CS 168
Introduction to the Internet: Architecture and Protocols

Fall 2014

Sylvia Ratnasamy

http://inst.eecs.berkeley.edu/~cs168/fa14/

Today

- Introductions
- What is (this course on) networking about?

  5 minute break

- Class policies, administrivia and roadmap

Introductions

Teaching Assistants

- Peter Gao
- Sangjin Han (co-*head TA*)
- Anurag Khandelwal
- Chang Lan
- Murphy McCauley
- Akshay Narayan
- Radhika Mittal (*co-head TA*)
- Shoumik Palkar
- Qifan Pu

See the course website for TA office hours and sections

Peter Xiang Gao

- 2nd year PhD student
- Research focus
  - fault-tolerant network devices

Sangjin Han

- Co-Head TA
- 4th year PhD student
- Research focus:
  - high performance network software

Anurag Khandelwal

- 2nd year PhD student
- Research focus
  - cloud computing

Chang Lan

- 2nd year PhD student
- Research focus
  - programmable routers

Murphy McCauley

- 2nd year PhD student
- Research focus
  - Software-Defined Networking

Radhika Mittal

- Co-Head TA
- 3rd year PhD student
- Research focus:
  - network congestion control and queue management

Akshay Narayan

- Senior
- Research focus
  - Congestion control in datacenters

Shoumik Palkar

- Senior
- Research focus
  - Programmable routers

Qifan Pu

- 2nd year PhD student
- Research focus
  - Distributed systems

Instructor: Sylvia Ratnasamy

- Ph.D. in Computer Science from UCB in 2002
- Worked at Intel between 2002-2011; back at UCB since 2011
- Networking has been my research focus throughout

- My teaching style
  - I talk too fast
  - The more bored you look, the faster I talk
  - So, stop me with questions!!

- Office hours: 4:30-5:30pm Thursday in 413 Soda Hall
  - Always happy to chat if you have a problem (send email)

What is networking about?

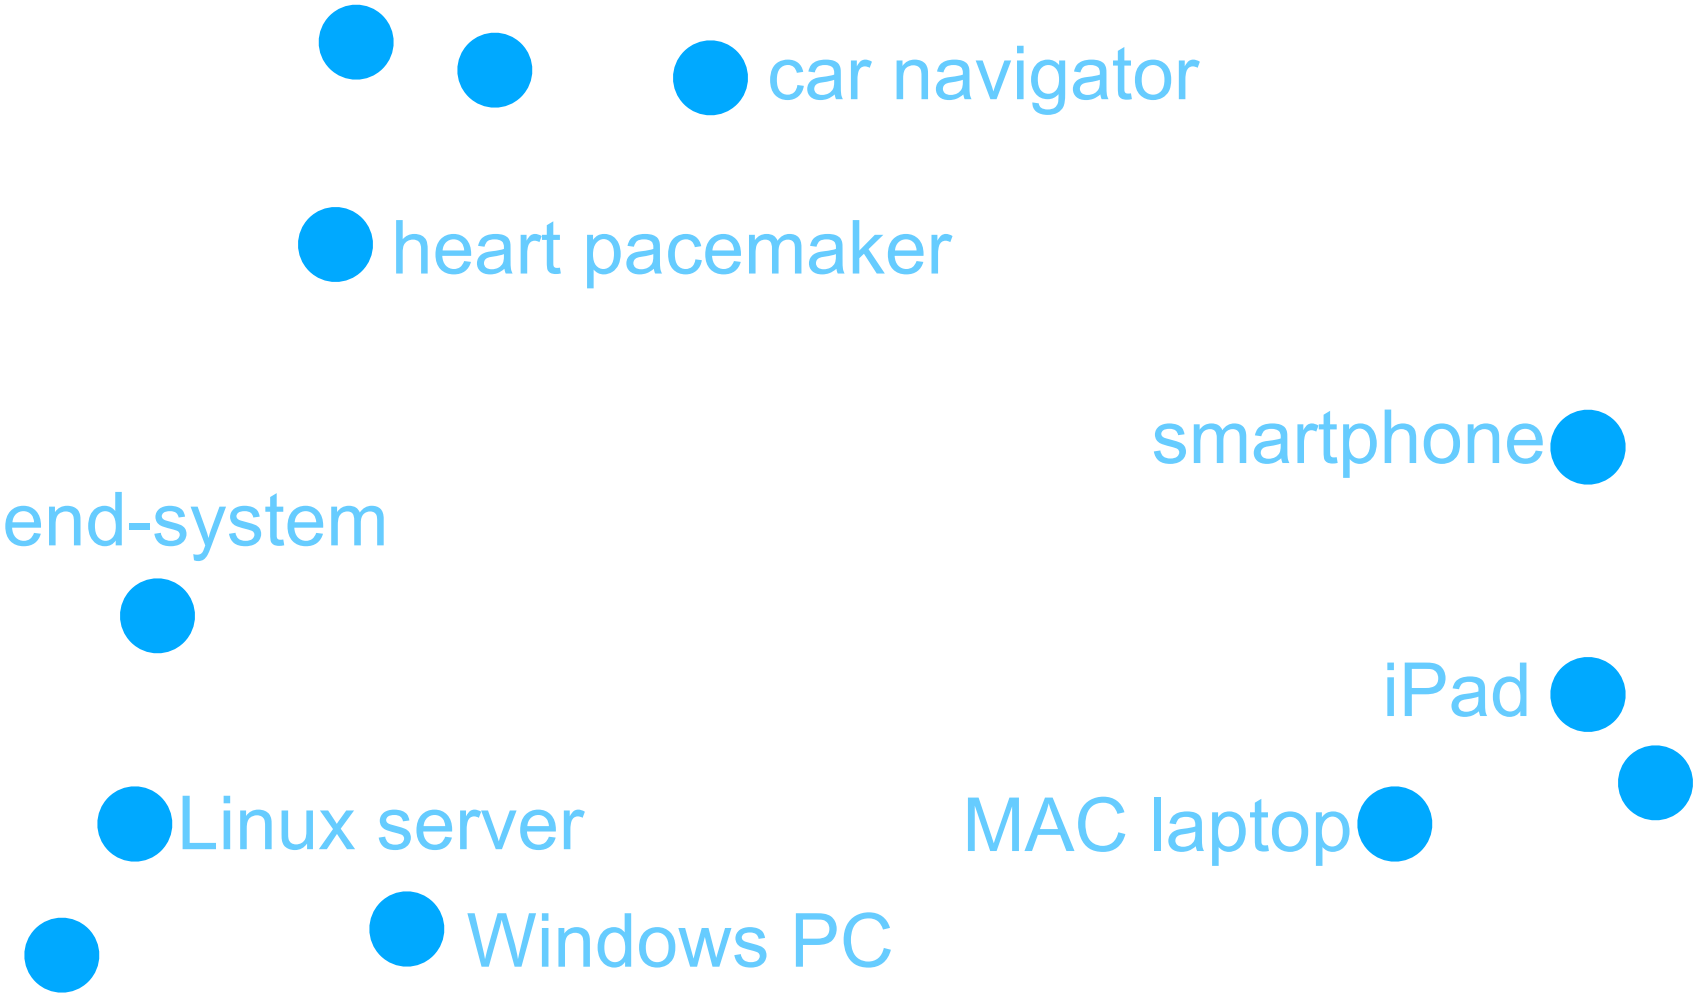
car navigator
heart pacemaker
smartphone
end-system
iPad
Linux server
MAC laptop
Windows PC

