CS244

Advanced Topics in Networking

Stuff you should know about CS244

Changhoon Kim

Spring 2022

In-person Class!

- Let’s take advantage of being in the same room:
  High-bandwidth, low-latency communication!
- Be active, engaging, and inquisitive
- The Zoom session is for restricted use only;
  let me know ahead if you have legitimate reasons for having to use it

Goals

1. To become familiar with the field of networking research: *Network architecture, systems and programming*.
2. To practice the art of reading and critiquing research papers.
3. To learn the art of reproducing research results.

It's a big field, so we have
to focus on just a few topics.

Basics

Class Time

- Each class we will discuss 1-2 papers
- You must read the papers before class (super important!)
- You will write a critique before class
- Plan to read the papers carefully and in depth
- Most of the lecture will be spent on discussion
- (35% of your grade comes from critiques, and in-class participation)

So….read the papers, come to class, and be ready to participate

How grading works in Spring 2022

How we grade (so you know what we think is important)

Reading and participation 35%

- Critiques before class: 15%
- In-class participation: 20%

Programming assignments 65%

- PA #1: 15% Reproducing a particular research result with Mininet
- PA #2: 15% Reproducing the same research result with NS3
- Final Project: 35% EITHER an original research project
  OR reproducing a paper of your choice
  (not previously reproduced in CS244)

(Optional) Final exam

- We might ask all or some students to take a short oral exam as part of their final projects
- We will let you know by the end of Week 8
- Oral exams will cover the subject matter related to each student's final project

Please participate!

- Join online prepared to discuss the main ideas of the paper(s)
- We will all learn from each other
- Attendance and participation are vital parts of this class

Critiques

What to submit?

- Short critique for each paper before the class (by midnight the night before the lecture)
- Submit online (use Canvas)

Questions to answer while writing your critique:

- What problem are the authors solving?
- Why are the problems important?
- What is the main idea and what do you think of it?
- What was the *status quo ante* before this paper, and what is the clearest way to explain this paper's contribution?
- How well is the paper written?

Grade: 0, 1, or 2 points

Two Assignments

Goals:

Learn the dynamics of TCP sawtooth and router buffer occupancy in a network

Learn why large router buffers can lead to poor performance (**bufferbloat**)

Implementation:

PA1 - Mininet (emulator)

Released on: March 30, 2022

Due: April 15, Fri, 2022

PA2 - NS3 (simulator)

Released on: April 18, 2022

Due: April 29, 2022

Logistics

Who will lead the discussions

- Chang Kim (Changhoon.kim@stanford.edu)
- Several guest experts

TAs

- Sundar Renganathan (rsundar@stanford.edu)
- Preey Shah (preey@stanford.edu)

The Honor Code

- We take it seriously and we expect you to take it seriously too.
- In Fall 2020 about ten CS144 students got into a lot of trouble ☹
- Some set out to cheat on exams.
- Others didn't set out to cheat: At the last minute, under stress, they copied an assignment off the web, then tried to modify it.
  It doesn't work!
- We use special tools to compare solutions against current and previous years and solutions we find on the web.
- The tools work really well.
- Please, let's have a zero-violation year.

The Honor Code

Permitted Collaboration: The following items are encouraged and allowed at all times for all students in this class:

- Discussion of material covered during lecture, problem sessions, or in handouts
- Discussion of the requirements of an assignment
- Discussion of the use of tools or development environments
- Discussion of general approaches to solving problems
- Discussion of general techniques of coding or debugging

The Honor Code

Unpermitted Collaboration: All submissions must represent **original, independent work**. Some examples of activities that do not represent original work include:

- Copying solutions from others or knowingly allowing others to copy your solution.
- Use of solutions posted to websites is prohibited.
- Placing your source code in a public repository where others can copy it is unpermitted collaboration.
- Debugging code for someone else.
- Collaborating on or discussing the online graded quizzes before you have completed them.

Contact

Whenever possible: ~~Piazza~~ → Ed

- Official Q&A channel
- Someone else probably has the same question
- Please don't send questions to class email list
- If private: Post a private ~~Piazza~~ Ed post

Quick and pseudo-real-time conversation (best effort): Slack

- Please send email to me and the TAs right after the first class, and we'll invite you to the Slack channel

All extension requests should go to Chang.

The Internet: An Exciting Time

One of the most influential inventions

- A research experiment that escaped from the lab
- … to be the global communications infrastructure

Ever increasing reach

- Today
  - Over 4 billion users
  - All sorts of services powered by O(10^8) servers hosted typically in datacenters
- Tomorrow: more users, smartphones, computers, sensors, cameras…

Constant innovation

- Apps: Web, P2P, social networks, virtual worlds
- Links: optics, WiFi, cellular, 5G, ...

Transforming Everything

The ways we do business

- E-commerce, advertising, cloud computing, ...

The way we have relationships

- E-mail, IM, Facebook friends, virtual worlds

How we think about law

- Interstate commerce and sales tax, National boundaries, Wikileaks

The way we govern

- E-voting and E-government and fake news
- Censorship and wiretapping

The way we fight

- Cyber-attacks, including nation-state attacks

But what *is* networking?

