# Lecture 4: IP Addresses, Sockets, and System Programming

# COMP 411, Fall 2022 Victoria Manfredi





Acknowledgements: materials adapted from Computer Networking: A Top Down Approach 7<sup>th</sup> edition: ©1996-2016, J.F Kurose and K.W. Ross, All Rights Reserved as well as from slides by Abraham Matta at Boston University and some material from Computer Networks by Tannenbaum and Wetherall.

# Today

## 1. Announcements

- homework 1 due today, homework 2 posted
  - tictactoe.py solution code will be posted once homework1 submitted
- 2. Network applications
- 3. Network programming
  - TCP sockets

## 4. Network tools

- Wireshark: looking at real traffic

# Internet Organization IP ADDRESSES

vumanfredi@wesleyan

# Every device on Internet has an IP address

## IPv4 addresses

- 4 bytes
  - space of addresses: 0-255 . 0-255 . 0-255 . 0-255
  - hostnames are human-readable, IP addresses are machine-readable
- Loopback address: send traffic to yourself
  - traffic sent here is "looped back" through network stack on machine on which sending process is running
  - 127 . \* .\* .\*
  - typically 127.0.0.1, also called localhost
- Private subnet addresses ]
  - 10 .\* .\* .\*
  - 172.16-31 .\* .\*
  - 192.168 .\* .\*

Subnet: shared prefix portion of addr

### IPv6 addresses

- 16 bytes: we're running out of 4 byte addresses ...

# Who owns what address ranges?

#### Amazon

- 50.19.\*.\* → 256 x 256 = 65536 addresses
- 54.239.98.\*  $\rightarrow$  256 addresses

#### Facebook

- 57.240.0.0/17
- 157.240.10.0/24
- 157.240.1.0/24

#### Google

. . .

- 64.233.160.0 to 64.233.191.255
- 66.102.0.0 to 66.102.15.255

Wesleyan

. . .

- 129.133.21.\*

— ...

# How are IP addresses assigned?

#### Your ISP or institution has block of IP addresses

- you are assigned one of those IP addresses
- (possible you will get NAT'd address …)

### Static IP address

- manual configuration: set in network settings

#### **Dynamic IP address**

- using Dynamic Host Configuration Protocol (DHCP) in network-layer
- client (you) broadcasts request for IP address
- DHCP server on network assigns you address from address pool
  - typically get IP address for fixed period of time
  - router can be configured to act as DHCP server

# Actually ...

## Many hosts have multiple IP addresses

### How?

- IP address associated with network interface not host
- network interface card (NIC): connects computer to network
- A host may have 1 or more network interfaces
  - my laptop has (at least) 2 NICs: 1 wireless and 1 wired (via USB)
  - router needs at least two interfaces
    - otherwise can't connect multiple networks together
  - Cisco core router: can have up to 10,000 interfaces!
    - one interface per link: router has many IP addresses

#### VirtualBox Virtual Machine (VM)

- you can set the number and type of network interfaces for VM

# What's my IP address?

## ifconfig

- what network interfaces does my machine have?
- what are my IP and MAC # addresses?
- configure/enable/disable an interface



# What's host's IP address?

#### Host

> host www.google.com												
www.google.com has addre	ss 74.125.141.99											
www.google.com has addre	ss 74.125.141.103											
www.google.com has addre	ss 74.125.141.105											
www.google.com has addre	ss 74.125.141.147											
www.google.com has addre	ss 74.125.141.104											
www.google.com has addre	ss 74.125.141.106											
www.google.com has IPv6	address 2607:f8b0:400c:c06::93											

#### What's host name for IP address?

> host 8.8.8.8 8.8.8.8.in-addr.arpa domain name pointer google-public-dns-a.google.com.

