Networking Introduction

CSE 333 Summer 2020

Instructor: Travis McGaha

Teaching Assistants:

Jeter Arellano
Ramya Challa
Kyrie Dowling
Ian Hsiao
Allen Jung
Sylvia Wang

Poll Everywhere

pollev.com/cse33320su

About how long did Exercise 13 take?

A. 0-1 Hours
B. 1-2 Hours
C. 2-3 Hours
D. 3-4 Hours
E. 4+ Hours
F. I didn't submit / I prefer not to say

Side question:
Best movie?

Administrivia

- hw3 is due Thursday (8/6)
- Rest of the quarter:
  - Topics: Networking, Concurrency, Processes
  - hw3: file search shell
  - hw4: file search shell over the network
- ex15 released on Friday
  - You will write a program that sends data over TCP.
  - After ex15, 2 remaining: TCP listener, concurrency

Lecture Goal

- Networking is a very common area to interact with. You will likely have to create a program that will read/write over the network at some point in your career.

- This course can't go over everything, but we want to give you some understanding of how the network works.
  - (Take CSE 461 if you want to know more, the course is pretty cool ☺)

- Lecture will be more "story-like". Please feel free to take notes, but no need to memorize everything that is said.

Lecture Outline

- **Introduction to Networks**
  - **Layers upon layers upon layers…**

Networks From 10,000 ft

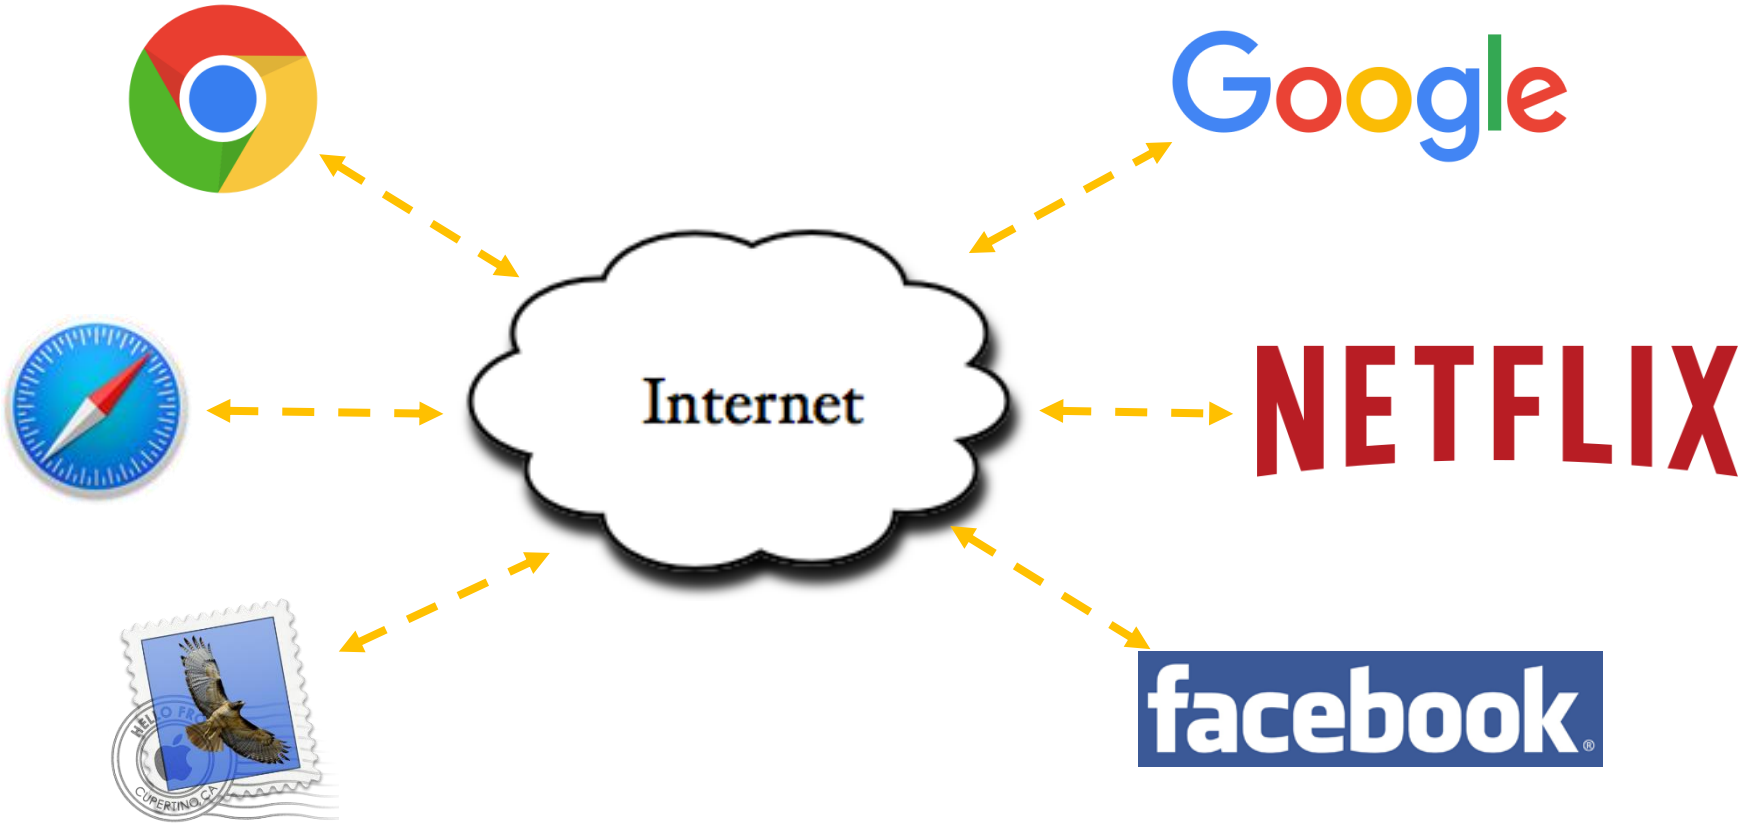

The Physical Layer

- Individual bits are modulated onto a wire or transmitted over radio
  - Physical layer specifies how bits are encoded at a signal level
  - Many choices, e.g., encode "1" as +1v, "0" as -0v; or "0"=+1v, "1"=-1v, …