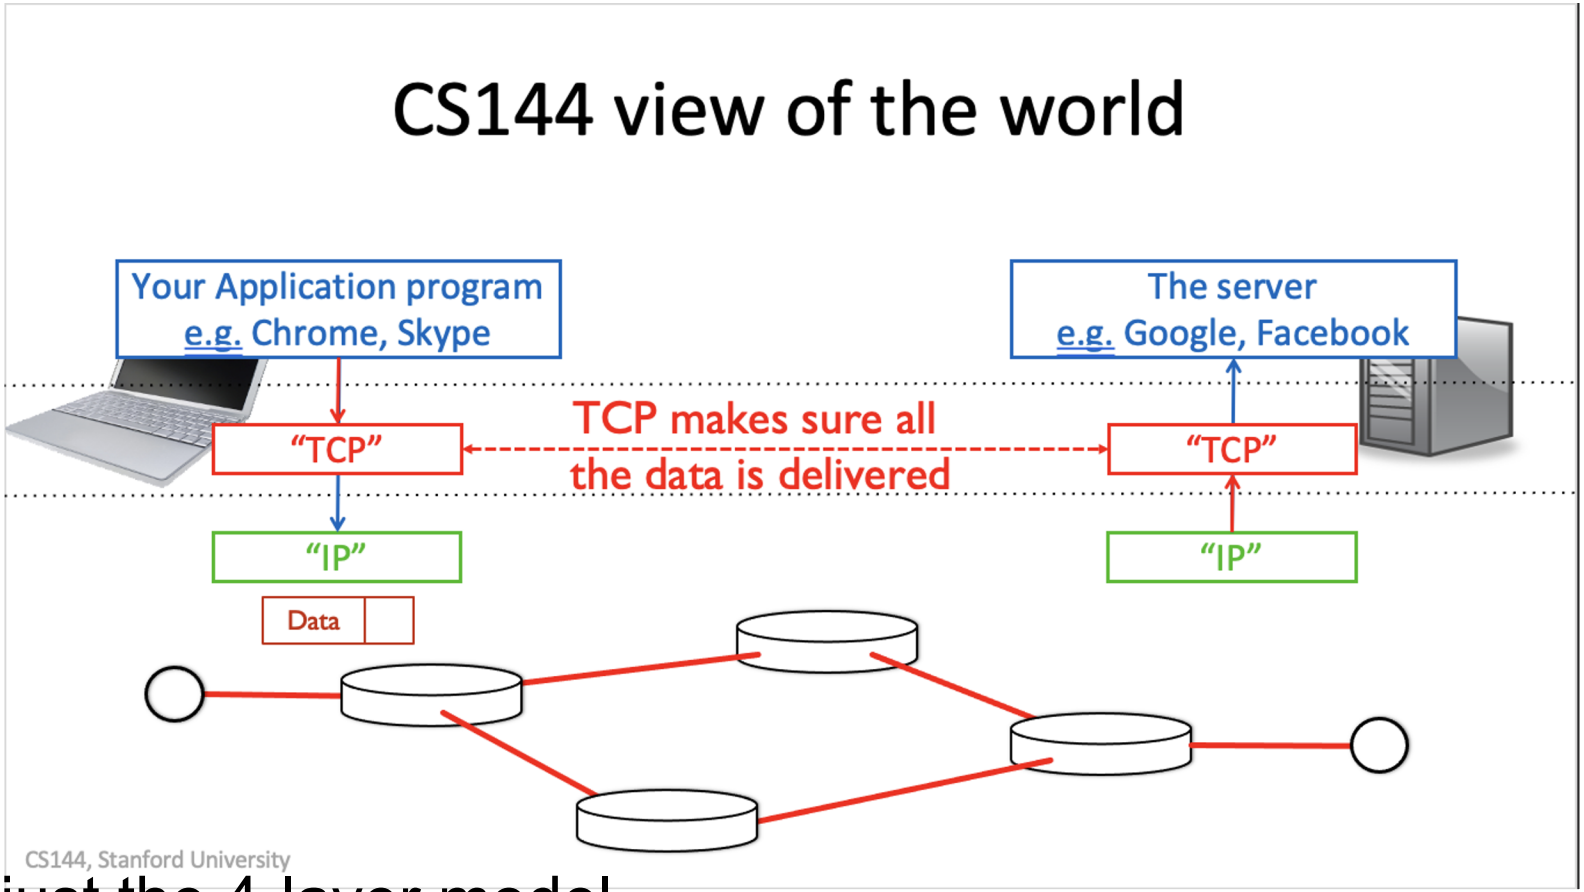

Is it just the 4-layer model we learn in CS144?

A Plethora of Protocol Acronyms?

