CS 268: Graduate Computer Networking

Sylvia Ratnasamy

(Spring, 2013)

Today

- Overview
  - What is networking research about?
  - What you will learn in CS268
- Logistics
- Recap networking basics
  - Give you a sense of what I'm assuming you know

The Internet is transforming everything

- The way we do business
  - E-commerce, advertising, cloud-computing
- The way we have relationships
  - Facebook friends, E-mail, IM, virtual worlds
- The way we learn
  - Wikipedia, MOOCs, search engines
- The way we govern and view law
  - E-voting, censorship, copyright, cyber-attacks

Started as a research experiment!

But what *is* networking about?

“What’s your formal model for the Internet?” -- *theorists*

“Aren’t you just writing software for networks” – *OS community*

“Isn’t it just another distributed system?” – *PODC community*

“You use simulation for evaluation????” – *all*

“What’s with all these TLA protocols?” – *all*

“But the Internet seems to be working now…” – *my parents*

A few defining characteristics of the Internet

A federated system

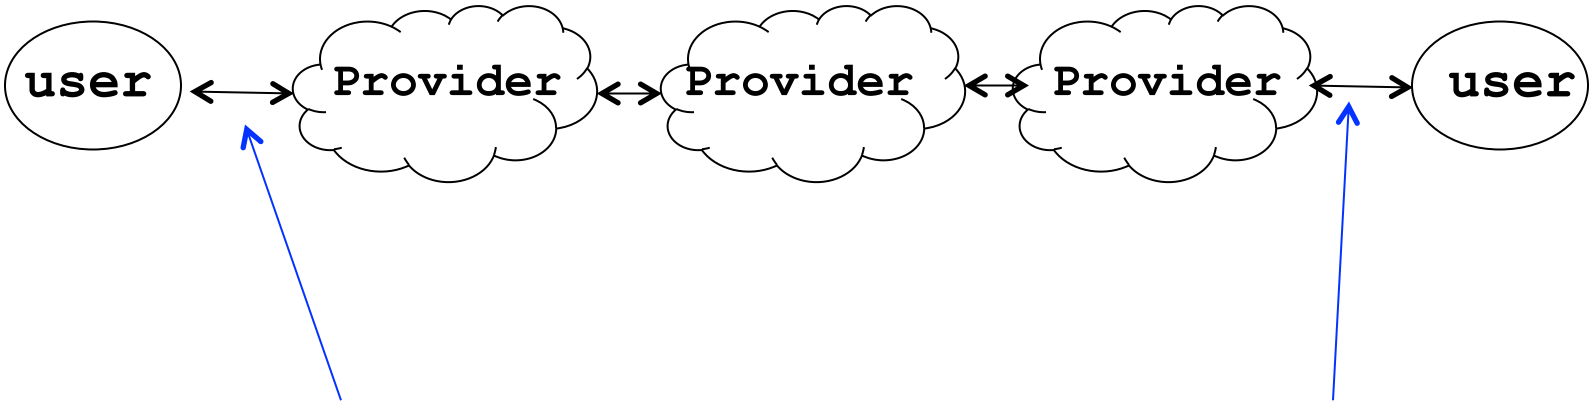

- IP as the common interface between users and the network
- IP hides the inner workings of the federation from the user

