CS168
Introduction to the Internet: Architecture and Protocols

Peyrin Kao, Sylvia Ratnasamy, Rob Shakir

Fall 2024

Today

- Introductions
- What is (this course on) the Internet about?
- Class logistics

NEW

Peyrin Kao (he/him)

- **Background**
  - 2022 – present: lecturer in EECS (CS161, 61b, 61c, 188)
  - 2021 – 2022: MS in EECS; Research focus: CS education
  - 2017–2021: BA in CS/Data Science, UC Berkeley

Rob Shakir (he/him)

- **Background**
  - Got into networking via a startup he founded in 2003
  - Learnt a lot through "just doing it”
  - Tech lead for multiple global networks, including British Telecom
  - Moved to the US to join Google and now a lead architect and engineer working on Google’s global WAN network

Sylvia Ratnasamy (she/her)

- **Background**
  - PhD from UC Berkeley
  - Joined the UCB faculty in 2011
  - Industry experience: ~10 years at Intel; co-founded startup; stints at Google
  - Networking has been my focus throughout

TAs (see class website for office hours and sections)

- TODO

Today

- Introductions
- What is (this course on) the Internet about?
- Class logistics (Peyrin)

- **Internet**
- **Protocols**
- **Architecture**

Two Meanings of “Internet”

- **The infrastructure that ties together computing devices**
  - TCP, IP, BGP, DNS, OSPF, …
- **The ecosystem of applications built on top of the above infrastructure**
  - amazon, facebook, google, twitter, ….
- **In this class, we use the first definition!**

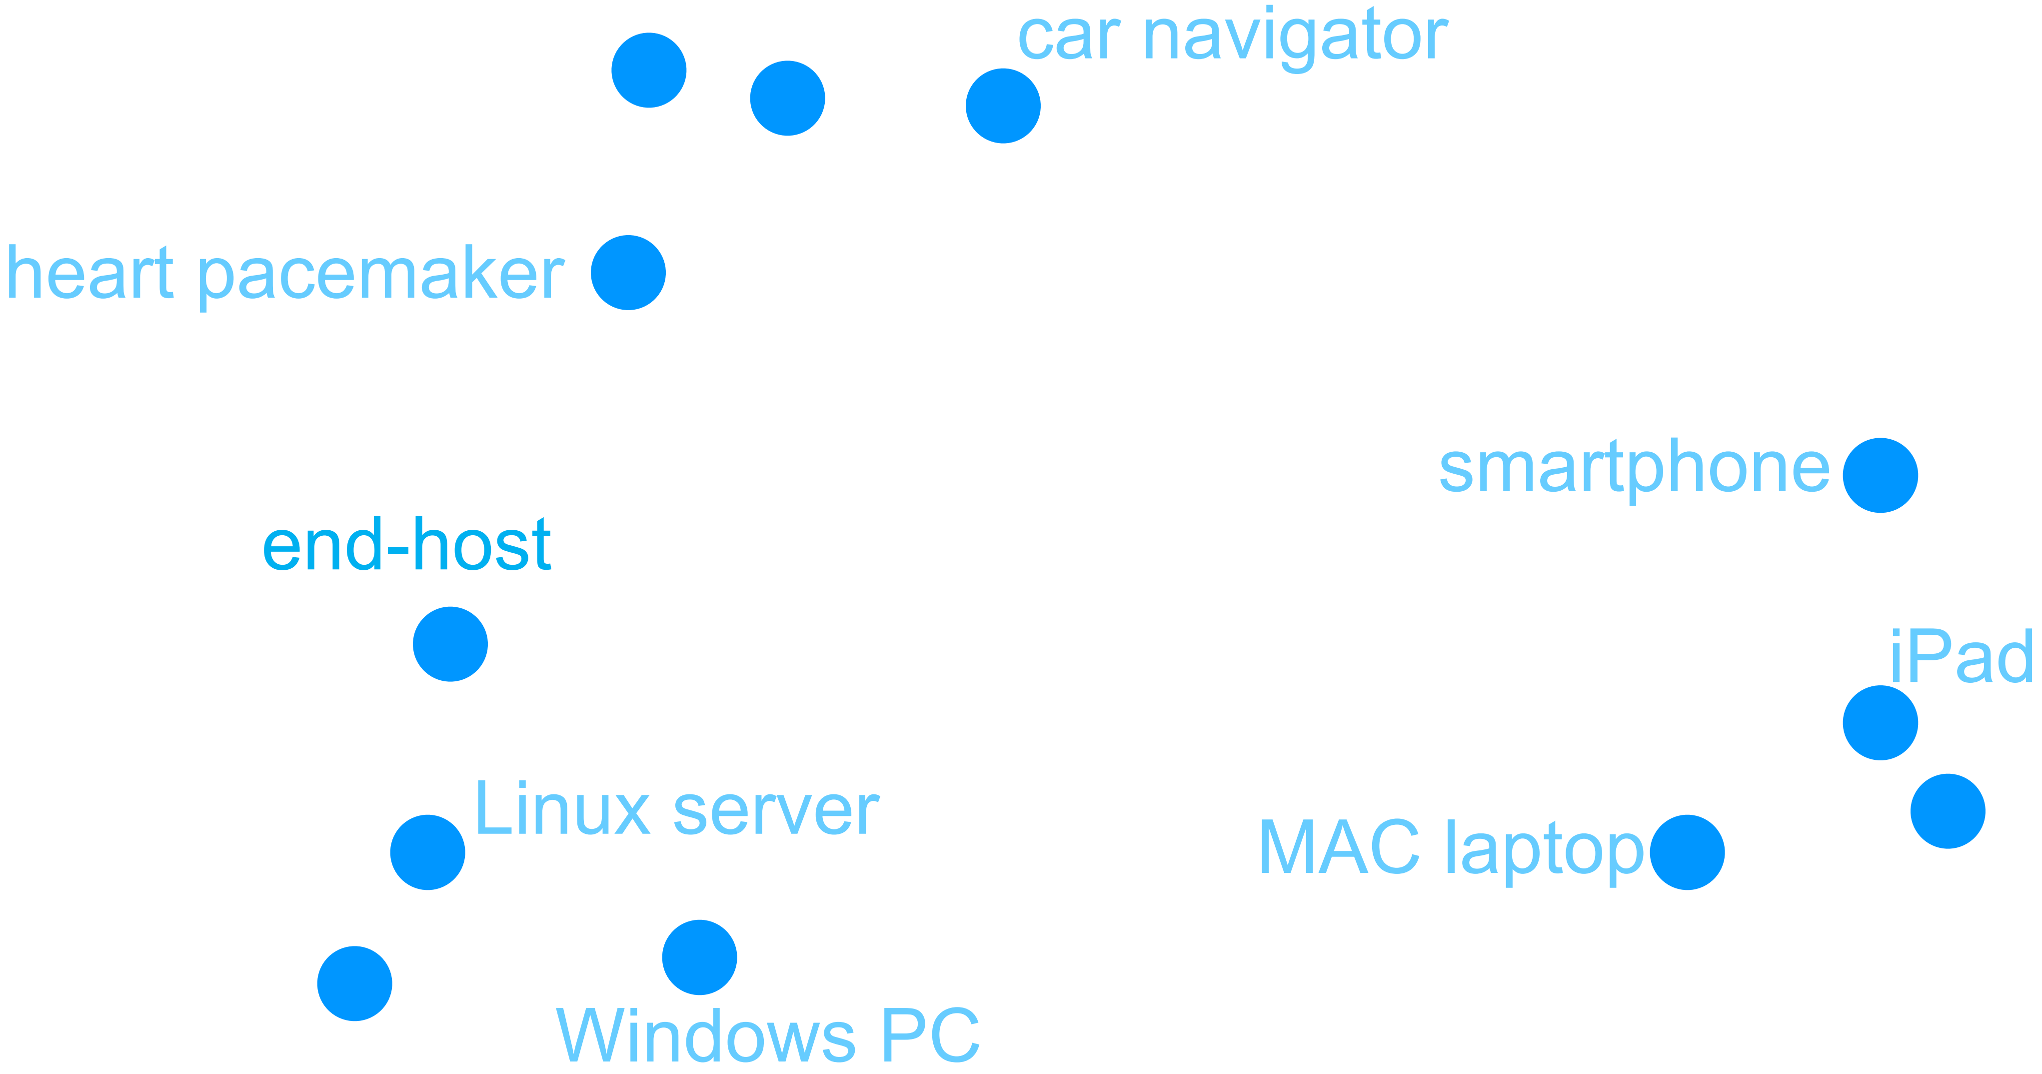
car navigator
heart pacemaker
smartphone
end-host
iPad
Linux server
MAC laptop
Windows PC