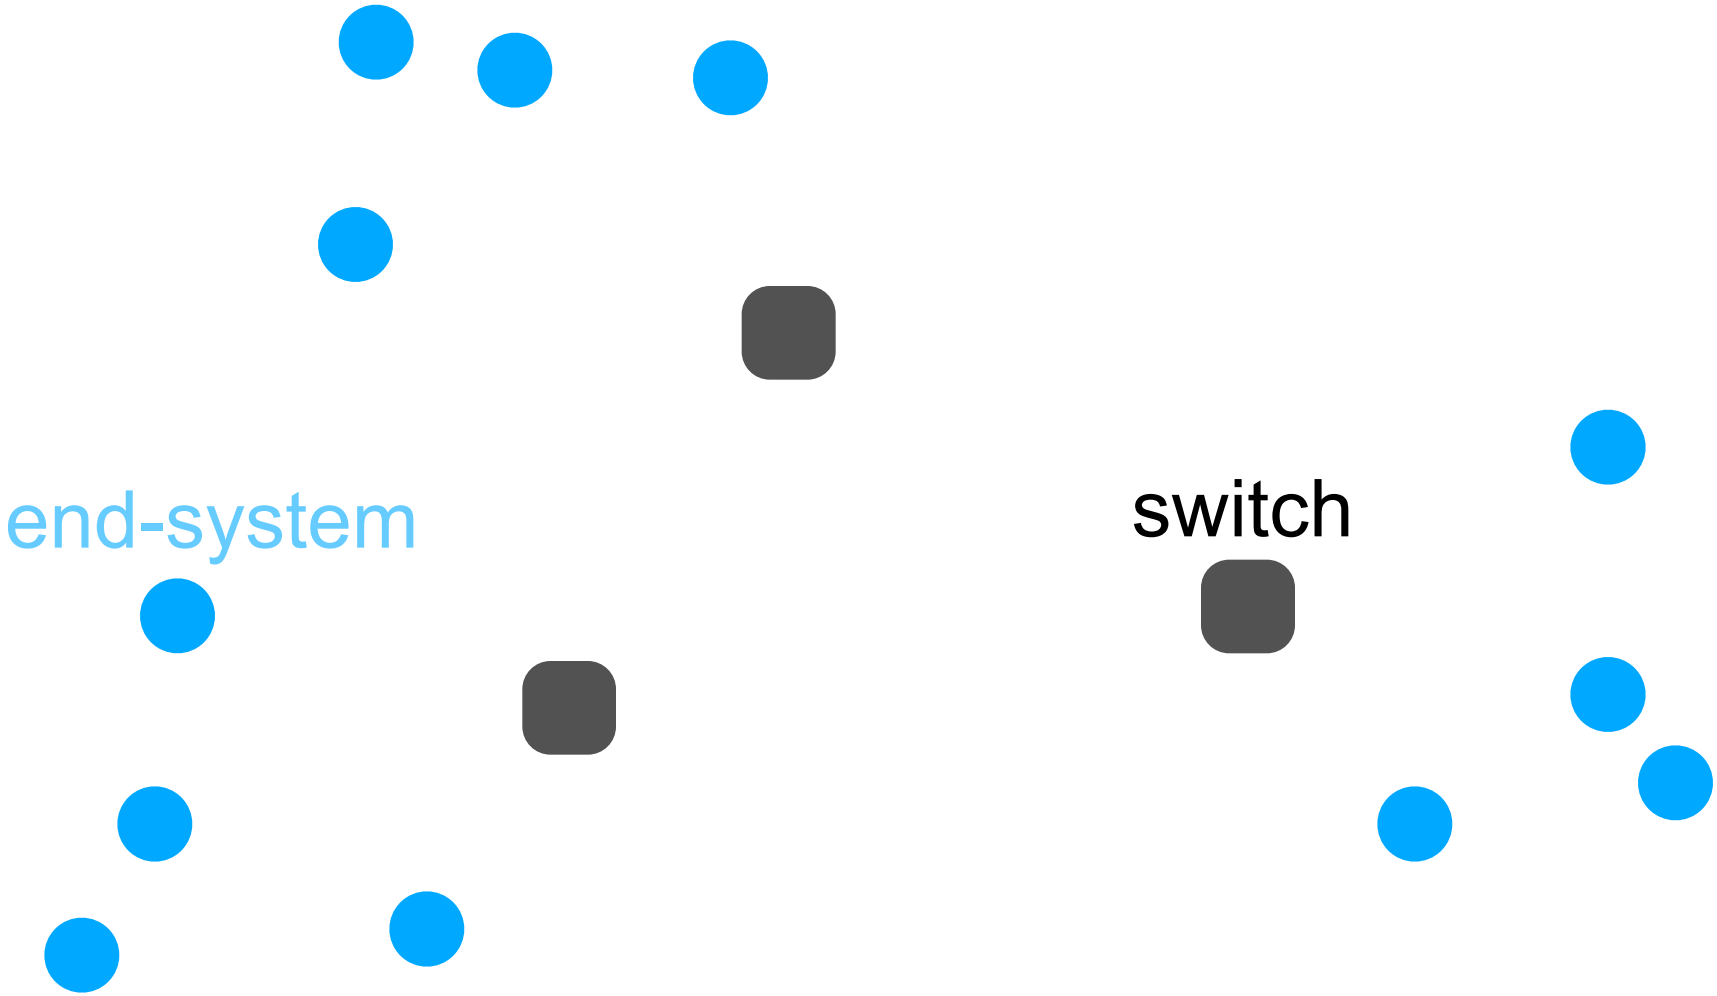
end-system
switch

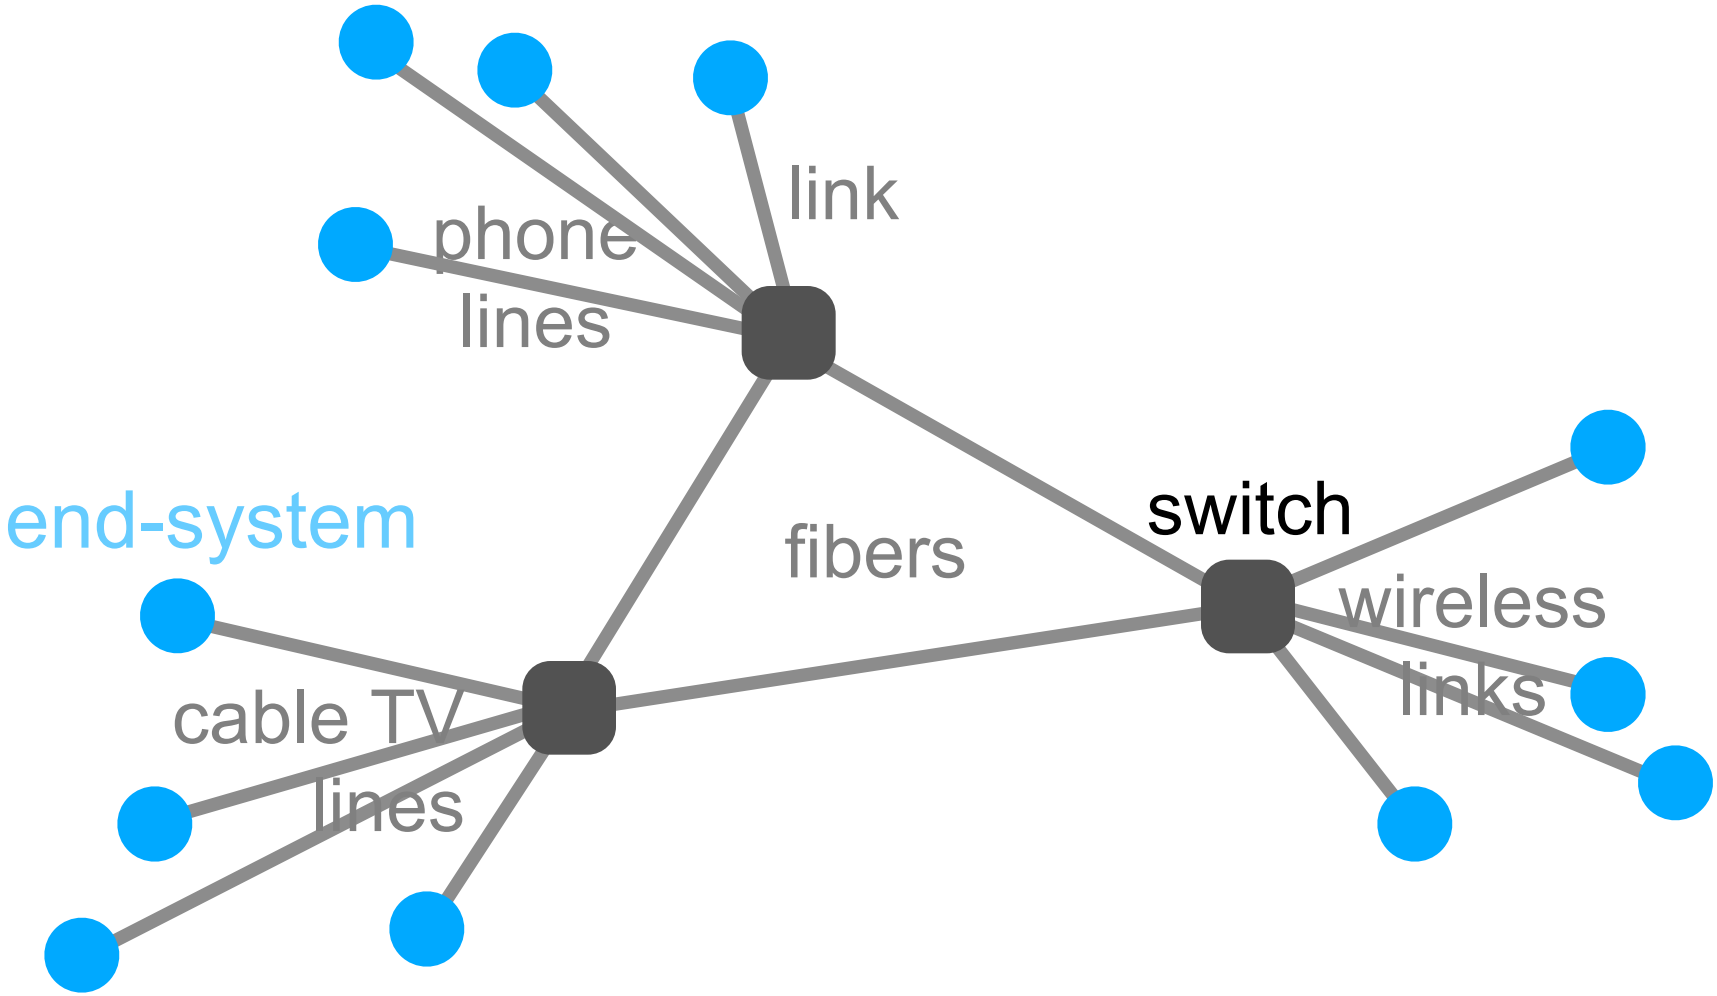
link
phone
lines
end-system
fibers
switch
wireless
