CSCI 325: Distributed Systems

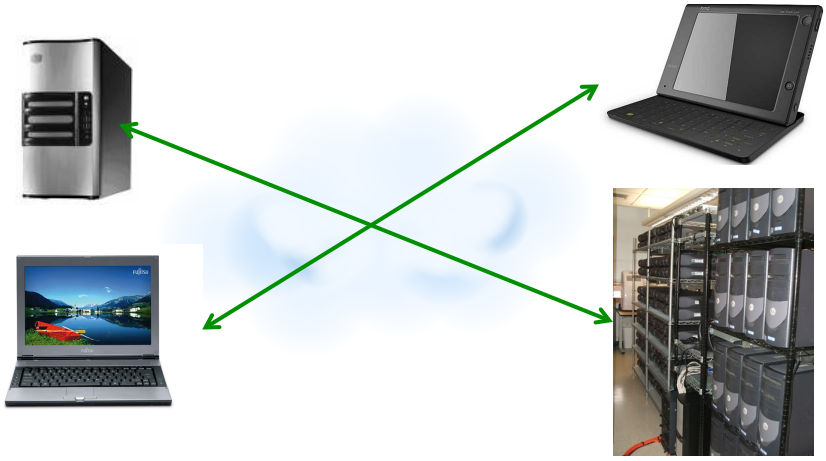

Objectives

- Course overview
- Overview of distributed systems
- Introduction to reading research papers

Distributed Systems?

- What is a distributed system?
- Know any examples of distributed systems?
  - Any used?

Distributed Systems?

- What is a distributed system?
  - Collections of independent, networked computers working together
- Examples of distributed systems
  - Networked printers, storage
  - Internet
  - Peer-to-peer systems
  - Grid computing
  - Games
  - Sensor networks

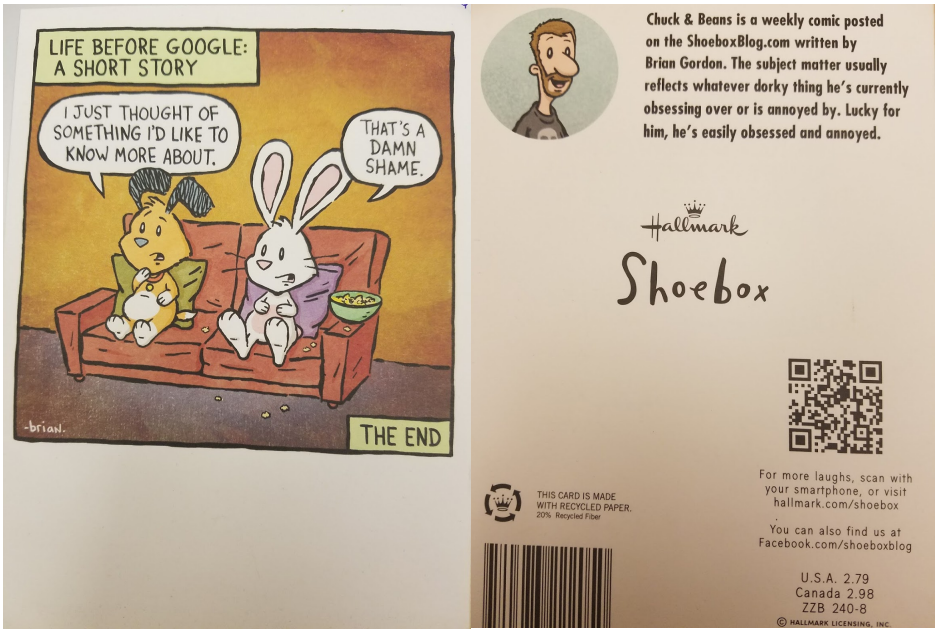

Distributed Systems Architectures

- Two main models
  - Client-server
    - Most common, and arguably the simplest
    - Multi-tier client-server
      - Variation on the simple client-server architecture
      - Multiple levels of communication
  - Peer-to-peer
    - All processes involved in a task or activity play similar roles