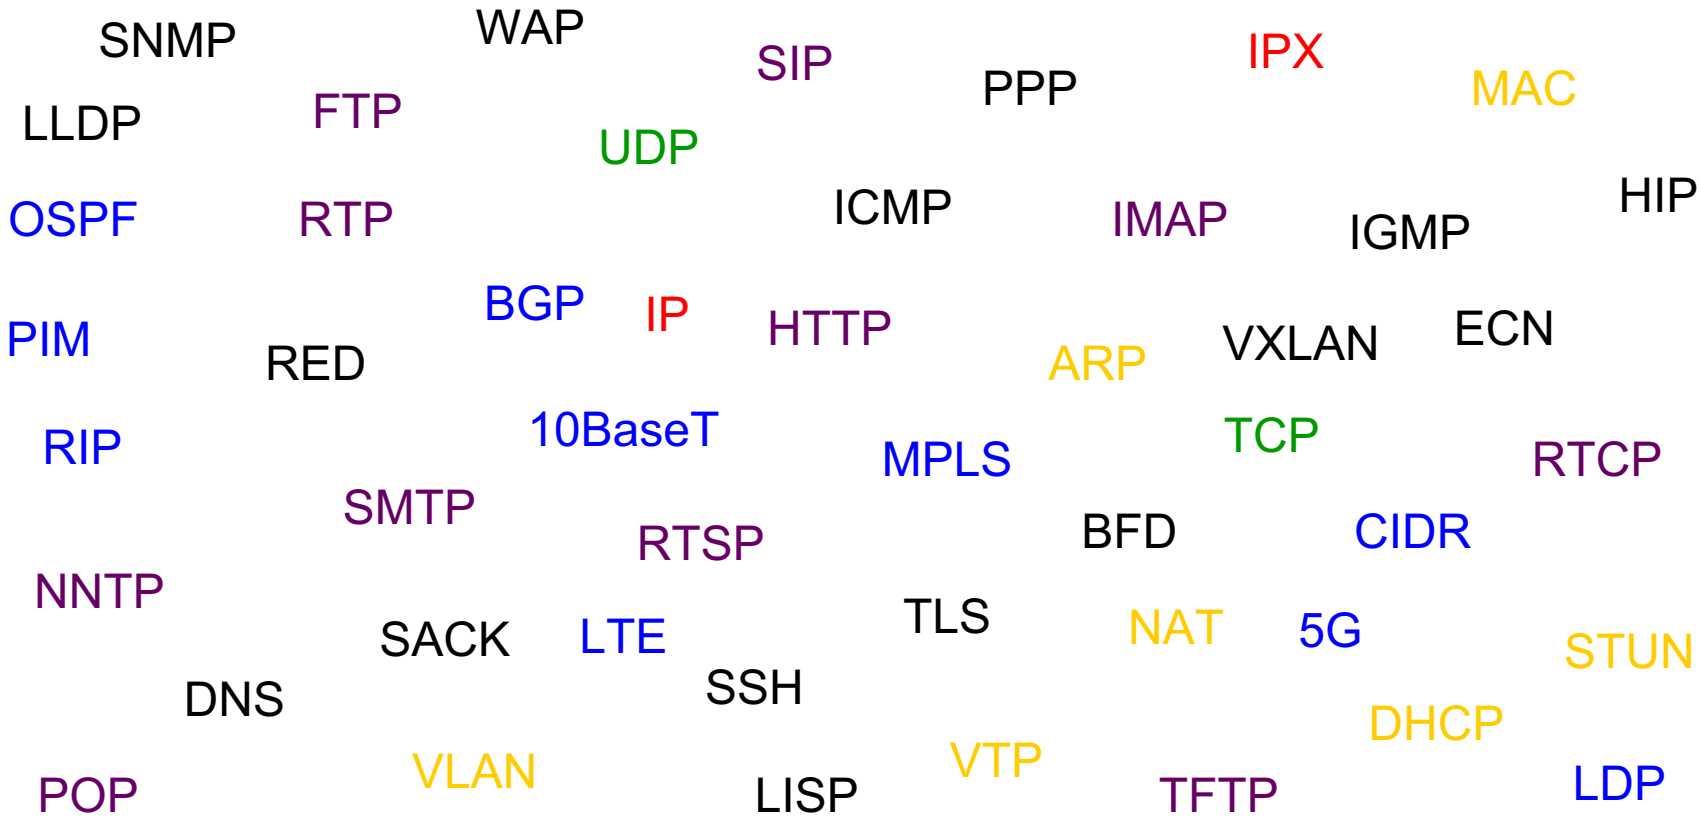

A Heap of Header Formats?