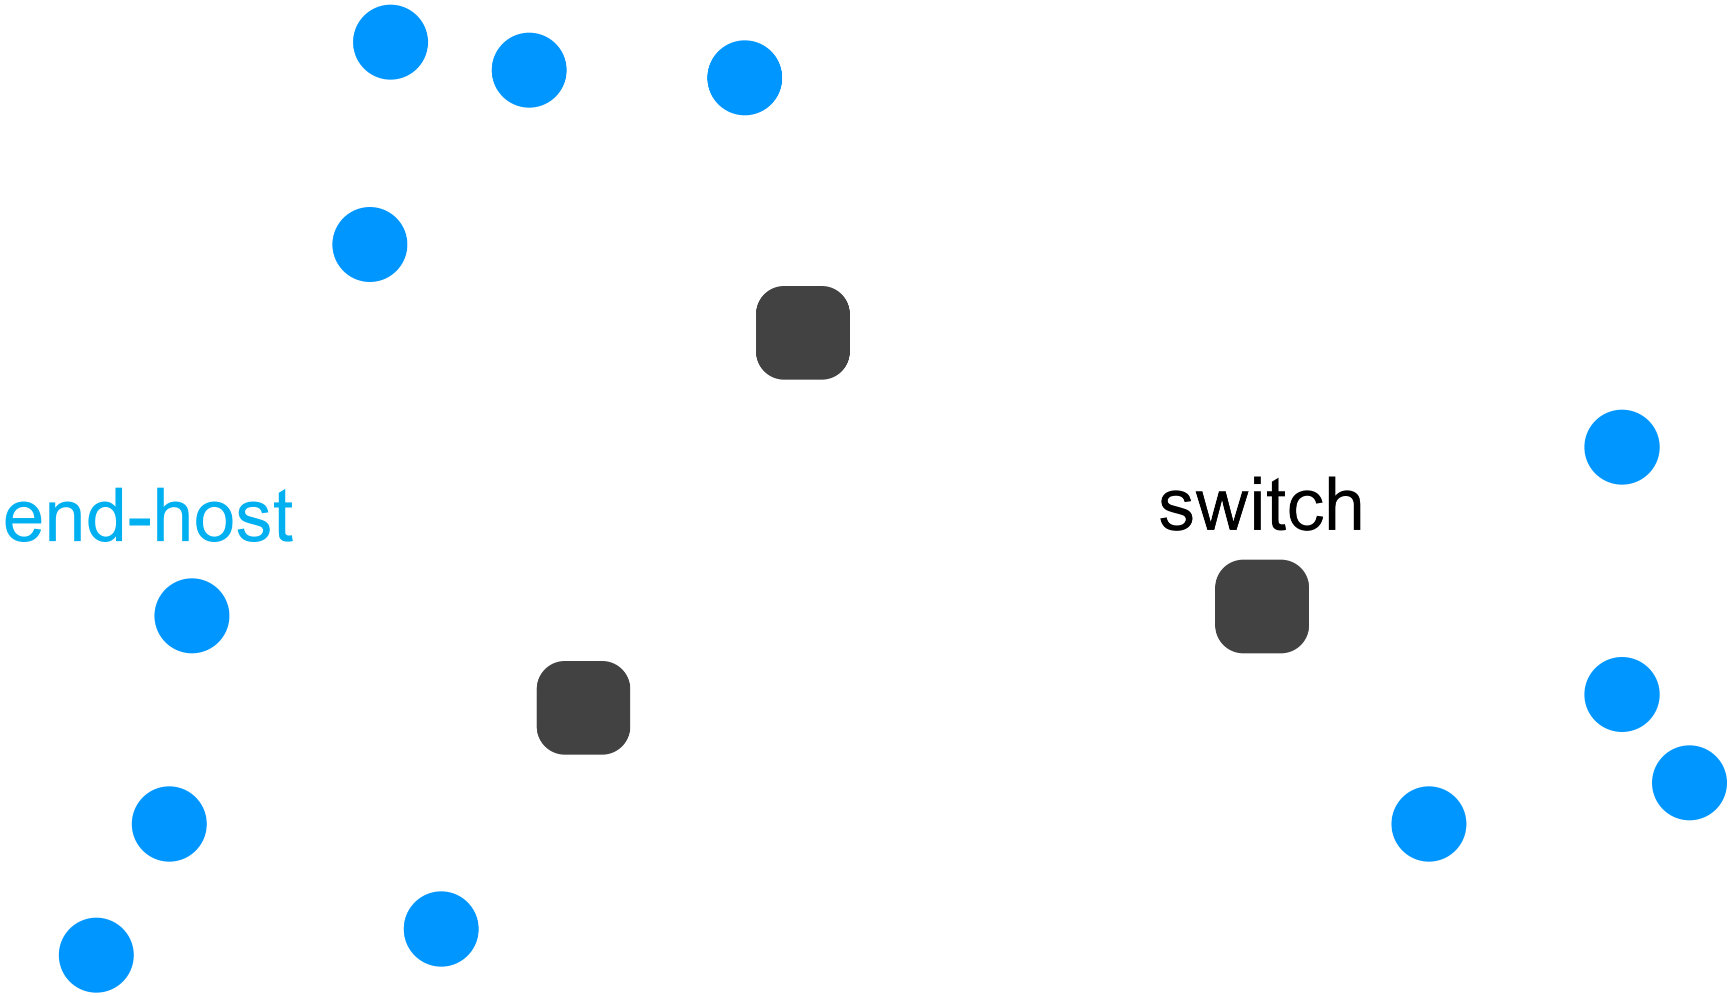
end-host
switch

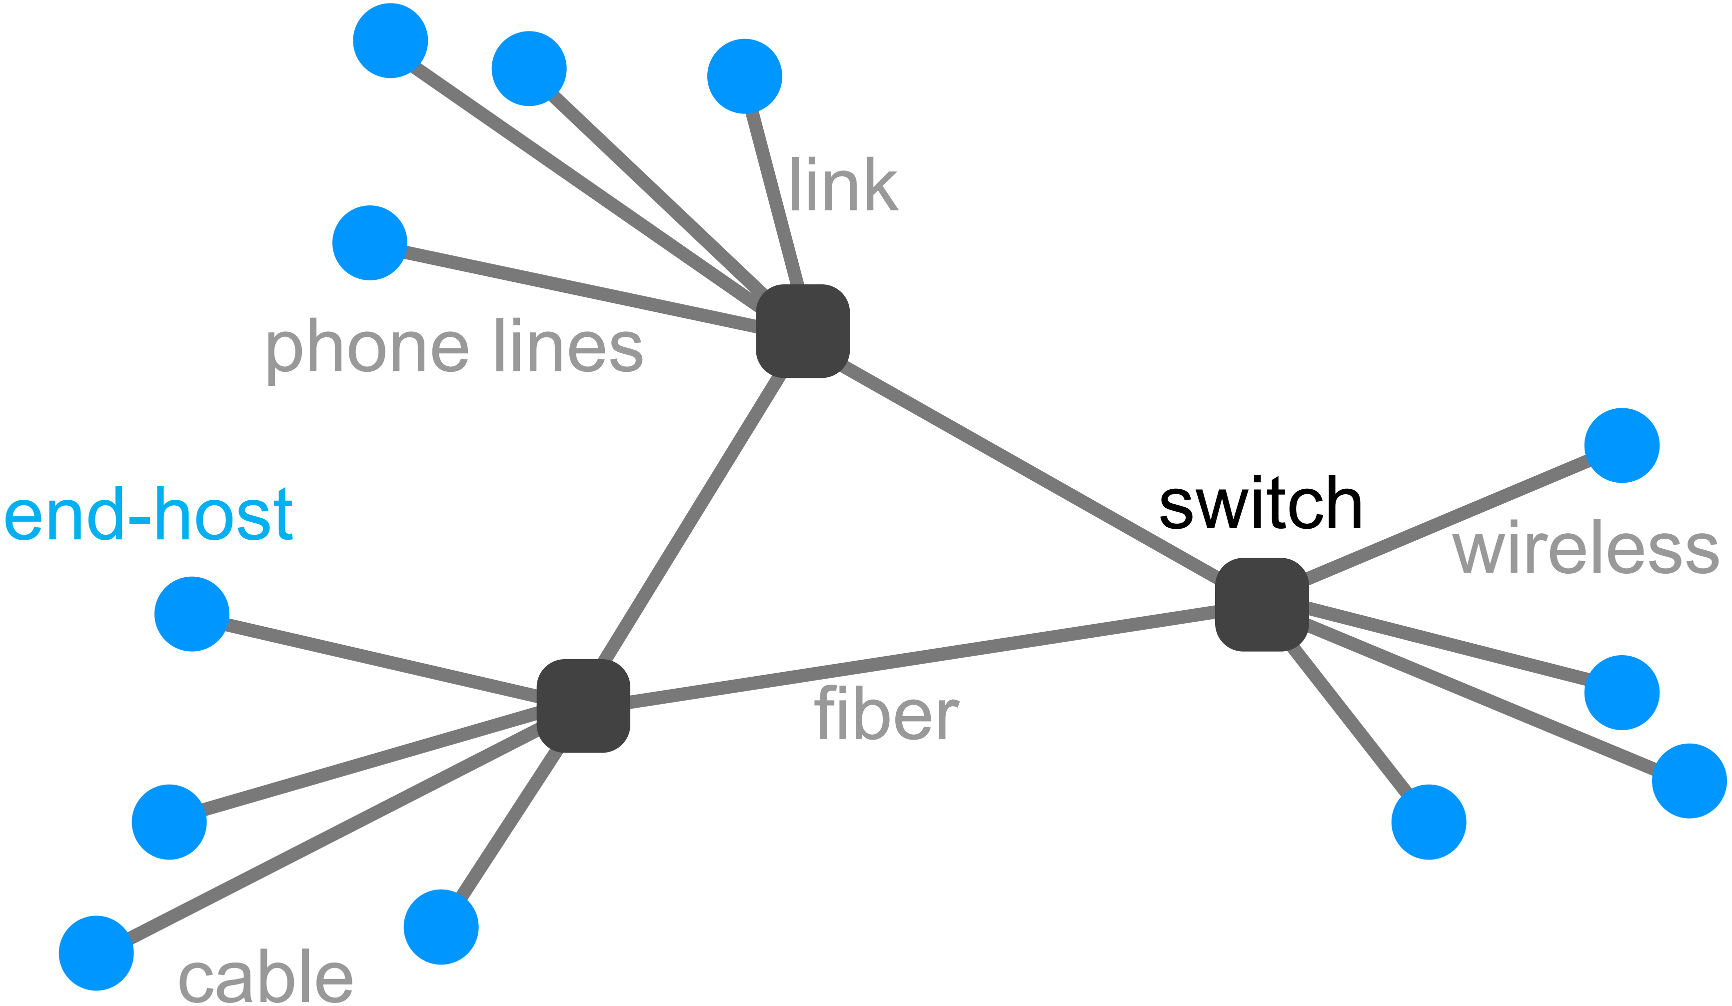
link
phone lines
end-host
switch
wireless
fiber
cable