Client-Server Model

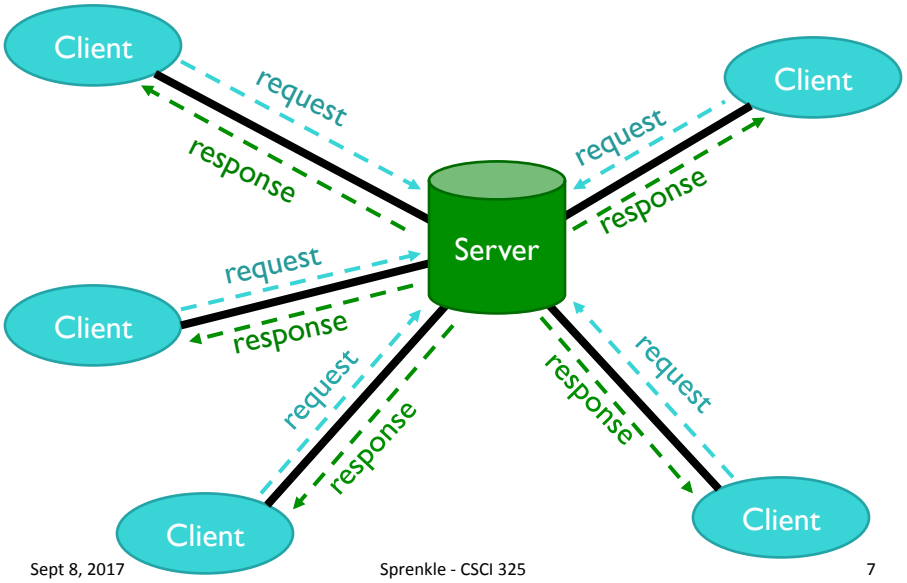

Peer-to-Peer Systems

Connections between peers

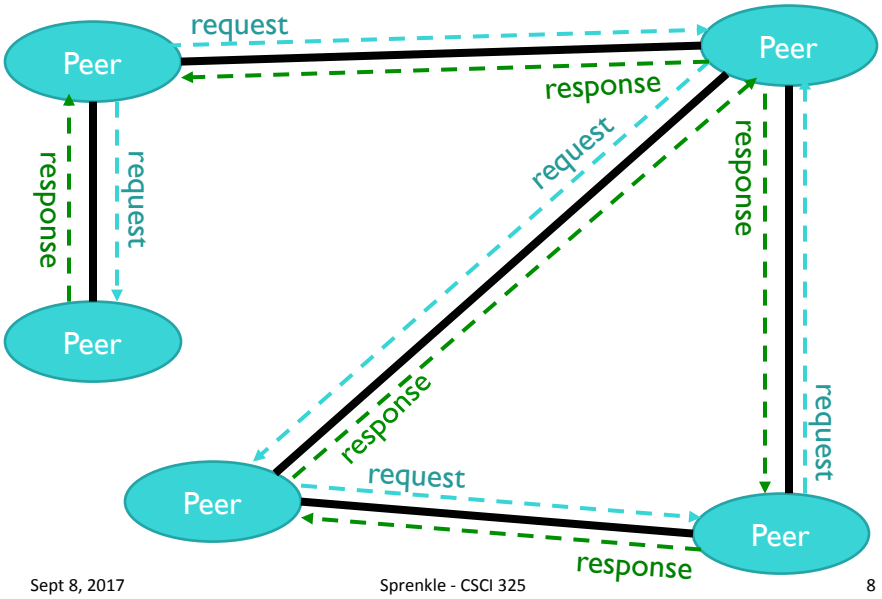

Challenges

- What are challenges in dealing with distributed systems?

Distributed Systems Challenges

- Communication
- Naming
- Distribution of workload
- Distribution transparency
- Consistency
- Handling failure
- Security
- Scaling

What This Course is About

- Networking fundamentals
- Distributed systems
  - Challenges of distributed systems
  - Design principles
  - Learn how to build large-scale distributed systems
    - Several programming projects
- Emerging research issues
  - Study fundamental research papers
- Life-skills
  - Reading, writing, discussion, presentation

Bonus: OS

Overall goal: Emphasize "*why*" and "*how*" over "*what*"

What made distributed systems possible?

A LITTLE BIT OF HISTORY

The Internet

- Connection of computer networks using the Internet Protocol (IP)
  - Allows network applications, e.g., email, file transfer, world wide web, remote login, …