A federated system

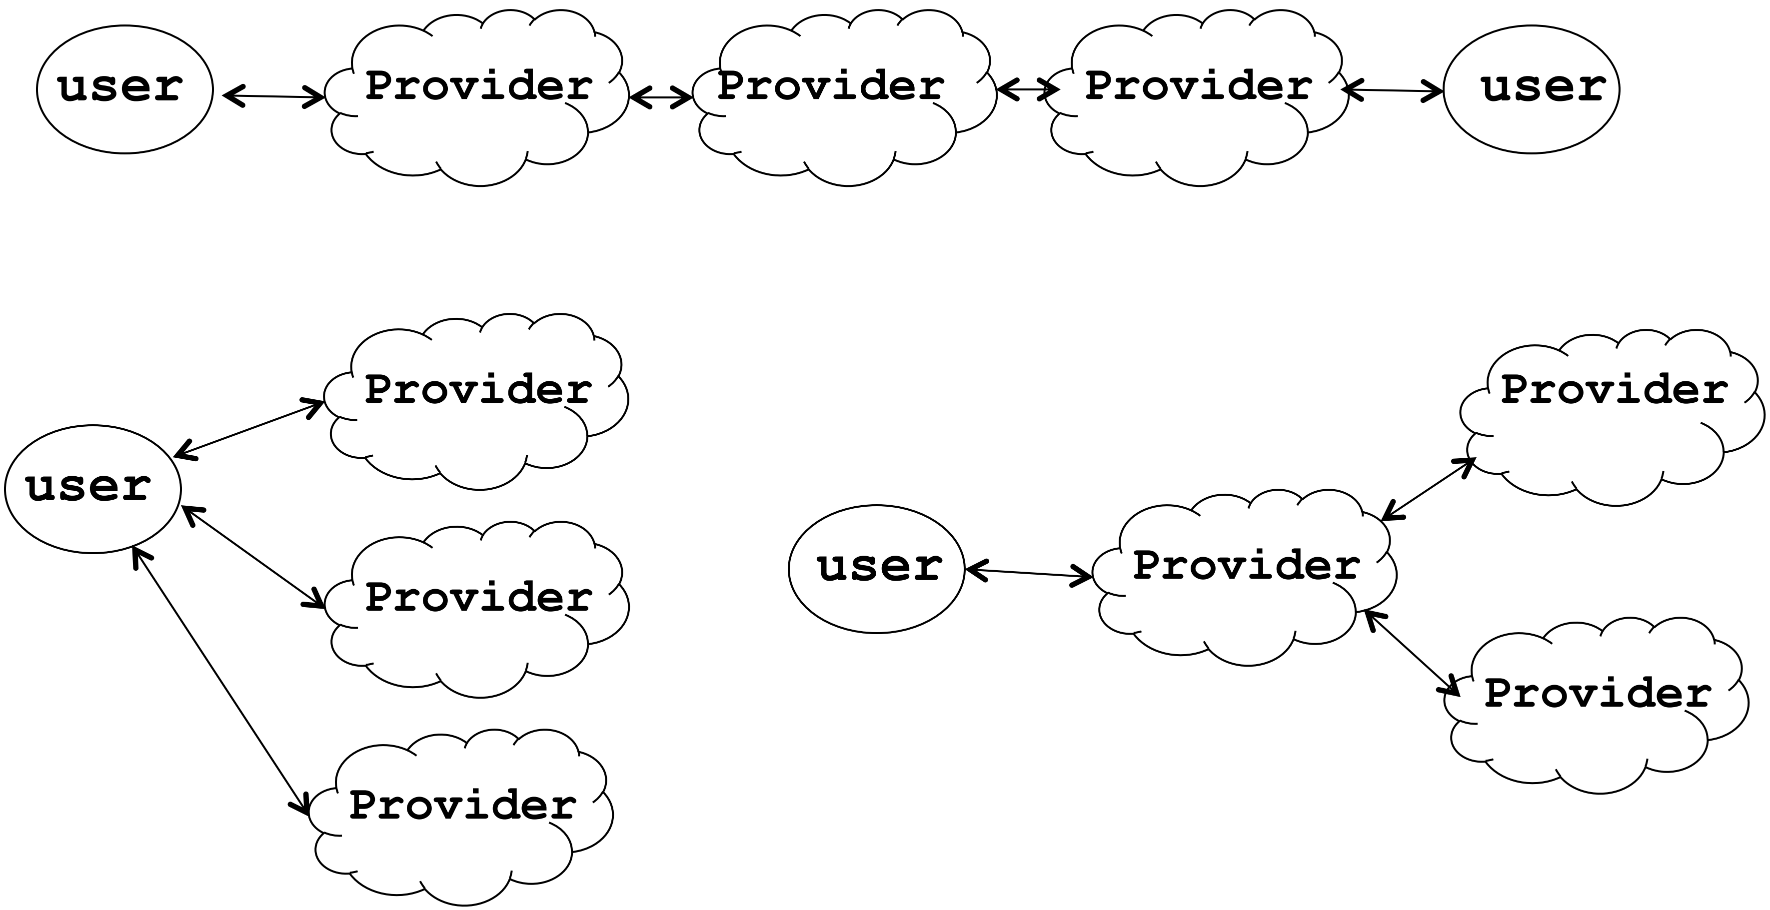

A federated system

- Parties with different goals, policies and constraints cooperate to provide an *identical* service interface

- Complicated by the fact that this federation is “opaque”
  - User doesn’t “see” the effect of individual providers
  - Provider A doesn’t see what Provider B does with A’s packets

- Practical realities of incentives, peering economics and real-world trust drive topology, route selection and service evolution

Tremendous scale

- 2.4B users (34% of world population)
- 1 Trillion unique URLs
- 294B emails sent per day
- 1B smartphones
- 937M Facebook users
- 2B YouTube videos watched per day
- Routers that switch 100Tbps
- Links that carry 100Gbps
- #Links = ???