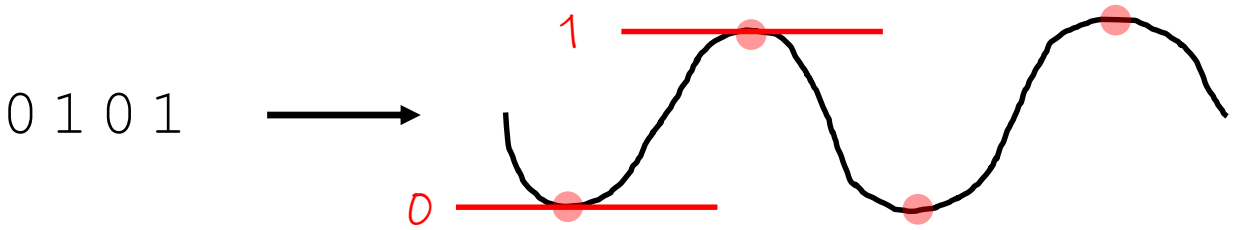

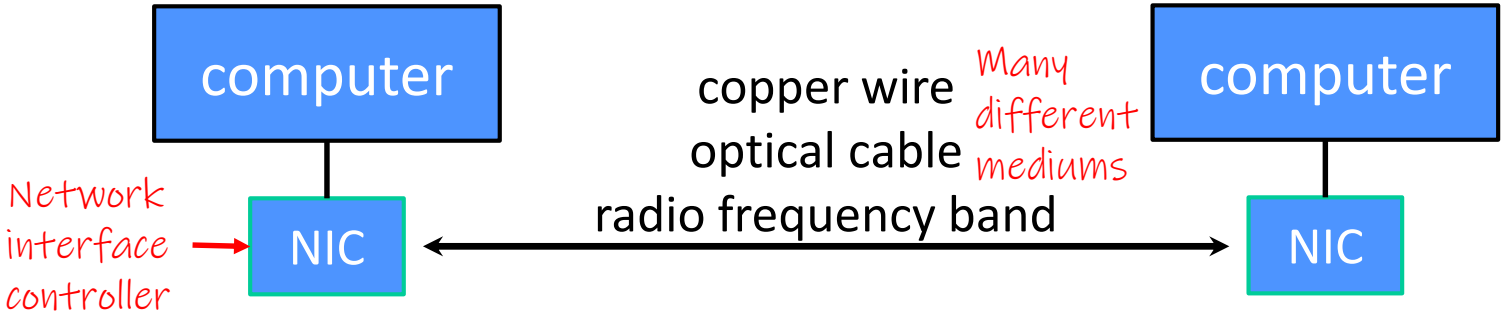

The Data Link Layer

- Multiple computers on a LAN contend for the network medium
  - Media access control (MAC) specifies how computers cooperate
  - Link layer also specifies how bits are "packetized" and network interface controllers (NICs) are addressed

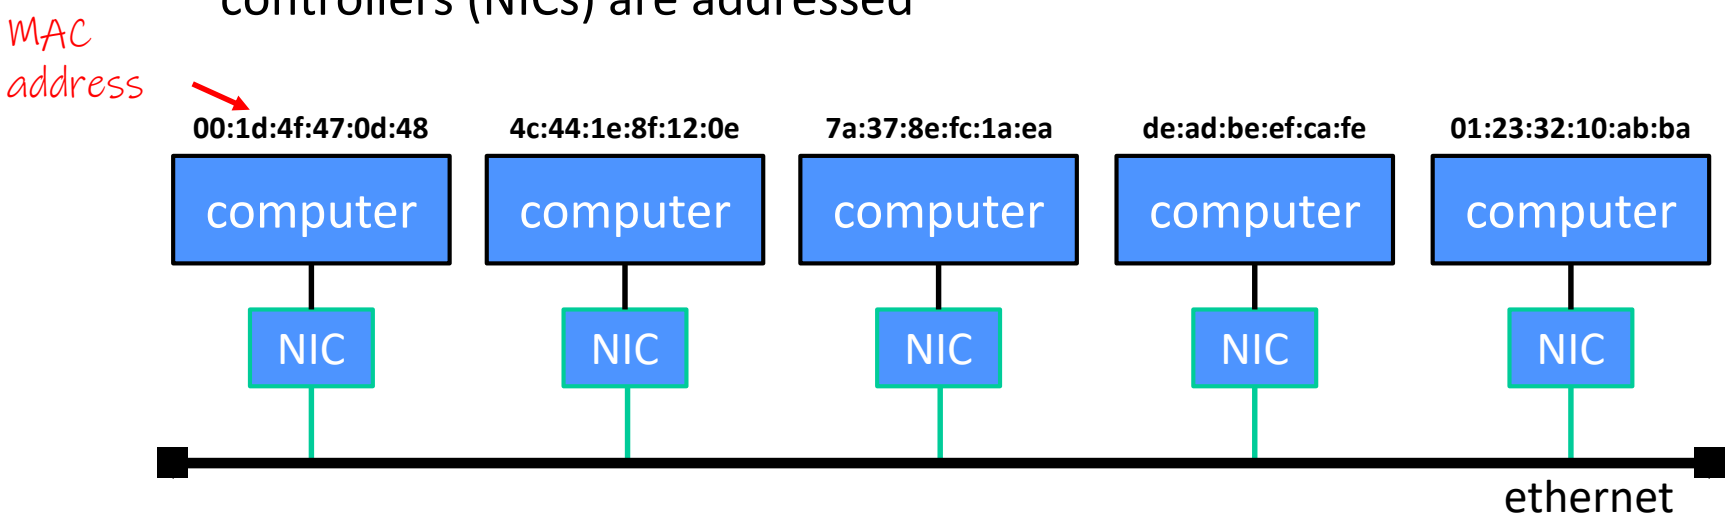

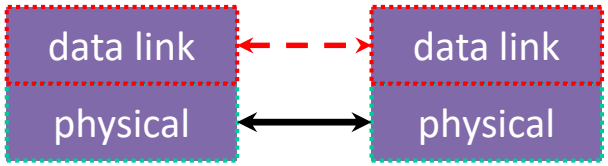

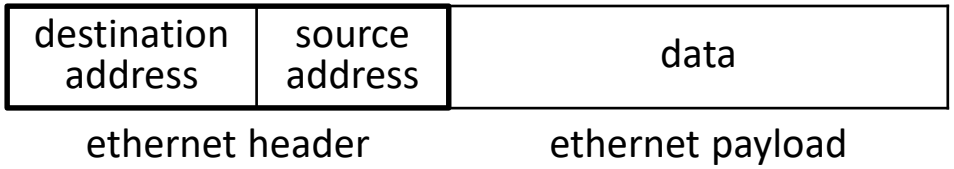

The Network Layer (IP)

- Internet Protocol (IP) routes packets across multiple networks
  - Every computer has a unique IP address
  - Individual networks are connected by routers that span networks

