CcompSci 356 Computer Network Architecture
Lecture 1: Introduction and course overview

Xiaowei Yang
xwy@cs.duke.edu

Roadmap

- Course introduction
  - Audience
  - Staff
  - Contents
  - Assignments
  - Grading policy
  - Other questions?
- A networking example

Who should be taking this class?

- Curious about how things work
  - A lot of internals of networking
- Loves to build things
  - Programming labs
- Outcome
  - Build networks, analyze packet traces, write software that uses the network

Why study networking?

Networking = "Plumbing"

Networking is the "plumbing" of computing

Almost all areas of computing are network-based.

- Distributed computing
- Big Data
- Cloud Computing
- Internet of Things
- Smart Cities

Networking is the backbone of computing.

Source: Raj Jain Keynote speech at SIGCOMM 2017

Networking is Fueling All Sectors of Economy

Networking companies are among the most valued companies: Apple, AT&T, Samsung, Verizon, Microsoft, China Mobile, Alphabet, Comcast, NTT, IBM, Intel, Cisco, Amazon, Facebook, …

All tech companies that are hiring currently are networking companies

Note: Apple became highly valued only after it switched from computing to communications (iPhone)

Networking = Economic Indicator

Source: Raj Jain Keynote speech at SIGCOMM 2017

Smart Everything

Smart Watch


Smart TV


Smart Car


Smart Health


Smart Home


Smart Kegs


Smart Space


Smart Industries


Smart Cities

Source: Raj Jain Keynote speech at SIGCOMM 2017

What's Smart?

Old: Smart = Can think Computation

= Can Recall Storage

Now: Smart = Can find quickly, Can Delegate

Communicate = Networking

Smart Grid, Smart Meters, Smart Cars, Smart homes, Smart Cities, Smart Factories, Smart Smoke Detectors, …

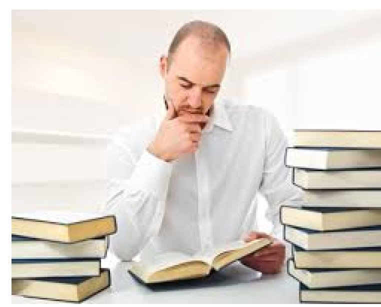

Not-Smart Smart

Networked Smart

Source: Raj Jain Keynote speech at SIGCOMM 2017

Networks change lives

MENU

Our Impact

Through our connectivity efforts we've brought more than 25 million people online who otherwise would not be and introduced them to the incredible value of the internet. They're doing better in school, building new businesses, and learning how to stay healthy. Here are a few stories of how Free Basics is having an impact.

https://info.internet.org/en/impact/

Networks foster innovations

- Google, Facebook, Internet of Things, online games, e-commerce, cloud computing …

- Fun examples: test of time paper awards
    1. "Ethane: Taking control of the Enterprise" by Martin Casado, Michael J. Freedman, Justin Pettit, Jianying Luo, Nick McKeown, Scott Shenker. SIGCOMM 2007.
       Ethane ushered in the age of Software-Defined Networking (SDN) and a new generation of research that inspired both academia and industry to design network control planes that we can reason about.

Martín Casado was born around 1976 in Cartagena, Spain.[1] He received his bachelor's degree from Northern Arizona University in 2000.[2] In 2017, he received an honorary doctorate from the same university.[3] He worked for Lawrence Livermore National Laboratory doing computational science[4] followed by work with the intelligence community from 2000 to 2003.[4] Casado attended Stanford University from 2003 to 2007,[*citation needed*] earning both his Masters and PhD in computer science.[5] While at Stanford, he began development of OpenFlow,[6] an open source protocol that enabled software-defined networking. During this period, he co-founded Illuminics Systems[7] with Michael J. Freedman.[7] Illuminics Systems was acquired by Quova, Inc. in November, 2006.[7] His PhD thesis, "Architectural Support for Security Management in Enterprise Networks," under advisors Nick McKeown, Scott Shenker and Dan Boneh, was published in 2008.[5]

Career [edit]

In 2007, Casado co-founded Nicira Networks along with McKeown and Shenker, a Palo Alto, California based company working on network virtualization. Along with McKeown and Shenker, Casado promoted software-defined networking.[6] His PhD work at Stanford University led to the development of the OpenFlow protocol, which was promoting using the term software-defined networking (SDN). McKeown and Shenker co-founded the Open Networking Foundation (ONF) in 2011 to transfer control of OpenFlow to a not-for-profit organization.[8]

In July 2012, VMware acquired Nicira for $1.26 billion.[9][10] At VMware he was made a fellow and held the positions chief technology officer (CTO) for networking and security and general manager of the Networking and Security Business Unit.[11]

Casado was named one of Business Insider's 50 most powerful people in enterprise tech in 2012,[12] and was featured in Silicon Valley's Business Journal's "Silicon Valley 40 Under 40" in 2013.[1] Casado was a 2012 recipient of the Association for Computing Machinery (ACM) Grace Murray Hopper Award as for helping create the Software Defined Networking movement.[13]

Networks foster innovations

- Google, Facebook, Internet of Things, online games, e-commerce, cloud computing

- Fun examples: test of time paper awards
    1. "Ethane: Taking control of the Enterprise" by Martin Casado, Michael J. Freedman, Justin Pettit, Jianying Luo, Nick McKeown, Scott Shenker. SIGCOMM 2007. (2017 award winner)
    Ethane ushered in the age of Software-Defined Networking (SDN) and a new generation of research that inspired both academia and industry to design network control planes that we can reason about.
    2. Link-level measurements from an 802.11b mesh network" by Daniel Aguayo, John Bicket, Sanjit Biswas, Glenn Judd, Robert Morris . SIGCOMM 2004. (2016 award winner)

Parent	Cisco Systems
Website	meraki.cisco.com

History [edit]

Meraki was founded by two MIT PhD students, Sanjit Biswas and John Bicket, along with Hans Robertson. The company was based in part on the MIT Roofnet project, an experimental 802.11b/g mesh network developed by the Computer Science and Artificial Intelligence Laboratory at the Massachusetts Institute of Technology.

Meraki was funded by Google and Sequoia Capital. The organization started in Mountain View, California in 2006, and before relocating to San Francisco. Meraki employed people who worked on the MIT Roofnet project.[2][3][4]

In 2007, Meraki selected San Francisco for their community-based Free the Net campaign.[why?] They started putting gateway devices in the Lower Haight neighborhood to provide Internet access and giving away repeaters. In the first year of the project, the growth of the network was primarily in the Mission District. By October 2007, they estimated 20,000 distinct users connected and about 5 terabytes of data transferred in this network. In July 2008, Meraki said 100,000 people in San Francisco used its "Free the Net" service. Since then, Meraki discontinued this public service, though many access points remain active, but with no connection to the Internet.