links
cable TV
lines

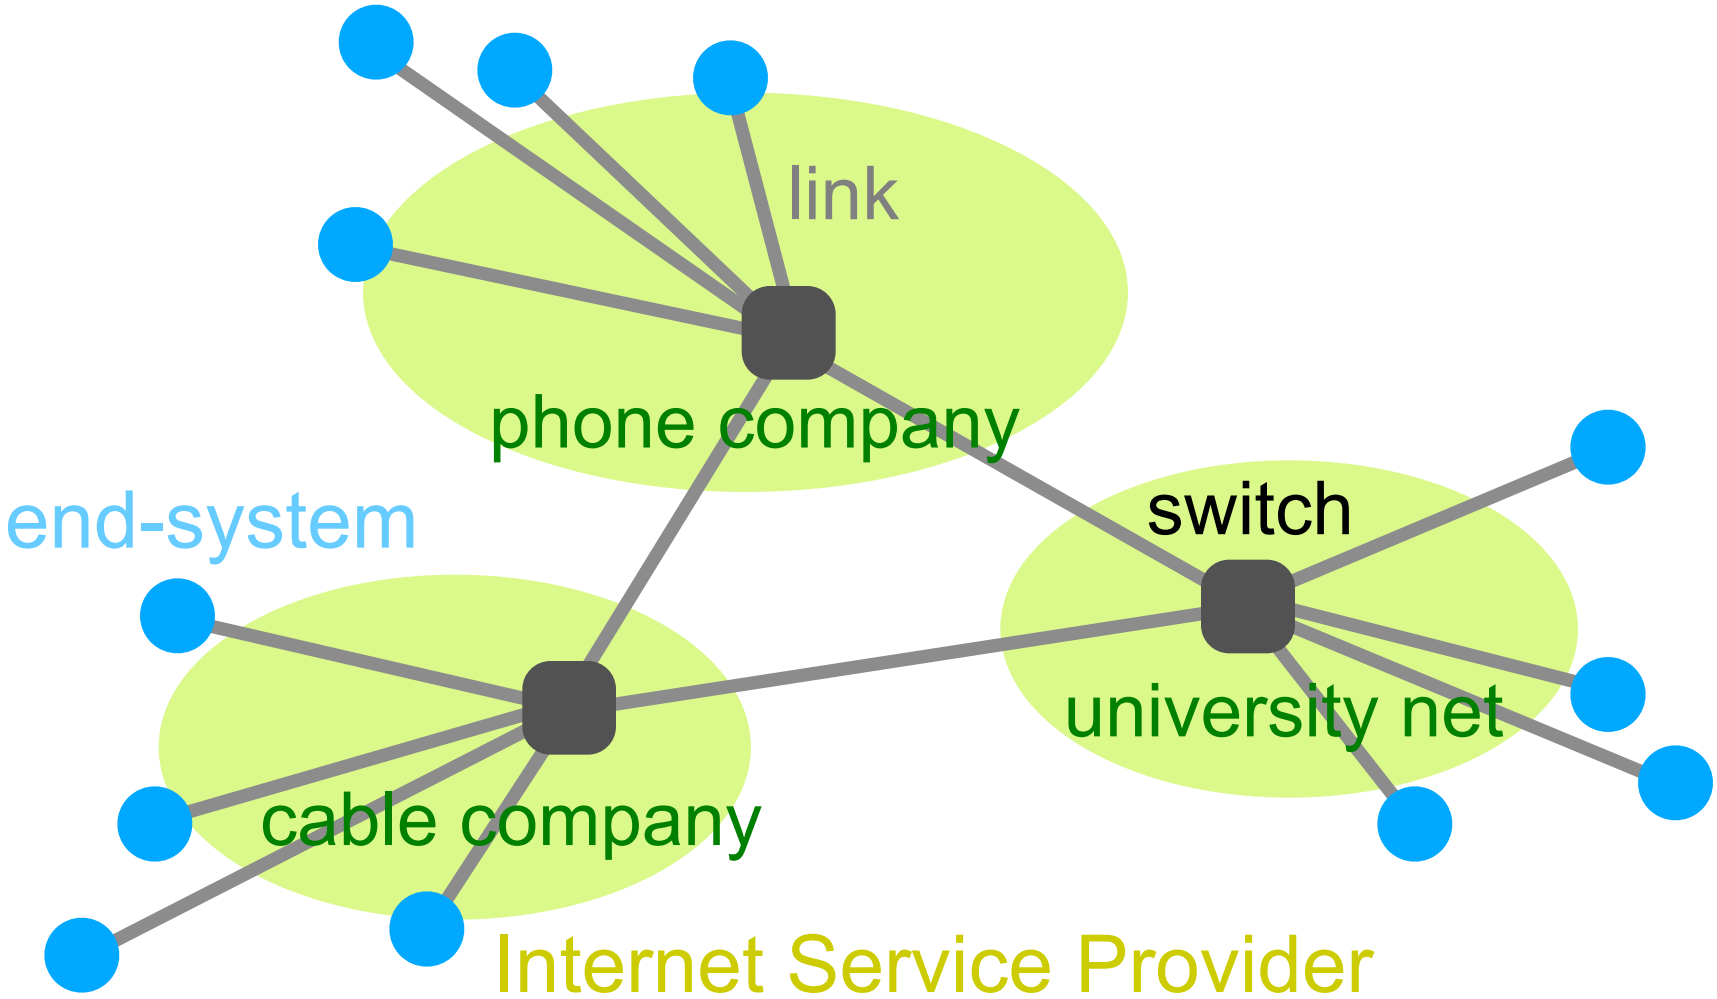
link
phone company
end-system
switch
university net
cable company
Internet Service Provider