# What's host's IP address?

dig

#### > dig www.google.com

- ; <<>> DiG 9.8.3-P1 <<>> www.google.com
- ;; global options: +cmd
- ;; Got answer:
- ;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 4619
- ;; flags: qr rd ra; QUERY: 1, ANSWER: 6, AUTHORITY: 0, ADDITIONAL: 0

;; QUESTION SECTION:

;www.google.com.	IN	Α				
;; ANSWER SECTION:						
www.google.com.	56	IN	Α	74.125.141.104		
www.google.com.	56	IN	Α	74.125.141.103		
www.google.com.	56	IN	Α	74.125.141.105		
www.google.com.	56	IN	Α	74.125.141.147		
www.google.com.	56	IN	A	74.125.141.99		
www.google.com.	56	IN	Α	74.125.141.106		

;; Query time: 7 msec
;; SERVER: 129.133.52.12#53(129.133.52.12)
;; WHEN: Mon Jan 22 14:06:38 2018

;; MSG SIZE rcvd: 128

# Is host up?

## Ping

- sends ICMP echo request to host
- host sends ICMP echo reply back
- If no reply within timeout period, packet deemed lost

```
> ping stanford.edu
PING stanford.edu (171.67.215.200): 56 data bytes
64 bytes from 171.67.215.200: icmp_seq=0 ttl=237 time=94.951 ms
64 bytes from 171.67.215.200: icmp_seq=1 ttl=237 time=94.738 ms
64 bytes from 171.67.215.200: icmp_seq=2 ttl=237 time=95.525 ms
64 bytes from 171.67.215.200: icmp_seq=3 ttl=237 time=194.993 ms
64 bytes from 171.67.215.200: icmp_seq=4 ttl=237 time=97.139 ms
64 bytes from 171.67.215.200: icmp_seq=5 ttl=237 time=95.878 ms
64 bytes from 171.67.215.200: icmp_seq=5 ttl=237 time=95.878 ms
64 bytes from 171.67.215.200: icmp_seq=6 ttl=237 time=95.667 ms
^C
--- stanford.edu ping statistics ---
7 packets transmitted, 7 packets received, 0.0% packet loss
round-trip min/avg/max/stddev = 94.738/109.842/194.993/34.770 ms
```

# Is one IP address per machine enough?

### What happens if you run multiple network applications?

- many processes running on computer
  - process is program in execution

How do messages received by computer get to right process?

- messages are addressed to (IP address, port #) pair
- different processes on computer will connect to network using same IP address but different port numbers

# 2 key functions of Internet core

How does Internet router determine outgoing link for packet?

## 1. Routing

- view Internet as giant graph
- run shortest path algorithms

### 2. Forwarding

- use paths to choose best output link for packet destination IP address
- if one link fails, chooses another



# Routing of packets across Internet

Each router uses its forwarding table to choose outbound link based on packet's destination



# Network Applications OVERVIEW

vumanfredi@wesleyan.edu

# Creating a network app

### Write programs that

- run on (different) end systems
- communicate over network
- e.g., web server software communicates with browser software

# Q: Do we need to write software for network-core devices?

- No, network-core devices do not run user applications
- applications on end systems allows for rapid app development, propagation



# **Client-server architecture**

# Client host requests and receives service from always on server host



#### Server

- always-on, dedicated host
  - e.g., web server
- permanent IP address
- data centers for scaling

#### Clients

- communicate with server
- may be intermittently connected
- may have dynamic IP addresses
- do not communicate directly with other clients

## Client and server devices are not equivalent

# Peer-to-peer (P2P) architecture

# Peers request service from other peers, provide service in return to other peers

#### End systems directly communicate

- self scalability new peers bring new service capacity, as well as new service demands
- minimal/no use of always-on server
- E.g., Skype, BitTorrent

#### **Complex management**

- peers are intermittently connected and change IP addresses
- Q: why is this complex?