On November 18, 2012, Cisco Systems announced it would acquire Meraki for an estimated $1.2 billion.[1]

- You could be the next!

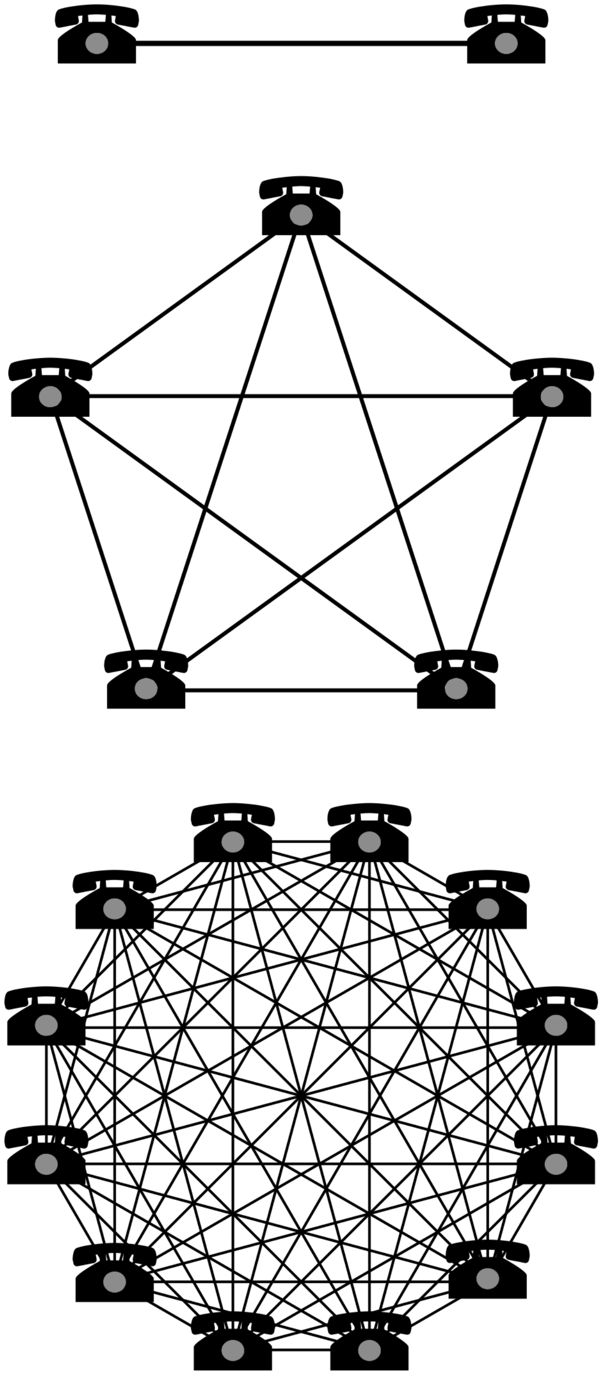

Metcalfe's Law

- The effect (value) of a telecommunications network is proportional to the square of the number of connected users of the system (n^2).

Course staff

- Instructor A: Xiaowei Yang
  - http://www.cs.duke.edu/~xwy/
  - Main research interests: network architectures and protocols, network security, distributed systems

- TA: Xiao Zhang
  - http://users.cs.duke.edu/~xzhang/

Prerequisites

- Know how to program in C/C++
  - Labs are done in C/C++
    - pointers, struct
  - Some knowledge of socket programming
    - Need to pick up the knowledge on your way

Course contents

- First half: the classic Internet
- Second half: current status and future
- Organization
  - Fundamental challenges
    - Routing, resource sharing, reliability, growth, evolution
  - Current solutions
  - Jargons and concepts

Recommended Textbook and resources

- Larry L. Peterson, Bruce S. Davie, Computer Networks: A Systems Approach, 5^{th} Edition, Morgan Kaufmann
  - Available from bookstore, amazon
  - A view as a system builder, not a user
  - 4^{th} is okay

- A great online book: https://intronetworks.cs.luc.edu/current/html/

Website: important

- http://www.cs.duke.edu/courses/spring19/compsci356/index.html

- Syllabus:
  - Book chapters related to lectures
  - Pointers to external papers for topics
  - Read before class for discussion
  - Assignments and due dates
  - Lectures

- All subject to change. Reload before checking

Discussion forum

- Piazza: link on website
- Useful discussions

Contact Us: We are here to help you

- Office hours
  - Xiao Zhang: Th 4pm - 6pm @ N003
  - Xiaowei Yang: Wed 2pm-4pm @ LSRC D336
- Email for meetings outside regular office hours
  - xwy@cs.duke.edu
  - xzhang@cs.duke.edu
- No recitation section
  - Do come to office hours
  - Do communicate with us via email, piazza

Your work

- Assigned readings in the Syllabus

- Class participation
  - Answering questions on Piazza counts

- Homework
  - Two planned, each 10 points (individual)

- Labs

Labs

- Lab 1: An echo server. 10 pts (individual)
- Lab 2: Simple router. 15 pts (group 1 or 2 students)
- Lab 3: Dynamic routing protocols. 15 pts (group 1 or 2 students)
- Some labs contain pre-lab questions that help you understand the basic concepts
- Labs are distributed with skeleton code and most of them have reference implementations for testing
- Turn-ins include answers to pre-lab questions, source code, lab reports if we ask for them

Exams

- Midterm
  – 3/7, in class (last lecture before the Spring break)

- Final
  – 5/3, 7-10pm

Grading policy

- Class participation: 5%
- Homework: 20%
- Labs: 40%
  - In a group assignment, all students get the same grade for the assignment
- Midterm: 15%
- Final: 20%

Final grade assignment

- No curving
- >= 90% A-/A/A+
- [80, 90) B-/B/B+
- [70, 80) -/C/C+
- [60, 70) D
- < 60% F
- May scale up, but not down

Collaboration policy

- Discussion is encouraged
- Individual assignments must be completed independently
- Group assignments only need to turn in one copy of the files with group members noted in the submission

Academic honesty policy

- Don't know if you are cheating? Please consult the description: http://www.cs.duke.edu/courses/spring19/compsci356/index.html
- If you are caught cheating, you will be reported to the Office of Student Conduct **and you will receive a failing grade in the class**

Late policy

- Due dates/times noted on the course website
- The deadline for any assignment can be extended with a 10% penalty per day
-
  - No deadline can be extended by more than two days.
  - Assignments will NOT be accepted 48 hours after the due date.
    - Tight schedule
    - Extension will delay next assignment
  - If you are ill: contact the instructor and provide a medical note.

Questions?

The first big question we study:

How to design a global computer network

What's a network?