Enormous diversity and dynamic range

- Latency: μseconds to seconds (10^6)
- Bandwidth: kbps to 10s of Gbps (10^7)
- Queuing delay: from ~0 to seconds
- Packet loss: 0 – 90%
- End-systems: from sensors and cell phones to datacenters and supercomputers
- Technology: optical, wireless, satellite, copper
- Applications: 40B to 5 GB transfers; live TV to backup; critical services to entertainment
- **Users**: the governing, governed, operators, malicious, naïve, savvy, embarrassed, paranoid, addicted, cheap …

Constant Evolution

1970s:

- 56kbps “backbone” links
- <100 computers, a handful of sites in the US
- Telnet and file transfer are the “killer” applications

Today

- 100+Gbps backbone links
- 5B+ devices, all over the globe
- Countless apps; 20M Facebook apps installed per day

Asynchronous Operation

- Fundamental constraint: **speed of light**
- Consider:
  - How many cycles does your 3GHz CPU in Berkeley execute before it can possibly get a response from a message it sends to a server in NY?
    - Berkeley to New York: 4,125 km
    - Traveling 300,000 km/s: 13.75 msecs
    - 3,000,000,000 cycles/sec * 0.028 = 84,000,000 cycles!
- Thus, communication feedback is always *dated*

Prone to Failure

- To send a message, **all** components along a path must function correctly
  - software, links, network interface cards, switch fabrics; at endpoints, wireless access points, modem, switches, routers, firewalls, proxies..
  - Including human operators
- Consider: 50 components, that work correctly 99% of time → 39.5% chance communication will fail
- Plus, recall
  - scale → lots of components
  - asynchrony → takes a long time to hear (bad) news

An Engineered System

- Constrained by what technology is practical
  - Link bandwidths
  - Switch port counts
  - Bit error rates
  - Cost
  - …

Recap: The Internet is…

- A complex federation
- Of enormous scale
- Dynamic range
- Diversity
- Constantly evolving
- Asynchronous in operation
- Failure prone
- Constrained by what's practical to engineer

Recap: The Internet is…

- Too complex for theoretical models
- “Working code” doesn’t mean much

So, what do we need?

We still don’t really know…

- No consensus on what constitutes the “correct” or “best” network architecture
- No consensus on “top 10 problems”
- No consensus on the right prioritization of goals

Before you flee…

What we do know

- The early Internet pioneers came up with a solution that was successful beyond all imagining
- Several enduring architectural principles and practices emerged from their work

Some key principles

- Packets
- Statistical multiplexing
- Best-effort service
- Protocol layers
- A “narrow waist” at the network layer
- Stateless (“dumb”) network and smart end-hosts
- Decentralization

What we do know

- The early Internet pioneers came up with a solution that was successful beyond all imagining
- Several enduring architectural principles and practices emerged from their work
- But it is just one design
- And numerous cracks have emerged over time
  - No accountability or verification
  - No hooks for traffic engineering, diagnostics
  - No visibility into network conditions
  - Complex and buggy protocols

What we do know

- The early Internet pioneers came up with a solution that was successful beyond all imagining
- Several enduring architectural principles and practices emerged from their work
- But it is just one design
- And numerous cracks have emerged over time
- As have new requirements
  - Mobility, security, manageability, reliability, low latency, advanced processing, scale, scale, scale…

Hence, Networking Research Today..

- Packets → McKeown "circuits"
- Statistical multiplexing → Vahdat "TDMA"
- Protocol layers
- A "narrow waist" at the network layer → Anderson "Pluralism"
- Best-effort service
- Stateless ("dumb") network and smart end-hosts
- Decentralization → 4D/SDN "centralize"

How does this list change?