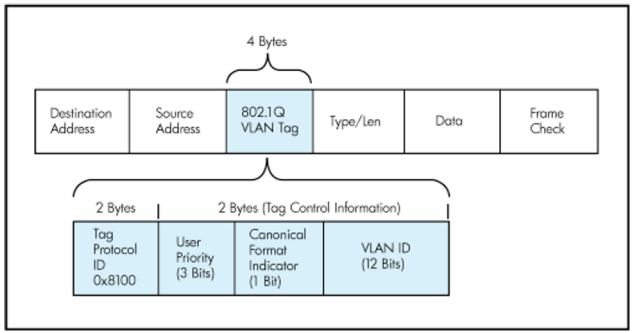

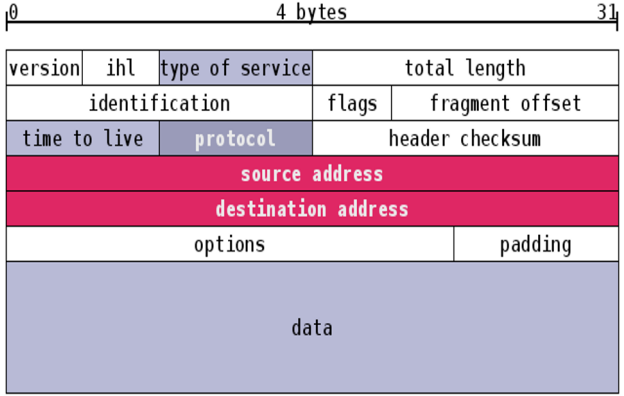

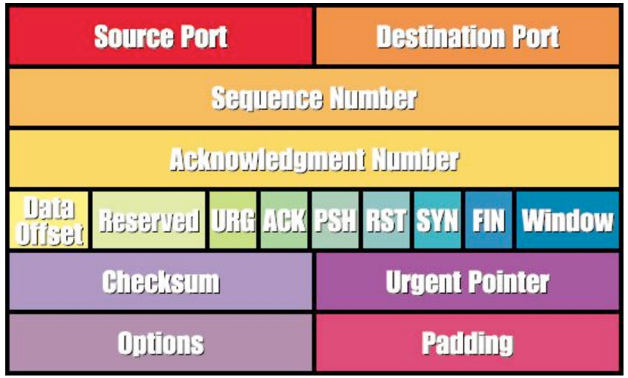

HTTP Response Header

Name	Value
HTTP Status Code: HTTP/1.1 200 OK	
Date:	Thu, 27 Mar 2008 13:37:17 GMT
Server:	Apache/2.0.55 (Ubuntu) PHP/5.1.2
Last-Modified:	Fri, 21 Mar 2008 13:57:30 GMT
ETag:	"358a4e4-56000-ddf5c680"
Accept-Ranges:	bytes
Content-Length:	352256
Connection:	close
Content-Type:	application/x-msdos-program

TCP/IP Header Formats in Lego

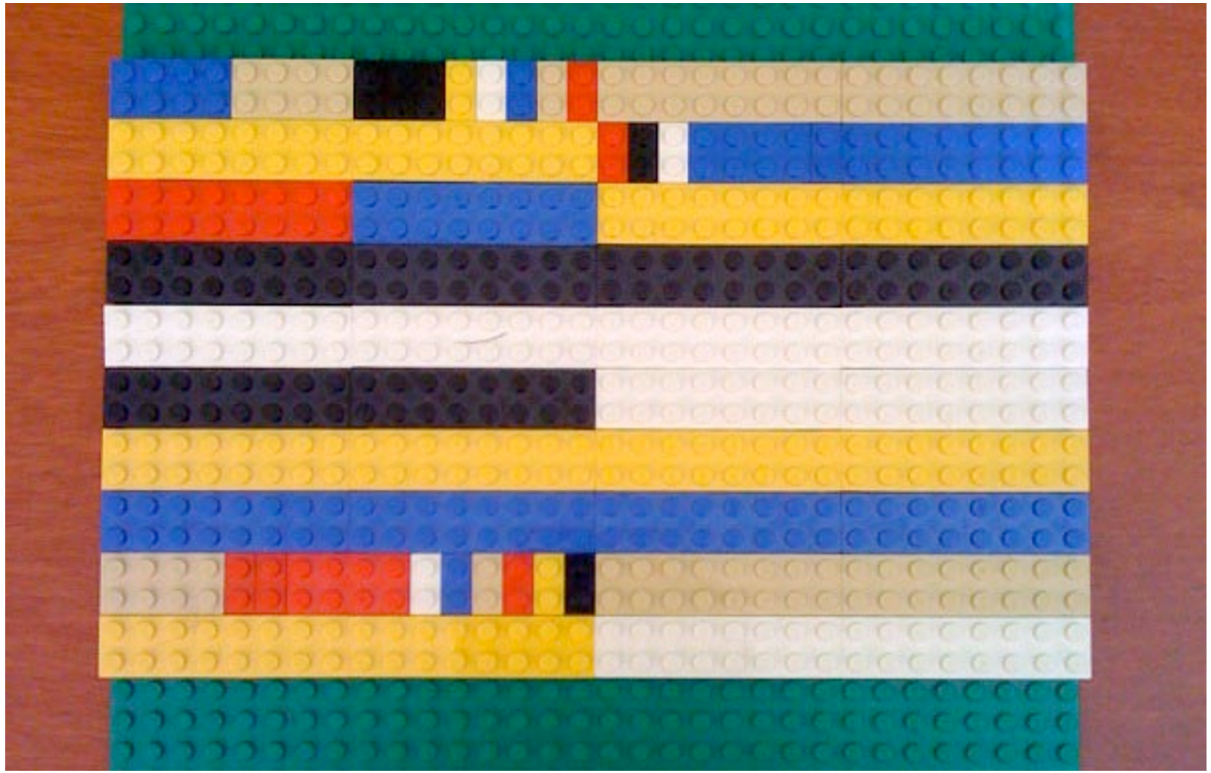

Lots of Different Equipment?

Router

Label Switched Router

Load balancer

Switch

Scrubber

Repeater

Gateway

Intrusion Detection System

Bridge

Route Reflector

Deep Packet Inspection

DHCP server

Firewall

Packet shaper

NAT

Hub

Packet sniffer

WAN accelerator

DNS server

Base station

Proxy

A place to apply theory?

- Algorithms and data structures
- Control theory
- Queueing theory
- Optimization theory
- Game theory and mechanism design
- Formal methods
- Formal verification
- Information theory
- Cryptography
- Programming languages
- Graph theory
- AI/ML

A place to build systems?

- Distributed systems
- Operating systems
- Computer architecture
- Software engineering
- …

What Peers in Other Fields Say

"What are the top ten classic problems in networking? I would like to solve one of them and submit a paper to SIGCOMM." After hearing that we don't have such a list: "Then how do you consider networking a discipline?"

"So, these networking research people today aren't doing theory, and yet they aren't the people who brought us the Internet. What exactly are they doing?"

"Networking papers are strange. They have a lot of text."

Is networking a problem domain or a scholarly discipline?

“There is a tendency in our field to believe that everything we currently use is a paragon of engineering, rather than a snapshot of our understanding at the time. We build great myths of spin about how what we have done is the only way to do it to the point that our universities now teach the flaws to students (and professors and textbook authors) who don't know better.”