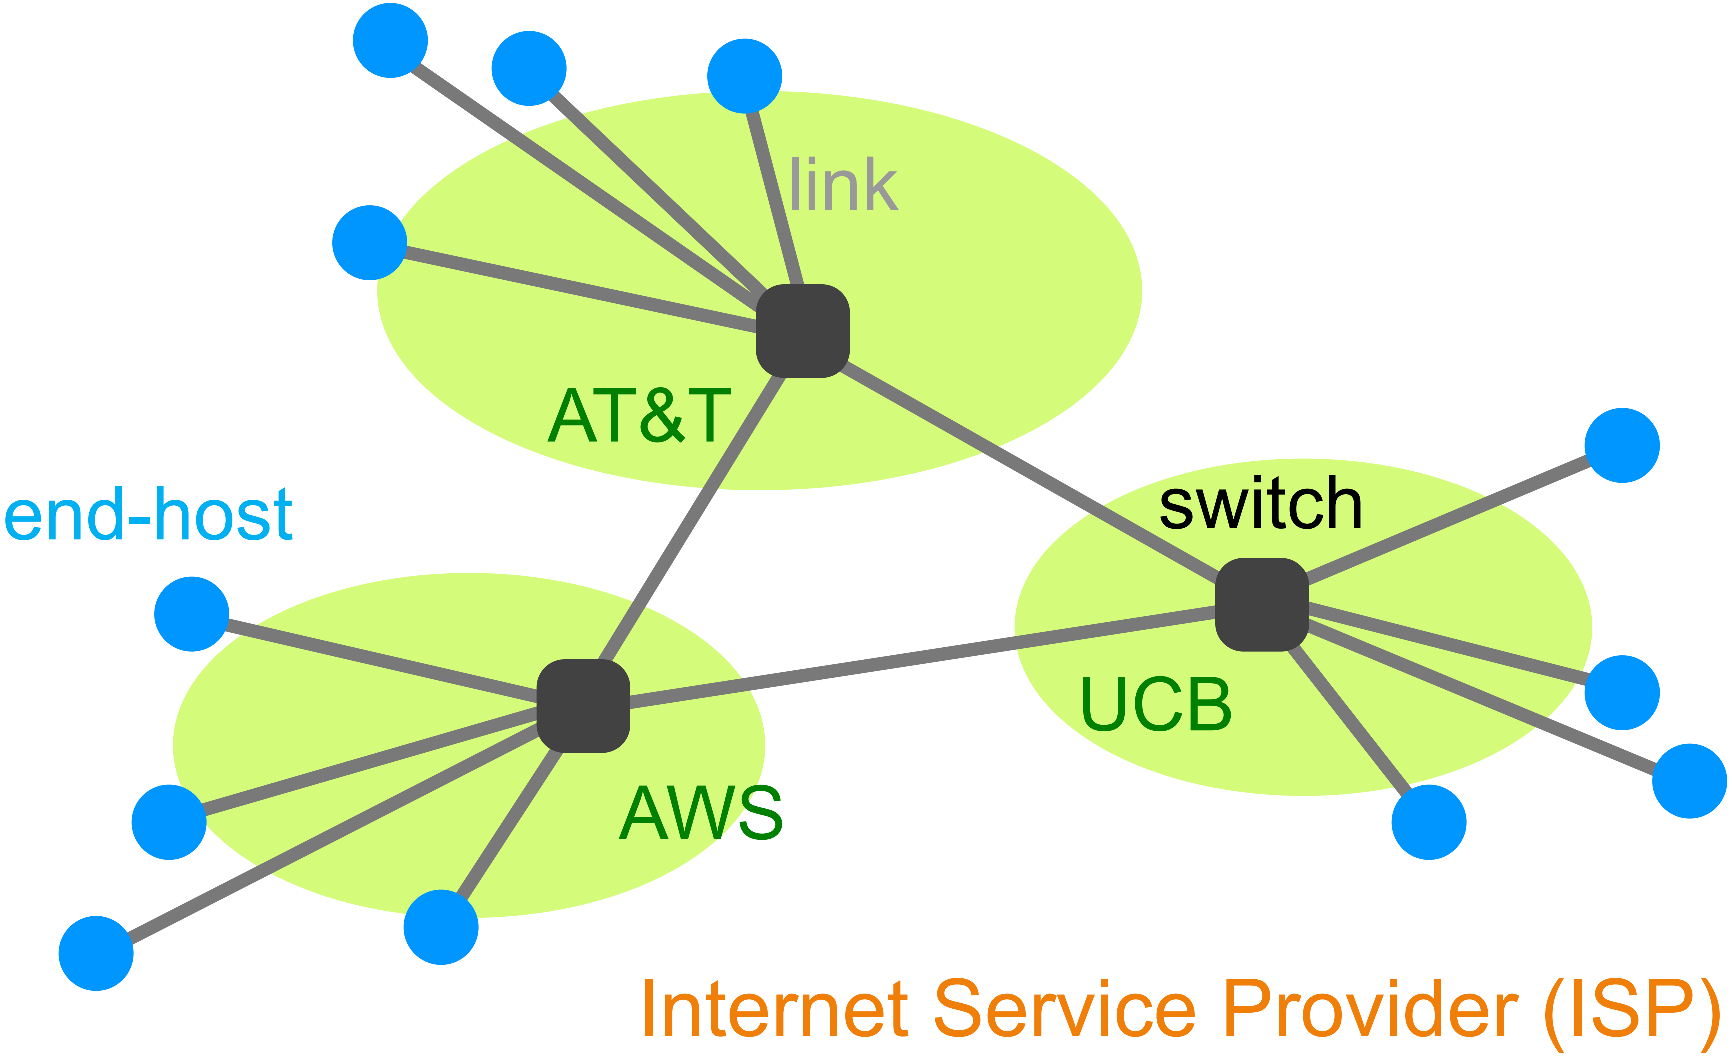
link
AT&T
end-host
switch
UCB
AWS
Internet Service Provider (ISP)