Internet Research

- Epicenter is SIGCOMM and HotNets
  - Also NSDI, SOSP/OSDI, Mobicom, CoNeXT, Infocom
- Architectural discussions tend(ed?) to dominate
  - But also new technologies (optical routers), applications (P2P), discoveries (topology), analysis (route stability) and models (economic, performance)
- Constant evolution means there's always ongoing debate
  - But also recognition that the need for interoperability and standardized protocols means we need to arrive at consensus

A Random Sample of Internet Researchers

- Point to note is the diversity in expertise and contributions

Internet Researchers

- **Closer to home**

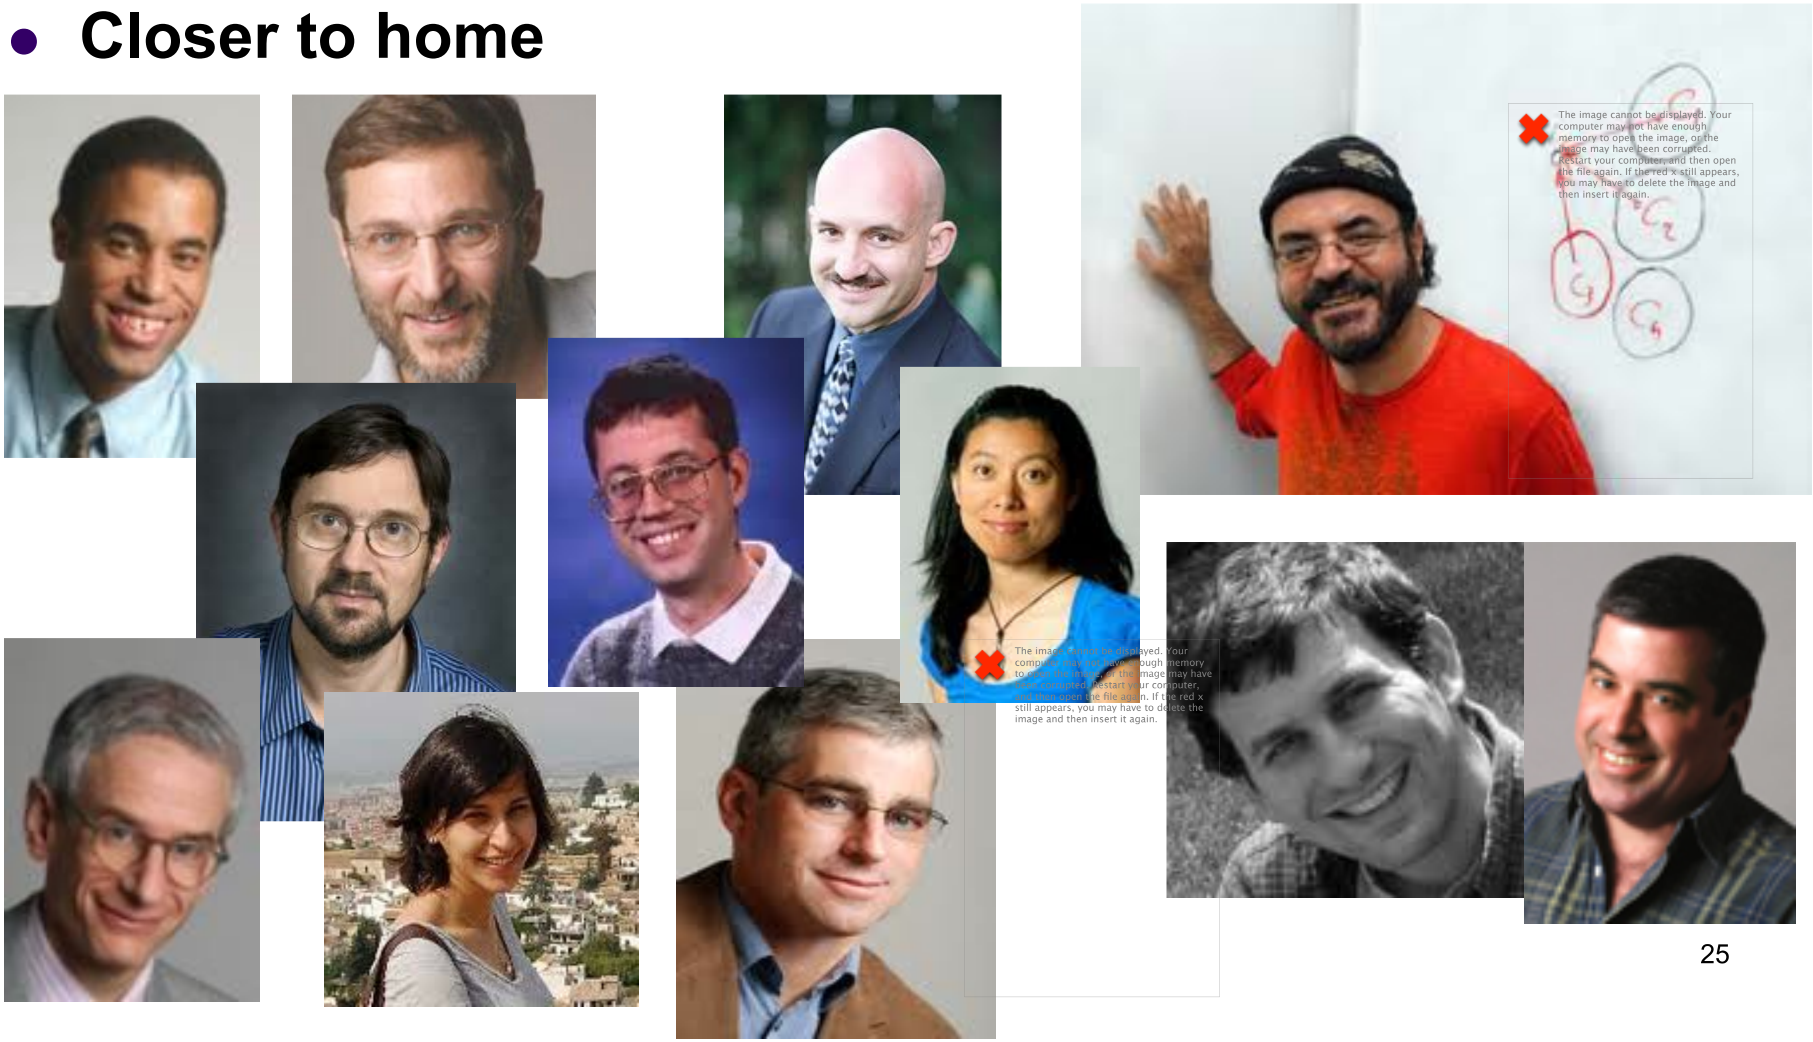

Can Internet Research Still Have Impact?

- My take: yes!
  - 201x: SDN
  - 201x: Multipath TCP
  - 200x: Fat trees, VL2
  - 200x: Nemo (Net memory)
  - 200x: NetSift (net fingerprints)
  - 200x: WAN optimizers (Spring-Wetherall de-duplication)
  - 200x: DHTs (BitTorrent, Azure Storage Tables, Amazon's Dynamo)
  - 199x: Lulea/TreeBitMap lookup
  - 199x: PIM (Core-Based Trees)
  - 199x: Akamai (distr. caching)
  - 199x: NAT
  - 199x: Router congestion signals (ECN, RED)
  - Earlier: BGP, DNS, TCP, IP…

Can Internet Research Still Have Impact?

- My take: yes!
  - A of B most cited papers in CS are Sigcomm papers
  - A of B most cited authors working in networking

Can Internet Research Still Have Impact?

- My take: yes!

- But to a degree
  - The core protocols – IP/BGP – have proven hard to change
    - A number of networking researchers lose sleep over this
  - My take: this isn't a bad thing
    - Provided change is *possible* (some lose sleep over this too)