-- John Day (Internet pioneer)

Before you all leave …

Tell me: *Why is Networking Cool?*

Young, relatively immature field

- Tremendous intellectual progress is still needed
- *You* can help decide what the Internet really is

Widely-read papers

- Many of the most cited papers in CS are in networking

Interdisciplinary

- CS, EE, MS&E, Policy, Economics, Law, Ethics, Physics

Lots of platforms for building your ideas

- Simulation & Emulation: NS2, NS3, Mininet
- Open source control software: ONOS, SONiC, ODL, NOX, POX
- Programmability: Click, NetFPGA, P4 language
- Routing software: Quagga, XORP, and Bird
- Testbeds: Emulab, GENI
- Measurements: RouteViews, traceroute, Internet2, Inband Network Telemetry

Architectural questions tend to dominate CS networking research

Decomposition of Function

Definition and placement of function

- What to do, and where to do it

The “division of labor”

- Between the host, network, and management systems
- Across multiple concurrent protocols and mechanisms

Software Defined Network (SDN)

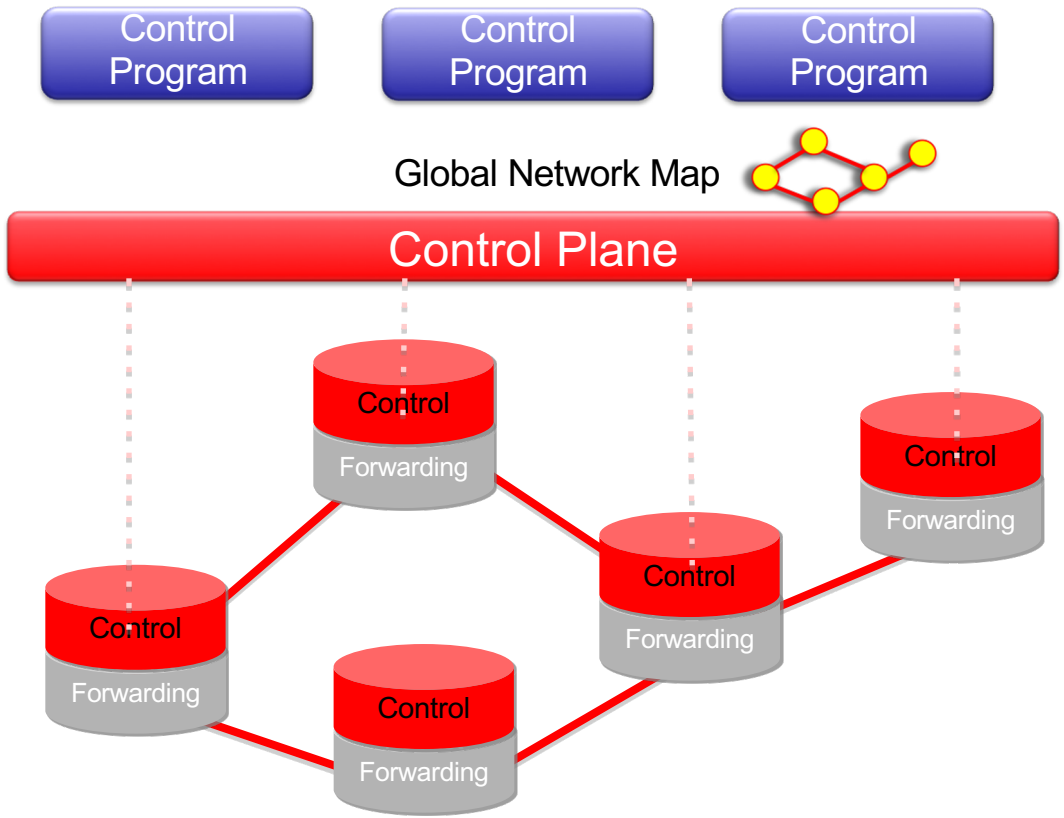

Network Function Virtualization (NFV)