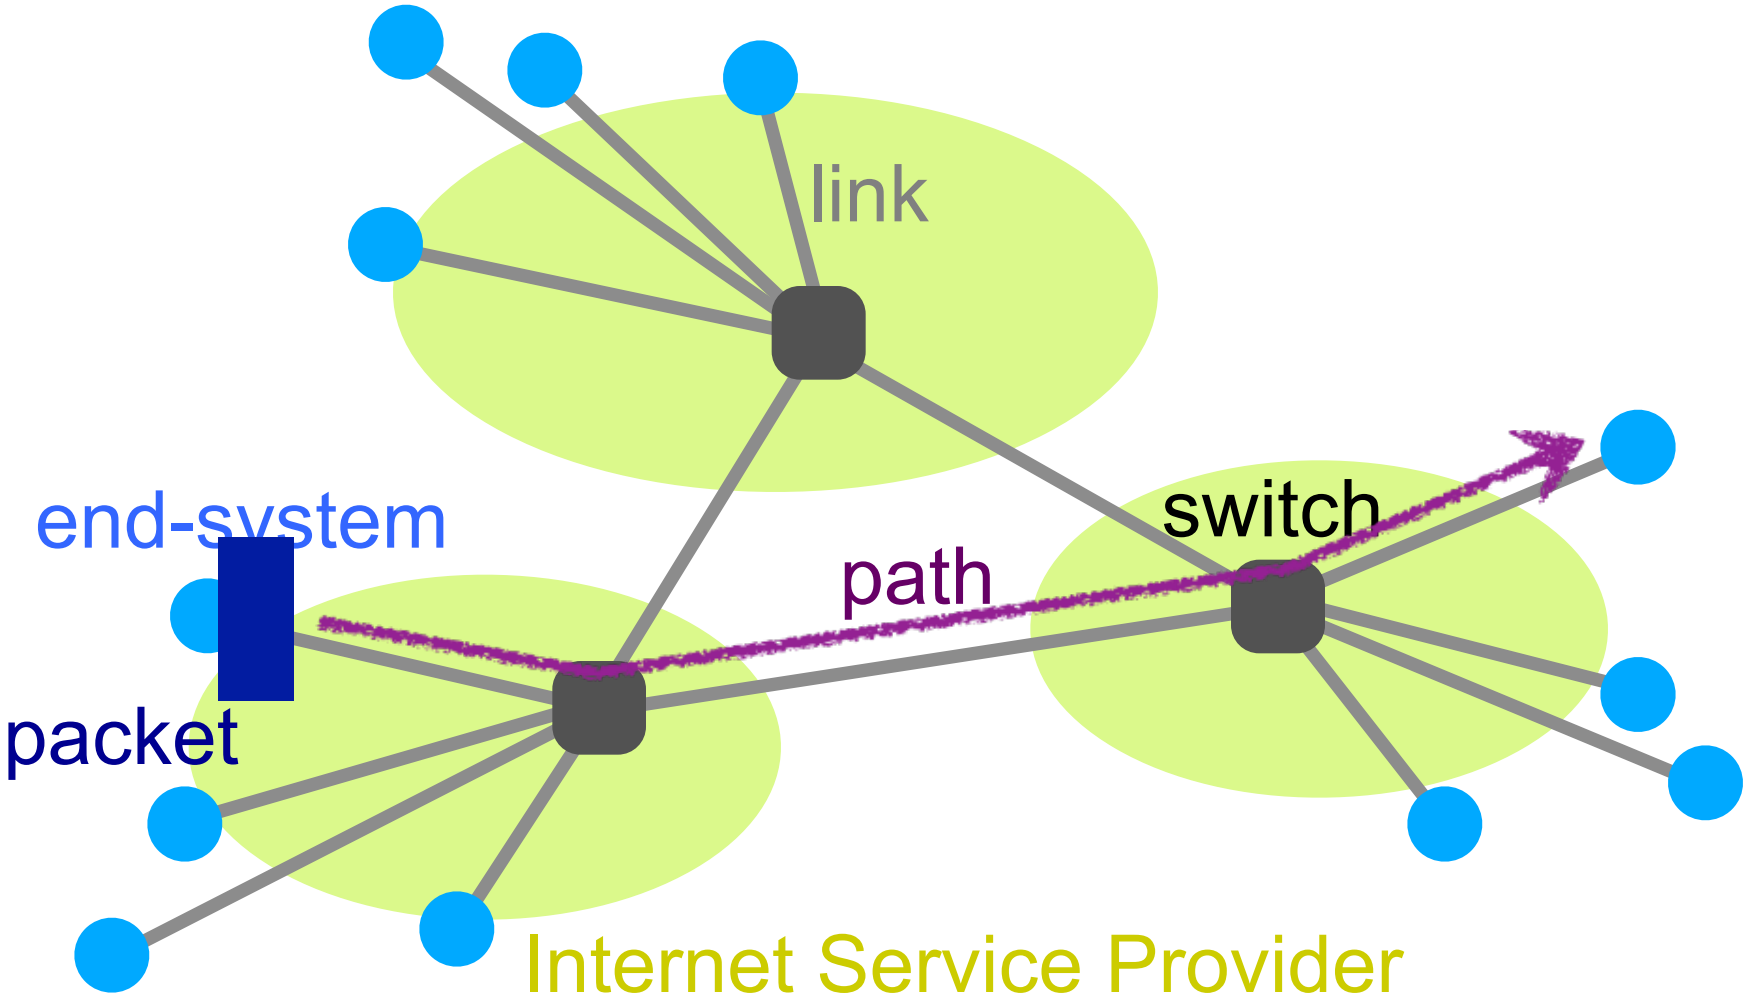
link
end-system
switch
path
packet
Internet Service Provider

