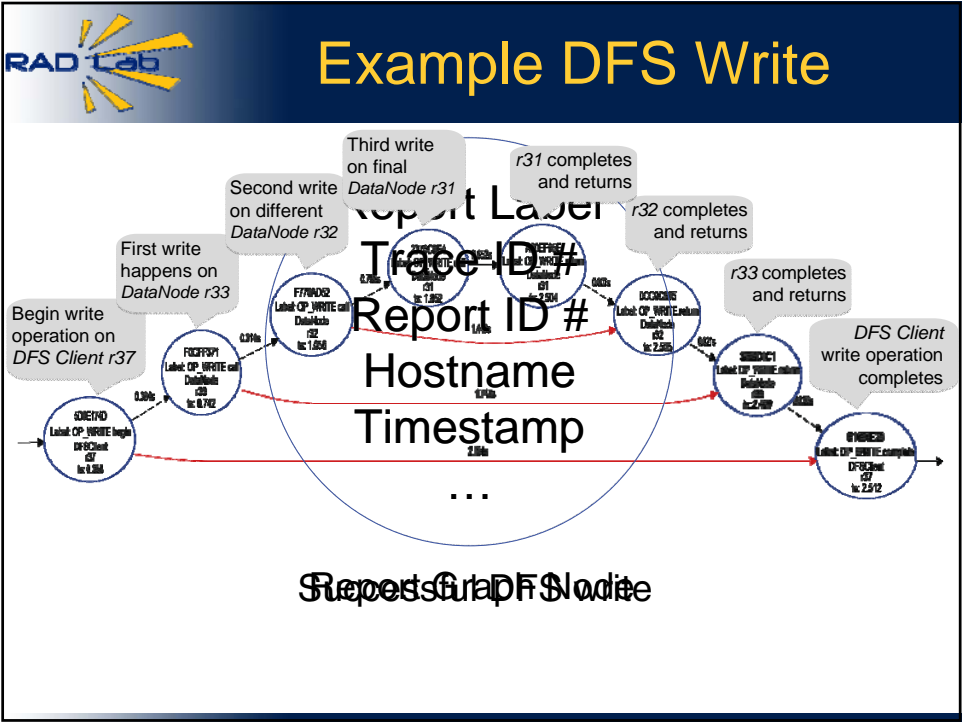
How X-Trace Works

- *Path based* tracing framework
- Generate event graph to capture causality of events across network
 - Examples of events: RPCs, HTTP requests
- Annotate *messages* with trace metadata (16 bytes) carried along execution path
 - Instrument Protocol APIs and RPC libraries




DFS Write Message Flow







MapReduce Event Graph

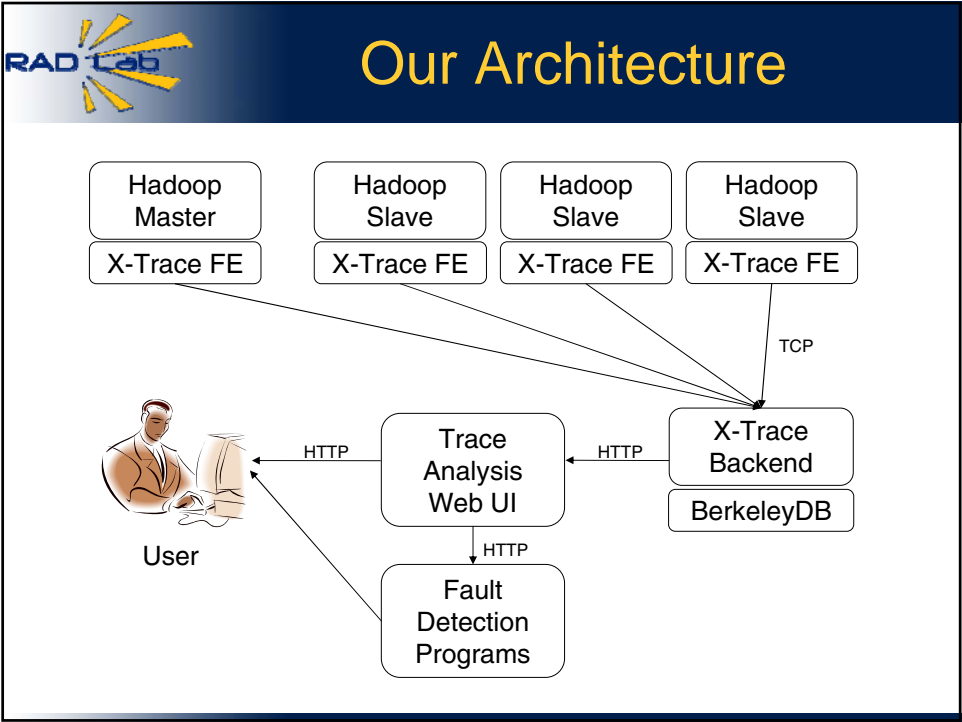


Example MapReduce graph



Properties of Path Tracing

- Deterministic causality and concurrency
- Control over which events get traced
- Cross-layer
- Low overhead
- Modest modification of app source code
 - Fewer than 500 lines for Hadoop