# All devices are equivalent: a client can also be a server



# **Processes communicating**

#### Process

 program in execution, running within a host

#### Processes within same host

 communicate by using interprocess communication (defined by OS)

#### Clients, servers

- client process
  - process that initiates communication

#### server process

 process that waits to be contacted

#### Processes on different hosts

communicate by exchanging messages

#### Aside

 applications with P2P architectures also have client & server processes

Our goal learn how to build client/server applications that use sockets to communicate

vumanfredi@wesleyan.edu

# Network Programming OVERVIEW

vumanfredi@wesleyan.edu

# How do two processes talk over a network?

### Via sockets

- interface transport layer provides to apps to access network
- connection endpoint with associated IP addr, port #



# Python socket module

#### import socket

- gives access to BSD (Berkeley Socket Distribution) socket interface
  - POSIX sockets <-> Berkeley sockets <-> BSD sockets
  - available on pretty much every modern operating system

#### Resources

- https://docs.python.org/3/howto/sockets.html
- <u>https://docs.python.org/3/library/socket.html</u>

#### Socket exceptions

- <u>https://docs.python.org/3/library/socket.html#exceptions</u>
- You must read/write bytes from/to a socket
  - encode string to bytes: string.encode('utf-8')
  - decode string from bytes: string.decode('utf-8')

# Sockets

### **Address families**

- AF\_UNIX
  - local, inter-process communication
- AF\_INET4
  - Internet protocol v4
- AF\_INET6
  - Internet v6

### Socket types

- SOCK\_DGRAM
  - UDP packets
- SOCK\_STREAM
  - TCP packets
- SOCK\_RAW
  - don't let OS process transport header on packet, have OS send/receive raw packet

#### Part of process identifier: e.g., <ip address, port>

To send HTTP message to wesleyan.edu web server

- IP address: 129.133.7.68
- port number: 80

Different types of service offered by different socket types

# 2 main socket types for 2 transport services

#### TCP (Transmission Control Protocol)

#### connection-oriented

 before data exchange takes place, a logical connection is first established

#### - reliable, byte stream-oriented

 delivery is in-order, error- and loss-free, no duplication

#### App reads in-order, error-free bytes from socket

#### UDP (User Datagram Protocol)

- connection-less
  - · data is sent directly in a best-effort way
- unreliable
  - data can arrive out-of-order, be lost, corrupted, duplicated

App reads whatever is currently at socket, whether out-of-order, missing etc.

Any reliability must be implemented by app

# Send data (from python reference)

#### socket.send(bytes) - TCP

 Send data to the socket. The socket must be connected to a remote socket. Returns the number of bytes sent. Applications are responsible for checking that all data has been sent; if only some of the data was transmitted, the application needs to attempt delivery of the remaining data

#### socket.sendall(bytes) - TCP

 Send data to the socket. The socket must be connected to a remote socket. Unlike send(), this method continues to send data from bytes until either all data has been sent or an error occurs. None is returned on success. On error, an exception is raised, and there is no way to determine how much data, if any, was successfully sent.

#### socket.sendto(bytes, address) - UDP

 Send data to the socket. The socket should not be connected to a remote socket, since the destination socket is specified by address.

# Receive data (from python reference)

### Socket.recv(num\_bytes)

 Receive data from the socket. The return value is a bytes object representing the data received. The maximum amount of data to be received at once is specified by *bufsize*.

# **Partial Send/Recv**

## socket.sendall()

- generally preferable to use to eliminate partial send

## socket.recv()

- app needs way to know whether it has read everything from socket
  - "end" flag
  - a priori knowledge of number of bytes to read

• ...

- typically put recv() in while loop
  - · keep reading until nothing left to read from socket

# Endianness

## **Big endian**

big end first: largest byte (containing most significant bit) first

### Little endian

- little end first: smallest byte (containing least significant bit) first

#### Network byte order

- big endian

### UTF-8 byte order

- stays the same regardless of endian-ness of machine
- i.e., you shouldn't need to worry about byte order

# Network Programming TCP SOCKETS

vumanfredi@wesleyan.edu

# Socket programming with TCP

# Client must first contact server before sending data

- server process must be running
  - creates socket (door) that welcomes client's contact

## How?

- create TCP socket
  - specify server IP addr, port #
- "handshake" occurs
  - TCP Syn/Synack/Ack exchanged
  - if succeeds, connection established, can send data

### When contacted by client

- server TCP creates new socket for server process to communicate with that particular client
  - allows server to talk with multiple clients
  - source port numbers used to distinguish clients