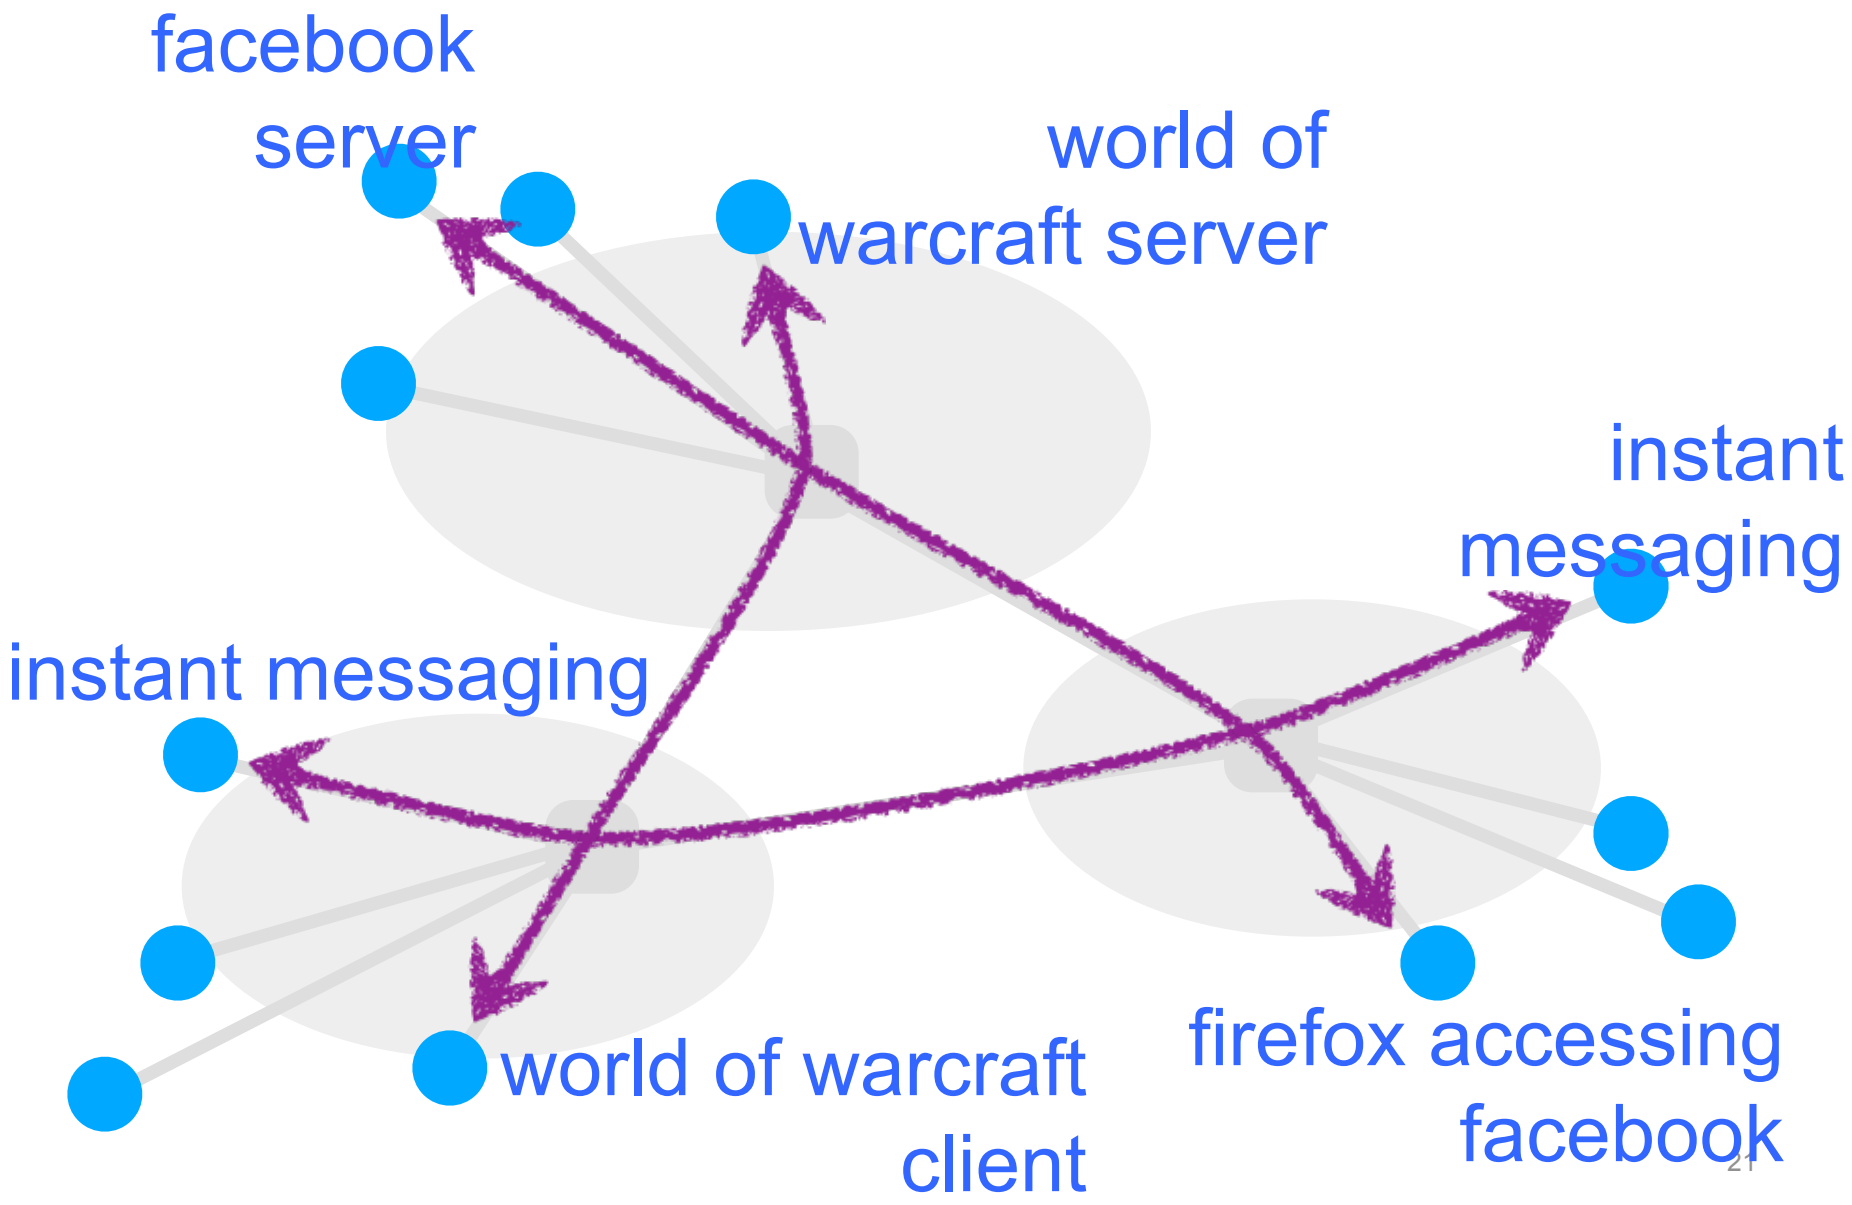
facebook server
world of warcraft server
instant messaging
instant messaging
world of warcraft client
firefox accessing facebook

