

CS162
Operating Systems and
Systems Programming
Lecture 25

Why Systems Fail and
What We Can Do About It

November 30, 2011
Anthony D. Joseph and Ion Stoica
<http://inst.eecs.berkeley.edu/~cs162>

Goals for Today

- Definitions for Fault Tolerance
- Causes of system failures
- Possible solutions
- Staying broad

"You know you have a distributed system when the crash of a computer you've never heard of stops you from getting any work done." —LESLIE LAMPART

Note: Some slides and/or pictures in the following are adapted from slides from a talk given by Jim Gray at UC Berkeley on November 9, 2000.

11/30/2011

Anthony D. Joseph and Ion Stoica CS162 ©UCB Spring 2011

Lec 25.2

Dependability: The 3 ITIES

- **Reliability / Integrity:**
does the right thing.
(Also large MTBF)
- **Availability:** does it now.
(Also small $\frac{MTTR}{MTBF+MTTR}$)
- **System Availability:**
if 90% of terminals up & 99% of DB up?
(=> 89% of transactions are serviced on time)



MTBF or MTTF = Mean Time Between (To) Failure
MTTR = Mean Time To Repair

11/30/2011

Anthony D. Joseph and Ion Stoica CS162 ©UCB Spring 2011

Lec 25.3

Fault Tolerance versus
Disaster Tolerance

- **Fault-Tolerance:** mask local faults
 - RAID disks
 - Uninterruptible Power Supplies
 - Cluster Failover
- **Disaster Tolerance:** masks site failures
 - Protects against fire, flood, sabotage,..
 - Redundant system and service at remote site.
 - Use design diversity



11/30/2011

Anthony D. Joseph and Ion Stoica CS162 ©UCB Spring 2011

Lec 25.4

High Availability System Classes Goal: Build Class 6 Systems

| System Type | Unavailable (min/year) | Availability | Availability Class |
|------------------------|------------------------|--------------|--------------------|
| Unmanaged | 50,000 | 90.0% | 1 |
| Managed | 5,000 | 99.0% | 2 |
| Well Managed | 500 | 99.99% | 3 |
| Fault Tolerant | 50 | 99.999% | 4 |
| High-Availability | 5 | 99.9999% | 5 |
| Very-High-Availability | .5 | 99.99999% | 6 |
| Ultra-Availability | .05 | 99.999999% | 7 |

UnAvailability = MTTR/MTBF
can cut it in 1/2 by cutting MTTR or MTBF


11/30/2011

Anthony D. Joseph and Ion Stoica CS162 ©UCB Spring 2011

Lec 25.5

Case Study – Japan

"Survey on Computer Security", Japan Info Dev Corp., March 1986. (trans: Eiichi Watanabe).



| | |
|--------------------------------|-----------------|
| Vendor (hardware and software) | 5 Months |
| Application software | 9 Months |
| Communications lines | 1.5 Years |
| Operations | 2 Years |
| Environment | 2 Years |
| Total | 10 Weeks |

1,383 institutions reported (6/84 - 7/85)

7,517 outages, MTBF ~ 10 weeks, avg duration ~ 90 MINUTES


To Get 10 Year MTBF, Must Attack All These Areas

11/30/2011

Anthony D. Joseph and Ion Stoica CS162 ©UCB Spring 2011

Lec 25.6

Case Studies - Tandem Trends Reported MTBF by Component



| | 1985 | 1987 | 1990 | Years |
|-------------|------|------|------|-------|
| SOFTWARE | 2 | 53 | 33 | Years |
| HARDWARE | 29 | 91 | 310 | Years |
| MAINTENANCE | 45 | 162 | 409 | Years |
| OPERATIONS | 99 | 171 | 136 | Years |
| ENVIRONMENT | 142 | 214 | 346 | Years |
| SYSTEM | 8 | 20 | 21 | Years |


Problem: Systematic Under-reporting

11/30/2011

Anthony D. Joseph and Ion Stoica CS162 ©UCB Spring 2011

Lec 25.7

Operations Failures



What went wrong??