#### **Application viewpoint**

TCP provides reliable, in-order byte-stream transfer ("pipe") between client and server

# **TCP Socket**



# **Client/server socket interaction: TCP**



# **Application example**

### 1. Client

 reads a line of characters (data) from its keyboard and sends data to server via socket

### 2. Server

- receives data from socket and converts characters to uppercase

#### 3. Server

- sends modified data to client

### 4. Client

- receives modified data and displays line on its screen

# **Application example: TCP server**

#### Python TCPServer



# Application example: TCP client

## **Python TCPClient**

create TCP socket for server, remote port \_\_\_\_\_ 12000

No need to attach server name, port

from socket import \* serverName = 'servername' serverPort = 12000clientSocket = socket(AF\_INET, SOCK\_STREAM) clientSocket.connect(serverName,serverPort) sentence = raw input('Input lowercase sentence:') clientSocket.send(sentence.encode()) modifiedSentence = clientSocket.recv(1024) print ('From Server:', modifiedSentence.decode()) clientSocket.close()

# echo\_client.py and echo\_server.py

# Look at code and run: available on class schedule

# Packet sniffing WIRESHARK

vumanfredi@wesleyan

# How can I look at network traffic?

### Packet sniffer

- passively observes messages transmitted and received on a particular network interface by processes running on your computer
- often requires root privileges to run

#### Popular packet sniffers

- Wireshark (also command-line version, tshark)
- tcpdump (Unix) and WinDump (Windows)
- use command line sniffers to analyze packet traces with bash script

# Packet sniffer operation



# Wireshark

## Install

- https://www.wireshark.org/download.html

Run

- type Wireshark in terminal, or double-click icon
- Wireshark display may look different for Linux vs. Mac vs. Windows

>>

• •	•	📕 The V	Wireshark Ne	twork Analy	zer		
	<ul> <li>2</li> <li>2</li> <li>3</li> <li>4</li> <li>4&lt;</li></ul>		۹ (=	۵ 🔮	T 1		Ð O
Арр	ly a display filter <発/>					•	Expression
	Welcome to Wireshark						
	Capture						
	using this filter: 📙 Enter a d	capture filter					•
Choose an	Wi-Fi: en0 awdl0	mm					
interface to	Thunderbolt Bridge: bridge0						
capture	Thunderbolt 2: en2						
traffic on	p2p0 Loopback: lo0	·					

# What do we see?

•	🖲 😑 🧧 Wi-Fi: enO													
(	1 🔳 🧕 🔘	🖿 🖹 🕅	Q 🍋 🔿 🖄	₹ 🛃 📘		⊕, ⊖, €, ፹								
	Apply a display filter	Expression +												
No	. Time 77 7.313771 <b>S</b>	ource IP	Dest IP	Protocols	ngth 166 194	Protocol State 24fc A in								
	79 7.315676	129.133.6.10	129.133.178.53	DNS	166	Standard query response 0xbd43 A in								
	80 7.374379	173.192.82.195	129.133.182.236	TLSv1.2	97	Application Data								
		29.133.182.236	173.192.82.195	TCP	66	62762 → 443 [ACK] Seq=1 Ack=32 Win=								
	Captured	29.133.182.236	173.192.82.195	TLSv1.2	101	Application Data								
		73.192.82.195	129.133.182.236	TCP	66	443 → 62762 [ACK] Seq=32 Ack=36 Win								
	packets	29.133.182.236	129.133.72.61	TCP	181	[TCP segment of a reassembled PDU]								
	03 0.01/203	129.133.72.61	129.133.182.236	TCP	181	[TCP segment of a reassembled PDU]								
	86 8.017283	129.133.182.236	129.133.72.61	TCP	66	62496 → 8009 [ACK] Seq=231 Ack=231								
	87 8.578356	JuniperN_1e:18:01	Broadcast	ARP	64	Gratuitous ARP for 129.133.176.1 (R								
	88 8.622793	129.133.182.236	216.58.219.229	TCP	54	63800 → 443 [ACK] Seq=1 Ack=1 Win=4								
	89 8.639661	216.58.219.229	129.133.182.236	тср	66	[TCP ACKed unseen segment] 443 → 63								
	90 9.602437	JuniperN_1e:18:01	Broadcast	ARP	64	Gratuitous ARP for 129.133.176.1 (R								
	91 9.848778	129.133.182.236	198.105.244.104	TCP	78	668 → 515 [SYN] Seq=0 Win=65535 Len								
►	Frame 77: 166 byt	es on wire (1328 bits	), 166 bytes captur	red (1328 bits) on	inter	rface Ø								
►	Ethernet II, Src:	JuniperN_1e:18:01 (3	c:8a:b0:1e:18:01),	Dst: Apple_c5:b4:	9a (78	3:31:c1:c5:b4:9a)								
►	Internet Protocol	Version 4, Src: 129.	133.6.11, Dst: 129.	133.178.53										
►	User Datagram Pro	otocol, Src Port: 53 (	53), Dst Port: 4400	55 (44065)		Packet								