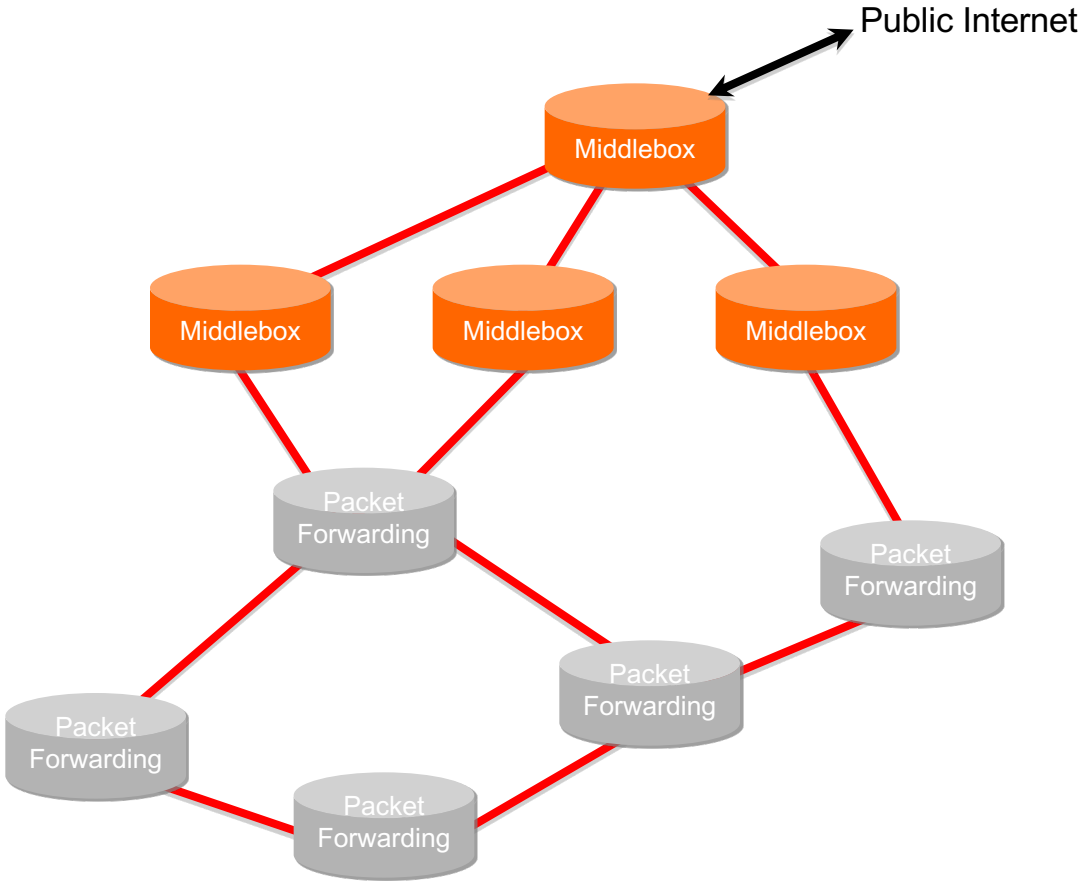

Network Function Virtualization (NFV)

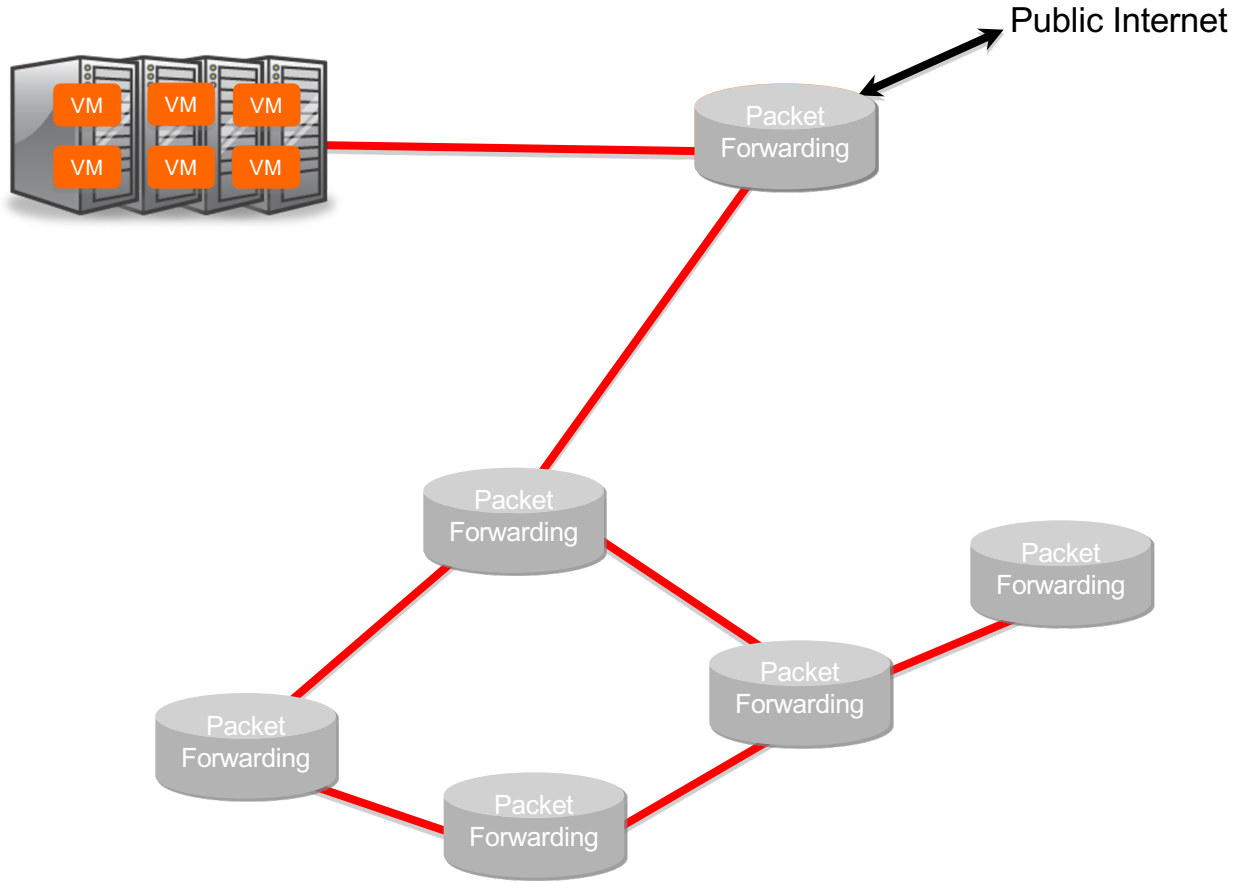

The Internet architecture is evolving faster than ever

“closed and proprietary”
“proliferation of standards”
“barrier to entry”
“stranglehold by vendors”