In Short

- Networking research is fun
  - Real, important problems
  - Opportunities for impact
  - Inherently interdisciplinary
- But the field is young
  - Often ad-hoc, often poorly defined
- And success has brought constraints

This Class

- A bit of architecture (principles, service models)
- A bit of technology (algorithms, infrastructure, software)
- A bit of how things work (BGP, congestion control)
- A bit of the classics (QoS, Active Networks)
- A bit of new directions (datacenters, homenets)

- Topics that are given short shrift
  - Wireless, security (see Vern's class), datacenters (Randy's class)

Let’s take a 5 minute break.

Logistics

- Sylvia Ratnasamy (sylvia@eecs.berkeley.edu)
  - 413 Soda Hall
  - Office Hours: Wednesdays 3-5pm

- Sorry, no Teaching Assistant!

- Course Info: http://www.cs.berkeley.edu/~sylvia/cs268/

Goals of this Course

- Understand
  - How does the Internet work?
  - What are the Internet's design principles?
  - How and why is the Internet changing?
- Get familiar with Internet research efforts
  - papers that changed practice; e.g., [J88, OFlow]
  - papers that changed how we think; e.g., [ESM, NoFair]
  - papers that changed what we think about; e.g., [Active, DDiff]
  - papers that brought clarity; e.g., [Design, E2E]

Goals of this Course (cont’d)

- Appreciate what is good research
  - Problem selection
  - Solution & research methodology
  - Presentation
- Identify gaps and flaws in the state-of-the-art
  - requires deep reading and active class participation!
- Apply what you learned in a class project

What Do You Need To Do?