11/30/2011

Anthony D. Joseph and Ion Stoica CS162 ©UCB Spring 2011

Lec 25.8

Operations Failures

RAID Drive 1 failed!
Replace immediately



1

i.9

Cloud Computing Outages 2011

| Vendor | When | Duration | What Happened & Why |
|---|----------------|----------|--|
| Apple iPhone 4S Siri | November 2011 | 1 Day | Siri loses even the most basic functionality when Apples servers are down. Because Siri depends on servers to do the heavy computing required for voice recognition, the service is useless without that connection. Network outages caused the disruption according to Apple. |
| Blackberry outage | October 2011 | 3 Days | Outage was caused by a hardware failure (core switch failure) that prompted a "ripple effect" in RIM's systems. Users in Europe, Middle East, Africa, India, Brazil, China and Argentina initially experienced email and message delays and complete outages and later the outages spread to North America too. Main problem is message backlogs and the downtime produced a huge queue of undelivered messages causing delays and traffic jams. |
| Google Docs | September 2011 | 1 Hour | Google Docs word collaboration application cramp, shutting out millions of users from their document lists, documents, drawings and Apps Scripts. Outage was caused by a memory management bug software engineers triggered in a change designed to "improve real time collaboration within the document list." |
| Microsoft's services - Hotmail & SkyDrive | 2011 | | Domain Name Service (DNS) Network traffic balancing tool had an update and the update did not work properly which caused the issue. |
| Amazon's EC2 cloud & | August 2011 | 1-2 days | Transformer exploded and caught fire near datacenter that resulted in power outage due to generator failure. Power back up systems at both the data centers failed causing power outages. Transformer explosion was caused by lightning strike but disputed by local utility provider. |
| Microsoft's BPOS | August 2011 | 1-2 days | Transformer exploded and caught fire near datacenter that resulted in power outage due to generator failure. Power back up systems at both the data centers failed causing power outages. Transformer explosion was caused by lightning strike but disputed by local utility provider. |

From: <http://analysiscasesstudy.blogspot.com/>

11/30/2011

Anthony D. Joseph and Ion Stoica CS162 @UCB Spring 2011

Lec 25.10

Cloud Computing Outages 2011

| Vendor | When | Duration | What Happened & Why |
|-----------------------------|------------------|-----------|--|
| Amazon Web Services | April, 2011 | 4 Days | During the upgrade, the traffic shift was executed incorrectly and rather than routing the traffic to the other router on the primary network, the traffic was routed onto the lower capacity redundant EBS network. This led to Amazon Elastic Block Store ("EBS") volumes in a single Availability Zone within the US East Region that became unable to service read and write operations. It also impacted the Relational Database Service ("RDS"). RDS depends upon EBS for database and log storage, and as a result a portion of the RDS databases hosted in the primary affected Availability Zone became inaccessible. |
| Microsoft BPOS Outages | May 2011 | 2 Hours | During customer email messages delayed by network errors hours. Database outgoing messages started getting stuck in the pipeline. |
| Twitter Outages | March & Feb 2011 | 1-4 Hours | Outages due to over capacity and moving operations to new data center. |
| Intuit Quick Books Online | March 2011 | 2 Days | Service failures on human error during scheduled maintenance operations. Intuit changed its network configuration and inadvertently blocked customer access to a portion of the company's servers. A surge in traffic overloaded the servers when connectivity was restored, so the company opted to restore service. |
| Google Mail and Apps Outage | February 2011 | 2 Days | Google mail and Google Apps users experienced login errors and empty mailboxes. Google Engineering determined that the root cause was a bug inadvertently introduced in a Gmail storage software update. The bug caused the affected users' messages and account settings to become temporarily unavailable from the datacenters. |