User Datagram Protocol, Src Port: 53 (53), Dst Port: 44065 (44065)

Domain Name System (response) 2 hex digits = 1 byte= 1 ascii char ►

c1 c5 b4 9a 3c 8a b	0 1e 18 01 08	00 45 00	x1 <e< td=""></e<>
20 98 00 00 3e 11 a	0 72 81 85 06	0b 81 85	>r
If you click on okt o	r boodor field	00 01	.5.5.!\$
IT you click of pkt c		33 63	i nt.nyt.
will highlight hex/a	scii fields and	00 01	om
vice ver	79 74	".wild card.ny	
	30	vนโกสติท	rethe worsledgeke

#### Packet contents in hex and ascii: can match bytes to header

details

wireshark\_pcapng\_en0\_20160824155218\_HN8Ru3

0000

0010

0020

0030

0040

0050

0060

1

78 31

00 98

b2

00

6f

ad

69

Packets: 48516 · Displayed: 48516 (100.0%) · Dropped: 0 (0.0%) Profile: Default

# What do we see?

	87	8.5/	8350	,	1	Jun	ıpe	rN_	1e:	18:0	01	l	oroa	aca	IST			AK	٢		64
	88	8.62	2793	1		129	.13	3.1	82.2	236		2	216.	58.	219	9.229	)	TC	Р		54
	89	8.63	9661	l.		216	.58	.21	9.22	29		1	129.	133	3.18	82.23	16	тс	P		66
Layers	90	9.60	2437			Jun	ipe	rN_	1e:1	18:0	01	[	Broa	dca	st			AR	Р		64
	91	9.84	8778		3	129	.13	3.1	82.2	236		1	198.	105	5.24	44.10	14	TC	Р		78
Physical —	Eram	0 77	166	; h	100	0.0		re	(13	28	hit	c)	160	5 h	vte	c (2)	atured	(13	28 hite)	on i	nter
Link — — — — — — — — — — — — — — — — — — —	► Ethe	rnet	II,	Sro	: J	luni	per	N_1	le:1	8:0	1 (	3c:	8a:1	00:	le:	18:0	1), Ds	t: A	pple_c5:t	04:9a	a (78
Network —	► User	Data	agran	n Pr	roto	col	, S	arc	Por	t:	53	(53	), [	)st	Po	rt:	44065	(440	65)		
Transport	▶ Doma	in Na	ame s	syst	cem	(re	espo	onse	2)												
Application	0000	78 31	c1	c5	b4	9a 1	3c	8a	bØ	1e	18	01	08	00	45	00	x1	.<.	E.	1	
	0010	00 98	20	98 35	ac	21	3e 00 1	11 84	a0 ee	/2 d2	24	85 fc	06 81	80	81	85 01	.5.5	.>.	.r		
	0030	00 03	00	00	00	00	03	69	6e	74	03	6e	79	74	03	63		i	nt.nyt.c		
	0040	6f 6d	00	00	01	00	01	c0	0c	00	05	00	01	00	00	01	om				
	0050	ad 00	22	08	77	69	6c	64	63	61	72	64	07	6e	79	74	·· · · ·	vild	card.nyt		
	0060	69 6d	65	73	03	63	6f	6d	07	65	64	67	65	6b	65	79	imes.	com	.edgekey		
	0 7	wire	shark_	pcap	ong_e	en0_2	2016	082	41552	218_	HN8	Ru3						Pa	ckets: 48516	· Disp	layed: 4