- Wikipedia: A wide variety of systems of interconnected components are called **networks**.
- Examples of networks: what components are connected?
  - The Internet
  - Telephone networks
  - TV networks
  - Power networks
  - Sewage networks
  - Water networks
  - ….
- Why do we build networks?
  - To distribute/transfer something

Features of computer networks

- Generality
- Carry many different types of data
- End points are computers or devices that speak certain network protocols
- Support an unlimited range of applications
  - Can you name several Internet applications?

What’s the Internet?

- The **Internet** is a large-scale general-purpose computer network.
  - Run more than one application
- The Internet transfers information between computers.
- The Internet is a **network of networks**.

What the Internet looks like

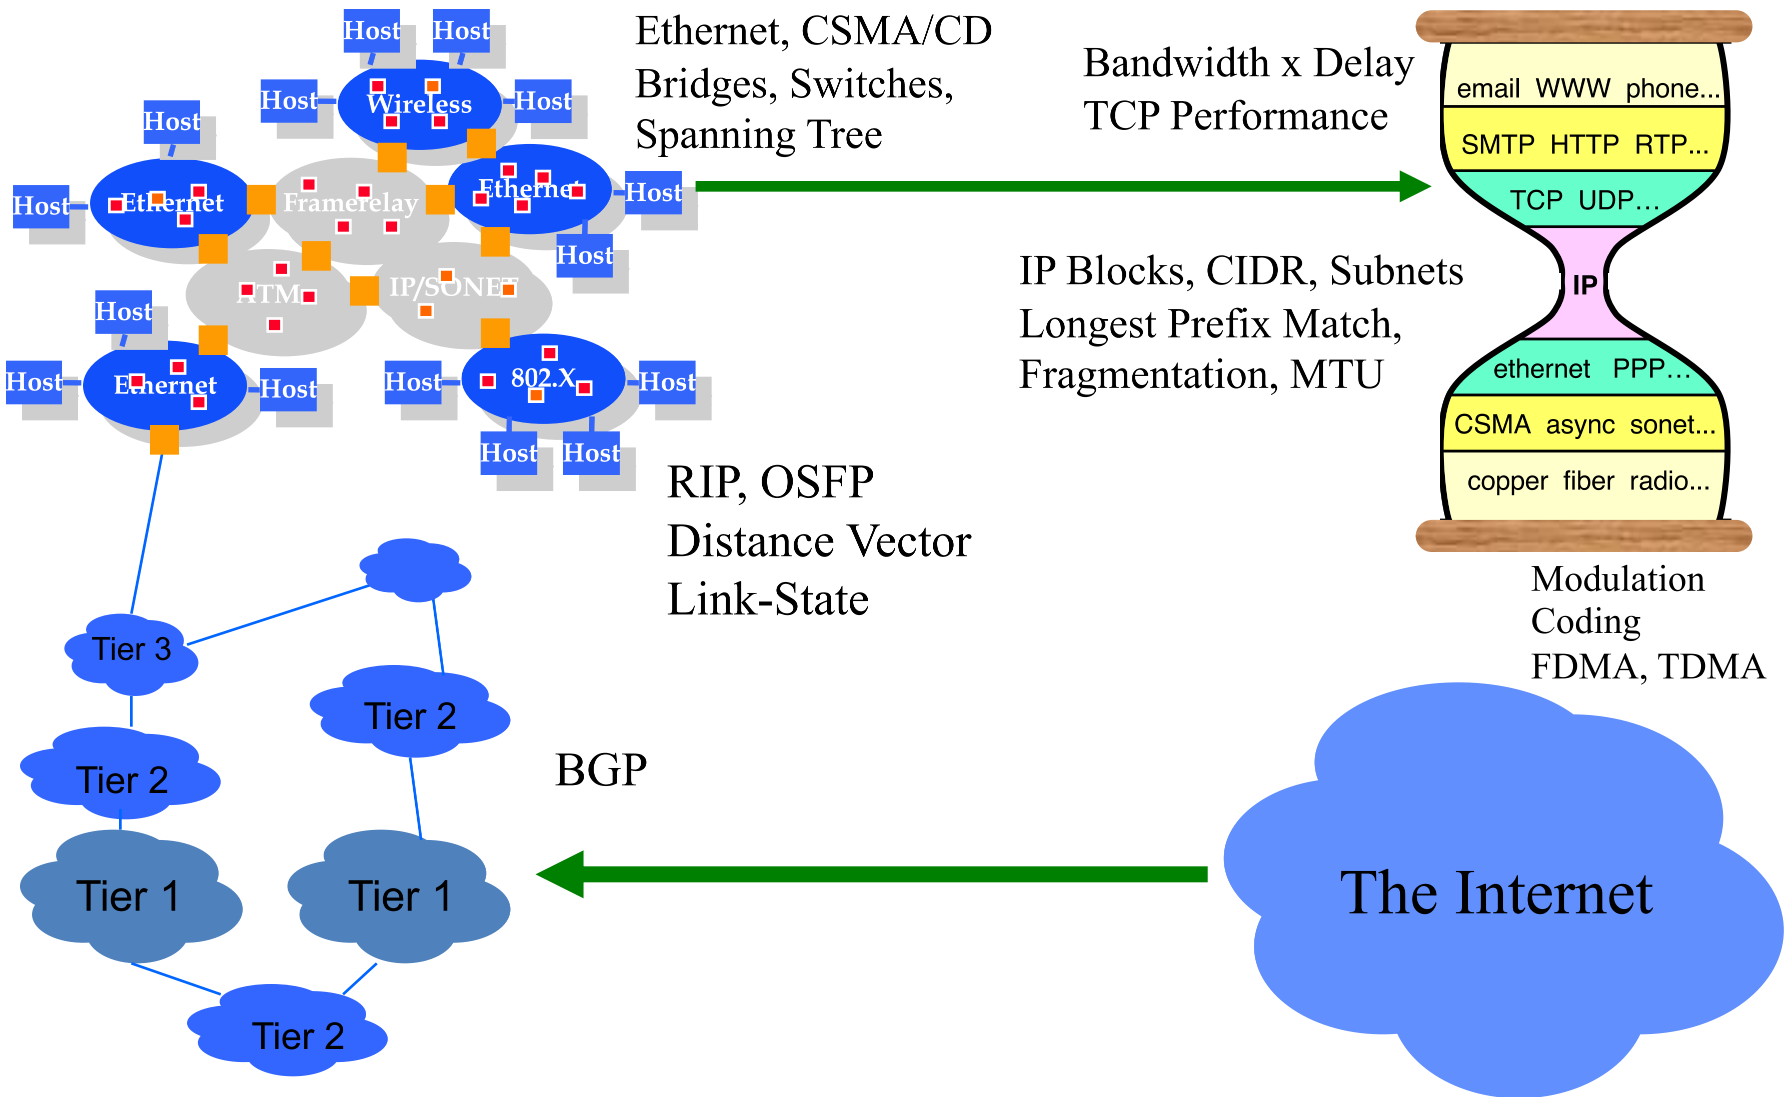

An Example

A simple TCP/IP Example

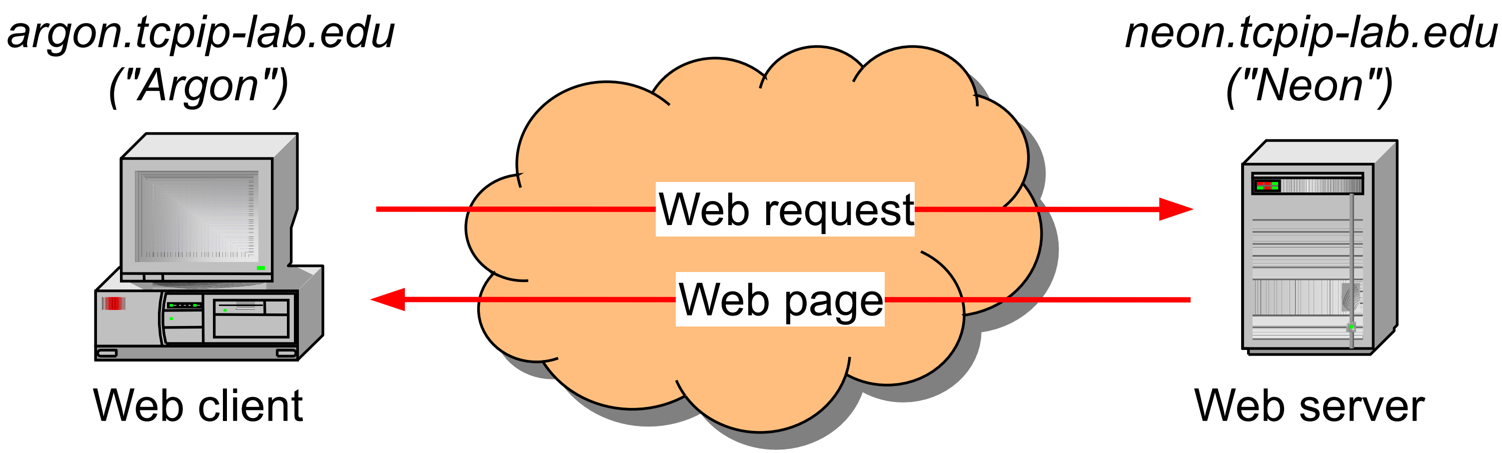

- A user on host *argon.tcpip-lab.edu* ("*Argon*") makes web access to URL

  http://neon. tcpip-lab.edu/index.html.

- What actually happens in the network?

HTTP Request and HTTP response

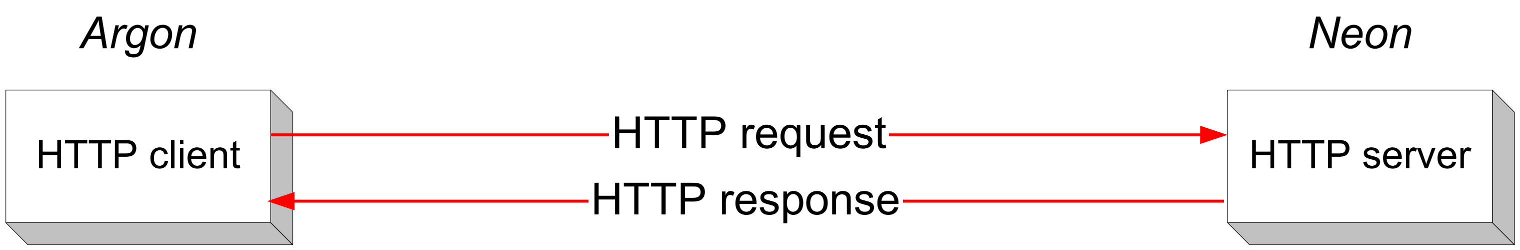

- Web server runs an HTTP server program
- HTTP client Web browser runs an HTTP client program