From: <http://analysiscasesstudy.blogspot.com/>

11/30/2011

Anthony D. Joseph and Ion Stoica CS162 @UCB Spring 2011

Lec 25.11

Cloud Computing Outages 2010

| Vendor | When | Duration | What Happened & Why |
|------------------------|-------------------|-------------|--|
| Hotmail Outage | December 2010 | 3 Days | A number of our users reported their email messages and folders were missing from their Hotmail accounts. Error occurred from a script that was meant to delete dummy accounts created for automated testing and that inadvertently targeted 47,000 real accounts. |
| Skype Outage | December 2010 | 1 Day | Cluster of support servers responsible for offline instant messaging became overloaded and the P2P network became unstable and suffered a critical failure. A supernode is important to the P2P network acting like a directory, supporting other Skype clients, helping to establish connections between them etc. The failure of 25-30% of supernodes in the P2P network resulted in an increased load on the remaining supernodes. |
| Paypal Outage | November 2010 | 3 Hours | A network hardware failure was the trigger for an outage. The hardware failure was worsened by problems in shifting traffic to another data center, resulting in about 90 minutes of downtime. |
| Facebook Outage | September 2010 | 2 1/2 Hours | Outage due to an error condition. An automated system for verifying configuration values ended up causing much more damage than it fixed. Every single client saw the invalid value and attempted to fix it that led to a query to a database cluster and cluster was overloaded with thousand of queries per second. Even after fixing problem stream of queries continued. |
| Microsoft BPOS Outages | September 2010 | 2 Hours | A design issue in the upgrade to Hotmail caused messages not to get sent, but the issue occurred in a 2 hour period of intermittent access for BPOS organizations served from North America. |
| Wikipedia Outage | July & March 2010 | 2-3 Hours | In July, the power failure is understood to have affected Wikimedia's 'prmpa' cluster. Due to the temporary unavailability of several critical systems and the large impact on the available systems capacity, all Wikimedia projects went down. In March, Wikimedia servers overheated in the organization's European data center and shut themselves off automatically. Wikimedia then switched all its traffic to its server cluster in Florida, but the failover process, which involves changing servers' DNS entries, malfunctioned, knocking the organization's sites offline around the world. |
| Hosting.com Outage | June 2010 | 2 Hours | Failure of a Cisco switch at the Newark, N.J., data center caused intermittent network connectivity. Dedicated switch had failed, the second failover switch had crashed as well and the problem was caused by a software bug. |
| Twitter.com outage | June 2010 | 5 hours | Increased activity on the site, combined with system enhancements and upgrades, have uncovered networking issues. Incidences of poor site performance and a high number of errors due to one of the internal sub-networks being over-capacity. |

11/30/2011

Anthony D. Joseph and Ion Stoica CS162 @UCB Spring 2011

Lec 25.12

| Cloud Computing Outages 2009 | | | |
|------------------------------|----------------------------|-----------------|---|
| Vendor | When | Duration | What Happened & Why |
| Salesforce.com Outage | January 2010, 2009 | 1-2 Hours | Outages were caused by server disruption, when a core network device failed, stopping all data from being processed in Japan, Europe, and North America. The technical reason for the outage: a core network device had failed, due to memory allocation errors. The backup plan, which was supposed to trigger a cut-over to a redundant system, also failed. |
| Amazon's EC2 | June 2009 | 1.5 Hours | A lightning storm caused damage to a single Power Distribution Unit (PDU) in a single Availability Zone |
| eBay Paypal | August 2009 | 1-4 Hours | Online payments system failed a couple of times led to non completion of transactions. Network hardware issue is blamed for outage. |
| Twitter | August 2009 | 1/2 Day | A denial-of-service attack was blamed for the problem |
| Google Gmail | September 2009 | 2 hours 2 times | Reasons from vendors include routing errors to server maintenance issues. |
| Microsoft Sidekick | October 2009 | 6 days | Microsoft's Danger server farm, that holds the cloud T-Mobile Sidekick subscriber's data crashed, depriving users of their calendar, address book, and other key data. Critical data was lost during outage. |
| Rackspace.com Outage | June 2009 December 2009 | 1 Day 1 Hour | Power outage and subsequent power generator failures that caused servers to fail. Company was forced to pay out between \$2.5 million and \$3.5 million in service credits to customers. The issues resulted from a problem with a router used for peering and backbone connectivity located outside the data center at a peering facility, which handles approximately 20% of Rackspace's Dallas traffic. The router configuration error was part of final testing for data center integration between the Chicago and Dallas facilities. |