# Add a filter

•	• •							-	Wi-Fi: er	0								
(		Only	ТСР			٩	۰	ا	2 7	5 👱 🛛		Ð	.Θ.	€,		_		
No	tcp	traf	fic			A Dest	ination		See	only	ТСР	fo			× → •	Exp	pression	+
140.	18	0.33001/	129.1	33.182.	236	129	.133.73	3.18		тср	1	81 IT	CP seam	ent o	f a rea	ssem	bled PDL	1
	20	0.362499	129.1	33.182.	236	129	.133.73	3.18		ТСР		66 629	919 → 8	009 [/	ACK] Se	q=11	6 Ack=11	16
	21	0.393788	129.1	33.182.	236	52.	209.21	. 15		ТСР	14	34 [T(	P seam	ent of	f a rea	ssemi	bled PDL	J]
	22	0.393789	129.1	33.182.	236	52.	209.21	. 15		TLSv1.2		bro	otoc	ol r	une			
	25	0.499503	129.1	33.182.	236	52.	209.21	. 15		ТСР	ILO	, bic			u115	=226	9 Ack=53	374
	30	1.725135	129.1	33.182.	236	129	.133.72	2.223		ТСР		OV	er T(	CP		semi	bled PDL	1] -
-												•••						-
•	Frame	e 18: 181 b	ytes on v	vire (14	448 bit	ts), 18	31 byte	s cap	tured (	1448 bit	s) on in	terfac	e 0					-
	Ether	rnet II, Sr	c: Apple	c5:b4:9	a (78:	:31:c1:	c5:b4:	9a),	Dst: Ju	niperN_1	le:18:01	(3c:8a	a:b0:1e	:18:01	1)			
	Inte	rnet Protoc	ol Versi	on 4, Sr	rc: 129	9.133.1	82.236	, Dst	: 129.1	33.73.18	3							
	Trans	smission Co	ntrol Pro	otocol,	Src Po	ort: 62	2919 (6	2919)	, Dst F	ort: 800	(8009)	, Seq:	1, Ac	k: 1,	Len: 1	15		
	So	urce Port:	62919															
	De	stination F	Port: 800	9														
	[S	tream index	(: 1]															
	[Т	CP Segment	Len: 115	]														
	Se	quence numb	ber: 1	(relat	ive se	quence	number	r)										
	[N	lext sequend	e number	: 116	(rel	ative	sequend	ce nui	mber)]									
	Ac	knowledgmer	nt number	: 1	(relat	ive ac	k numbe	er)										
	He	ader Length	n: 32 byt	es														
	▶ Fl	ags: 0x018	(PSH, AC	к)														
	Wi	ndow size w	alue: 40	96														
	[0	alculated w	vindow si	ze: 409	6]													
00	00 3	sc 8a b0 1e	18 01 78	31 c1	c5 b4	9a 08	00 45	00	<	×1	.E.							0
00	10 0	0 a7 71 c6	40 00 40	06 c5	81 81	85 b6	ec 81	85	q.@.(	a								
00	20 4	9 12 f5 c7	11 49 13	al 0e	17 4a	03 85	8e 80	18	11	·· ··J								
00	40 7	d e2 17 03	03 00 66	00 00	00 00	00 00	04 1e	15	}	n								
00	50 7	3 6f 3b 63	f0 86 d9	d3 bd	17 fc	04 3d	a9 43	8c	so;c	=	.c.							
00	60 4	le 63 ea d8	c0 b0 bf	f1 a1	d5 3b	6a a6	d5/92	ahfr	NGiODW	restevar	.K							
0	7	Transmission	Control Brote	ool: Brotor	col					Dackate: A	8516 Dienis	wed: 465	527 /05 09		anad: 0 (0	0%)	Profile: Det	fault