Trace Analysis UI

- Web-based
- Provides
 - Performance statistics
 - Graphs of utilization
 - Critical path analysis

The screenshot shows a browser window displaying statistics for task #AD000275174284. It includes three sections:

DFS Statistics

Operation	Count	Mean	Dev	Min	Max
OP_READ_BLOCK	61	0.01	0.01	0.00	0.07
OP_WRITE_BLOCK	6	0.28	0.21	0.03	0.75

Process Statistics

Method	Count	Mean	Dev	Min	Max
copyToLocalistic	13	0.05	0.03	0.01	0.07
saveTaskOutput	6	0.03	0.01	0.02	0.04
run MapRunner	10	70.76	5.90	66.12	87.67
run MapTask	10	76.92	9.10	71.22	103.67
run ReduceTask	3	116.31	0.28	115.97	116.65
saveTaskOutput	13	0.01	0.02	0.00	0.02

Utilization Statistics

The graph shows the number of machines in use over time. The y-axis is 'Machines in use' (0 to 24) and the x-axis is 'Time (s)' (0 to 200). The utilization starts at 24 machines, remains constant until approximately 100 seconds, then drops sharply to about 5 machines and continues to decrease slowly towards 200 seconds.



Checkpoint

- X-Trace overview
- Applications of trace analysis
- Conclusion




Optimizing Job Performance

- Examined performance of Apache Nutch web indexing engine on a Wikipedia crawl
- Time to creating an inverted link index of a 50 GB crawl
 - With default configuration, 2 hours
 - With optimized configuration, 7 minutes



Optimizing Job Performance

Machine Utilization over Time

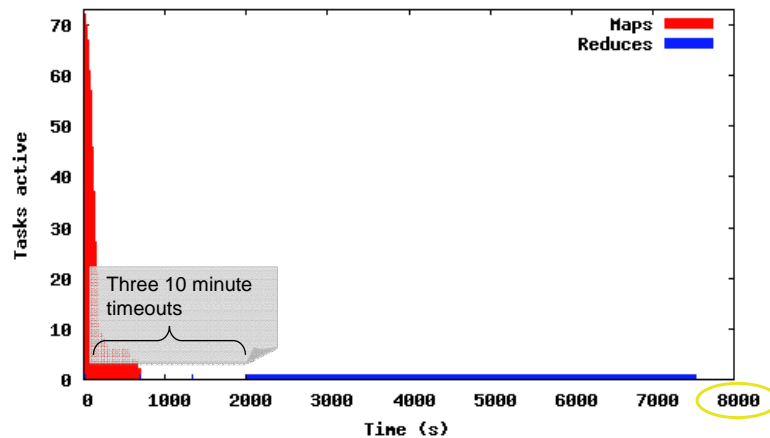


Machine utilization under default configuration



Optimizing Job Performance

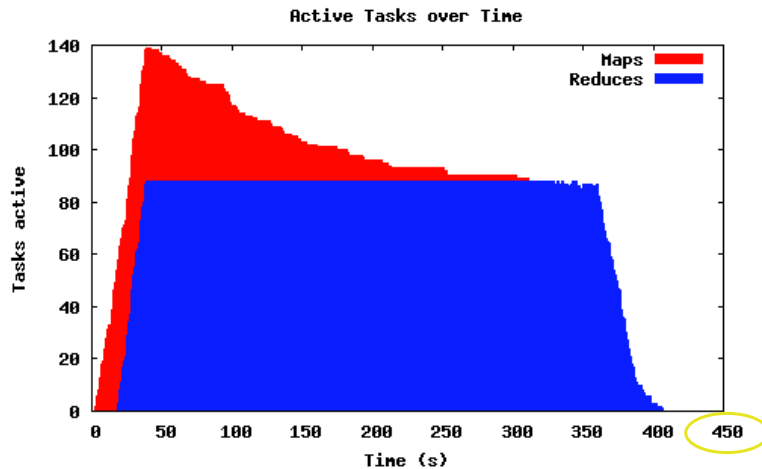
Active Tasks over Time



Problem: One single Reduce task, which actually fails several times at the beginning