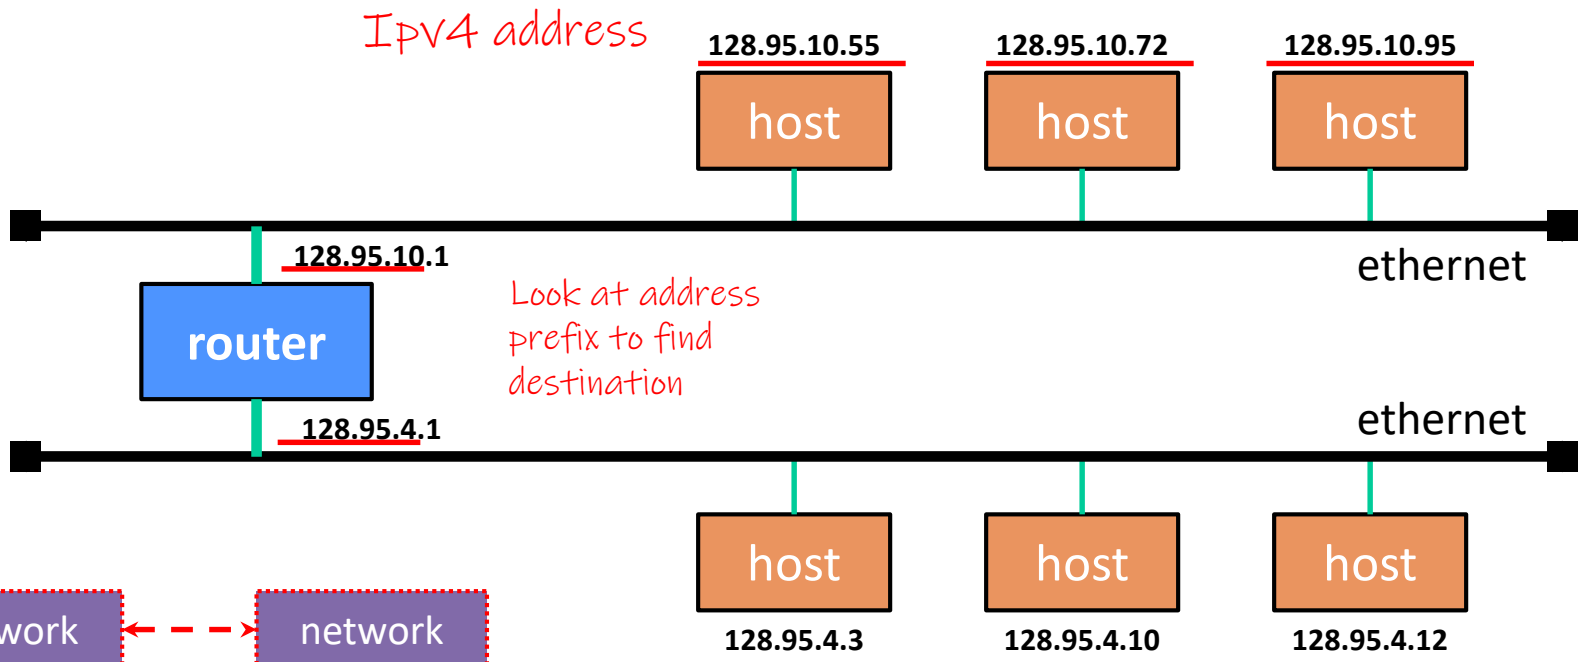

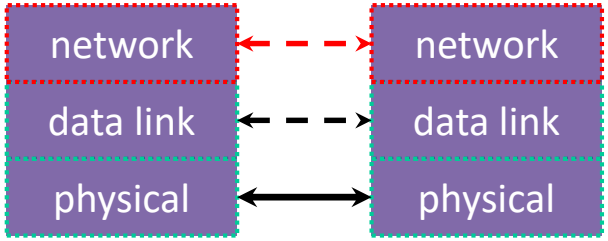

The Network Layer (IP)

- There are protocols to:
  - Let a host map an IP to MAC address on the same network
  - Let a router learn about other routers to get IP packets one step closer to their destination

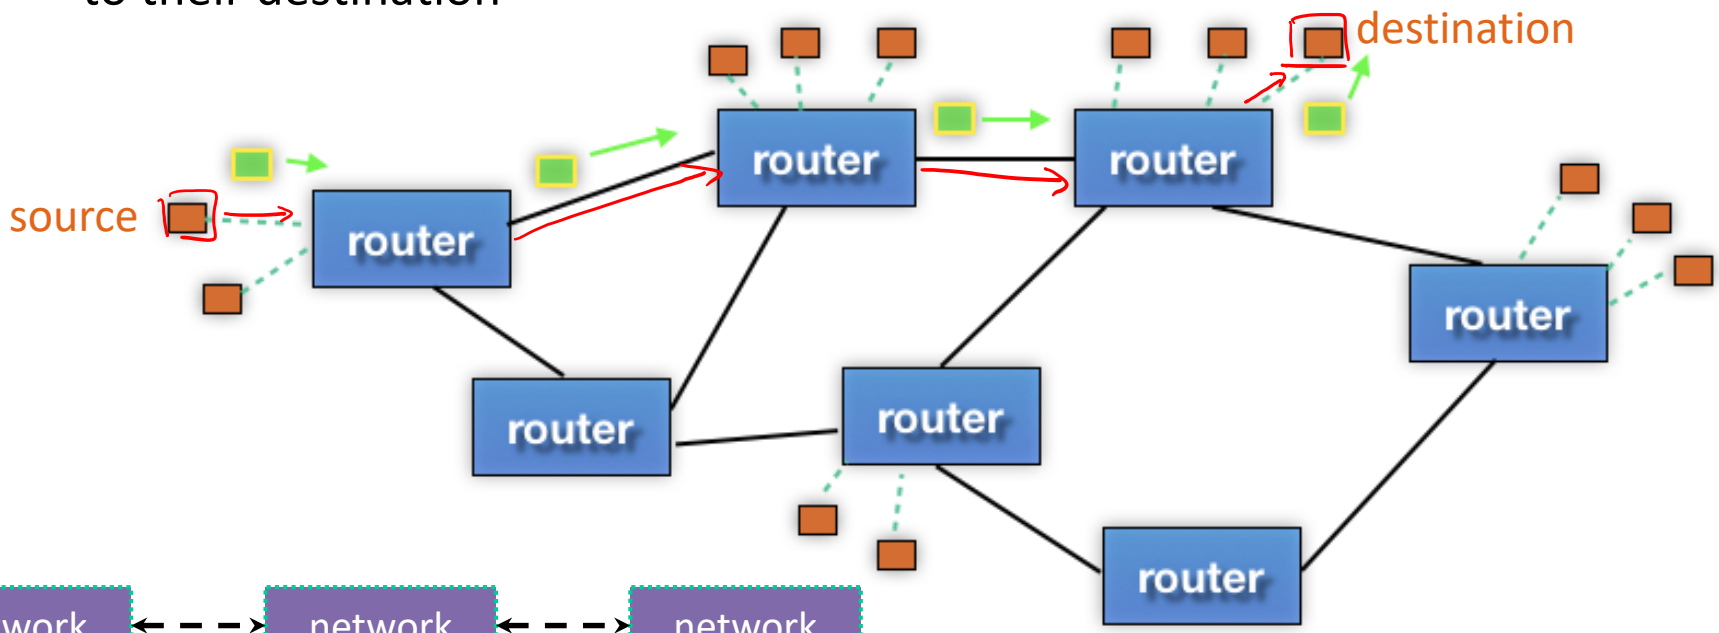

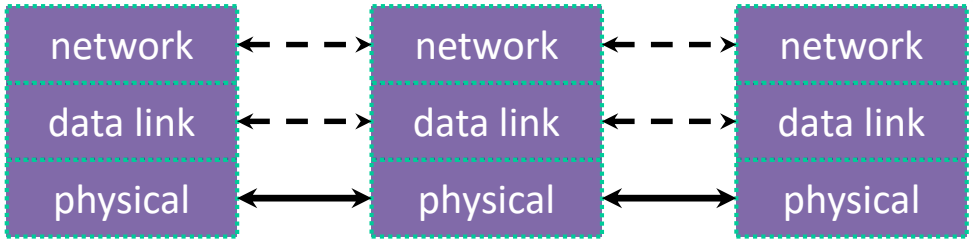

The Network Layer (IP)

