CS 294-42: Cloud Computing, Systems, Networking, and Frameworks

Fall 2011 (MW 1:00-2:30, **293 Cory Hall**)

Ion Stoica

(http://www.cs.berkeley.edu/~istoica/classes/cs294/11/)

Cloud Computing in Numbers

- Datacenter instance:
  - Cost in billion range
  - > 100,000 servers

- Number of servers (estimates*)
  - Google: ~1 mil servers
  - Microsoft, Yahoo!, IBM, HP: several 100,000s each
  - Amazon, Ebay, GoDaddy, Facebook, Akamai: > 50,000

(*http://www.datacenterknowledge.com/archives/2009/05/14/whos-got-the-most-web-servers/)

Why does Cloud Computing Matter?

- Fundamental change
  - The way applications are written and deployed
  - Internet traffic: Internet becomes last-hop between hosts and datacenter
  - Economics: pay as you go
- Opportunity to rethink:
  - Large scale distributed systems
  - Network architectures
  - Tradeoffs in computer systems

How is Cloud Computing Different?

- How is different from distributed systems?
- How is different from parallel systems?
- Axis:
  - Environment constraints
  - Scale
  - Type of failures
  - Application requirements
  - …

Academia Challenges

- Rapid evolving field
  - Need to be in the avantgarde of understanding challenges and trends
  - Unfortunately, academia trails industry
    - Very hard to achieve the scale and generate the workload that reveal main challenges
- How to address above challenges?

5

Rising to the Challenge

- A thriving open source ecosystem
  - Hadoop/HDF, Hive, Pig, …
  - Distributions: Cloudera, Hortonworks
  - Many users: Facebook, Twitter, Yahoo!, …
- Strong collaboration between Berkeley and industry
  - Google, Microsoft, Facebook, Twitter, Yahoo!, …
- Academia enables experimentation
  - Hard to explore alternatives in industry: usually pick a design and go with it for years

6

Research Trust

- Multiprogramming for datacenters: multiple frameworks sharing same datacenter
- Why?
  - Improve utilization, performance
  - Sharing data
  - Enable applications using different frameworks
  - Multiple frameworks sharing same datacenter
- Current project: Mesos, Orchestra, Memento, …
- Future: towards an OS for clusters

7

Research Trust (con'd)

- Real-time, interactive processing of big data
- Why?
  - Need to process huge volumes of unstructure data
  - Batch processing doesn't allow interactive data processing
- Current projects: Spark, Memento
- Future: trade between response time and accuracy

8

Grading

- Project: 60%
- Class presentations: 40%
  - Around 2 papers per student
  - See Randy's guidelines for leading discussion on papers
    - http://bnrg.eecs.berkeley.edu/~randy/Courses/CS294.F07/LeadingPapers.pdf

9

Administrative Information

- Class website: http://www.cs.berkeley.edu/~istoica/classes/cs294/11/
- Office Hours (Soda 465d):
  - Monday, 3-4pm
  - Wednesday, 11-noon
- Create a (anonymized) blog account for paper reviews if you don't have one yet (e.g., www.blogger.com)
  - Sent me an e-mail by Wednesday, August 31, with your blog url
  - Preferred e-mail for the class e-mail list

10

Papers

- Is the problem real?
- What is the solution's main idea (nugget)?
- Why is solution different from previous work?
  - Are system assumptions different?
  - Is workload different?
  - Is problem new?
- Does the paper (or do you) identify any fundamental/hard trade-offs?

11

Papers (cont'd)

- Do you think the work will be influential in 10 years?
  - Why or why not?
- Predicting the future always hard, but worth a try
  - Look at past examples for inspiration

12

Streaming Over TCP

- Countless papers:
  - Why cannot be done…
  - New protocols to do it…
- Today
  - Virtually all streaming over TCP
  - Trend to stream over HTTP!

13

Why did it Succeed?

14

Multicast

- Countless papers:
  - Why world will come to a standstill without multicast…
  - New protocols to do it…
- Today
  - Multicast is used only in enterprise settings at best
  - Overlay multicast widely used in the Internet
    - CDN based, e.g., WorldCup, March Madness, Iinagurations, ...
    - P2P, mostly popular outside US (e.g., China)

15

Why Did it Fail?

16

Consistency Everywhere

- Many papers & systems:
  - Group synchronous communication
  - Causally ordered message delivery
  - …
- Today:
  - Almost never used in WANs, and rarely used in LANs

17

Why Did it Fail?

18

Shared Memory

- Countless papers:
  - How shared memory simplifies programming parallel computers
  - Many, many systems proposed and build
- Today:
  - Message passing (MPI) took over as the de facto standard for writing parallel applications

19

Why Did it Fail?

20

Network Computer

- Big in 90s
  - Promoted by an alliance of Sun, Oracle, Acorn
- Promise: many of advantages of cloud computing
  - Easy to manage
  - Application sharing
  - …
- Failed miserably

21

Why Did it Fail?

22

Coming Back: ChromeOS

- Will it succeed this time?

23

What are Hard/Fundamental Tradeoffs?

- Brewer's CAP conjecture: "Consistency, Availability, Partition-tolerance", you can have only two in a distributed system
- In a *in-order, reliable* communication protocol cannot minimize overhead and latency simultaneously
- Hard to simultaneously maximize evolvability and performance

24