The Internet transfers data between end hosts

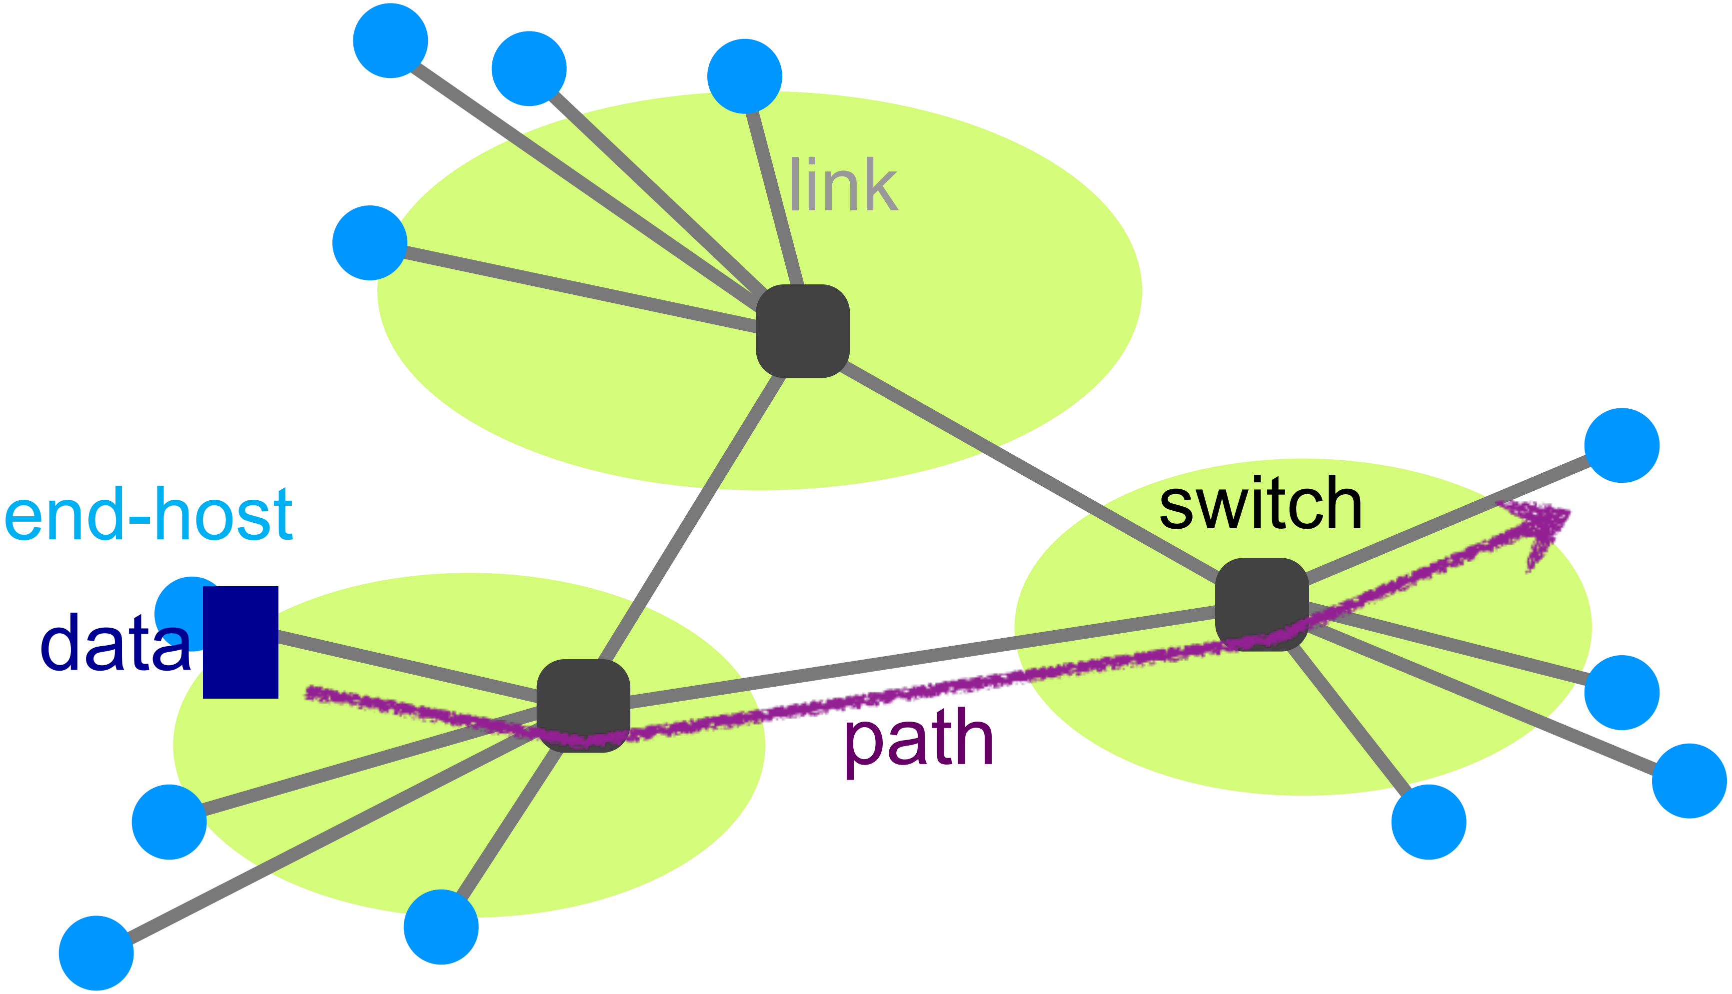

- Internet
- **Protocols**
- **Architecture**