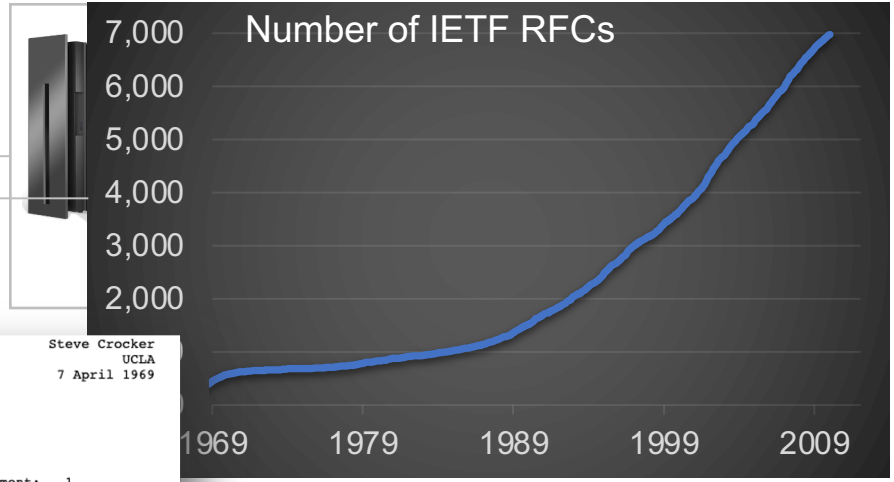

Network Working Group — Steve Crocker
Request for Comments: 1 — UCLA
7 April 1969

Title: Host Software
Author: Steve Crocker
Installation: UCLA
Date: 7 April 1969
Network Working Group Request for Comment: 1

2010 **2020** **2030**

-10 years

+10 years

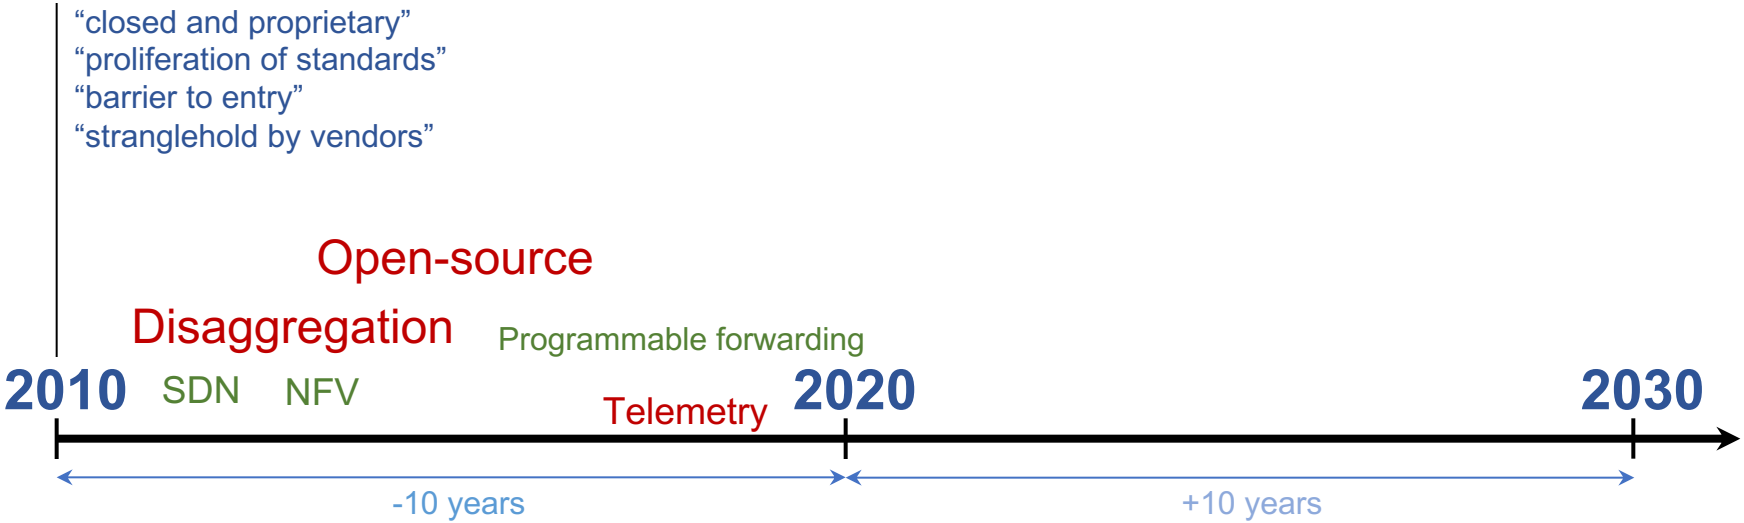
"closed and proprietary"
"proliferation of standards"
"barrier to entry"
"stranglehold by vendors"
Open-source
Disaggregation
Programmable forwarding
2010
SDN
NFV
Telemetry
2020
2030
-10 years
+10 years