From: <http://analysiscasestudy.blogspot.com/>

11/30/2011 Anthony D. Joseph and Ion Stoica CS162 ©UCB Spring 2011 Lec 25.13

Fault Model

- Failures are independent*
So, single fault tolerance is a big win
- Hardware fails fast (blue-screen, panic, ...)
- Software fails-fast (or stops responding/hangs)
- Software often repaired by reboot:
 - Heisenbugs – Works On Retry
 - Bohrbugs – Faults Again On Retry
- Operations tasks: major source of outage
 - Utility operations – UPS/generator maintenance
 - Software upgrades, configuration changes

11/30/2011 Anthony D. Joseph and Ion Stoica CS162 ©UCB Spring 2011 Lec 25.14

Some Fault Tolerance Techniques

- Fail fast modules: work or stop
- Spare modules: yield instant repair time
- Process/System pairs: Mask HW and SW faults
- Transactions: yields ACID semantics (simple fault model)

11/30/2011 Anthony D. Joseph and Ion Stoica CS162 ©UCB Spring 2011 Lec 25.15

Fail-Fast is Good, Repair is Needed

Lifecycle of a module
fail-fast gives short fault latency


High Availability
is low UN-Availability

$$\text{Unavailability} \sim \frac{\text{MTTR}}{\text{MTBF}}$$

Improving either MTTR or MTBF gives benefit
Simple redundancy does not help much

11/30/2011 Anthony D. Joseph and Ion Stoica CS162 ©UCB Spring 2011 Lec 25.16

Hardware Reliability/Availability (how to make HW fail fast)

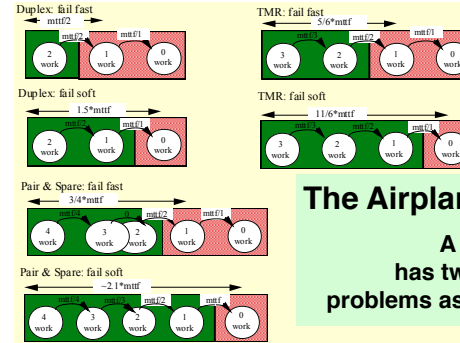


Comparator Strategies: (in recursive pairs, parent knows which is bad)

- Duplex: Fail-Fast: fail if either fails (e.g. duplexed CPUs)
 vs
 Fail-Soft: fail if both fail (e.g. disc, network,...)
- Triplex: Fail-Fast: fail if 2 fail (triplexed cpus)
 Fail-Soft: fail if 3 fail (triplexed FailFast CPUs)

11/30/2011 Anthony D. Joseph and Ion Stoica CS162 ©UCB Spring 2011 Lec 25.17

Redundant Designs have Worse MTBF!




The Airplane Rule:

A two-engine airplane has twice as many engine problems as a one engine plane

THIS IS NOT GOOD: Variance is lower but MTBF is worse
 Simple redundancy does not improve MTBF (sometimes hurts)

11/30/2011 Anthony D. Joseph and Ion Stoica CS162 ©UCB Spring 2011 Lec 25.18

Add Repair: Get 10^4 Improvement



Availability estimates
 1 year MTTF modules
 12-hour MTTR

| | MTTF | EQUATION | COST |
|------------------------------|------------|----------------|------|
| SIMPLEX | 1 year | - MTTF | 1 |
| DUPLEX: FAIL FAST | ~0.5 years | - MTTF/2 | 2+E |
| DUPLEX: FAIL SOFT | ~1.5 years | - MTTF(3/2) | 2+E |
| TRIPLEX: FAIL FAST | .8 year | - MTTF(5/6) | 3+E |
| TRIPLEX: FAIL SOFT | 1.8 year | - 1.8MTTF | 3+E |
| Pair and spare: FAIL-FAST | ~.7 year | - MTTF(3/4) | 4+E |
| TRIPLEX WITH REPAIR | >105 years | $MTTF^3/3MTTR$ | 3+E |
| Duplex fail soft + REPAIR | >104 years | $MTTF^2/2MTTR$ | 4+E |