- Paper readings, reviews & discussion
- Paper presentation
- A research-oriented class project

Grading

Term project	50%
Paper reviews and class participation	25%
Paper presentation	25%

- Can miss up to 3 paper reviews with no penalty
- Frequent absenteeism will affect your grade

Reading papers

- Plenty of advice out there (see website for links)
- My take: don't overthink it
  - Read the paper start to end. Carefully
  - Set the paper aside and think
    - Replay their motivating arguments – do they make sense?
    - Replay how their solution works for a simple example -- can you?
    - Once you're sure you understand, start critiquing
      - Is the problem important? ambitious? hard? have a long shelf-life?
      - Is the solution effective? Under what conditions does it break?
      - What other approaches are possible? Etc.

Reviewing papers

- Write a short review for each paper
  - Length: 2-3 paragraphs
  - Be honest
- Email your reviews to me by 9pm the night before class
  - See the course website for detailed instructions

Five Questions to Answer in your Review*

- What is the problem?
- Do you believe the problem is/was important?
  - Explain your thinking
  - Consider context
- What is the solution's main idea (nugget)?
- Do you think the solution is a good one?
  - Explain. Note: "yes" is a perfectly acceptable answer!
- Did you enjoy the paper?

Class Discussion

- Attendance is necessary but not sufficient
- Come prepared to discuss the main ideas
- Come prepared with your questions on things you didn’t understand or believe
- We will all learn from each other
- Let’s have no open laptops in class

Presenting Papers

- Each of you will be responsible for presenting and leading the discussion on one paper
- Go through the reading list (http://www.cs.berkeley.edu/~sylvia/cs268/) and send me the list of your 10 most preferred papers* by midnight Thu, January 24
- I'll do the match and send you the assignment on Friday, January 25
  - Presentations will start next week, or 02/04

Lecture Format

- 25-30 minutes per assigned reading (you lead)
  - ~10-15 minutes recap
  - ~10-15 minutes discussion
- 20 minutes to fill in the gaps (I lead)
  - key ideas from papers we didn't read*
  - context/background for next lecture's readings
  - project ideas, etc.
- 4 guest lectures (no student presentations)
- The above are guidelines, not rules
- Discussion style and participation will dominate!

Presenting papers (cont'd)

- Organize presentation to suit your style
  - Summarize-then-discuss, discuss-as-you-go, ...
- Some tips/expectations
  - Assume the class has read the papers →recap, don't explain
  - Go beyond the assigned paper
    - optional readings (for sure) and more
    - 1-2 papers don't represent the body of work on a topic; your job is to give the class a more complete picture
  - Come prepared with a set of questions to initiate discussions
- recall: this is 25% of your grade

Research Project

- Investigate new ideas and solutions in a class research project
  - Define the problem
  - Execute the research
  - Work with your partner
  - Write up and present your research

- Aim high! Ideally, best projects will become conference papers (e.g., SIGCOMM, NSDI, SOSP)

Research Project: Steps

- I will distribute a list of projects
  - Choose one of these or come up with your own
- Pick your project, partner, and submit a 1-2 page proposal describing: due 02/25 (5pm)
  - The problem you propose to solve
  - The type of results you hope to obtain
  - Some discussion of related work
- A midterm presentation of your progress (8-10 minutes): due 03/20
- A final project presentation: 04/29 and 05/01
- Final project writeup: due 05/14 (5pm)
- Start early!

Recap: This Course

- We will read and discuss ~40 papers
  - Typically, 2 papers/class
- Three components to your grade
  - Paper reading, reviews and discussion (25% of grade)
  - Paper presentation (25%)
  - Project (50%)

Recap: Your Action Items

- By this Thursday (January 24) send me the top 10 papers you'd like to present

- If you haven't yet registered and plan to take this course, let me know today

Any questions?

Overview

- Internetworking
- Packets
- Broadcast vs. Switched
- Packet vs. Circuit switched
- Layers
- Protocols
- Best-effort delivery
- Routing
  - Distance Vector
  - Link State
- Congestion control
- Reliable data transfer
- Flow control
- DNS

Acknowledgments

These slides are inspired by, or borrow directly from, the lecture notes of various folks including Ion Stoica, Scott Shenker, Jennifer Rexford, Nick McKeown, Brad Karp and Kyle Jamieson, Arvind Krishnamurthy and Hari Balakrishnan.