

# Distributed Systems Intro

## Course Staff Intro

- Instructor:
  - David Andersen
- TAs:
  - Alex Katkova
  - Jonathan Harbuck



- Administrative: Angela Miller

## Logistics

- Course Policies
  - see web page...
  - <http://www.cs.cmu.edu/~dga/15-440/F10>
  - obligatory discussion of {late days, cheating, etc.}
- Waitlist?
- No recitations this year
- Office hours (hands): Earlier, 3:30 - 5:30, 4:00 - 6:00, Sunday?

## Course Goals

- Systems requirement:
  - Learn something about distributed systems in particular;
  - Learn general systems principles (modularity, layering, naming, security, ...)
  - Practice implementing real, larger systems; in teams; must run in nasty environment;
- One consequence: Must pass homeworks, exams, and projects independently as well as in total.

# Course Format

- ~30 lectures
- Recitations: Practical issues for implementing projects; general questions and discussion
- 3 projects; 1 solo, 2 team
  - Distributed (internet-wide) password cracker
  - Distributed filesystem
  - Data-intensive cluster applications with MapReduce/Hadoop

# No Books?

- No official text
- Several useful references on web page
- We'll be compiling notes (and these slides) for your use over the course of the semester; based on, but not identical to, prior 15-440 instance

# Recitations

- Systems programming somewhat different from what you've done before
  - Low-level (C)
  - Often designed to run indefinitely (error handling must be rock solid)
  - Must be secure - horrible environment
  - Concurrency
  - Interfaces specified by documented protocols
- Recitations & "System Hacker's View of Software Engineering"
  - Practical techniques designed to save you time & pain

# Collaboration

- Working together important
  - Discuss course material
  - Work on problem debugging
- Parts *must* be your own work
  - Homeworks, midterm, final, solo proj
- Team projects: both students should understand entire project
- What we hate to say: we run cheat checkers...

# Late Work

- 10% penalty per day
- Can't be more than 2 days late
- Usual exceptions: documented medical, emergency, etc.
  - *Talk to us early if there's a problem!*
- Two "late points" to use - one day each (still can't be more than 2 days late)
- Regrade requests in writing to course admin