11/30/2011 Anthony D. Joseph an

Software Techniques: Learning from Hardware

Recall that most outages are not hardware
 Most outages in Fault Tolerant Systems are SOFTWARE
 Fault Avoidance Techniques: Good & Correct design

After that: Software Fault Tolerance Techniques:

Modularity (isolation, fault containment)

N-Version Programming: N-different implementations

Programming for Failures: Programming paradigms that assume failures are common and hide them

Defensive Programming: Check parameters and data

Auditors: Check data structures in background

Transactions: to clean up state after a failure

Paradox: Need Fail-Fast Software

11/30/2011 Anthony D. Joseph and Ion Stoica CS162 ©UCB Spring 2011 Lec 25.20

Fail-Fast and High-Availability Execution

Process Pairs: Instant restart (repair)


Use Defensive programming to make a process fail-fast
 Have restarted process ready in separate environment
 Second process “takes over” if primary faults

If software fault (bug) is a Bohrbug, then there is no repair
 “wait for the next release” or “get an emergency bug fix” or
 “get a new vendor”

If software fault is a Heisenbug, then repair is
 “reboot and retry” or “switch to backup process (instant restart)”

Tolerates HW faults too!


Repair time is seconds,
 could be ms if time critical



11/30/2011

Anthony D. Joseph and Ion

SYSTEM PAIRS FOR HIGH AVAILABILITY



- Programs, Data, Processes Replicated at 2+ sites
 - Pair looks like a single system
- System becomes logical concept
 - Like Process Pairs: System Pairs.
- Backup receives transaction log (spooled if backup down)
- If primary fails or operator switches, backup offers service

11/30/2011

Anthony D. Joseph and Ion Stoica CS162 ©UCB Spring 2011

Lec 25.22

Add Geographic Diversity to Reduce Single Points of Failure*



11/30/2011

Anthony D. Joseph and Ion Stoica CS162 ©UCB Spring 2011

Lec 25.23

Administrivia

- Final:
 - Thursday, December 15, 8-11am, 155 Dwinelle
 - Closed book, **two** pages of hand-written notes (both sides)
- Topics:
 - 30% first part
 - 70% second part
- Review session: **Next week TBA**
- Office hours:
 - Anthony: Tuesday Dec 6, 10:30-11:30 am and Wednesday, Dec 7, 12:30-1:30pm
 - Ion: Tuesday, Dec 6, 9:30-10:30am and Wednesday, 11:30-12:30pm
- We will post example questions for the final
 - Look at previous years' second/third midterms

11/30/2011

Anthony D. Joseph and Ion Stoica CS162 ©UCB Spring 2011

Lec 25.24

5min Break

11/30/2011

Anthony D. Joseph and Ion Stoica CS162 ©UCB Spring 2011

Lec 25.25

MapReduce

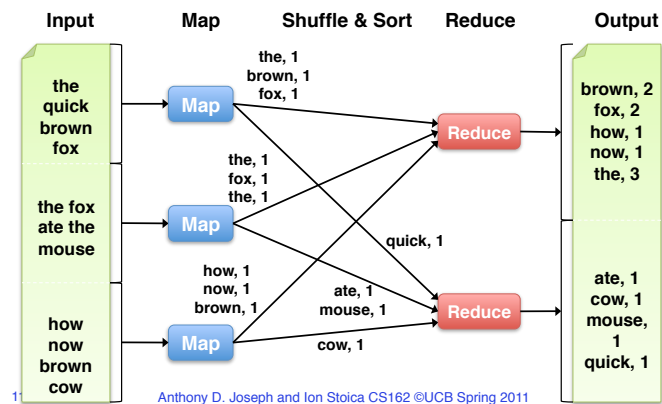
- First widely popular programming model for data-intensive apps on commodity clusters
- Published by Google in 2004
 - Processes 20 PB of data / day
- Popularized by open-source Hadoop project
 - 40,000 nodes at Yahoo!, 70 PB at Facebook
- Programming model
 - Data type: key-value *records*
 - » Map function: $(K_{in}, V_{in}) \rightarrow list(K_{inter}, V_{inter})$
 - » Reduce function: $(K_{inter}, list(V_{inter})) \rightarrow list(K_{out}, V_{out})$