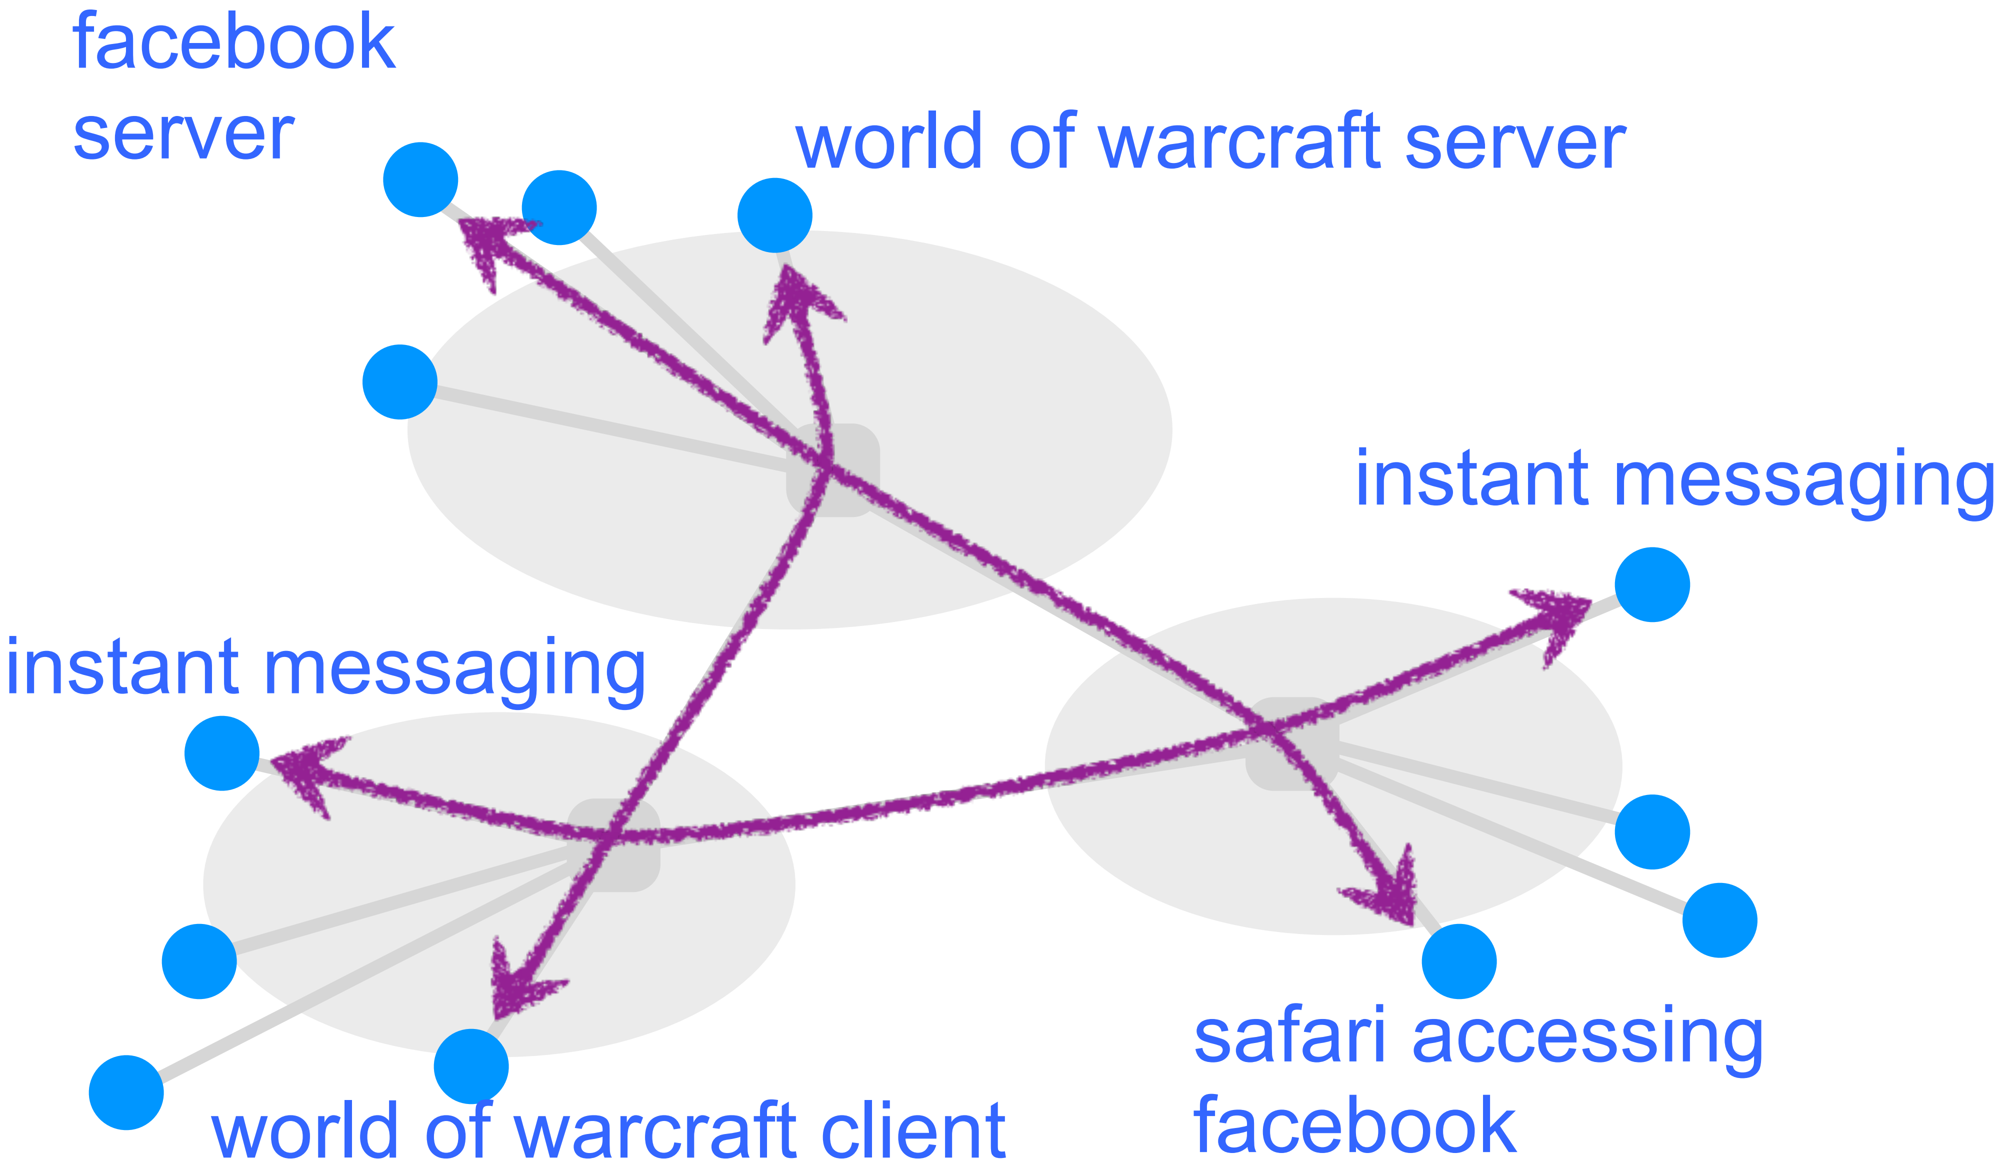
facebook
server
world of warcraft server
instant messaging
instant messaging
world of warcraft client
safari accessing
facebook