# Why take this course?

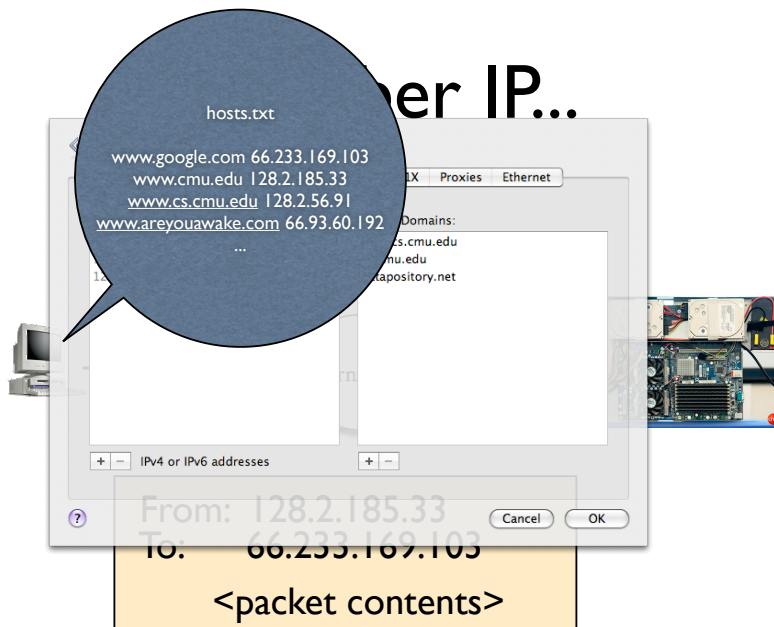
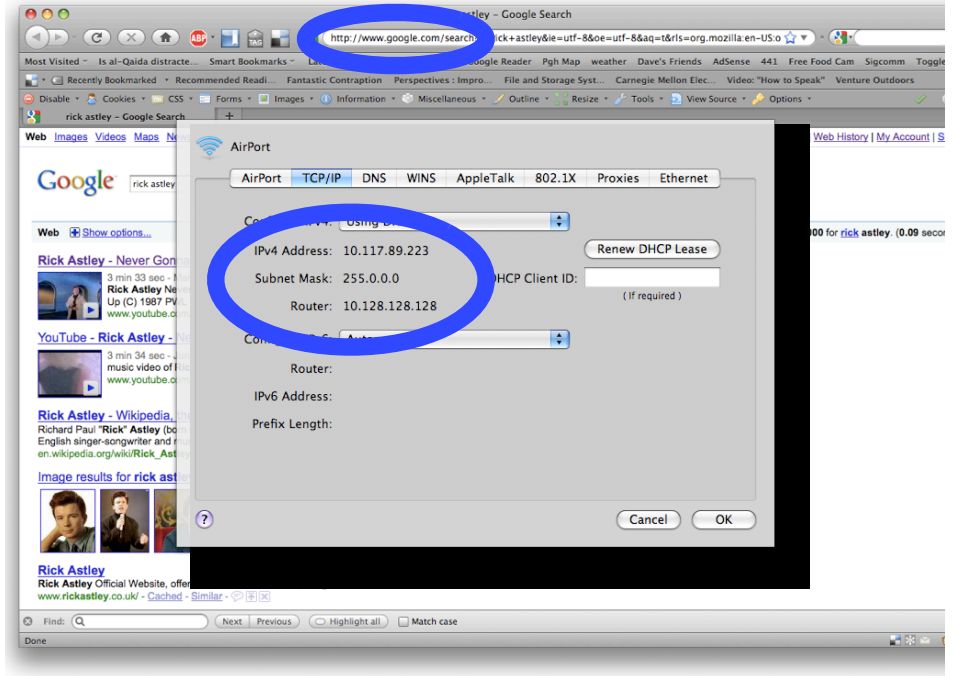
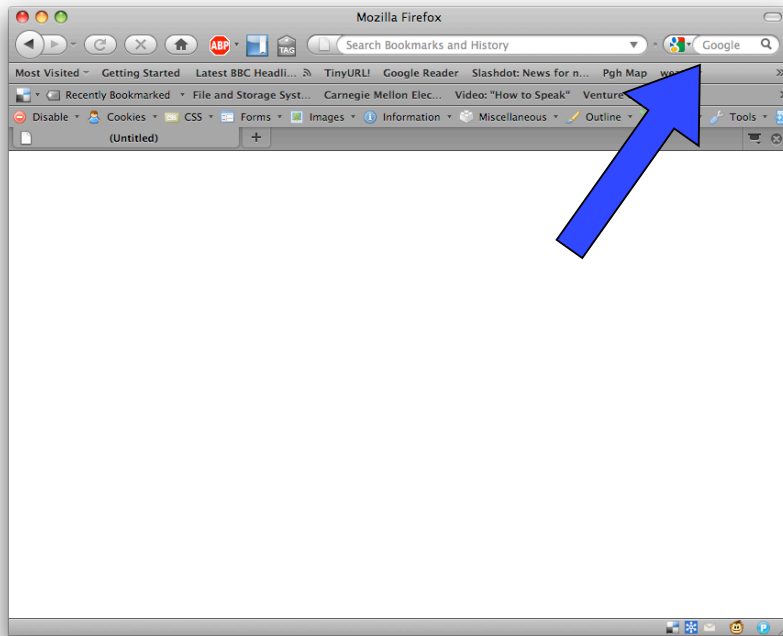
- Huge amounts of computing are now distributed...
  - A few years ago, Intel threw its hands up in the air: couldn't increase GHz much more without CPU temperatures reaching solar levels
  - But we can still stuff more transistors (Moore's Law)
  - Result: Multi-core.
  - Result 2: Your computer has become a parallel/distributed system. In a decade, it may have 128 cores.
- Oh, yeah, and that whole Internet thing..
  - my phone syncs its calendar with google, which i can get on my desktop with a web browser, ...
    - (That phone has the computing power of a desktop from 10 years ago and communicates wirelessly at a rate 5x faster than the average american home could in 1999.)
  - Stunningly impressive capabilities now seem mundane. But *lots* of great stuff going on under the hood...
  - Most things are distributed, and more each day