11/30/2011

Anthony D. Joseph and Ion Stoica CS162 ©UCB Spring 2011

Lec 25.26

Word Count Execution



11

Anthony D. Joseph and Ion Stoica CS162 ©UCB Spring 2011

Fault Tolerance in MapReduce

1. If a task crashes:
 - Retry on another node
 - » OK for a map because it had no dependencies
 - » OK for reduce because map outputs are on disk
 - If the same task repeatedly fails, fail the job

➤ **Note:** For the fault tolerance to work, tasks must be *deterministic and side-effect-free*

11/30/2011

Anthony D. Joseph and Ion Stoica CS162 ©UCB Spring 2011

Lec 25.28

Fault Tolerance in MapReduce

2. If a node crashes:
 - Relaunch its current tasks on other nodes
 - Relaunch any maps the node previously ran
 - » Necessary because their output files are lost

11/30/2011

Anthony D. Joseph and Ion Stoica CS162 ©UCB Spring 2011

Lec 25.29

Fault Tolerance in MapReduce

3. If a task is going slowly (straggler):
 - Launch second copy of task on another node
 - Take output of whichever copy finishes first
- Critical for performance in large clusters

11/30/2011

Anthony D. Joseph and Ion Stoica CS162 ©UCB Spring 2011

Lec 25.30

Takeaways

- By providing a data-parallel programming model, MapReduce can control job execution in useful ways:
 - Automatic division of job into tasks
 - Placement of computation near data
 - Load balancing
 - Recovery from failures & stragglers

11/30/2011

Anthony D. Joseph and Ion Stoica CS162 ©UCB Spring 2011

Lec 25.31

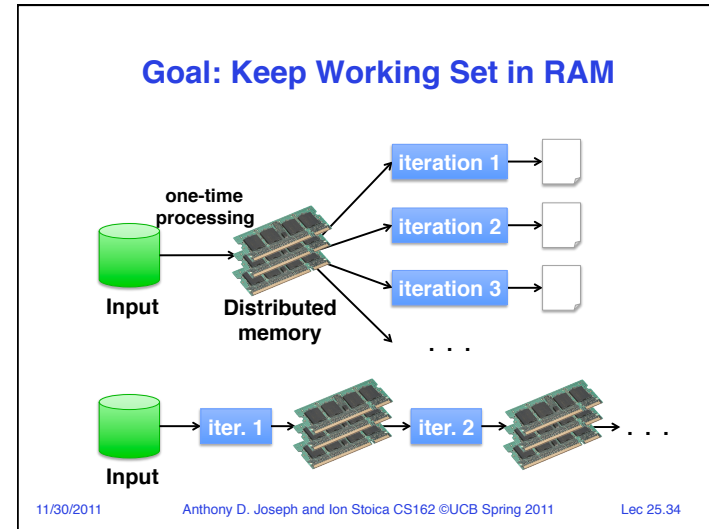
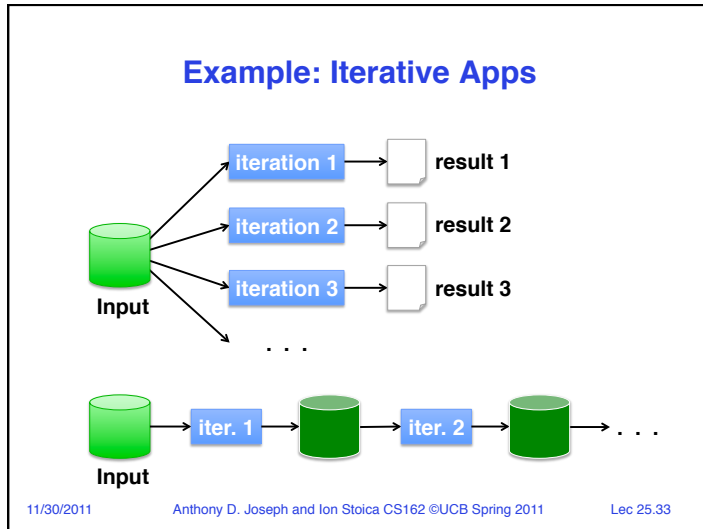
Issues with MapReduce