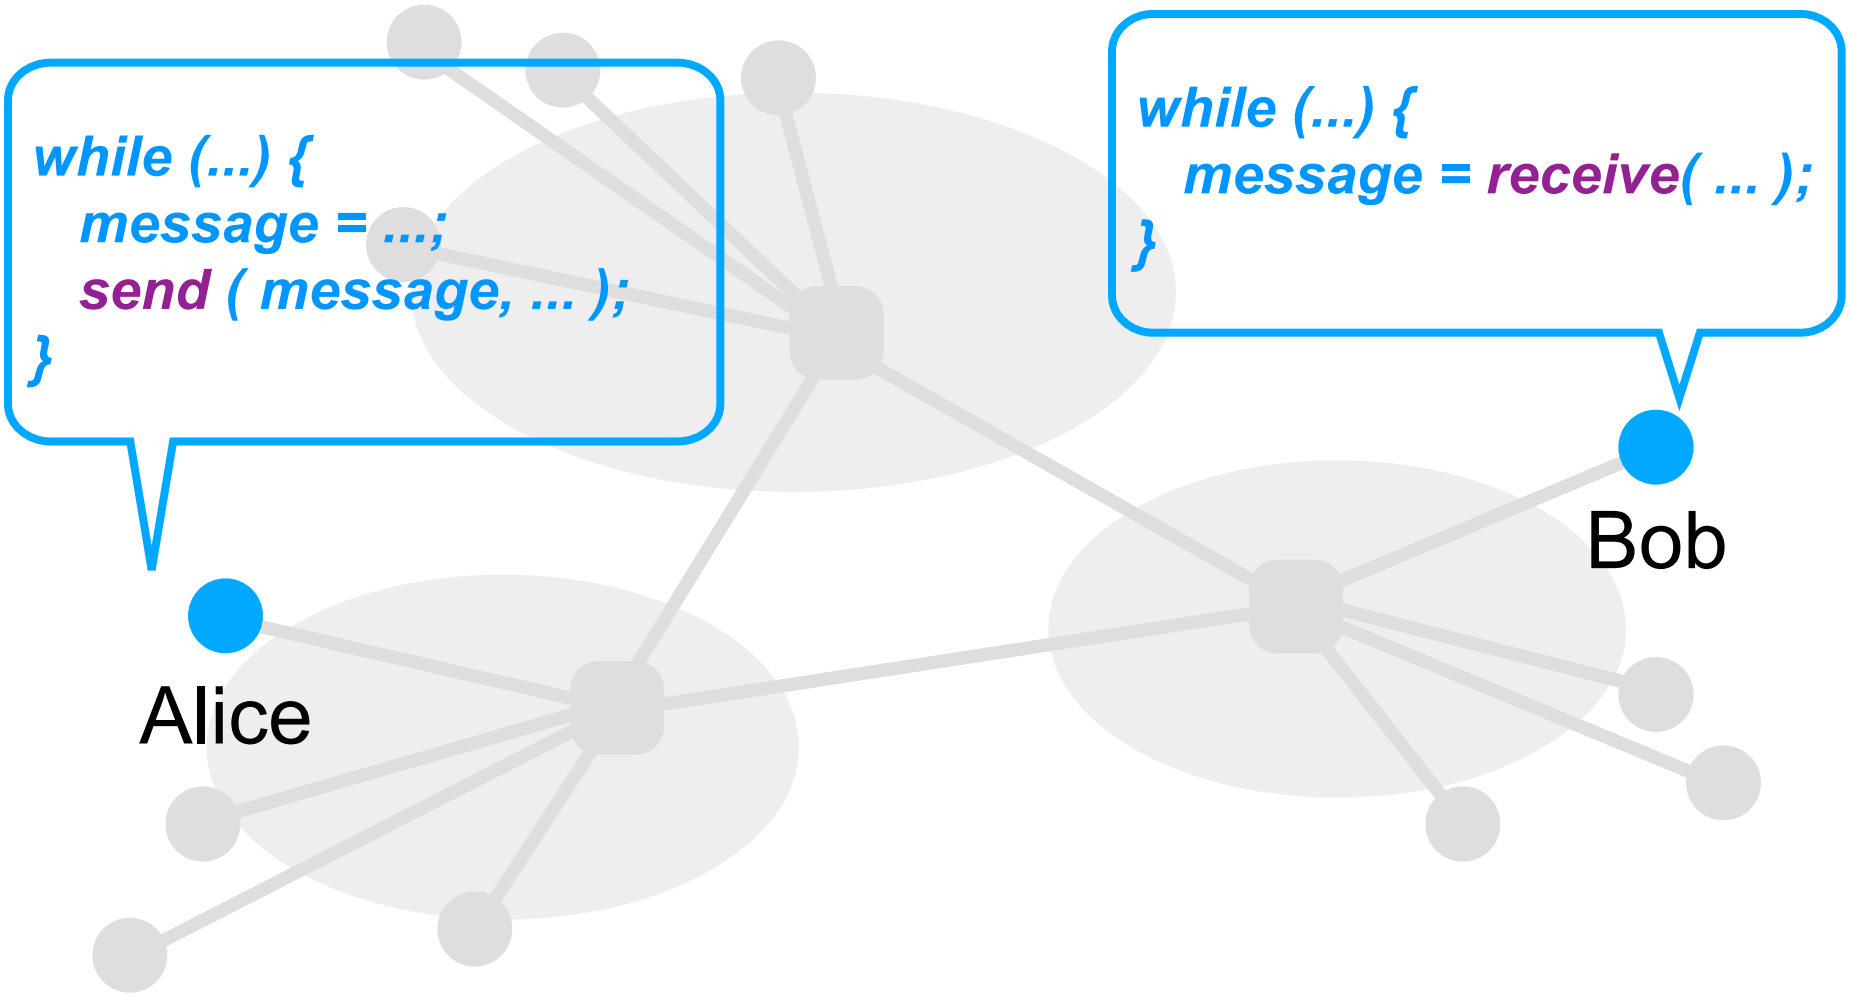
while (...) {
message = ...;
send (message, ...);
}
while (...) {
message = receive(...);
}
Alice
Bob

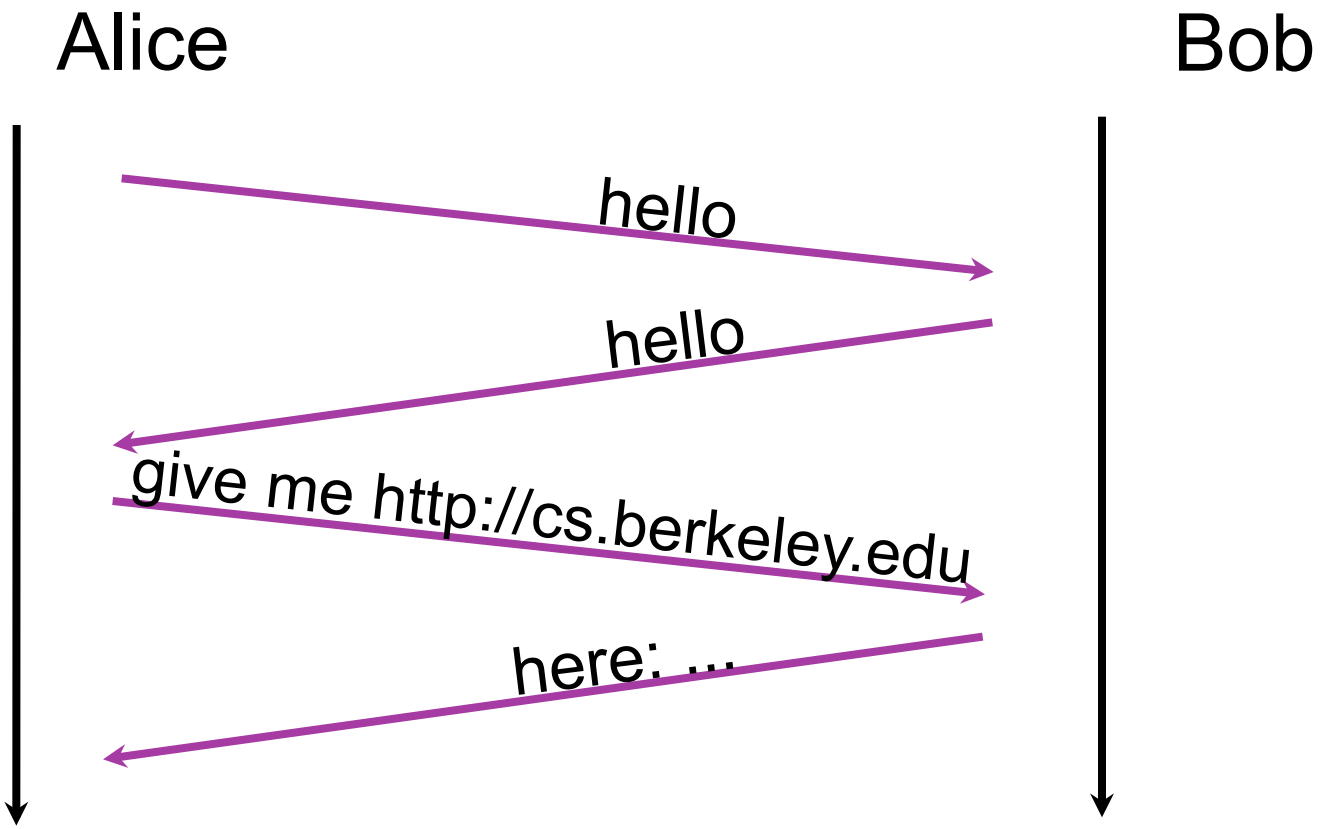
Alice
Bob
hello
hello
give me http://cs.berkeley.edu
here: ...