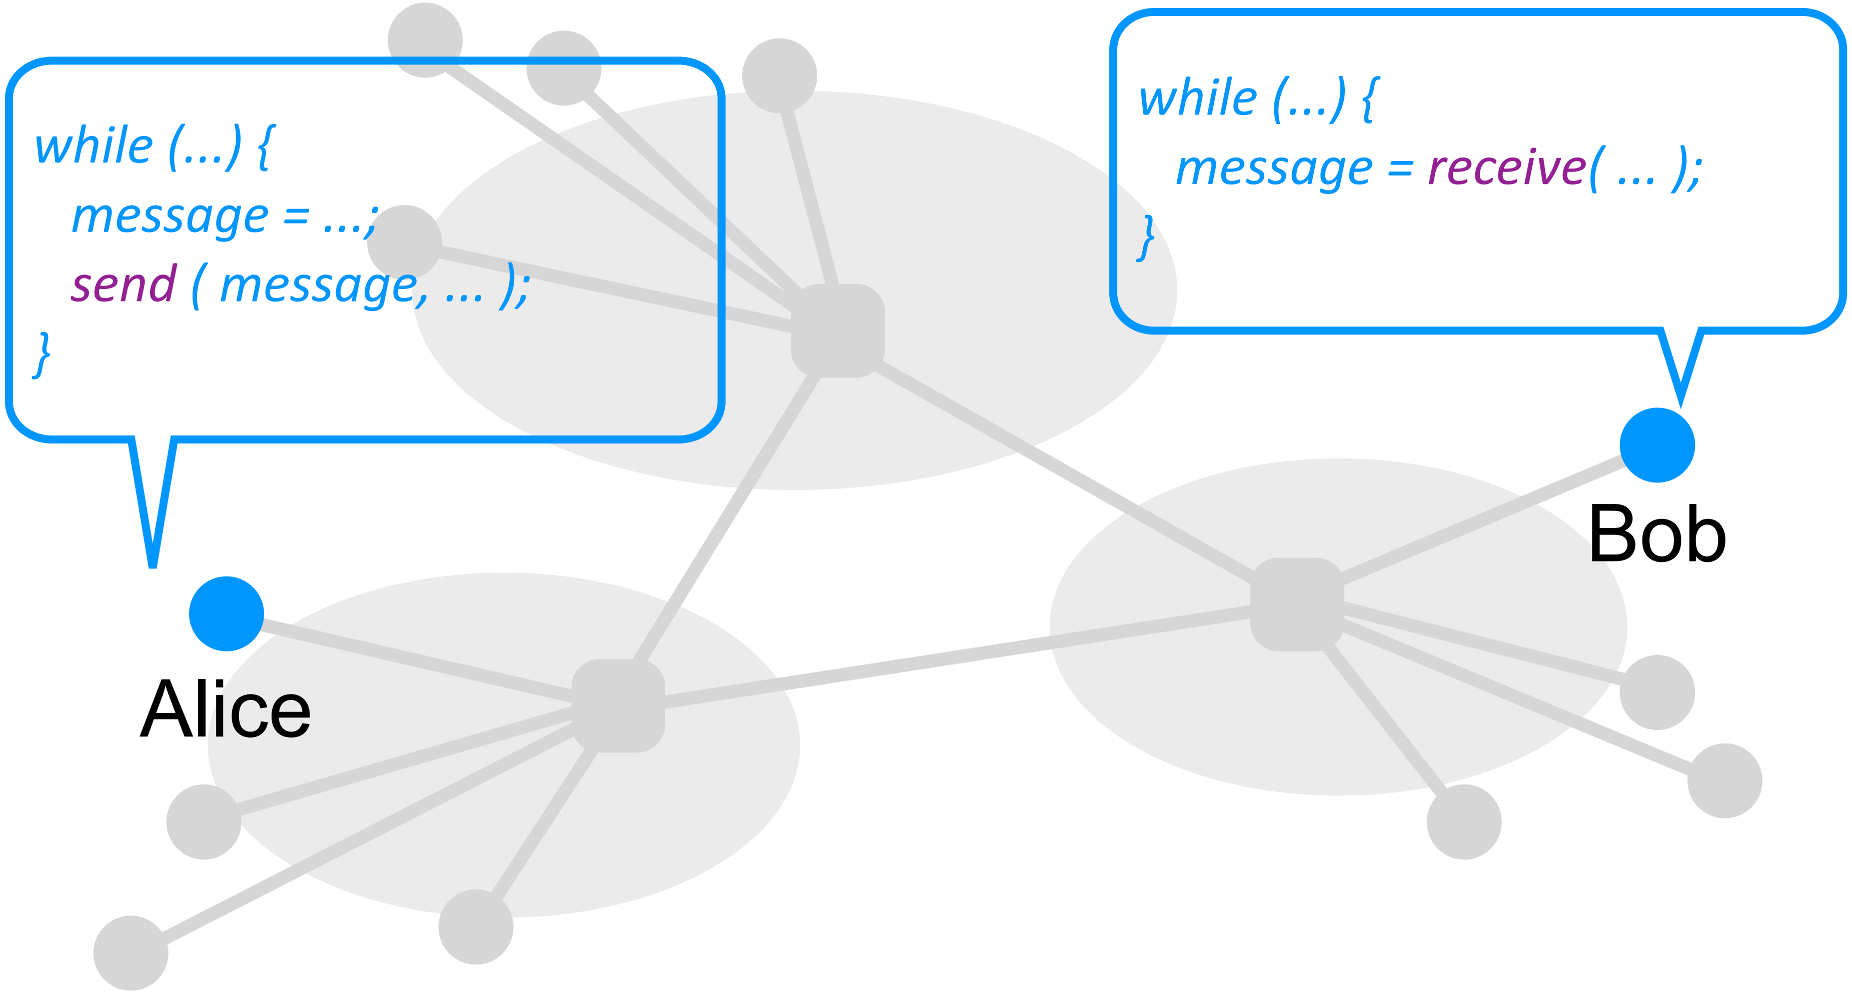
while (...) {
message = ...;
send (message, ...);
}
Alice
while (...) {
message = receive(...);
}
Bob