- Hard to express more complex programs
 - E.g. word count + a sort to find the top words
 - Need to write many different map and reduce functions that are split up all over the program
 - Must write complex operators (e.g. join) by hand
- Acyclic data flow -> poor support for applications that need to *reuse* pieces of data
 - Iterative algorithms (e.g. machine learning, graphs)
 - Interactive data mining (e.g. Matlab, Python, SQL)

11/30/2011

Anthony D. Joseph and Ion Stoica CS162 ©UCB Spring 2011


Lec 25.32



- ### Spark Goals
- Support apps with data reuse efficiently:
 - Let them keep data in memory
 - Retain the fault tolerance and automatic scheduling benefits of MapReduce
 - Enhance programmability:
 - Integrate into Scala programming language
 - Allow interactive use from Scala interpreter
- 11/30/2011 Anthony D. Joseph and Ion Stoica CS162 ©UCB Spring 2011 Lec 25.35

- ### Key Idea: Resilient Distributed Datasets (RDDs)
- *Restricted* form of distributed shared memory
 - Read-only, partitioned collections of records
 - Can only be created through deterministic transformations (map, group-by, join, ...)
 - Allows efficient implementation & recovery
 - Key idea: rebuild lost data using *lineage*
- 11/30/2011 Anthony D. Joseph and Ion Stoica CS162 ©UCB Spring 2011 Lec 25.36

RDD Recovery



11/30/2011

Anthony D. Joseph and Ion Stoica CS162 ©UCB Spring 2011

Lec 25.37

Programming Model

- Resilient distributed datasets (RDDs)
 - Immutable, partitioned collections of objects
 - Created through parallel *transformations* (map, filter, groupBy, join, ...) on data in stable storage
 - Can be *cached* for efficient reuse
- *Actions* on RDDs
 - Count, reduce, collect, save, ...

11/30/2011

Anthony D. Joseph and Ion Stoica CS162 ©UCB Spring 2011

Lec 25.38


Example: Log Mining

Load error messages from a log into memory, then interactively search for various patterns

```
lines = spark.textFile("hdfs://...")
errors = lines.filter(_.startsWith("ERROR"))
messages = errors.map(_.split('\t')(2))
cachedMsgs = messages.cache()
```

```
cachedMsgs.filter(_.contains("foo")).count
cachedMsgs.filter(_.contains("bar")).count
...
```

**Result: scaled to 1 TB data in 5-7 sec
(vs 170 sec for on-disk data)**



11/30/2011

Anthony D. Joseph and Ion Stoica CS162 ©UCB Spring 2011


Lec 25.39

RDD Fault Tolerance

RDDs maintain *lineage* information that can be used to reconstruct lost partitions

Ex:

```
cachedMsgs = textFile(...).filter(_.contains("error"))
    .map(_.split('\t')(2))
    .cache()
```



11/30/2011

Anthony D. Joseph and Ion Stoica CS162 ©UCB Spring 2011

Lec 25.40

Apache ZooKeeper


- Highly available, scalable, distributed coordination kernel
 - Leader Election, Group Membership, Work Queues, Sharding
 - Event Notifications/workflow, Config, and Cluster Mgmt
- Provides:
 - File API without partial reads/writes and no renames
 - Ordered updates and strong persistence guarantees
 - Conditional updates (version), Watches for data changes
- API:
 - String create(path, data, acl, flags)
 - void delete(path, expectedVersion)
 - Stat setData(path, data, expectedVersion)
 - (data, Stat) getData(path, watch)
 - Stat exists(path, watch)
 - String[] getChildren(path, watch)