- sends an HTTP request to HTTP server
- HTTP server responds with HTTP response

HTTP Request

```
GET /example.html HTTP/1.1
Accept: image/gif, */*
Accept-Language: en-us
Accept-Encoding: gzip, deflate
User-Agent: Mozilla/4.0
Host: 192.168.123.144
Connection: Keep-Alive
```

HTTP Response

```
HTTP/1.1 200 OK
Date: Sat, 25 May 2002 21:10:32 GMT
Server: Apache/1.3.19 (Unix)
Last-Modified: Sat, 25 May 2002 20:51:33 GMT
ETag: "56497-51-3ceff955"
Accept-Ranges: bytes
Content-Length: 81
Keep-Alive: timeout=15, max=100
Connection: Keep-Alive
Content-Type: text/html

<HTML>
<BODY>
<H1>Internet Lab</H1>
Click <a href="http://www.tcpip-lab.net/index.html">here</a> for the Internet Lab webpage.
</BODY>
</HTML>
```

- How does the HTTP request get from Argon to Neon ?

From HTTP to TCP

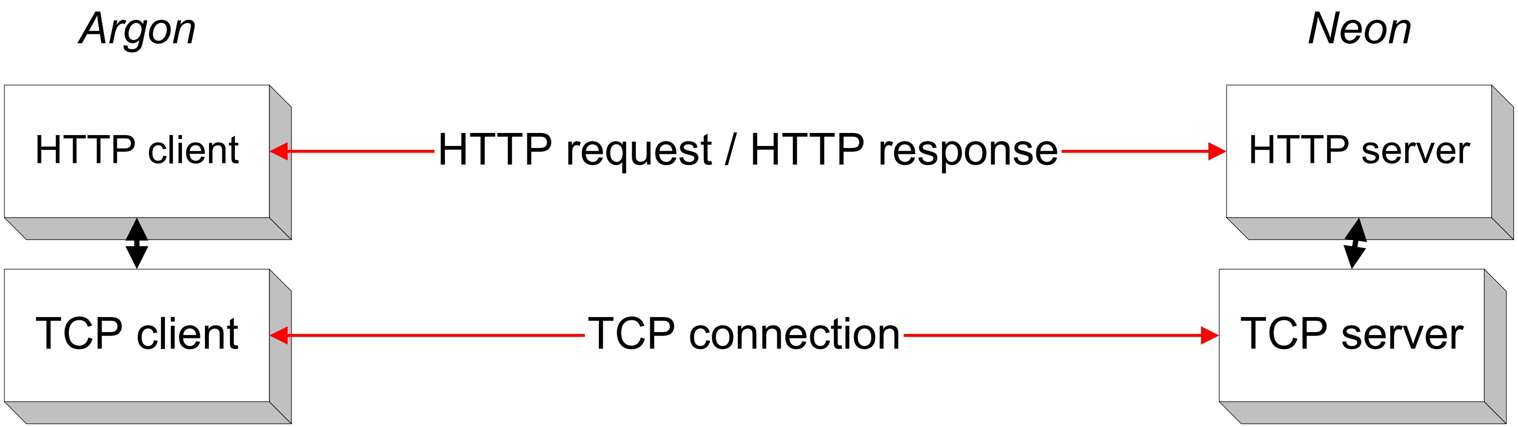

- To send request, HTTP client program establishes an TCP connection to the HTTP server Neon.
- The HTTP server at Neon has a TCP server running

Resolving hostnames and port numbers

- Since TCP does not work with hostnames and also would not know how to find the HTTP server program at Neon, two things must happen:
  1. The name “neon.tcpip-lab.edu” must be translated into an **IP address.**
  2. The HTTP server at Neon must be identified by a 16-bit **port number**.