# If you find yourself ...

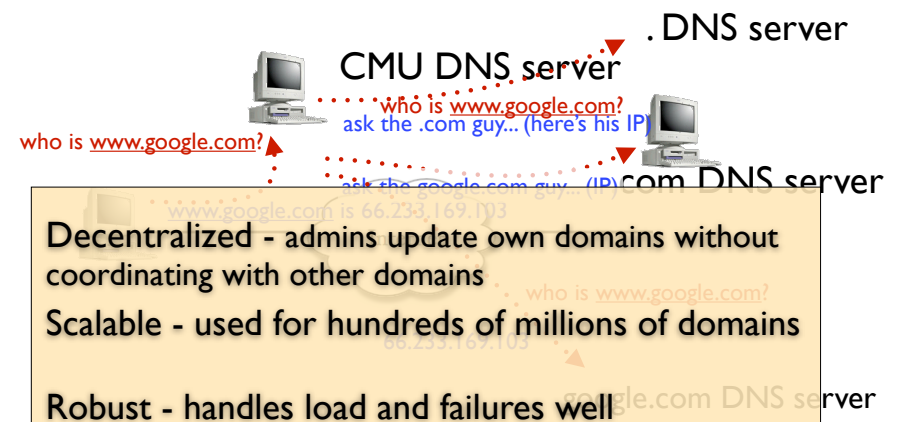
- In hollywood....
  - ... rendering videos on clusters of 10s of 1000s of nodes?
  - Or getting terabytes of digital footage from on-location to post-processing?
- On wall street...
  - tanking our economy with powerful simulations running on large clusters of machines
  - For 11 years, the NYSE ran software from cornell systems folks to update trade data
- In biochem...
  - using protein folding models that require supercomputers to run
- In gaming...
  - Writing really bad distributed systems to enable MMOs to crash on a regular basis
- not to mention the obvious places

# Enough advertising

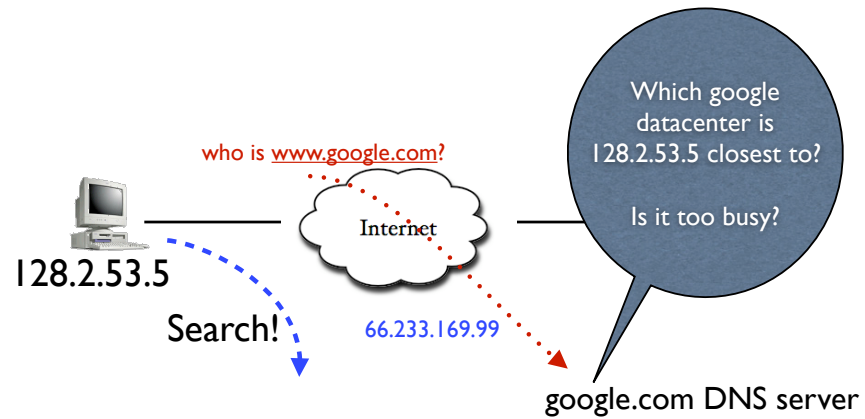
- Let's look at one real distributed system
- That's drastically more complex than it might seem from the web browser...



# Domain Name System



# But there's more...



## A Google Datacenter



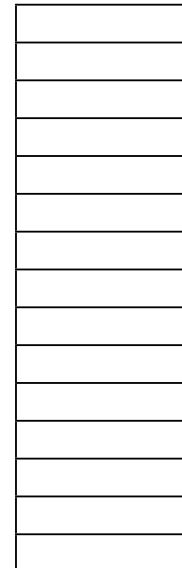
How big? Estimates range from 200-500k machines

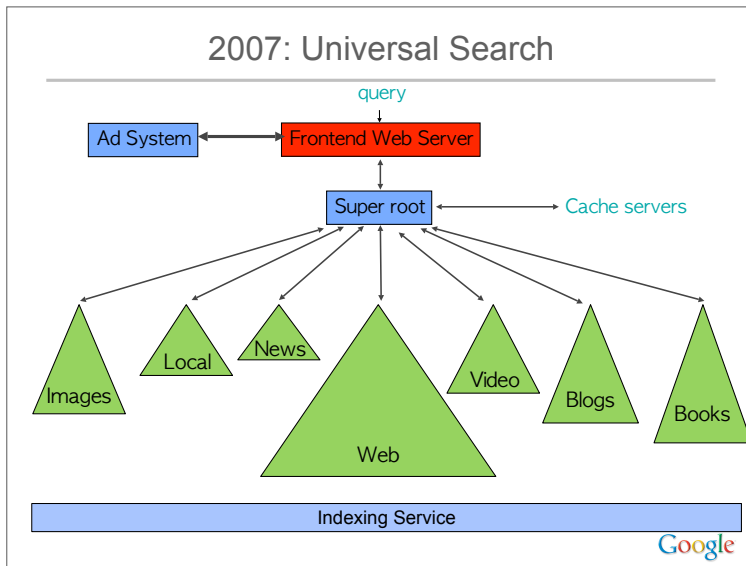
but it's not that bad...

usually don't use more than **20,000** machines to accomplish a single task.