11/30/2011

Anthony D. Joseph and Ion Stoica CS162 ©UCB Spring 2011

Lec 25.41

ZooKeeper Service



- All servers store a copy of the data (in memory)
- A leader is elected at startup, or upon current leader failure
- Followers service clients, all updates go through leader
- Update responses are sent when a majority of servers have persisted the change

11/30/2011

Anthony D. Joseph and Ion Stoica CS162 ©UCB Spring 2011

Lec 25.42

New CS162

- Different emphasis from CS162 in previous years; beginning in 2011 we shifted to give students a broad view on how today's systems and services
 - End-to-end system design, rather than OS only
 - More networking, database, and security concepts
 - New projects to reflect this emphasis
 - Better prepare students to design/develop such services
- Long term plan: make CS 162 a gateway course for
 - Database class (CS 186)
 - Networking class (EE 122)
 - Security class (CS 161)
 - Software engineering class (CS 169)
 - New OS class (cs16x in Fall 2012 with real OS)

11/30/2011

Anthony D. Joseph and Ion Stoica CS162 ©UCB Spring 2011

Lec 25.43

New vs. Old CS162

- Curriculum: 70% overlap
 - File systems, queuing theory, slightly fewer lectures on concurrency, caching, and distributed systems
 - + More networking, database transactions, p2p, and cloud computing
 - On-going analysis of what worked and didn't work*
- Different project: emphasize on how a system works end-to-end rather than focusing on implementing OS concepts in Nachos
- What if you want to do an OS project?
 - CS 16x in Fall 2012
 - Undergraduate research projects in the AMP Lab
 - » Akaros, Spartk, or Mesos projects

11/30/2011

Anthony D. Joseph and Ion Stoica CS162 ©UCB Spring 2011

Lec 25.44

Stay Broad!

- Very few influential nerds – never will be
 - Why?
 - » They're too narrow
 - » Real breakthroughs tend to come from people with breadth.
- One of the most important things you should do is to force yourself to stay broad

11/30/2011

Anthony D. Joseph and Ion Stoica CS162 ©UCB Spring 2011

Lec 25.45

Reasons to Stay Broad

- Breadth helps depth
 - The more you understand about many different things, the more you'll understand about each individual thing
 - Seeing how things in different areas are similar or different is a very good way of seeing what's important
- Breakthroughs often occur when people can cross traditional boundaries: compilers and architecture, graphics and VLSI, etc.
- Computers are tools: they serve people
 - To create effective tools, must understand the capabilities of computers, the capabilities of people, and the needs of the application areas where they'll be used

11/30/2011

Anthony D. Joseph and Ion Stoica CS162 ©UCB Spring 2011

Lec 25.46

More Reasons to Stay Broad

- Technology is changing fast
 - If you get cubby-holed, you'll get left behind.
- Why is there a shortage of 25-year-old engineers and a surplus of 45-year-old ones?
 - Companies encourage new graduates to get so narrow (short-term focus) that they're instantly obsolete
- *If you don't keep up, you'll be left behind...*

11/30/2011

Anthony D. Joseph and Ion Stoica CS162 ©UCB Spring 2011

Lec 25.47

Solutions

- Continuing education
 - Try always to be learning in whatever you do
 - Don't let your education stop when you leave Berkeley.
 - Consider getting a Masters degree or a PhD
- Explore new areas, both inside and outside Computer Science
 - Amazon AWS makes it easy to experiment
 - Everything you learn will someday be helpful, no matter how unlikely it seems – English, art, hobbies, all things are helpful
- *Bottom line: you are going to change the world!*
 - *You might not realize it yet, but the people in this classroom are going to do it*

11/30/2011

Anthony D. Joseph and Ion Stoica CS162 ©UCB Spring 2011

Lec 25.48