Alice
Bob
hello
give me http://...
give me http://...
give me http://...
give me http://...
give me http://...
give me http://...

Why study the Internet?

The Internet is transforming everything

- The way we do business
- sing, cl
- lations
- mail, IM
- search
- and vi
- copyri
- sease
  - Digital health, remote diagnostics

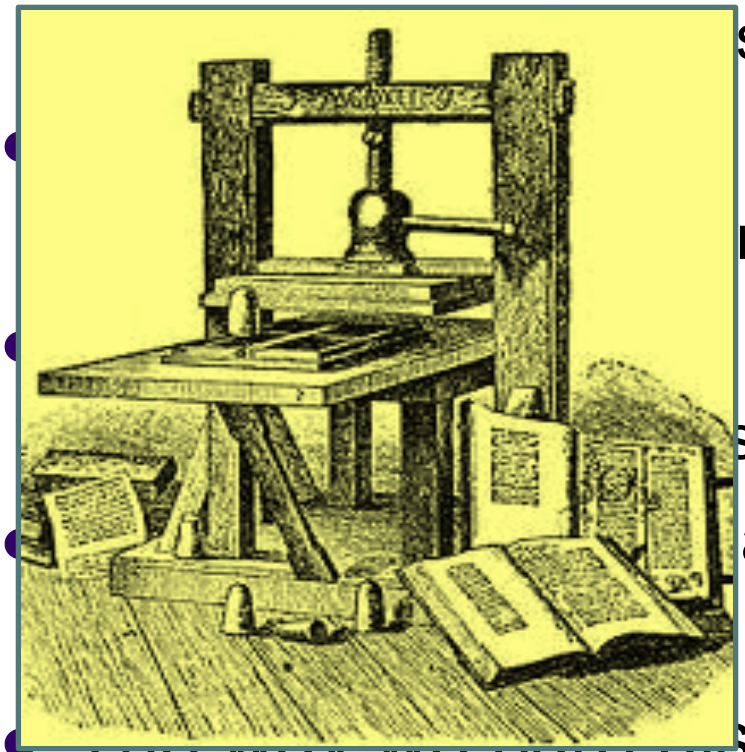

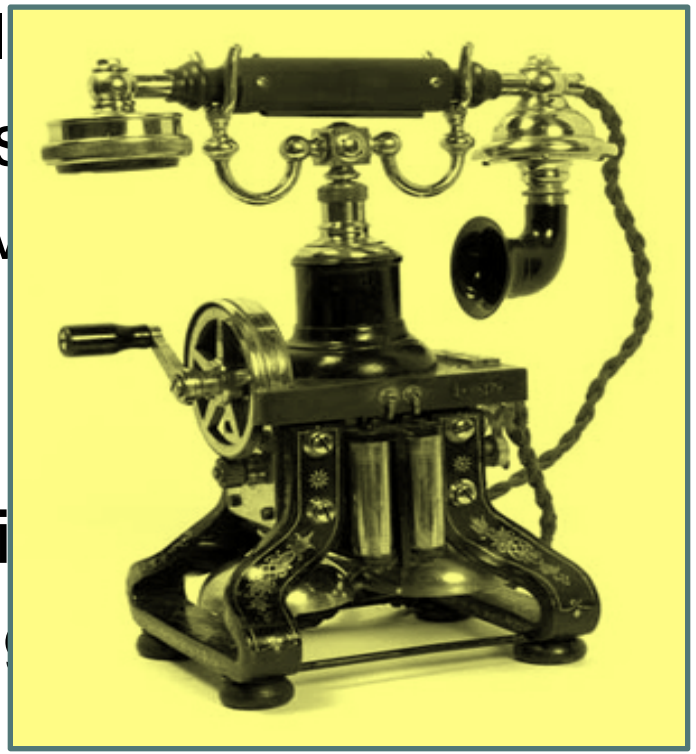

The Internet is big business

- Many large and influential networking companies
  - Cisco, Broadcom, AT&T, Verizon, Akamai, Huawei, …
  - $200B+ industry (carrier and enterprise alone)

- Networking central to most technology companies
  - Google, Facebook, Microsoft, HP, Dell, VMware, …

Internet research has impact

- The Internet started as a research experiment!
- 4 of 10 most cited authors work in networking
- *Many* successful companies have emerged from networking research(ers)

But why is the Internet *interesting*?

"What's your formal model for the Internet?" -- *theorists*

"Aren't you just writing software for networks" – *OS community*

"You don't have performance benchmarks???" – *hardware folks*

"It's just another communication network!" – *old timers at AT&T*

"What's with all these TLA protocols?" – *all*

"But the Internet seems to be working…" – *my parents*

A few defining characteristics of the Internet

A federated system

The Internet interconnects different networks (>18,000 ISPs)