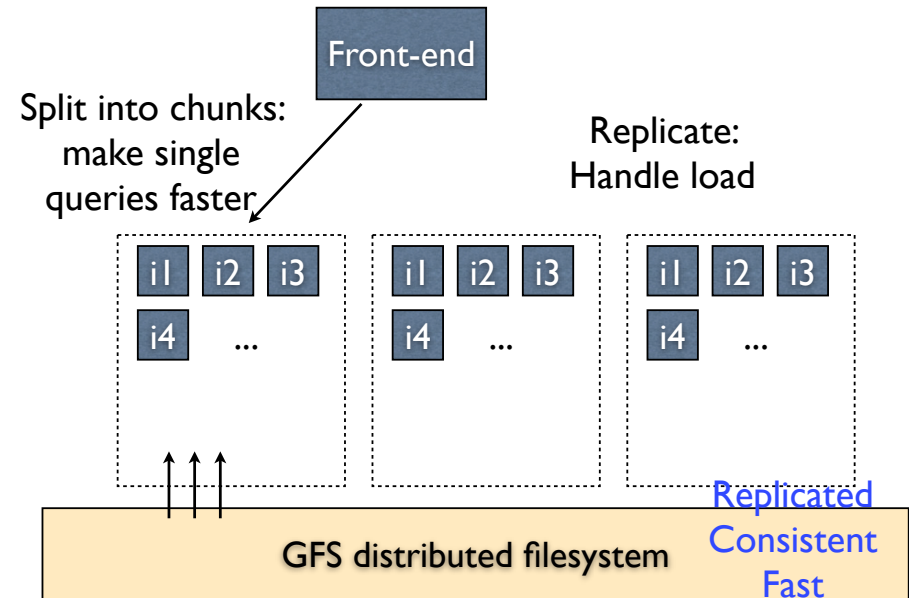
Rick Astley?

Front-end





slide from Jeff Dean, Google



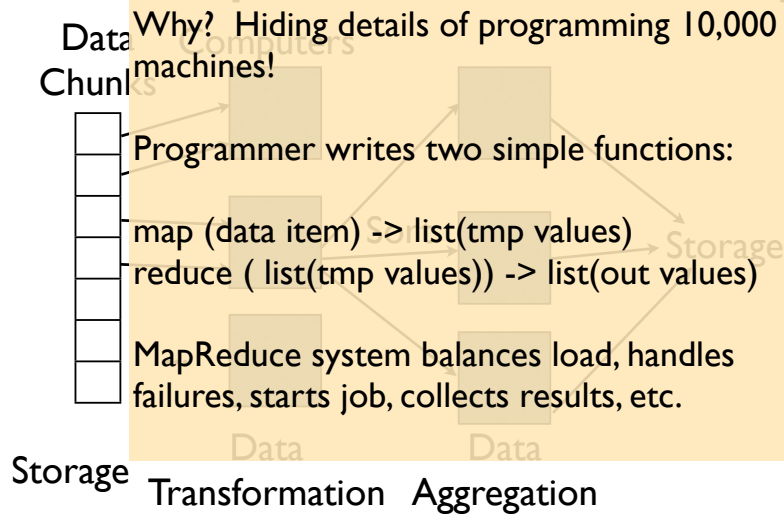
# How do you index the web?

1. Get a copy of the web.  
There are over 1 trillion unique URLs
2. Build an index.  
Billions of unique web pages
3. Profit.  
Hundreds of millions of websites  
30?? terabytes of text

=

- *Crawling* -- download those web pages
- *Indexing* -- harness 10s of thousands of machines to do it
- *Profiting* -- we leave that to you.
- “Data-Intensive Computing”

# MapReduce / Hadoop



## All that...

- Hundreds of DNS servers
- Protocols on protocols on protocols
- Distributed network of Internet routers to get packets around the globe
- Hundreds of thousands of servers
- ... to find one bad video in under 1/2 second

## How do you index the web?