Translating a hostname into an IP address

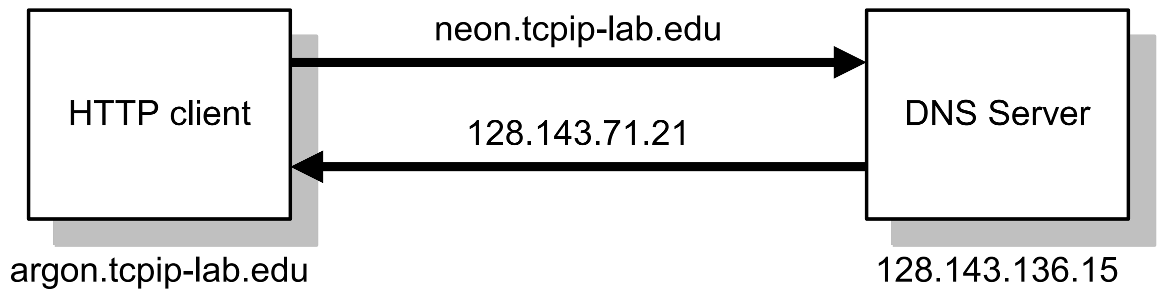

- The translation of the hostname *neon.tcpip-lab.edu* into an IP address is done via a database lookup
  - gethostbyname(host)
- The distributed database used is called the ***Domain Name System (DNS)***
- All machines on the Internet have an IP address:

argon.tcpip-lab.edu	*128.143.137.144*
neon.tcpip-lab.edu	*128.143.71.21*

Finding the port number

- **Note:** Most services on the Internet are reachable via well-known ports. E.g. All HTTP servers on the Internet can be reached at port number "80".
- **So:** Argon simply knows the port number of the HTTP server at a remote machine.
- On most Unix systems, the well-known ports are listed in a file with name /etc/services. The well-known port numbers of some of the most popular services are:

ftp	21	finger	79
telnet	23	http	80
smtp	25	nntp	119

Requesting a TCP Connection

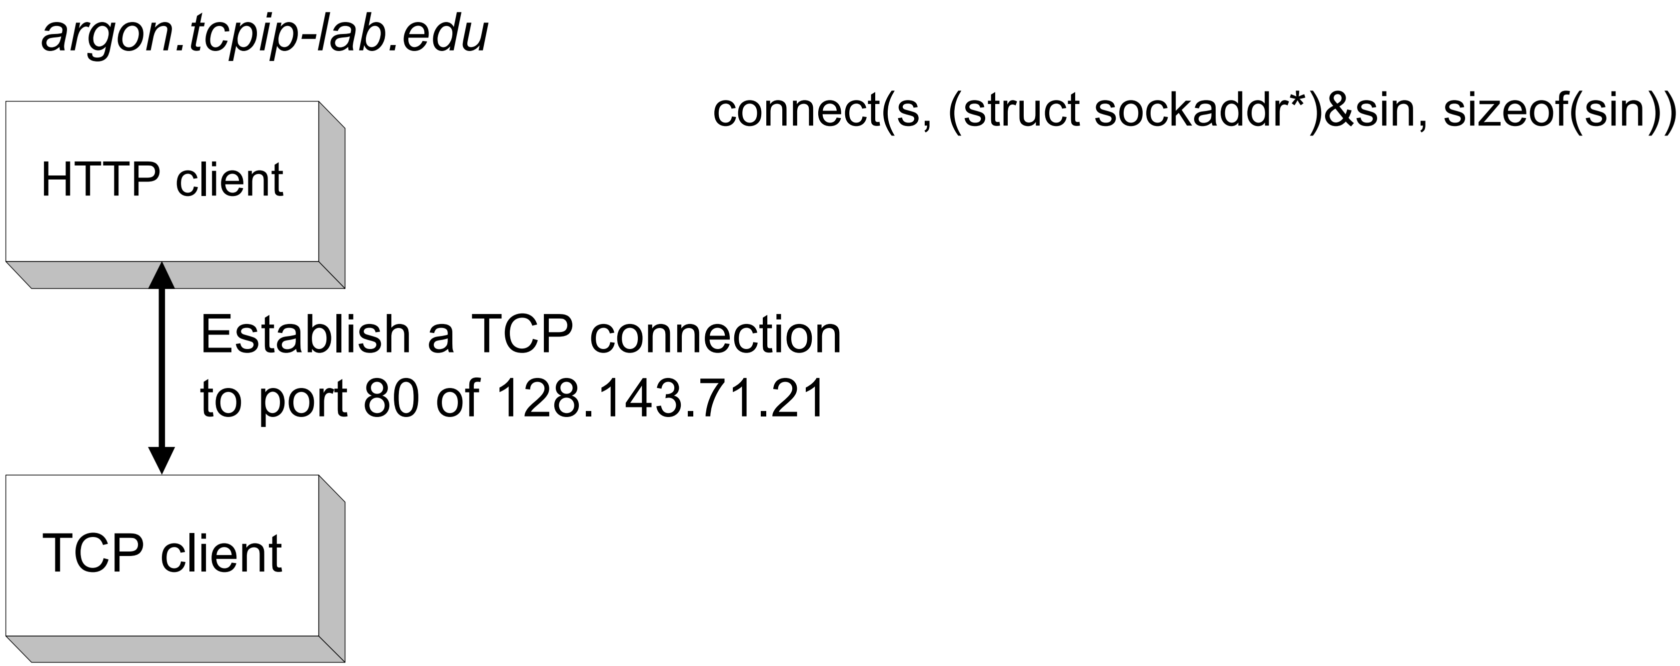

- The HTTP client at *argon.tcpip-lab.edu* requests the TCP client to establish a connection to port 80 of the machine with address 128.141.71.21