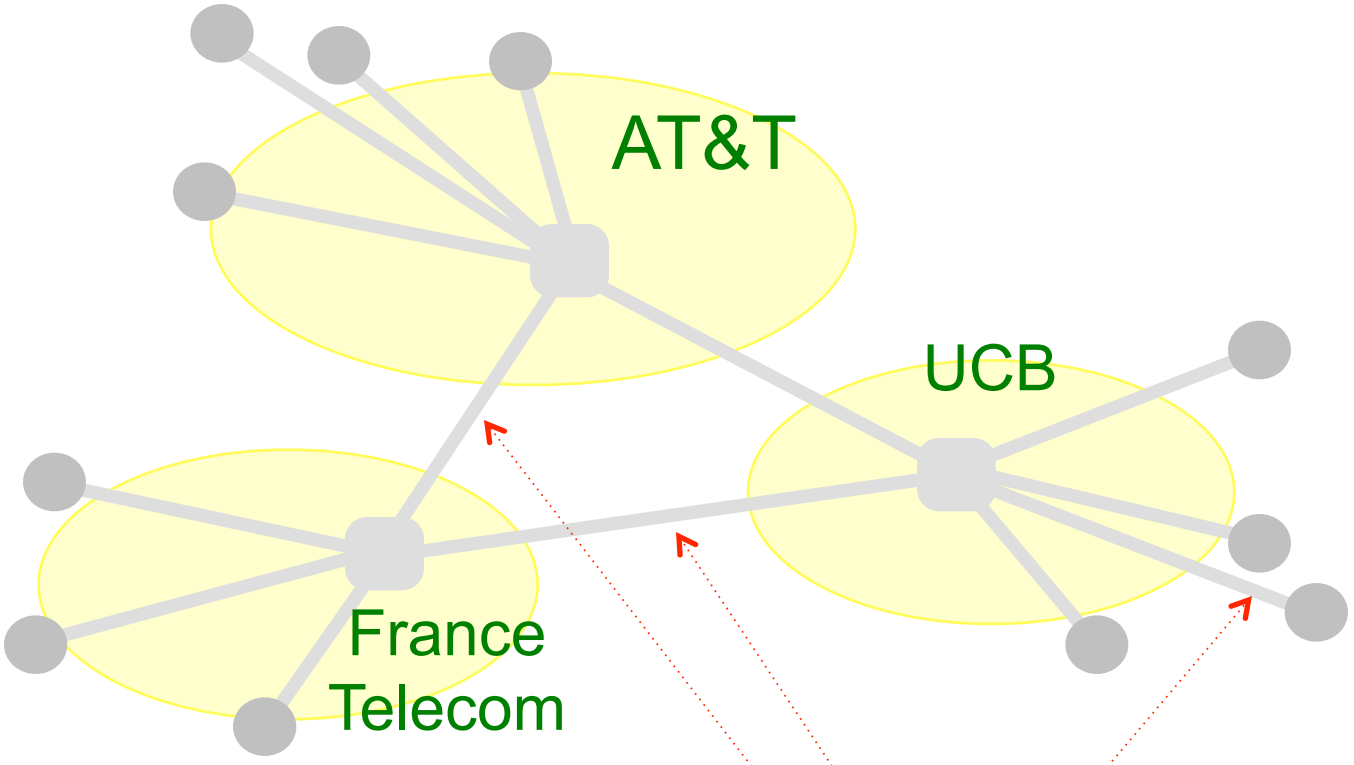

One common protocol -- the "Internet Protocol (IP) -- between users and the network and between networks

A federated system

- Interoperability is the Internet's most important goal
- Leads to a constant tussle between business and technical factors
  - competing ISPs must cooperate to serve their customers
  - practical realities of incentives, economics and real-world trust determine physical topology and path selection
  - a common protocol is great for interoperability …
  - … but complicates innovation

Tremendous scale

- 2.92 Billion users (41% of world population)
- 1 Trillion unique URLs (in 2008)
- 294 Billion emails sent per day
- 1.75 Billion smartphones
- 1.24 Billion Facebook users
- 100 hours of video uploaded to YouTube every minute
- Switches that move 300Terabits/second (10^{14})
- Links that carry 100Gigabits/second

Enormous diversity and dynamic range

- Communication latency: microseconds to seconds (10^6)
- Bandwidth: 1Kbits/second to 100 Gigabits/second (10^7)
- Packet loss: 0 – 90%
- Technology: optical, wireless, satellite, copper
- Endpoint devices: sensors, cell phones, datacenters
- Applications: skype, live video, gaming, remote medicine,
- Users: the governing, governed, operators, selfish, malicious, naïve, savvy, embarrassed, paranoid, …

Constant Evolution

1970s:

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

Today

- 100+Gigabits/second backbone links
- 5B+ devices, all over the globe
- 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 at 300,000 km/s: 13.75 milliseconds
    - Then back to Berkeley: 2 x 13.75 = 27.5 milliseconds
    - 3,000,000,000 cycles/sec * 0.0275 = 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, modem, wireless access point, firewall, links, network interface cards, switches,…
  - 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 limits of available technology
  - Link bandwidths
  - Switch port counts
  - Bit error rates
  - Cost
  - …

Recap: The Internet is…

- A federated system
- 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” needn’t mean much
- Performance benchmarks are too narrow

So, what do we need?

We still don't really know…

- No consensus on what constitutes the "correct" or "best" network design
- 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

Architectural principles

- Decentralization [lectures: all]
- Packets [lecture# 2]
- Statistical multiplexing [lecture 2]
- Best effort service [lecture 3]
- The “end to end” design principle [lecture 8+]
- “Layered” decomposition [lectures: all]
- IP as “narrow waist” interface [lecture 8]

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
  - want to diagnose problems but federation hides inner workings
  - want to block unwanted traffic but the network doesn't authenticate
  - can't optimize for different applications or customers
  - upgrading protocols is deeply painful

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, reliability, data centers, sensors, …

Hence, networking today is still debating the big questions...

- Packets →"circuits"
- Statistical multiplexing → "reservations"
- Protocol layers
- A "narrow waist" at the network layer
- Best-effort service →"Quality of Service (QoS)"
- The "end to end" design principle → "middleboxes"
- Decentralization → "centralize"

Backing up a level

- The Internet offers us a lesson on how to reason through the design of a complex system
  - What are our goals and constraints?
  - What's the right prioritization of goals?
  - How do we decompose a problem?
  - Who does what? How?
  - What are the tradeoffs between design options?

- In short: a lesson in how to architect a system

Network Architecture

- More about thinking rigorously than doing rigorous math
- More about understanding tradeoffs than running benchmarks
- More about practicality than optimality

Done right, can be a powerful thing

What (I hope) CS 168 will teach you

- How the Internet works
- Why it works the way it does
- How to reason through a complicated (networking) design problem

Let’s take a 5 minute break

Today

- Introductions
- What is (this course on) networking about?

5 minute break

- Class policies, roadmap, administrivia

Class Workload

- Three projects
- Three homeworks
- Exams:
  - midterm: October 20 in class
  - final: December 18, 8-11am, location TBA
  - closed book, open crib sheet

- No lecture on November 26 (for Thanksgiving)

Grading

3 Homeworks	15% (3x 5% each)
3 Projects	40% (10+10+20)
Midterm exam	20%
Final exam	25%

- Course graded to mean of B

Topics we will cover

- Basic concepts [Lectures 2, 3]
  - packets, circuits, delay, loss, protocols
- How the “insides” of the Internet work [Lectures 3-8]
  - IP, DV/LS routing, BGP
- How endpoints use the network [Lectures 9-16]
  - TCP, DNS, HTTP
- Crucial lower-level technologies [Lectures 17-20]
  - Ethernet, wireless
- Important new(er) topics [Lectures 21-26]
  - management, security, datacenters

Three projects

- Project 1: Routing (in simple simulator)
- Project 2: Reliable Transport (in simple simulator)
- Project 3: Build a network firewall
  - Larger project, in two phases

TAs will handle all project-related questions!

Administrivia: Textbook

- J. Kurose and K. Ross, *Computer Networking: A Top-Down Approach,* 6th Edition, 2012.
  - 5th Edition ok, but translate the reading assignments
- **You will not be tested on material we didn't cover in lecture or section**
  - Use as a reference and a source of examples

Enrollment and wait list

- Class size will not increase
- Wait-listed students will be admitted as and when registered students drop the class
  - If you're planning to drop, please do so soon!
  - Waitlist will be processed in order; seniors get priority

Class communications

- Web site: http://inst.eecs.berkeley.edu/~cs168/fa14/
  - Assignments, lecture slides, announcements
- Use your instructional account to hand in assignments
  - Accounts will be handed out next week
- Use Piazza for all other intra-class communication
  - You should all be signed up by now
- Copy Radhika (radhika@cs) and Sangjin (sangjin@cs) on any emails sent directly to me (sylvia@cs)

Policy on late submissions, re-grade requests, cheating

- Detailed description is on the class website
- Summary version:
  - You may submit assignments late or request re-grades but to a point, and it will cost you
  - The policy on re-grades for projects will be announced by the lead TA on the project and may vary across projects
  - Your responsibility to keep your code private!
  - When in doubt about the policy, ask us!

Class Participation

- We will post slides ~10minutes before class
- Ask and answer questions!!
  - it helps you understand
  - it helps others understand
  - it helps you stay awake
  - it helps me stay awake
  - it's just more fun for all of us
- Sit towards the front
- Limit electronic access for < 90 minutes
  - you will have a 5 minute break in the middle to get online

Next Steps

- For our next lecture
  - read 1.1 and 1.3 of K&R
- Make sure you are registered with the correct email address and on piazza
- Discussion sections will start on September 10
  - Sections on Wednesday and the following Monday will cover the same material
  - OK to swap sections unless we hit room capacity limits

Any questions?