- Packet encapsulation:
  - An IP packet is encapsulated as the payload of an Ethernet frame
  - As IP packets traverse networks, routers pull out the IP packet from an Ethernet frame and plunk it into a new one on the next network

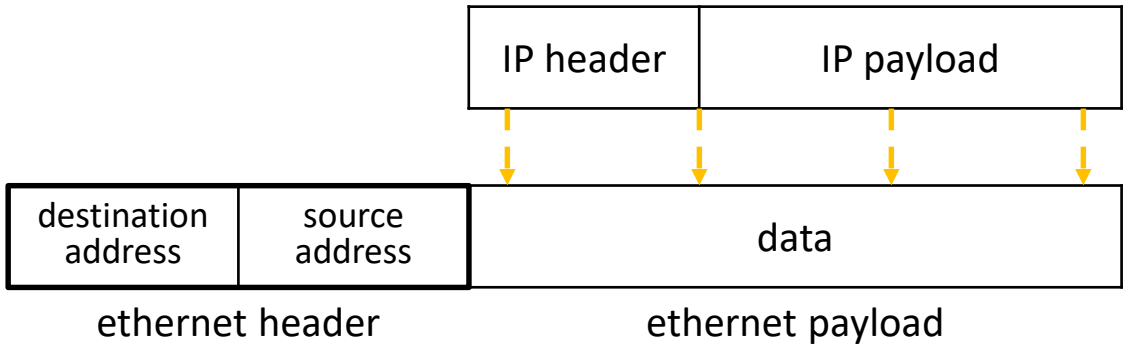

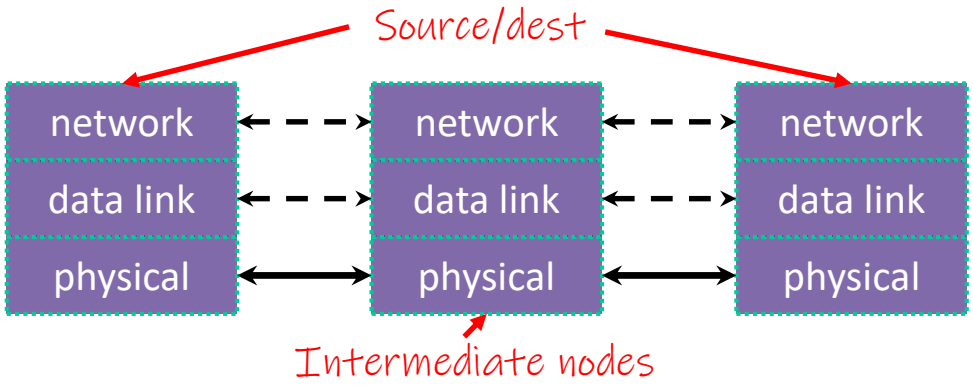

Reliability

- Packet loss?
- Physical Layer interference?
- Link going down?

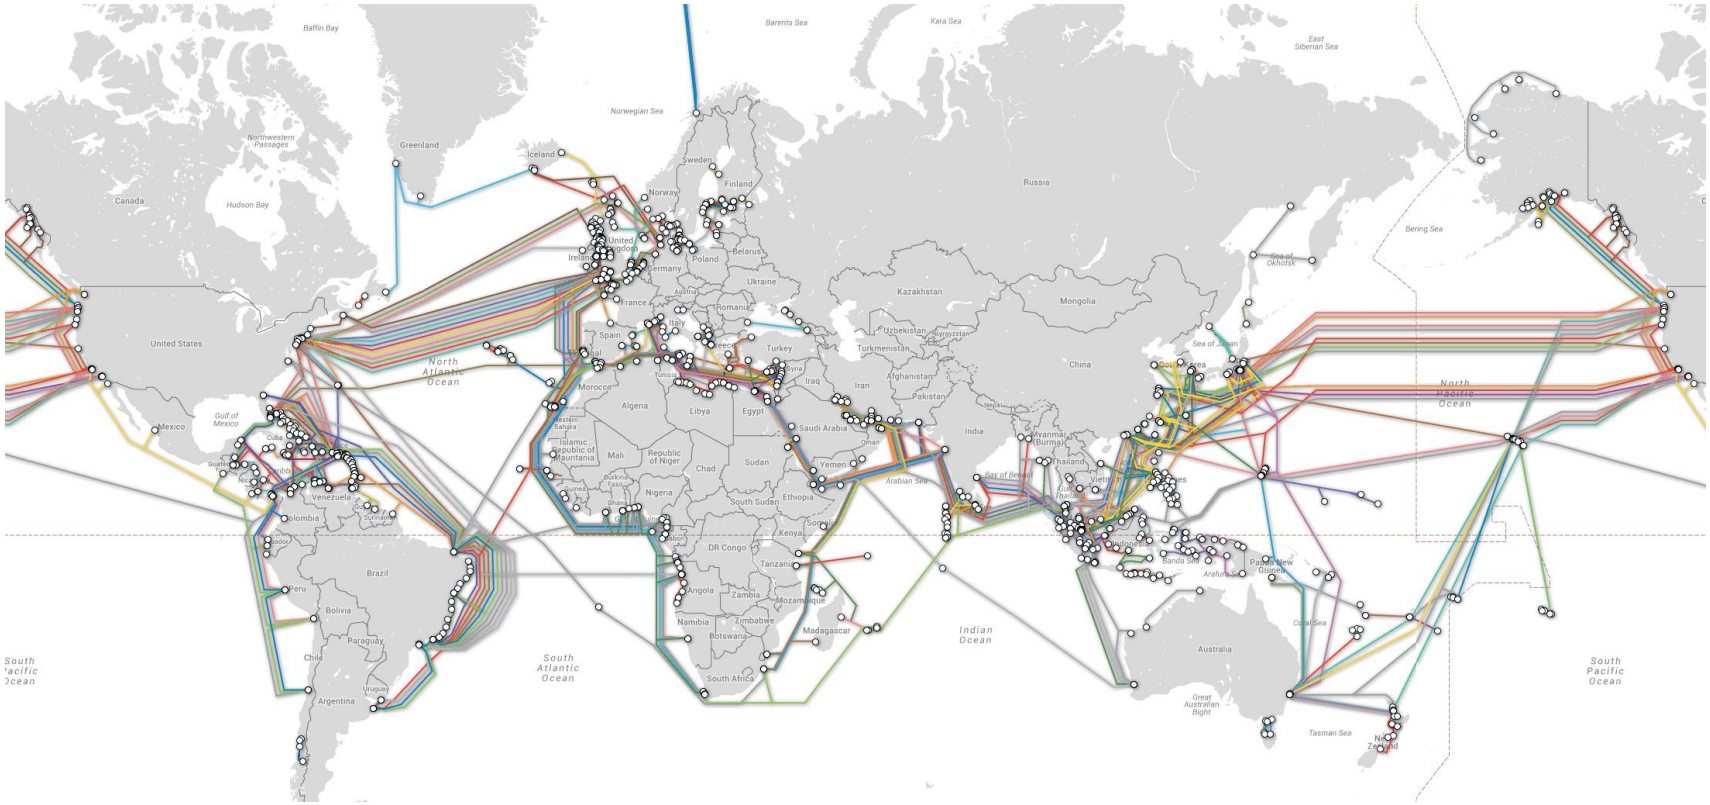

The Transport Layer

- Provides an interface to treat the network as a data stream
- Provides different protocols to interface between source and destination:
  - TCP - Transmission Control Protocol
  - UDP - User Datagram Protocol
- These protocols still work with packets, but manages their order, reliability, multiple applications using the network…