Optimizing Job Performance

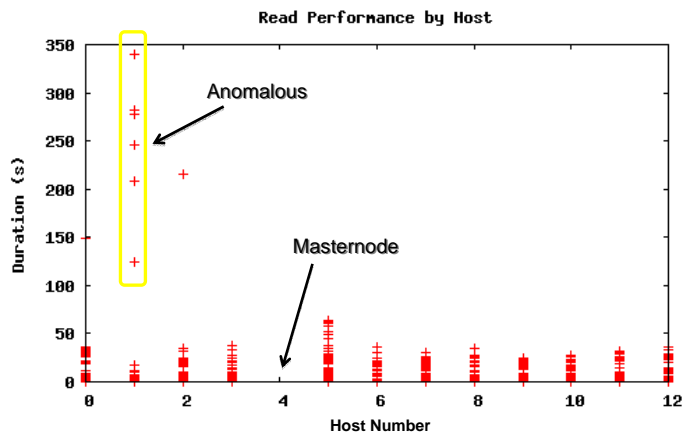


Active tasks vs. time with improved configuration
(50 reduce tasks instead of one)



Faulty Machine Detection

- Motivated by observing slow machine



Diagnosed to be failing hard drive



Statistical Analysis

- Off-line Machine Learning
 - Faulty machine detection
 - Buggy software detection
- Current Work on graph processing and analysis



Future Work

- Tracing more production MapReduce applications
 - Larger clusters + real workloads
- More advanced trace processing tools
- Migrating our code into Hadoop codebase



Conclusion

- Efficient, low overhead tracing of Hadoop
- Exposed multiple interesting behaviors


andyk@cs.berkeley.edu

matei@cs.berkeley.edu

<http://x-trace.net>



Questions?





Andy's ToDo Slide:

- **Hadoop Summit prep TODO:**
 - A section on telemetry collection/correlation
 - Add url where they can go to check out x-trace (more about "official release")
 - ~~get feedback~~
 - ~~Practice talk~~
 - ~~Update all outline slides~~
 - Change green to blue on all graphs (color blind bastards!)
 - Email yahoo and joydeep about referencing our work with them



Overhead of X-Trace

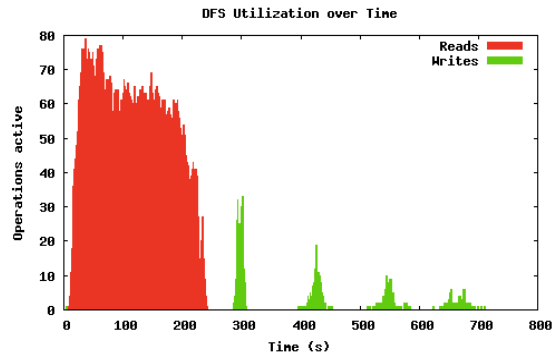
- Negligible overhead in our 40-node test cluster
- Because Hadoop operations are large
- E.g. 18-minute Apache Nutch indexing job with 390 tasks generates 50,000 reports (~20 MB of text data)



Optimizing Job Performance

Init	0.0s (0.1%)
Start JVM	1.6s (4.8%)
Init runner	4.5s (13.4%)
Run map	13.1s (39.3%)
Cleanup	14.1s (42.4%)

Breakdown of longest map
with 3x more mappers



Tracing also illustrates the inefficiency of having
too many mappers and reducers



Related Work

- Per-machine logs
 - Active tracing: log4j, etc
 - Passive tracing: dtrace, strace
- Log aggregation tools: Sawzall, Pig
- Cluster monitoring tools: Ganglia



Event Graphs

- Events are nodes in a “causal” directed graph
 - Captures causality (graph edges)
 - Captures concurrency (fan-out, fan-in, parallel edges)
- Spans layers, applications, administrative boundaries




Observations about Hadoop

- Hardcoded timeouts may be inappropriate in some cases (e.g. long reduce task)
 - Solution: expose timeouts in configuration file?
- Highly variable DFS performance under load (on *our* cluster), which can slow the entire job
 - Solution: multiple hard disks per node?




Hardware Provisioning

*Node with 1 Disk
(high variance typical)*



*Node with 4 Disks
(variance typically lower)*



Faulty Machine Detection

- Approach: Compare performance of each machine vs. others in that trace
- Statistical Anomaly detection tests