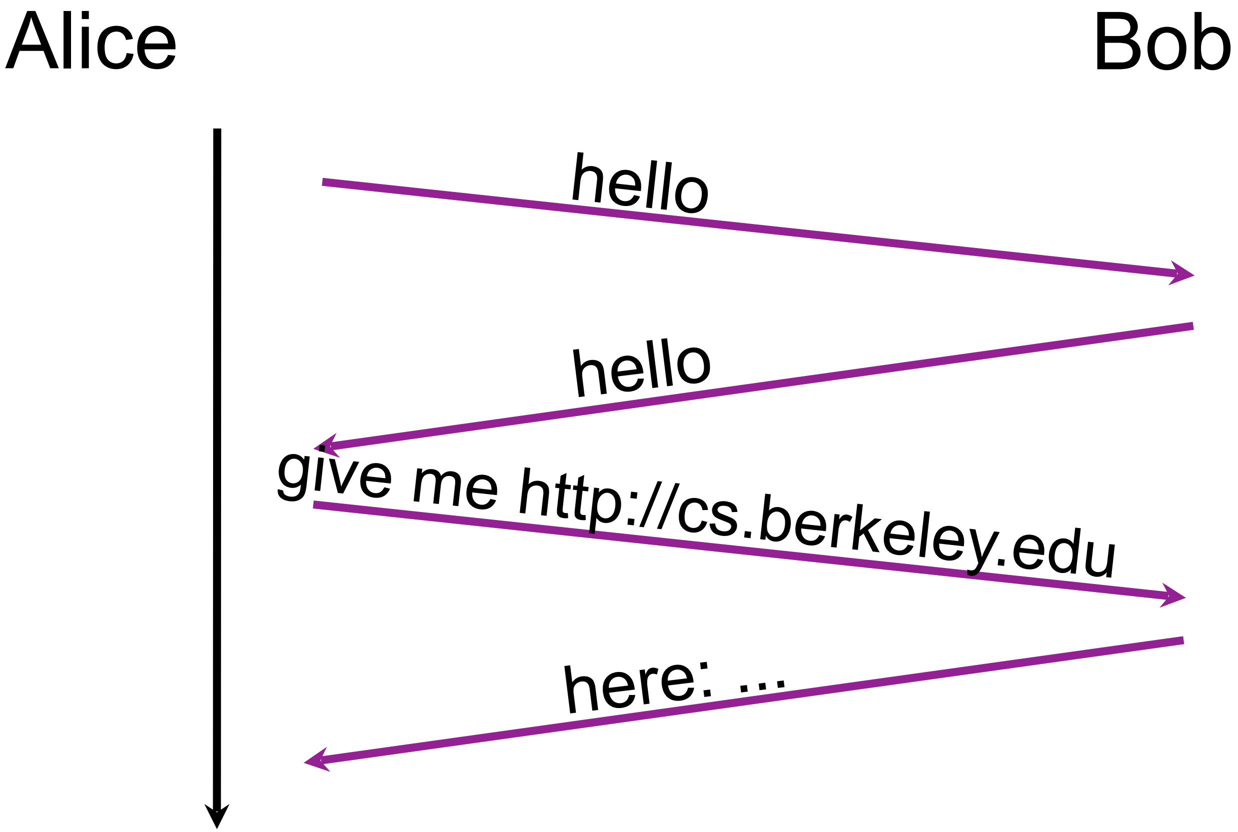
Alice
Bob
hello
hello
give me http://cs.berkeley.edu
here: ...

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

Protocol

- A specification of the messages that communicating entities exchange
    - their syntax and semantics
- Very much like conversational conventions … determining who should talk next and how they should respond

- Designing a good protocol is harder than it first seems!

- Internet
- Protocols
- **Architecture**

Why study the Internet?

The Internet has and is transforming everything

- **The way we do business …**
  - retail, advertising, cloud computing
- **The way we have relationships**
  - Twitter, chat
- **The way we learn**
  - Wikipedia, ChatGPT, AR/VR
- **The way we govern**
  - E-voting, censorship, cyber-warfare
- **The way we cure disease**
  - digital health, remote surgery

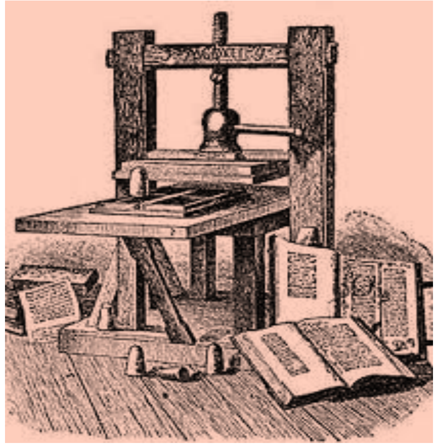

What’s your formal model for the Internet? -- theorists

Aren’t you just writing software for networks? – OS community

But why is the Internet *interesting*?

You don’t have performance benchmarks??? – hardware folks

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

A few defining characteristics of the Internet…

Network versus “The Internet”

- **There are many kinds of network technologies (switches and links)**
  - Ethernet, optical, wifi access points, DSL modems, Infiniband switches, …

- **The Internet is not a new/particular kind of network technology**

- **Instead, the Internet ties different networks together**
  - The <u>Inter</u>net

A federated system

Interoperability is the Internet's most important goal!