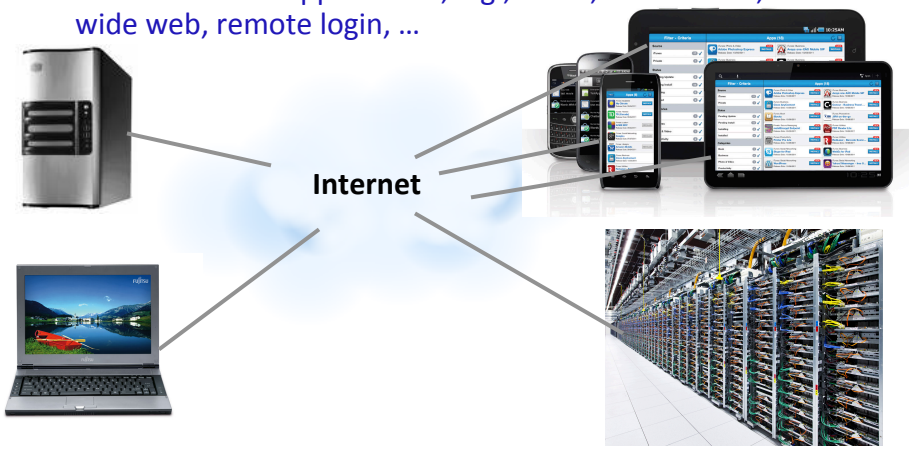

Vannevar Bush

- Established the U.S. military/university research partnership that developed ARPANET
- Wrote 1st visionary description of potential use for information technology
  - inspired many of Internet's creators

Source: Livinginternet.com

Could you envision the WWW years before it existed?

"Consider a future device for individual use, which is a sort of mechanized private file and library. It needs a name, and to coin one at random, 'memex' will do. A memex is a device in which an individual stores all his books, records, and communications, and which is mechanized so that it may be consulted with exceeding speed and flexibility. It is an enlarged intimate supplement to his memory.

"It consists of a desk, and while it can presumably be operated from a distance, it is primarily the piece of furniture at which he works. On the top are slanting translucent screens, on which material can be projected for convenient reading. There is a keyboard, and sets of buttons and levers. Otherwise it looks like an ordinary desk."

- Vannevar Bush, "As We May Think," *Atlantic Monthly*, July 1945

J. C. R. Licklider

Source: Livinginternet.com

- Joseph Carl Robnett "Lick" Licklider developed idea of universal network
- Spread his vision throughout the IPTO (Information Processing Techniques Office)
- Inspired his successors to realize his dream by creating ARPANET

"It seems reasonable to envision, for a time 10 or 15 years hence, a 'thinking center' that will incorporate the functions of present-day libraries together with anticipated advances in information storage and retrieval.

"The picture readily enlarges itself into a network of such centers, connected to one another by wide-band communication lines and to individual users by leased-wire services. In such a system, the speed of the computers would be balanced, and the cost of the gigantic memories and the sophisticated programs would be divided by the number of users."

- J.C.R. Licklider, *Man-Computer Symbiosis*, 1960.

Background

- 1957: USSR launches Sputnik, first artificial earth satellite
  - U.S. responds by forming Advanced Research Projects Agency (ARPA)
- 1962: Licklider's Galactic Network
- 1966: Marill and Roberts (MIT) paper: "Toward a Cooperative Network of Time-Shared Computers"
  - http://dl.acm.org/citation.cfm?id=1464336
- 1967: Roberts (MIT): ACM SOSP "Multiple Computer Networks and Intercomputer Communication"
  - http://dl.acm.org/citation.cfm?id=811680

1969 Internet Map: ARPANET

1st assignment:
draw today's Internet

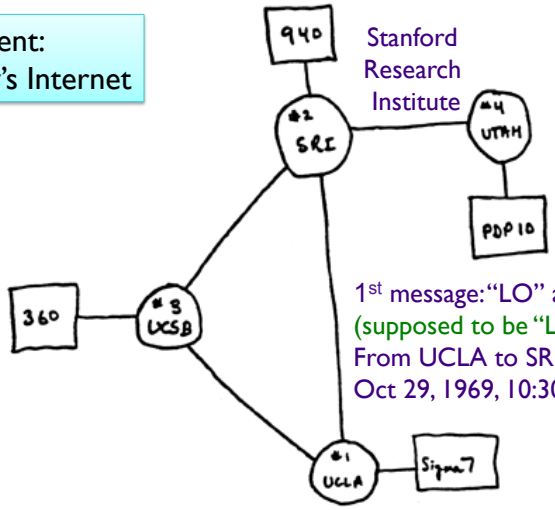

1st message: "LO" as in "Lo and Behold"
(supposed to be "LOG" but failure!)
From UCLA to SRI
Oct 29, 1969, 10:30 p.m.