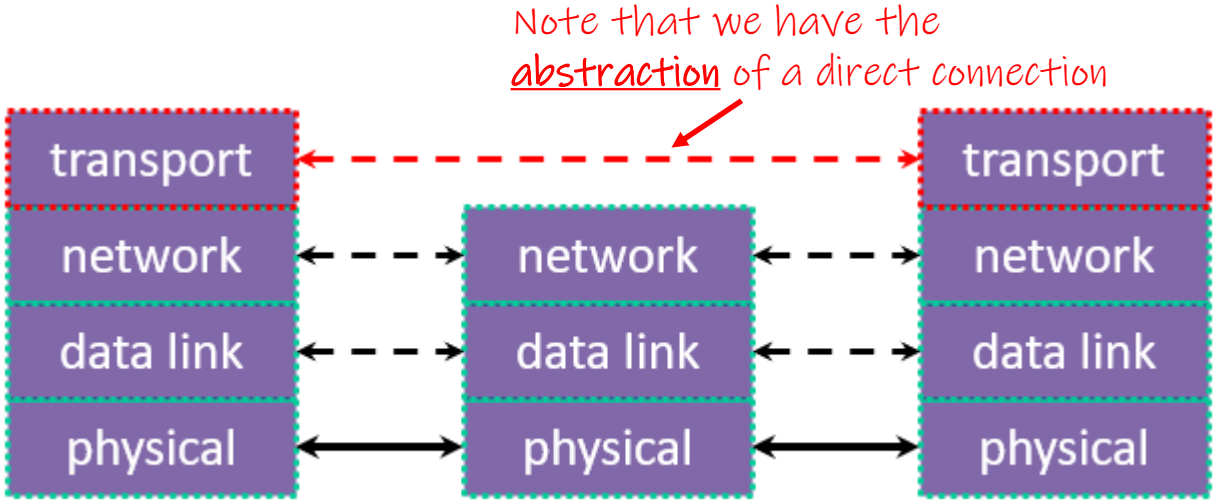

The Transport Layer (TCP)

- Transmission Control Protocol (TCP):
  - Provides applications with reliable, ordered, congestion-controlled byte streams
    - Sends stream data as multiple IP packets (differentiated by sequence numbers) and retransmits them as necessary
    - When receiving, puts packets back in order and detects missing packets
  - A single host (IP address) can have up to 2^{16} = 65,535 "ports"
    - Kind of like an apartment number at a postal address (your applications are the residents who get mail sent to an apt. #)

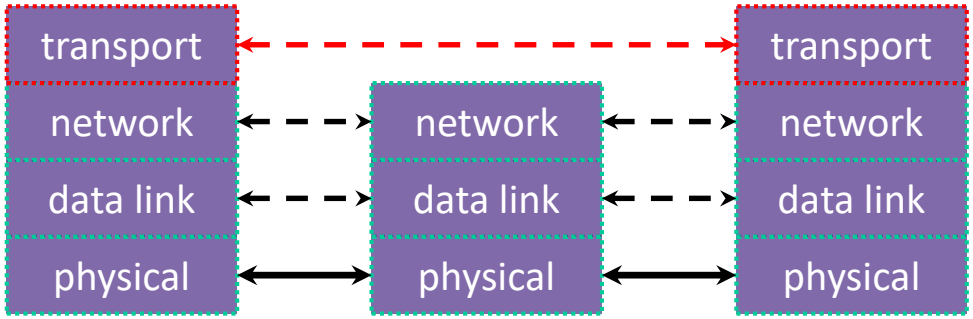

The Transport Layer (TCP)

- Packet encapsulation – one more nested layer!