Invoking the IP Protocol

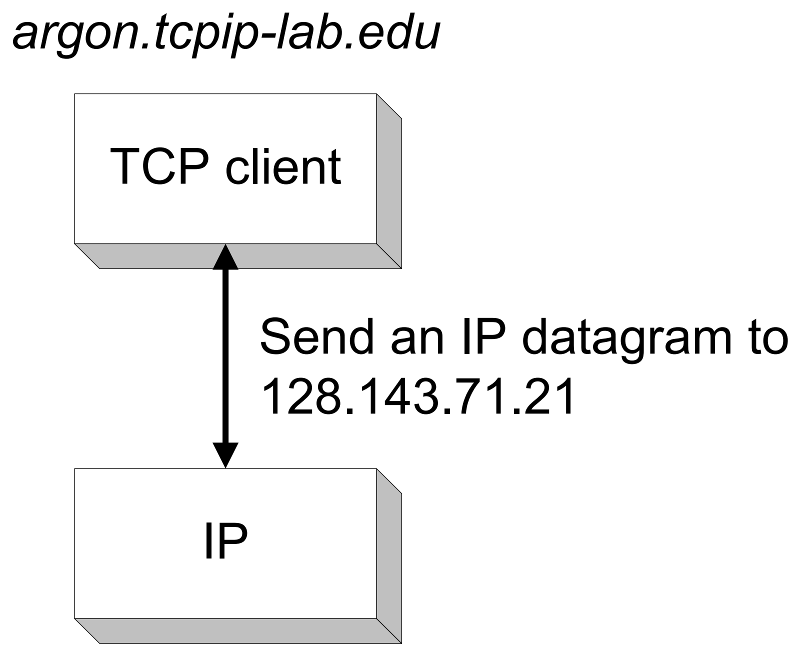

- The TCP client at *Argon* sends a request to establish a connection to port 80 at *Neon*
- This is done by asking its local IP module to send an IP datagram to *128.143.71.21*
- *(The data portion of the IP datagram contains the request to open a connection)*

Sending the IP datagram to the default router

- *Argon* sends the IP datagram to its default router
- The default gateway is an IP router
- The default gateway for *Argon* is *Router137.tcpip-lab.edu* (128.143.137.1).

Invoking the device driver

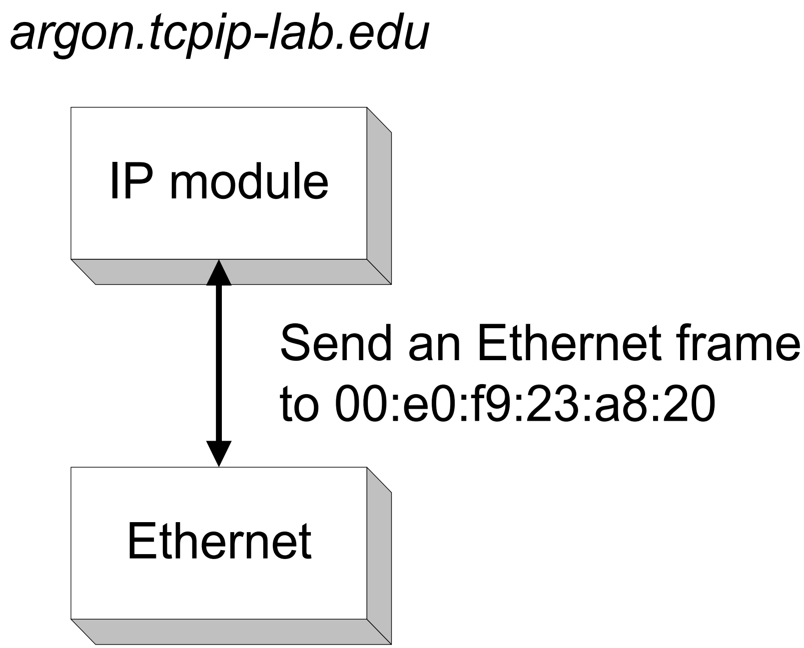

- The IP module at *Argon*, tells its Ethernet device driver to send an Ethernet frame to address *00:e0:f9:23:a8:20*
- Ethernet address of the default router is found out via ARP