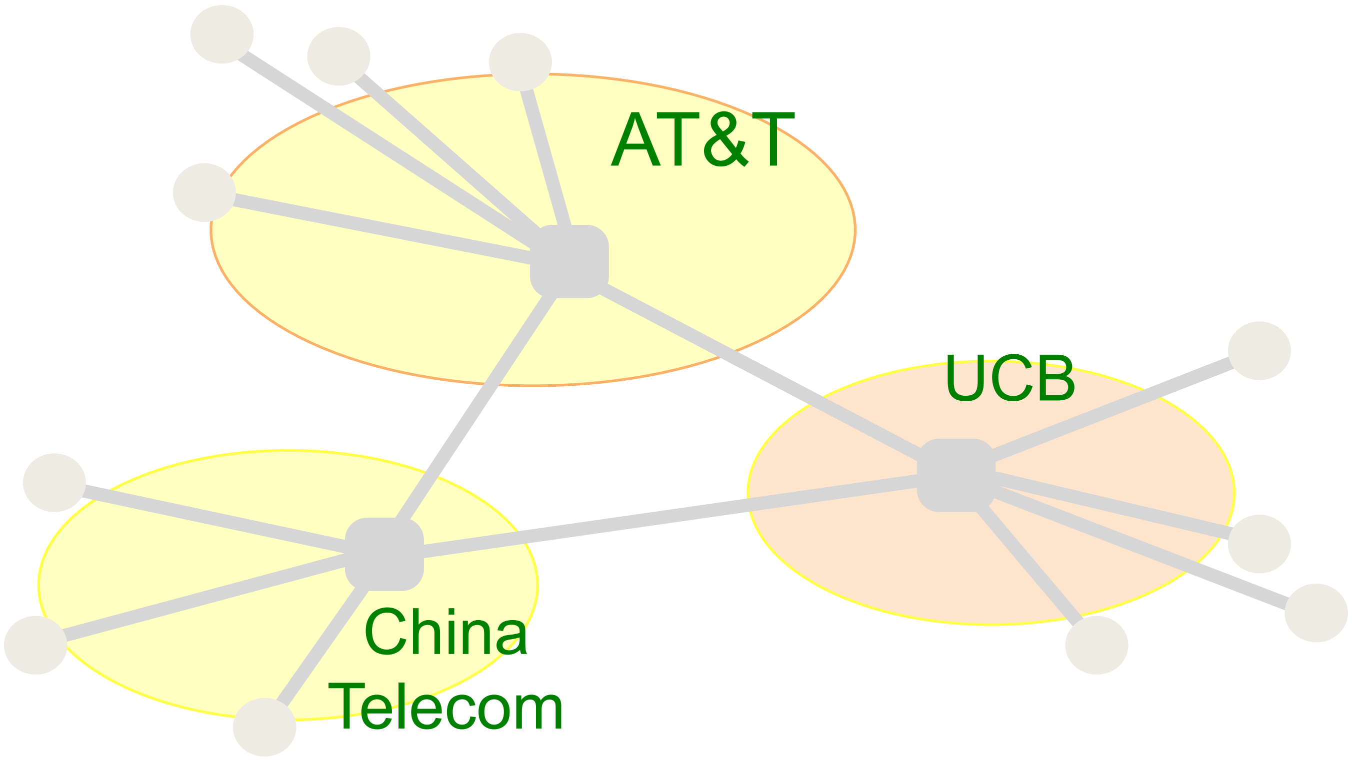

The Internet interconnects over 100,000 independently operated networks

A federated system

- **Fundamental challenge: how do you interconnect competing entities?**
  - Competing network providers must cooperate to serve their customers!
- **Leads to a constant tussle between business and technical factors**
  - Real-world incentives determine topology, path selection, diagnostics, and more
- **And complicates innovation**
  - How do you differentiate when interoperability relies on supporting a common protocol?
  - Upgrading “the Internet” is not an option

Tremendous scale

- > 5 Billion users (> 50% of world population)
- 1.24 Trillion unique URLs (web pages)
- Every second, we generate >10000 tweets, >100,000 Google queries, >3M emails

Enormous diversity and dynamic range

- **Technologies**: optical, wireless, satellite, copper, …
- **Communication latency:** microseconds to seconds (10^6 operating range)
- **Bandwidth:** 1Kbits/second to 1 Terabit/second (10^8 operating range)
- **Reliability**: 0 – 90%
- **Devices:** sensors, cell phones, datacenters, …
- **Users**: the governing, governed, operators, malicious, ...
- **Applications:** skype, live video, gaming, remote medicine, ...

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 for a message it sends to a server in NY?
  - Berkeley to New York: 4,125 km
  - Traveling to NY and back at 300,000 km/s: 27. 5 milliseconds
  - 3,000,000,000 cycles/sec * 0.0275 = 84,000,000 cycles!
- Thus, communication feedback is always **dated**

Prone to Failure

- Many components along a path
  - software, switches, links, network interface cards, wireless access points, modem,…
- Consider: 50 components, that work correctly 99% of time → 39.5% chance communication fail
  - Plus asynchrony → takes a long time to hear (bad) news

Constant evolution

1970s:

- 10^4 bits/second links
- < 100 computers in the US
- File transfer is the “killer” app

Today

- 10^{14} bits/second links
- 10B+ devices, all over the globe
- 3B+ facebook users; self-driving ca

Yet change must be backward compatible, incremental, and “in place”

Recap: The Internet is …

- A federated system …
- of enormous scale …
- with tremendous dynamic range and diversity …
- that is asynchronous in operation …
- failure prone ...
- **and constantly evolving**

Recap: The Internet is …

- Too complex for theoretical models
- “Working code” needn’t mean much
- Performance benchmarks are too narrow

The creation of the Internet required a new design paradigm

The Internet design paradigm

- Decentralized control
- A best-effort service model
- “Route around trouble”
- Dumb infrastructure (w/ smart endhosts)
- The end-to-end design principle
- Layering
- Federation via a “narrow waist” interface

A radical departure from systems at the time

Example: a best-effort service model

- **Fundamental question: what's the right service model that a network should support?**
  - "contract" between network and its users/end-hosts

- **Some possibilities:**
  - "guarantee that data will be delivered"
  - "guarantee that data will be delivered within X time"
  - "return a confirmation of successful delivery or an error"

- **Instead, what the Internet supports: "best effort" delivery of data**
  - No guarantee on whether or when data will be delivered
  - No notification of outcome!

The Internet design paradigm

- Decentralized control
- A best-effort service model
- “Route around trouble”
- Dumb infrastructure (w/ smart endpoints)
- The end-to-end design principle
- Layering
- Federation via a “narrow waist” interface

A radical departure from systems at the time

Now the de-facto blueprint for scalable services

The Internet design paradigm

- Decentralized control → **SDN: centralize?** → **dSDN: (re)decentralize?**
- A best-effort service model → **"quality of service" guarantees?** → **Nvidia's Infiniband**
- "Route around trouble"
- Dumb infrastructure (w/ smart endpoints) → **in-network attack detection?**
- The end-to-end design principle → **Edge computing?**
- Layering → **cross-layer optimizations**
- Federation via a "narrow waist" interface

But it is just one design …

… that is *constantly* being questioned

Backing up a level

- **The Internet poses a design challenge like no other**
- **From its creation emerged a new design paradigm**
- **That shaped how we reason about the design of scalable systems**
  - What's the right prioritization of goals?
  - What are fundamental constraints?
  - How do we decompose a problem?
  - What abstractions do we need?
  - What are the tradeoffs?
- **In short, a lesson in how to <u>architect</u> a (networked) system**

- Internet
- Protocols
- **Architecture**

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!

Class topics, more concretely

Reflect three broad phases in the Internet's evolution

1. Building a global data communication network
2. Scaling communication; and the emergence of a commercial ecosystem
3. (Networks that enable) scaling data; and a shifting commercial ecosystem

Phase 1: Building a global data communication network

Military interest in a communication infrastructure to withstand a nuclear attack

Scientists want to share data

1960

Concepts

Building a testbed

Scaling the testbed

Commercialization

1995

packet switching

statistical multiplexing

layering

routing

reliability

naming (DNS)

congestion control

Impact: transformed how humans communicate

Phase 2: Scaling & the emergence of a commercial ecosystem

"Content is king" – *Bill Gates*

Proliferation of PCs

1995

Commercialization

Exponential scaling

2005

Policy (BGP)

Content retrieval (HTTP)

Security (firewalls, DoS)

Caching (CDNs, load balancers)

Scalable addressing (NA

High-speed routers

Impact: everything moves to the Internet (content, brick-and-mortar businesses, banks, *etc.*)

Phase 3: Data and a shifting ecosystem

2005

2020

Big Data & AI

Always-on users

Datacenters & cloud

Mobile access

Fat-tree networks

Network virtualization

Low-latency commn. (RDMA)

Software-defined networking (SDN)

WiFi

Cellular networks

Impact: The Internet is everywhere, transforming everything

To recap, what we hope CS 168 will teach you

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

Today

- Introductions
- What is (this course on) the Internet about?
- **Class logistics**