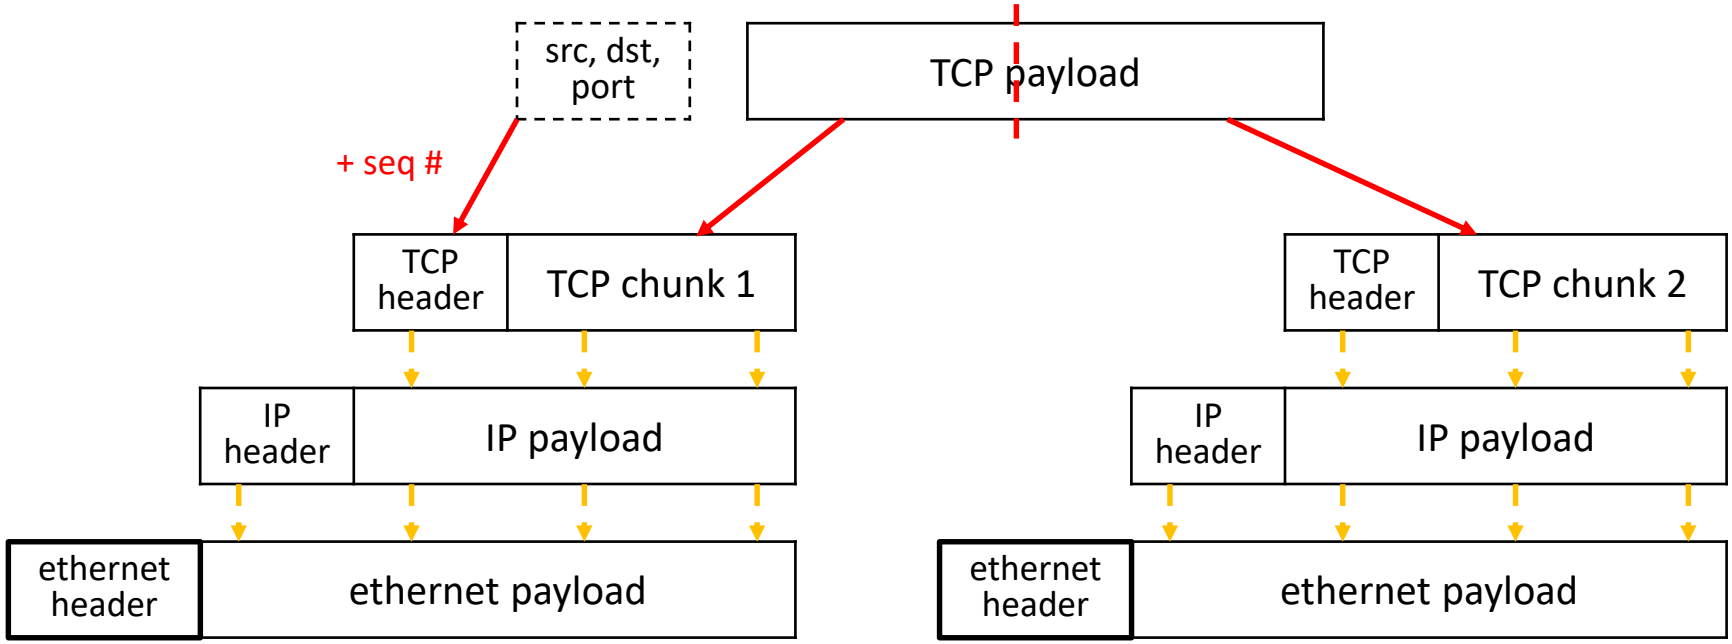

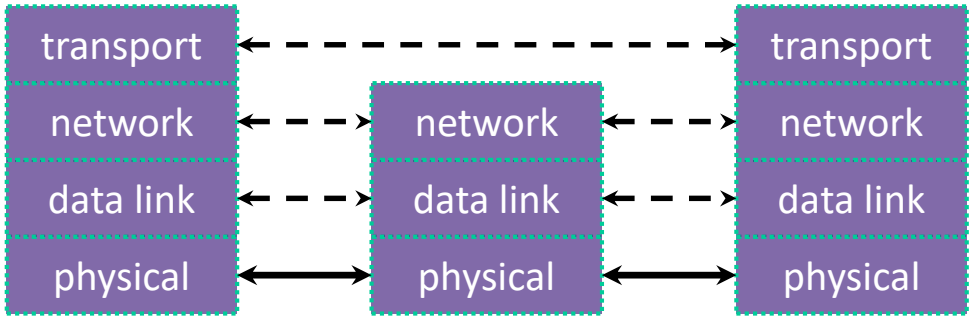

The Transport Layer (TCP)

- Applications use OS services to establish TCP streams:
  - The "Berkeley sockets" API
    - A set of OS system calls (Part of POSIX on linux)
  - Clients `connect()` to a server IP address + application port number
  - Servers `listen()` for and `accept()` client connections
  - Clients and servers `read()` and `write()` data to each other

Used same as in File I/O

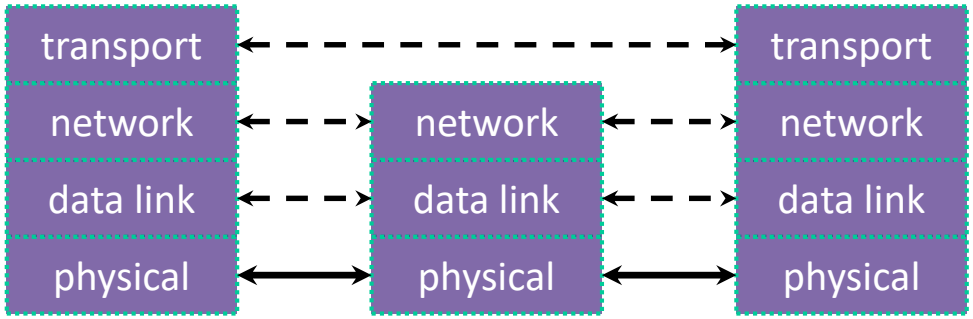

The Transport Layer (UDP)

- User Datagram Protocol (UDP):
  - Provides applications with *unreliable* packet delivery
  - UDP is a really thin, simple layer on top of IP
    - Datagrams still are fragmented into multiple IP packets

Ok when we want speed. (VOIP or ZOOM)

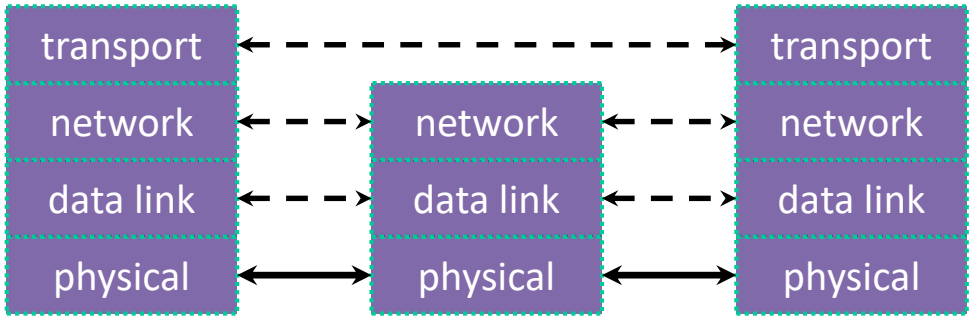

The (Mostly Missing) Layers 5 & 6

- Layer 5: Session Layer
  - Supposedly handles establishing and terminating application sessions
  - Remote Procedure Call (RPC) kind of fits in here
- Layer 6: Presentation Layer
  - Supposedly maps application-specific data units into a more network-neutral representation
  - Encryption (SSL) kind of fits in here

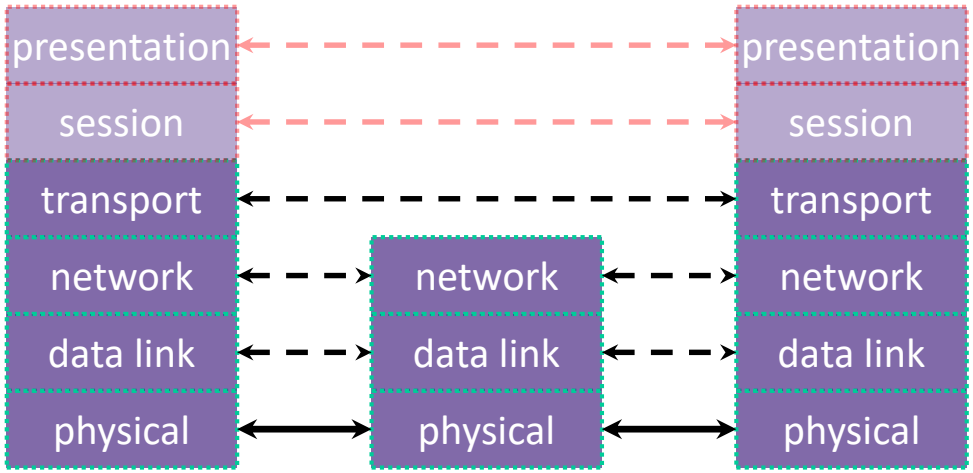

The Application Layer

- Application protocols
  - The format and meaning of messages between application entities
    - Example: HTTP is an application-level protocol that dictates how web browsers and web servers communicate
      - HTTP is implemented *on top of* TCP streams

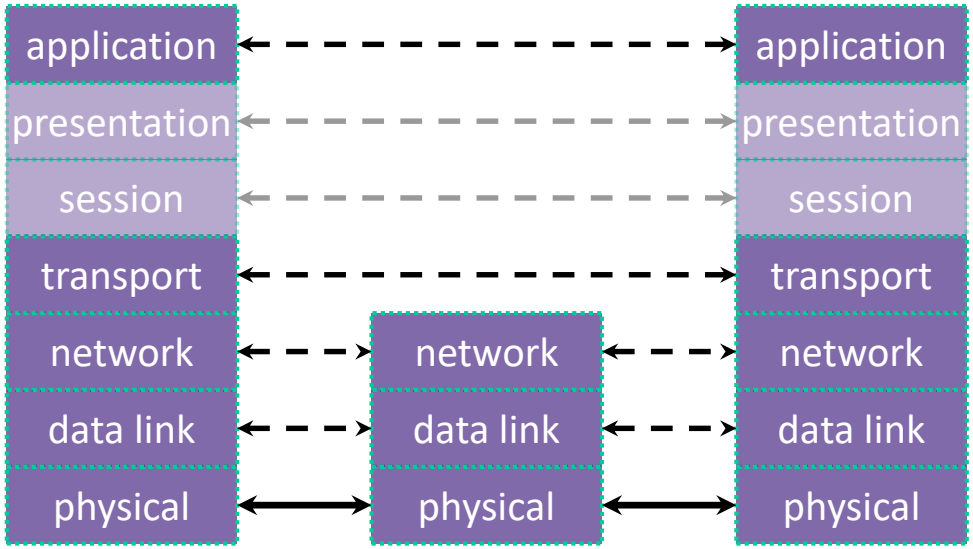

The Application Layer

- Packet encapsulation:

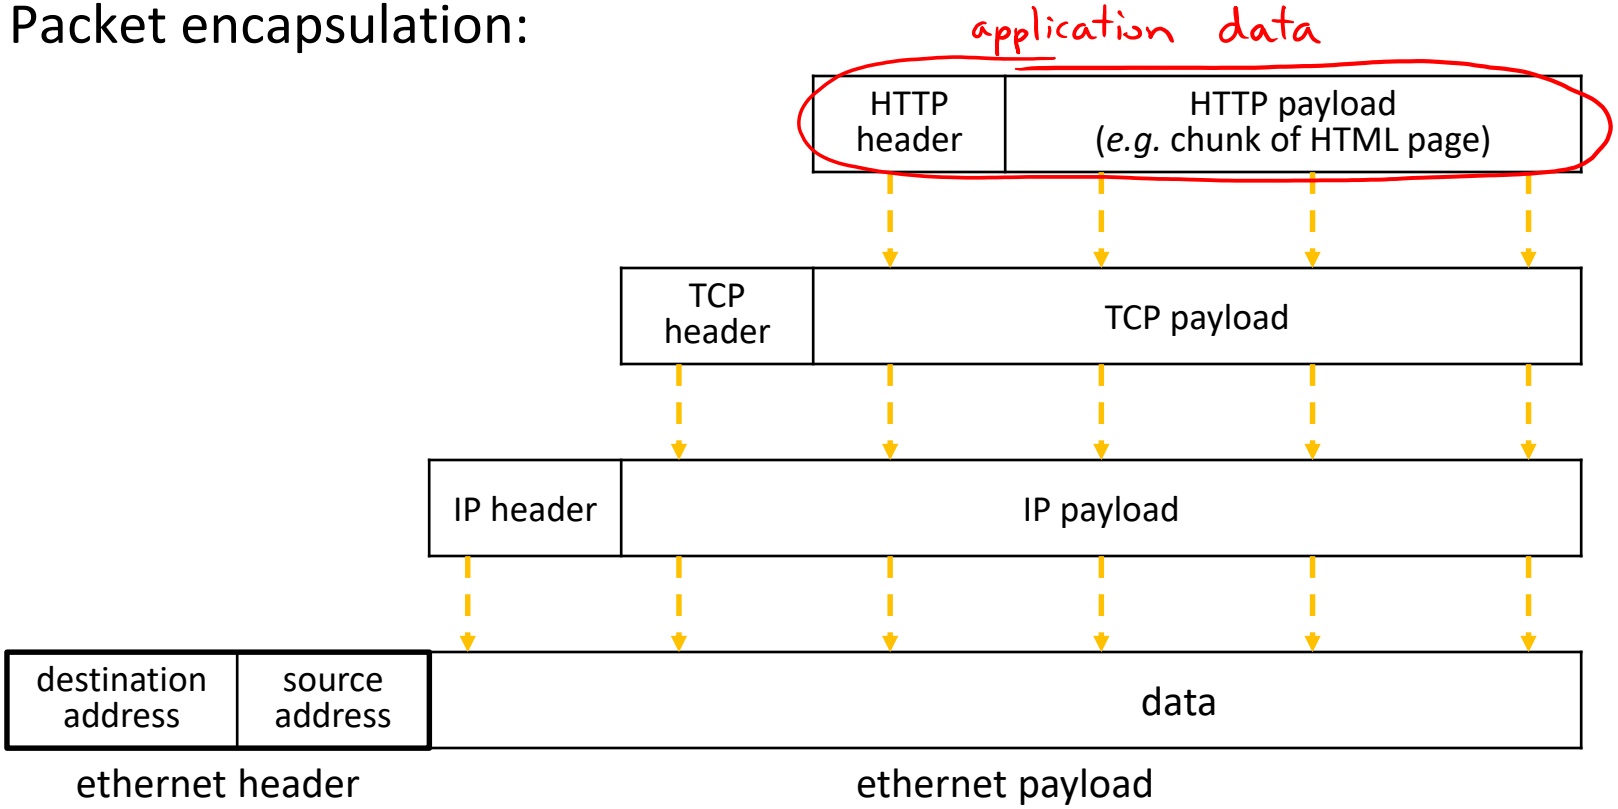

The Application Layer

- Packet encapsulation:

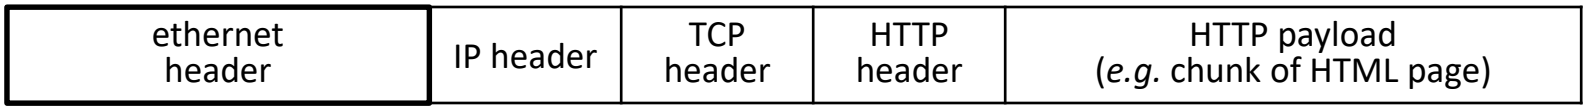

The Application Layer

- Popular application-level protocols:
  - **DNS:** translates a domain name (*e.g.* www.google.com) into one or more IP addresses (*e.g.* 74.125.197.106)
    - Domain Name System
    - An hierarchy of DNS servers cooperate to do this
  - **HTTP:** web protocols
    - Hypertext Transfer Protocol
  - **SMTP, IMAP, POP:** mail delivery and access protocols
    - Secure Mail Transfer Protocol, Internet Message Access Protocol, Post Office Protocol
  - **SSH:** secure remote login protocol
    - Secure Shell
  - **bittorrent:** peer-to-peer, swarming file sharing protocol