The route from ***Argon*** to ***Neon***

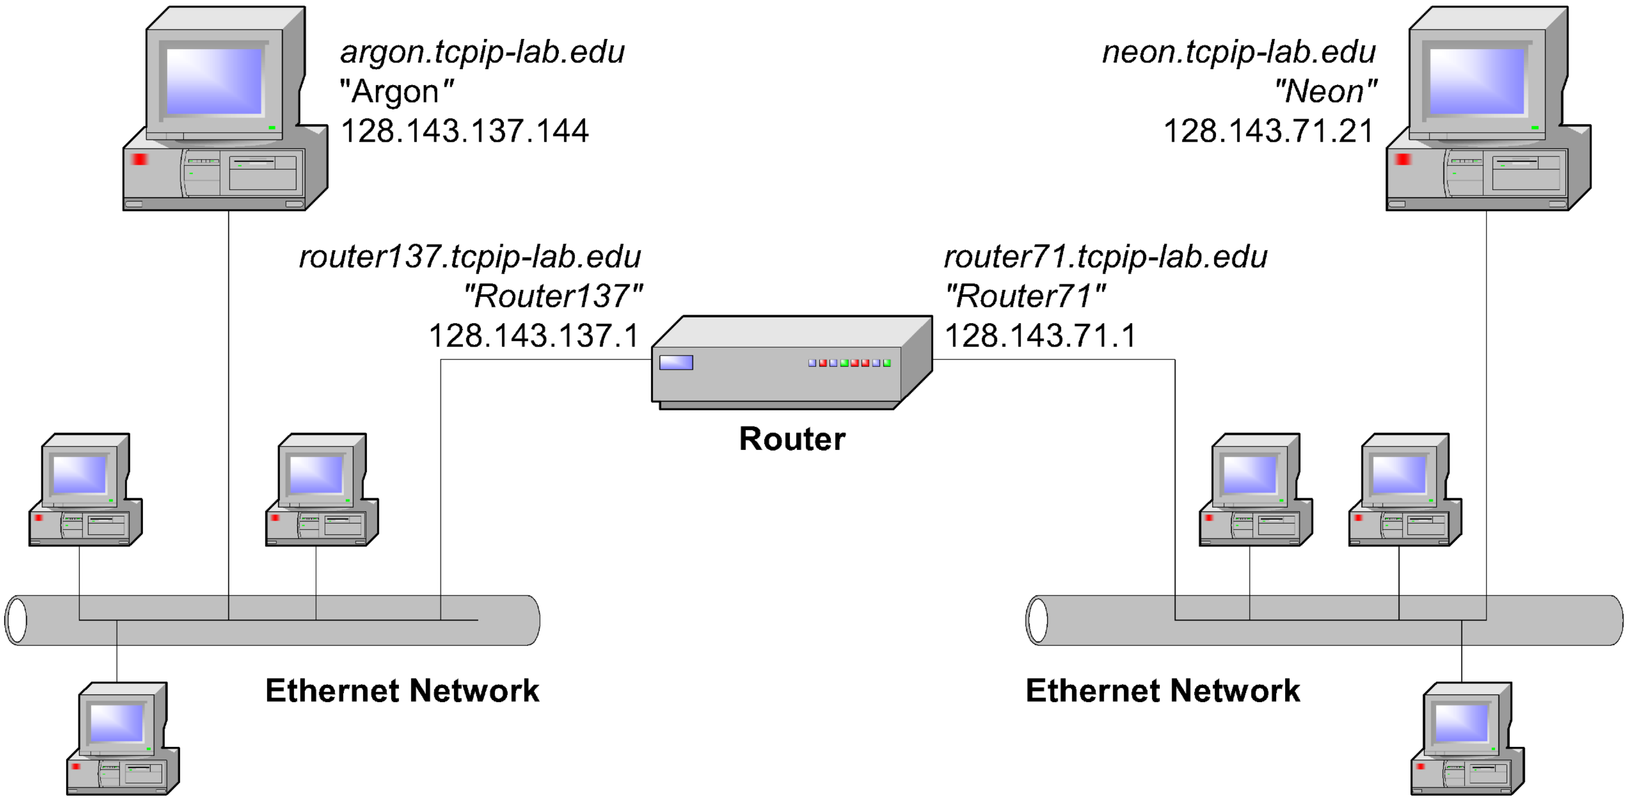

- Note that the router has a different name for each of its interfaces.

Sending an Ethernet frame

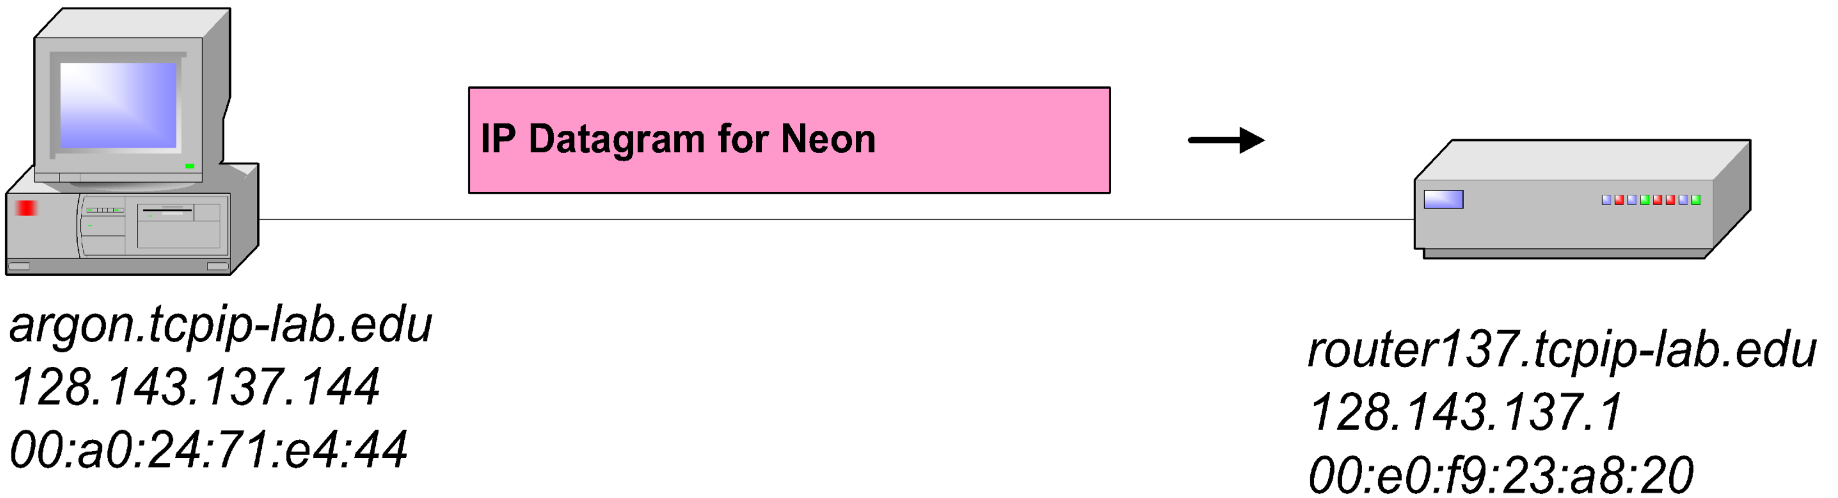

- The Ethernet device driver of *Argon* sends the Ethernet frame to the Ethernet network interface card (NIC)
- The NIC sends the frame onto the wire
  - *Turning bits into physical signals*

Forwarding the IP datagram

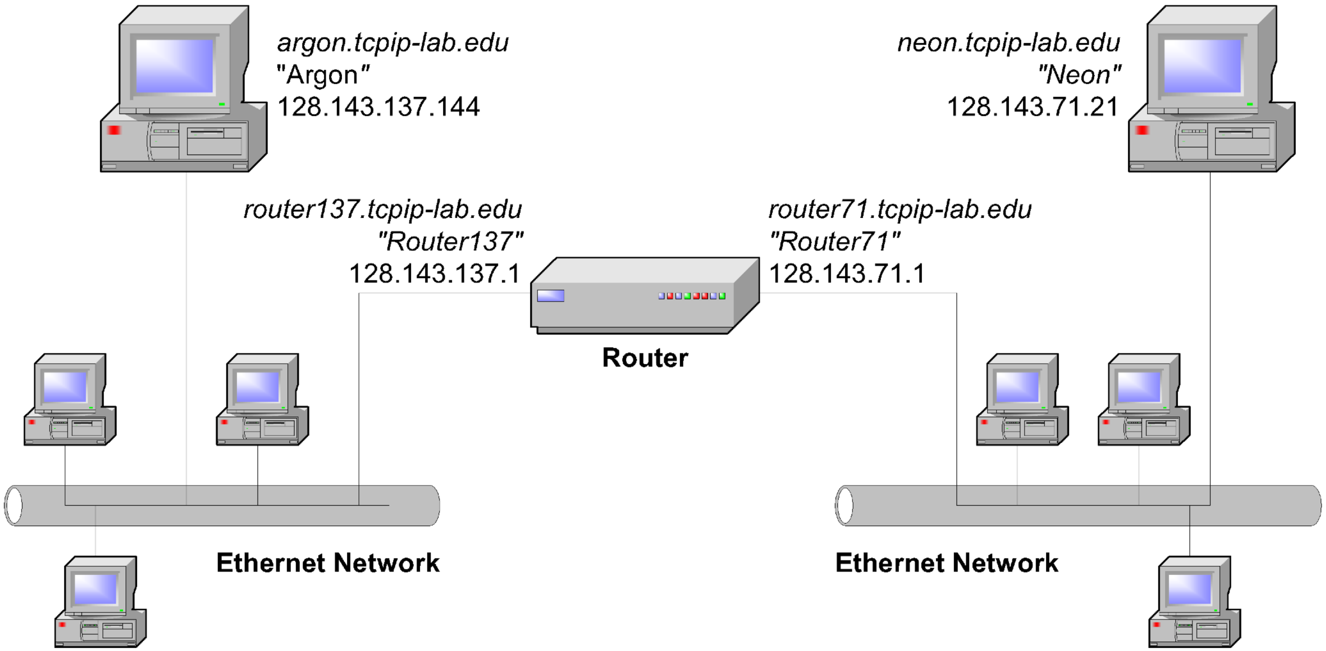

- The IP router receives the Ethernet frame at interface 128.143.137.1
  1. recovers the IP datagram
  2. determines that the IP datagram should be forwarded to the interface with name 128.143.71.1
- The IP router determines that it can deliver the IP datagram directly