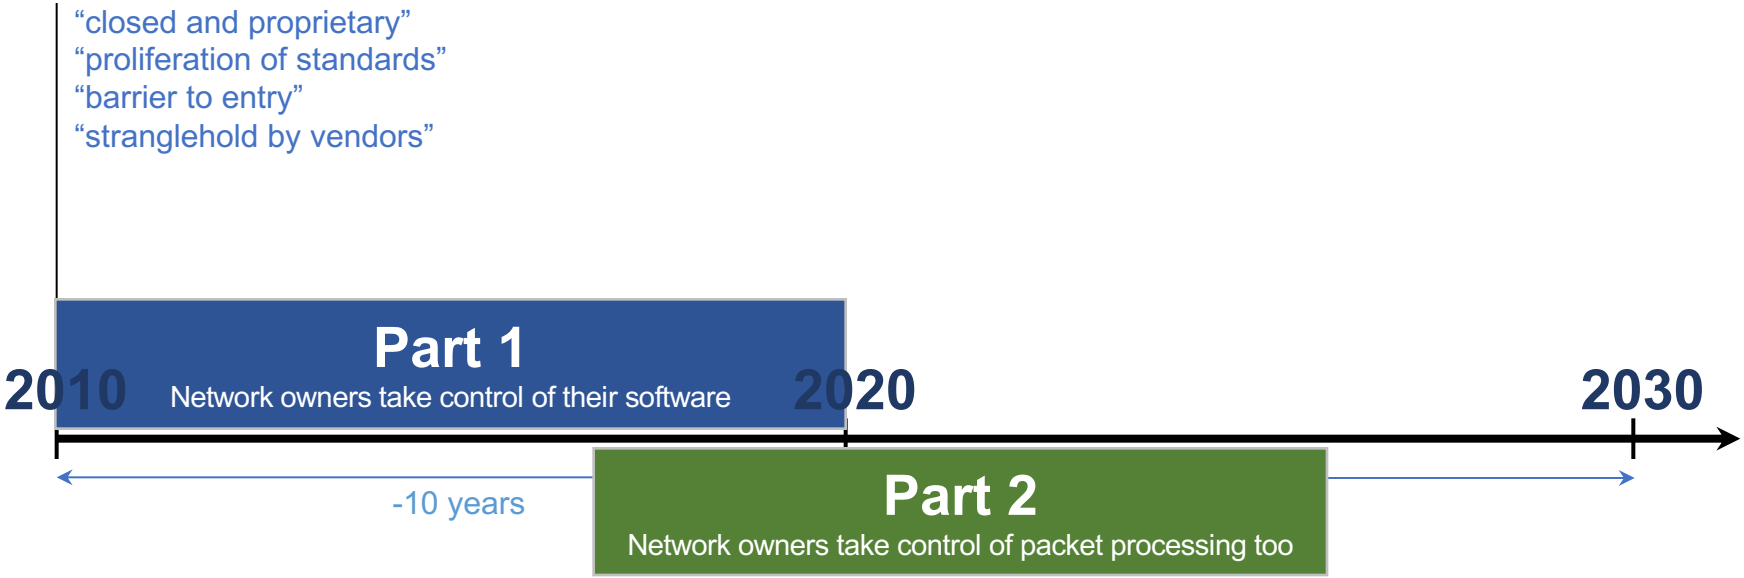
“closed and proprietary”
“proliferation of standards”
“barrier to entry”
“stranglehold by vendors”
Part 1
Network owners take control of their software
2010
2020
2030
-10 years
Part 2
Network owners take control of packet processing too

“Software will eat the world”

-- Marc Andreessen

Conclusion

Networking is extremely cool right now

- Real, important problems
- The Internet is evolving more rapidly than ever
- You can influence its future
- There are real opportunities for impact
- Inherently interdisciplinary

But the field is immature

- More of a “domain” than a “discipline”

Hints on reading a paper

Keshav: “How to Read a paper”, CCR 2007

Three stage approach

1. Read quickly in 5-10minutes
2. Read with greater care; ignore proofs
3. Deconstruct paper; question all assumptions

First papers

Read thoroughly, submit critiques

For Wed, Mar/30:

The Design Philosophy of the DARPA Internet Protocols
Clark, 1988

For Mon, Apr/4:

End-to-End Arguments in System Design
Saltzer, Reed and Clark, 1984

First optional papers

Read through briefly

Optional (easy reads)

1. A Brief History of the Internet
   Leiner et al., 2003
2. On Distributed Communication Networks
   Paul Baran, 1963

Our first paper

The Design Philosophy of the DARPA Internet Protocols

David D. Clark*
Massachusetts Institute of Technology
Laboratory for Computer Science
Cambridge, MA. 02139

(Originally published in Proc. SIGCOMM '88, Computer Communication Review Vol. 18, No. 4, August 1988, pp. 106–114)

Abstract

The Internet protocol suite, TCP/IP, was first proposed fifteen years ago. It was developed by the Defense Advanced Research Projects Agency (DARPA), and has been used widely in military and commercial systems. While there have been papers and specifications that describe how the protocols work, it is sometimes difficult to deduce from these why the protocol is as it is. For example, the Internet protocol is based on a connectionless or datagram mode of service. The motivation for this has been greatly misunderstood. This paper attempts to capture some of the early reasoning which shaped the Internet protocols.

1. Introduction

For the last 15 years[1], the Ad

architecture into the IP and TCP layers. This seems basic to the design, but was also not a part of the original proposal. These changes in the Internet design arose through the repeated pattern of implementation and testing that occurred before the standards were set.

The Internet architecture is still evolving. Sometimes a new extension challenges one of the design principles, but in any case an understanding of the history of the design provides a necessary context for current design extensions. The connectionless configuration of ISO protocols has also been colored by the history of the Internet suite, so an understanding of the Internet design philosophy may be helpful to those working with ISO.

This paper catalogs one view of th