In Other Words…

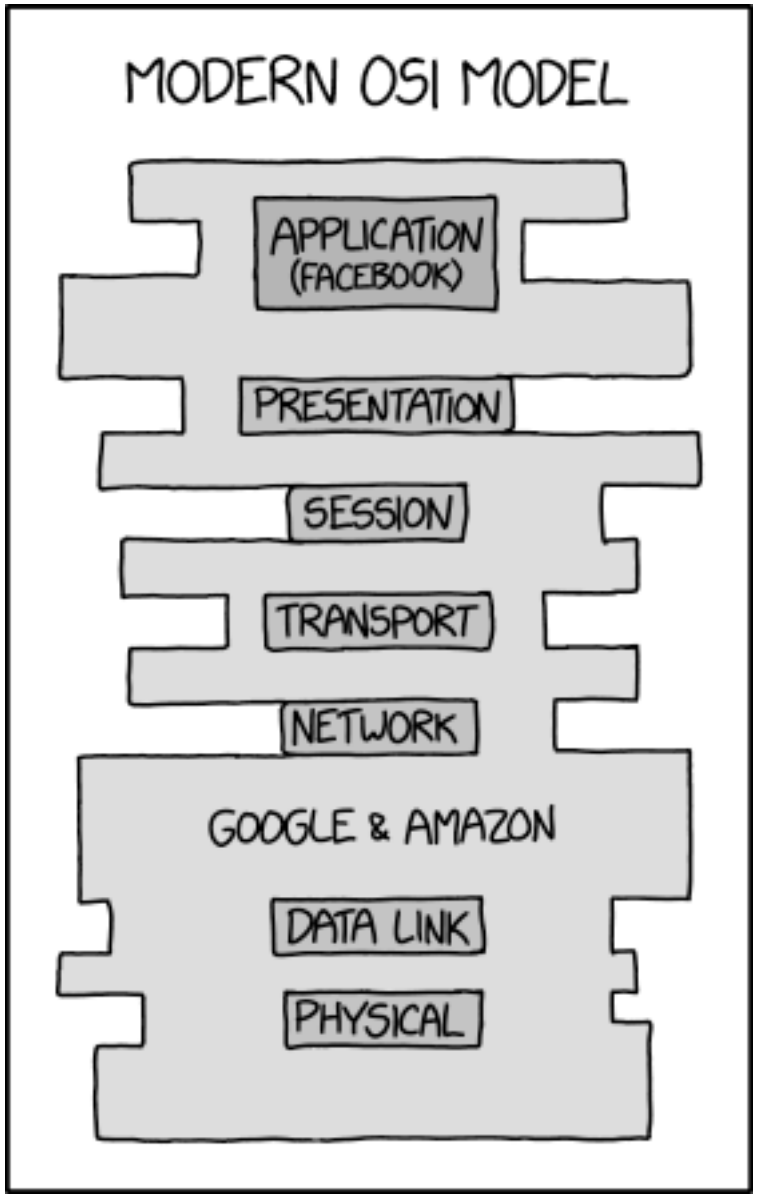

https://xkcd.com/2105/

"Network" Latency is Highly Variable

- Jeff Dean's "Numbers Everyone Should Know" (LADIS '09)

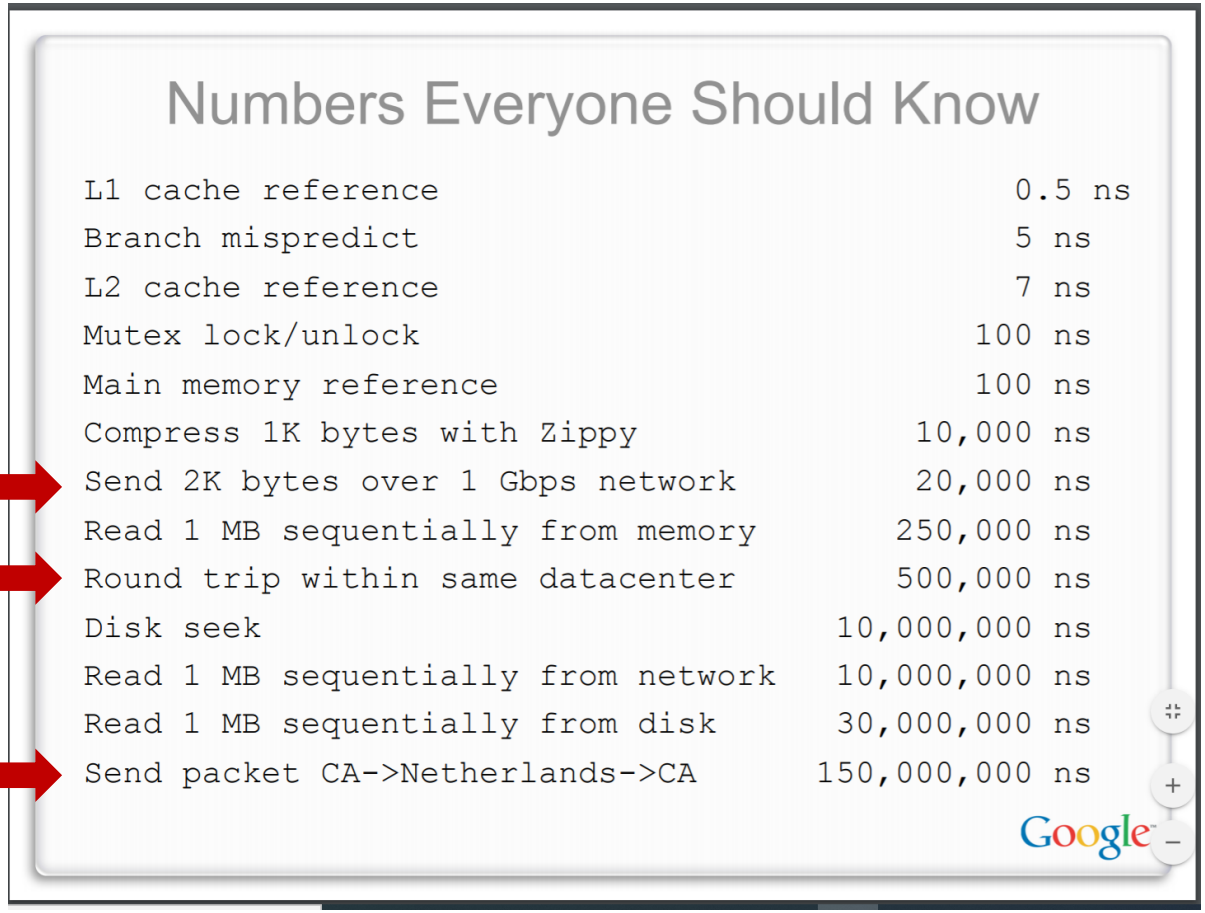

Numbers Everyone Should Know

L1 cache reference	0.5 ns
Branch mispredict	5 ns
L2 cache reference	7 ns
Mutex lock/unlock	100 ns
Main memory reference	100 ns
Compress 1K bytes with Zippy	10,000 ns
Send 2K bytes over 1 Gbps network	20,000 ns
Read 1 MB sequentially from memory	250,000 ns
Round trip within same datacenter	500,000 ns
Disk seek	10,000,000 ns
Read 1 MB sequentially from network	10,000,000 ns
Read 1 MB sequentially from disk	30,000,000 ns
Send packet CA->Netherlands->CA	150,000,000 ns

Google

Latency: Materials Matter

- Fiber optic cables are lower-latency and higher-bandwidth than traditional copper wiring
  - Much of the internet's "long haul" data is transmitted on these
  - (signal attenuation is much better too)

- Is it faster to send 1 person from UW to …
  - Downtown Seattle?
  - Downtown Ballard?

Affected by not only distance, speed limit & number of lanes.
Using I-5 vs a local road

Latency: Distance Matters

- Distances within a single datacenter are smaller than distances across continents
- Even within a datacenter, distances can sometimes matter

123Net Data Center, Wikimedia

Latency: Topology Matters

- Some places are surprisingly well- or poorly-connected to "backbone" infrastructure like fiber optic cables

- Unintuitive topology creates interesting failures
  - *e.g.*, 2006 Hengchun Earthquake disrupted communications to Singapore, Philippines, Thailand, China, etc. for a month