SDS Sigma 7

http://www.nsf.gov/news/special_reports/nsf-net/1960s.jsp
https://www.nsf.gov/news/special_reports/nsf-net/kleinrockvideopop.html

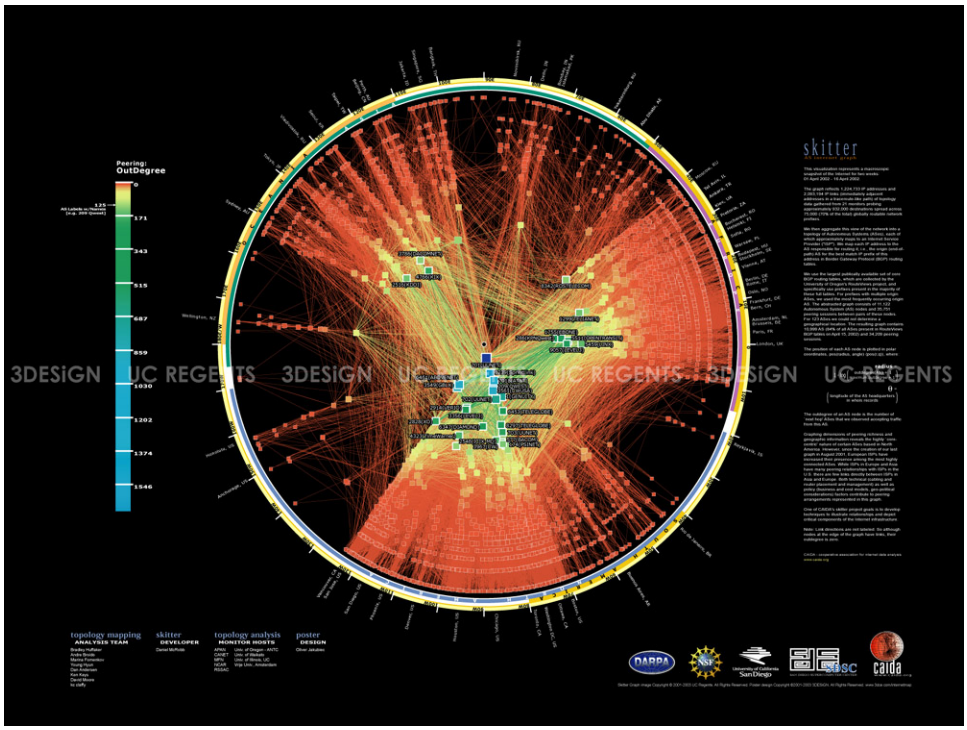

Internet Timeline

Year	Milestone
1971	Tomlinson develops email program, big hit
1972	Telnet
1973	File Transfer Protocol (FTP)
1974	Transmission Control Protocol (TCP)
1978	TCP split into TCP and IP (Internet Protocol)
1979	USENET (newsgroup) established
1984	1000 hosts connected to Internet, DNS introduced
1988	Internet worm brings down 10% of Internet
1991	WAIS, Gopher, WWW released

Internet Growth Trends

Year	Hosts on Internet
1977	111
1981	213
1983	562
1984	1000
1986	5000
1987	10,000
1989	100,000
1992	1,000,000
2001	151-175 million
2002	Over 200 million
2014	1.01 billion

# of computers connected directly to the Internet increased at a yearly rate >37% across 21 years

https://www.internetsociety.org/sites/default/files/Global_Internet_Report_2014_0.pdf

Statistics from the IITF Report "The Emerging Digital Economy" *

- To get a market of 50 Million people participating:
  - Radio: 38 years
  - TV: 13 years
  - Internet: 4 years
    - After open to general public

- http://govinfo.library.unt.edu/ecommerce/EDEreprt.pdf
  - Released on April 15, 1998

* Delivered to the President and the U.S. Public on April 15,1998 by Bill Daley, Secretary of Commerce and Chairman of the Information Infrastructure Task Force

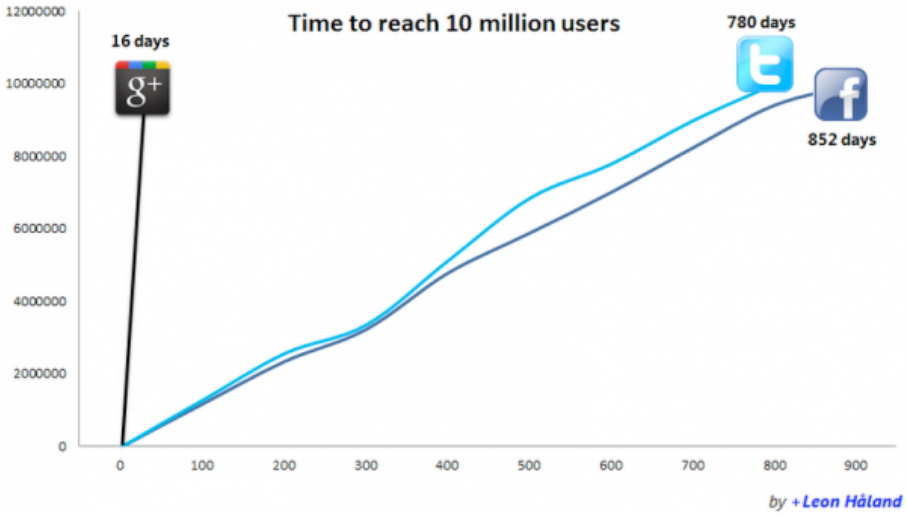

COURSE INFO

My Responsibilities

- Prepare useful, interesting knowledge
- Come to class prepared, on time
- Interesting, relevant, and challenging assignments
- Prompt feedback on assignments