Invoking the Device Driver at the Router

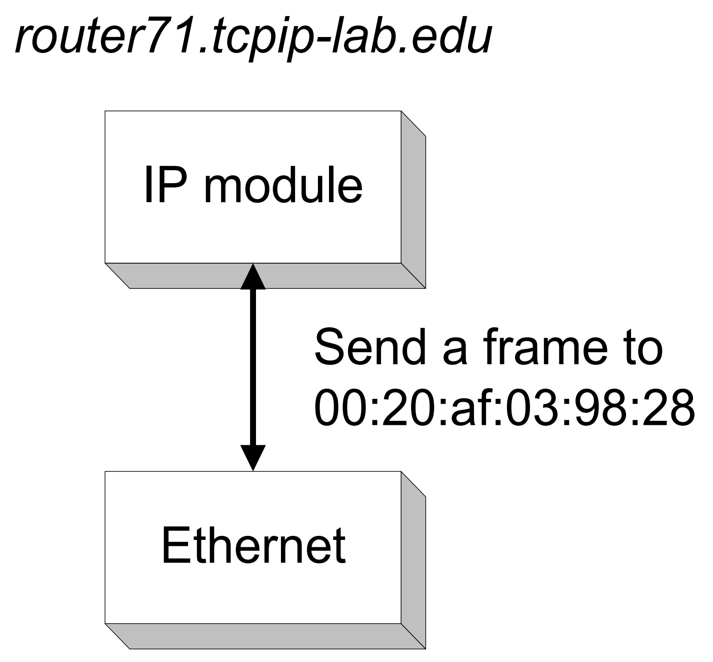

- The IP protocol at *Router71*, tells its Ethernet device driver to send an Ethernet frame to address *00:20:af:03:98:28*

Sending another Ethernet frame

- The Ethernet device driver of *Router71* sends the Ethernet frame to the Ethernet NIC, which transmits the frame onto the wire.

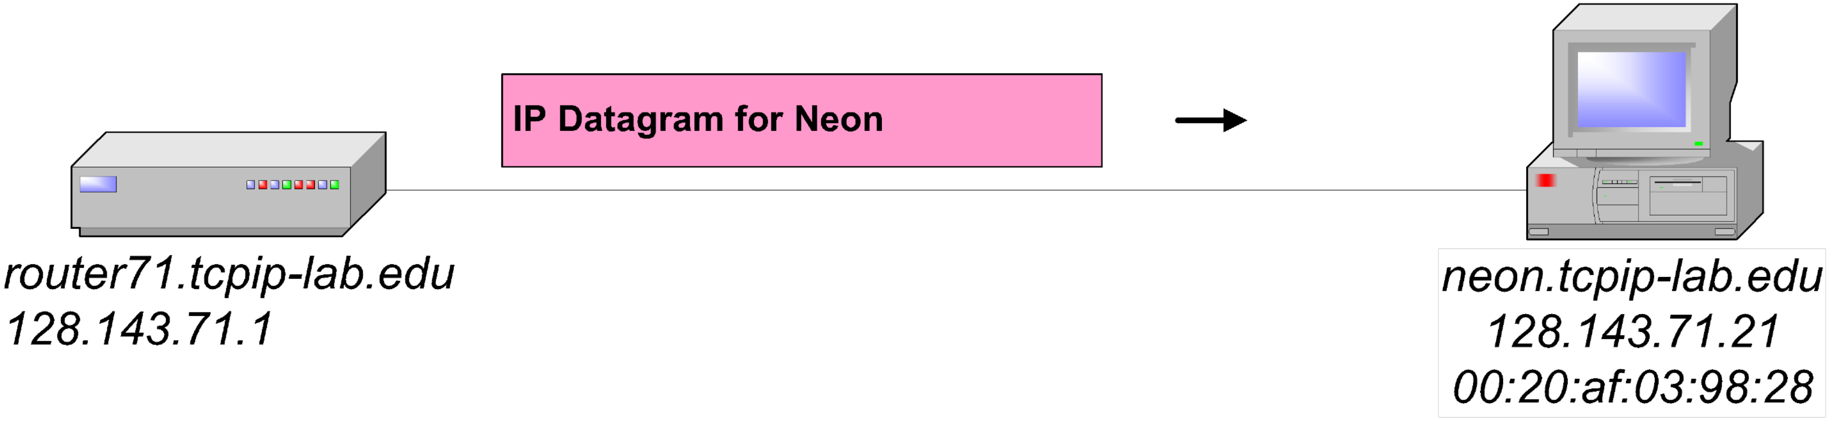

Data has arrived at Neon

- *Neon* receives the Ethernet frame
- The payload of the Ethernet frame is an IP datagram which is passed to the IP protocol.
- The payload of the IP datagram is a TCP segment, which is passed to the TCP server

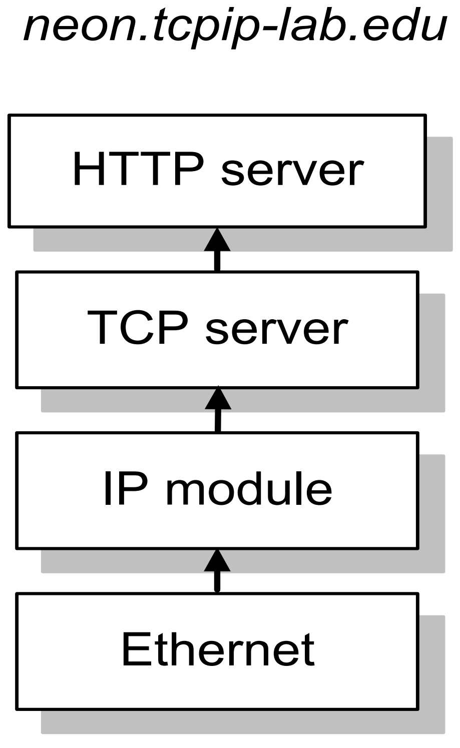

Summary

- Course introduction
  - Administrivia
  - Course outline
- An example introducing the layering concept
- Next
  - Network architectures: different styles of building networks