Your Responsibilities

- Come to class prepared, on time, and PARTICIPATE
- Turn in assignments *on time*
- When you're having trouble
  - Look for help on the Web
    - Find, adapt solutions
    - Give credit to where you found solution, if novel enough
  - Ask me for help!
- Learn, absorb, synthesize
  - Extra Credit: take it to the next level

Textbook

- **Required**: Distributed Systems, by van Steen and Tanenbaum, 3rd ed.
  - Provides background for class discussions and projects
  - Available online
- Optional: Distributed Systems, Concepts and Designs, by Courlouis, Dollimore, Kindberg, 5th ed.

Grading

- 17% Individual programming, reading, writing assignments
- 20% Midterm exam
- 33% Programming projects
- 25% Final Project
  - Including paper and presentation
  - Start thinking about possible topics
- 5% Participation and attendance
  - Success of class depends on student participation

Programming Projects

- 3 projects spanning the semester
  - Hands-on construction of interesting distributed services
  - Approximately 2.5 weeks to complete
  - Work in teams of 2 or 3
    - Use version control
  - Start early!

READING RESEARCH PAPERS

What to Look For While Reading

- Overall problem
  - How large/important is the problem?
- Goals
- Contributions
  - Keywords: new, novel
- Technical approach
  - Key insights ("leverage", "utilize")
- Evaluation
  - Answers all your questions about approach?
- Limitations
  - May not be a general-purpose solution
  - Check assumptions

Some Concrete Questions

- Statement of the Problem/Goals
  - Try to state succinctly the overall problem being addressed in this paper.
  - What particular goals do these researchers have in addressing this problem?
  - What contribution are they seeking to make to the state-of-the-art?
- Technical Approach
  - What is the key insight of this group's approach to tackling the stated problem? What is their overall approach/strategy to solving the problem?
- Discussion/Critique
  - How did the researchers evaluate their efforts?
  - What conclusions did they make from their evaluation results?
  - What application/useful benefit do the researchers/you see for this work?
  - What limitations do the researchers mention with their approach?
  - What additional limitations do you think there are?
  - Write one interesting question to ponder with regard to this paper beyond content understanding.

SEDA

We propose a new design for highly concurrent Internet services, which we call the staged event-driven architecture (SEDA). SEDA is intended to support massive concurrency demands and simplify the construction of well-conditioned services. In SEDA, applications consist of a network of event-driven stages connected by explicit queues. This architecture allows services to be well-conditioned to load, preventing resources from being overcommitted when demand exceeds service capacity. SEDA makes use of a set of dynamic resource controllers to keep stages within their operating regime despite large fluctuations in load. We describe several control mechanisms for automatic tuning and load conditioning, including thread pool sizing, event batching, and adaptive load shedding. We present the SEDA design and an implementation of an Internet services platform based on this architecture. We evaluate the use of SEDA through two applications: a high-performance HTTP server and a packet router for the Gnutella peer-to-peer file sharing network. These results show that SEDA applications exhibit higher performance than traditional service designs, and are robust to huge variations in load.

SEDA

- Problem/Goals
  - Highly concurrent internet systems
    - Goal: well-behaved under load
- Technical Approach
  - Staged, event-driven architecture (SEDA)
  - Automatic tuning, load conditioning
- Discussion
  - Evaluation: Used SEDA architecture for web server, P2P packet router
    - Measured performance, robustness to load variation

Reading Feedback: Annotations

- Perusall: Application accessible through Sakai
  - Allows you to comment on an article such that all students and professor can view them
- You will be expected to make a certain number of annotations on each article
  - "Certain number" = 5, typically
- Annotations can be questions or comments
  - must be substantive
- Each annotation will be graded as
  - 2: thoughtful; full-credit
  - 1: partial-credit
  - 0: thoughtless or not complete; no credit

TODO

- Set up Perusall, through Sakai
- Explore Course Web Page
- Check out "Welcome to the Machine"
  - Reviewing some terms from CSCI210 (plus maybe more)
- Read E2E Argument paper for Friday
  - Skim through once, review section headings
  - 3 hours max
  - Review paper
    - Write 5 annotations in Perusall
  - Wed: Discuss paper and questions