

## Chapter 1 Review Questions

1. There is no difference. Throughout this text, the words “host” and “end system” are used interchangeably. End systems include PCs, workstations, Web servers, mail servers, PDAs, Internet-connected game consoles, etc.
2. From Wikipedia: Diplomatic protocol is commonly described as a set of international courtesy rules. These well-established and time-honored rules have made it easier for nations and people to live and work together. Part of protocol has always been the acknowledgment of the hierarchical standing of all present. Protocol rules are based on the principles of civility.
3. Standards are important for protocols so that people can create networking systems and products that interoperate.
4. 1. Dial-up modem over telephone line: home; 2. DSL over telephone line: home or small office; 3. Cable to HFC: home; 4. 100 Mbps switched Ethernet: enterprise; 5. Wifi (802.11): home and enterprise; 6. 3G and 4G: wide-area wireless.
5. HFC bandwidth is shared among the users. On the downstream channel, all packets emanate from a single source, namely, the head end. Thus, there are no collisions in the downstream channel.
6. In most American cities, the current possibilities include: dial-up; DSL; cable modem; fiber-to-the-home.
7. Ethernet LANs have transmission rates of 10 Mbps, 100 Mbps, 1 Gbps and 10 Gbps.
8. Today, Ethernet most commonly runs over twisted-pair copper wire. It also can run over fibers optic links.
9. Dial up modems: up to 56 Kbps, bandwidth is dedicated; ADSL: up to 24 Mbps downstream and 2.5 Mbps upstream, bandwidth is dedicated; HFC, rates up to 42.8 Mbps and upstream rates of up to 30.7 Mbps, bandwidth is shared. FTTH: 2-10Mbps upload; 10-20 Mbps download; bandwidth is not shared.
10. There are two popular wireless Internet access technologies today:
  - a) Wifi (802.11) In a wireless LAN, wireless users transmit/receive packets to/from an base station (i.e., wireless access point) within a radius of few tens of meters. The base station is typically connected to the wired Internet and thus serves to connect wireless users to the wired network.
  - b) 3G and 4G wide-area wireless access networks. In these systems, packets are transmitted over the same wireless infrastructure used for cellular telephony, with the

base station thus being managed by a telecommunications provider. This provides wireless access to users within a radius of tens of kilometers of the base station.

11. At time  $t_0$  the sending host begins to transmit. At time  $t_1 = L/R_1$ , the sending host completes transmission and the entire packet is received at the router (no propagation delay). Because the router has the entire packet at time  $t_1$ , it can begin to transmit the packet to the receiving host at time  $t_1$ . At time  $t_2 = t_1 + L/R_2$ , the router completes transmission and the entire packet is received at the receiving host (again, no propagation delay). Thus, the end-to-end delay is  $L/R_1 + L/R_2$ .
12. A circuit-switched network can guarantee a certain amount of end-to-end bandwidth for the duration of a call. Most packet-switched networks today (including the Internet) cannot make any end-to-end guarantees for bandwidth. FDM requires sophisticated analog hardware to shift signal into appropriate frequency bands.
13. a) 2 users can be supported because each user requires half of the link bandwidth.  
b) Since each user requires 1Mbps when transmitting, if two or fewer users transmit simultaneously, a maximum of 2Mbps will be required. Since the available bandwidth of the shared link is 2Mbps, there will be no queuing delay before the link. Whereas, if three users transmit simultaneously, the bandwidth required will be 3Mbps which is more than the available bandwidth of the shared link. In this case, there will be queuing delay before the link.  
c) Probability that a given user is transmitting = 0.2  
d) Probability that all three users are transmitting simultaneously =  $\binom{3}{3} p^3 (1-p)^{3-3}$   
=  $(0.2)^3 = 0.008$ . Since the queue grows when all the users are transmitting, the fraction of time during which the queue grows (which is equal to the probability that all three users are transmitting simultaneously) is 0.008.
14. If the two ISPs do not peer with each other, then when they send traffic to each other they have to send the traffic through a provider ISP (intermediary), to which they have to pay for carrying the traffic. By peering with each other directly, the two ISPs can reduce their payments to their provider ISPs. An Internet Exchange Points (IXP) (typically in a standalone building with its own switches) is a meeting point where multiple ISPs can connect and/or peer together. An ISP earns its money by charging each of the the ISPs that connect to the IXP a relatively small fee, which may depend on the amount of traffic sent to or received from the IXP.
15. Google's private network connects together all its data centers, big and small. Traffic between the Google data centers passes over its private network rather than over the public Internet. Many of these data centers are located in, or close to, lower tier ISPs. Therefore, when Google delivers content to a user, it often can bypass higher tier ISPs. What motivates content providers to create these networks? First, the content

provider has more control over the user experience, since it has to use few intermediary ISPs. Second, it can save money by sending less traffic into provider networks. Third, if ISPs decide to charge more money to highly profitable content providers (in countries where net neutrality doesn't apply), the content providers can avoid these extra payments.

16. The delay components are processing delays, transmission delays, propagation delays, and queuing delays. All of these delays are fixed, except for the queuing delays, which are variable.
17. a) 1000 km, 1 Mbps, 100 bytes  
b) 100 km, 1 Mbps, 100 bytes
18. 10msec; d/s; no; no
19. a) 500 kbps  
b) 64 seconds  
c) 100kbps; 320 seconds
20. End system A breaks the large file into chunks. It adds header to each chunk, thereby generating multiple packets from the file. The header in each packet includes the IP address of the destination (end system B). The packet switch uses the destination IP address in the packet to determine the outgoing link. Asking which road to take is analogous to a packet asking which outgoing link it should be forwarded on, given the packet's destination address.
21. The maximum emission rate is 500 packets/sec and the maximum transmission rate is 350 packets/sec. The corresponding traffic intensity is  $500/350 = 1.43 > 1$ . Loss will eventually occur for each experiment; but the time when loss first occurs will be different from one experiment to the next due to the randomness in the emission process.
22. Five generic tasks are error control, flow control, segmentation and reassembly, multiplexing, and connection setup. Yes, these tasks can be duplicated at different layers. For example, error control is often provided at more than one layer.
23. The five layers in the Internet protocol stack are – from top to bottom – the application layer, the transport layer, the network layer, the link layer, and the physical layer. The principal responsibilities are outlined in Section 1.5.1.
24. Application-layer message: data which an application wants to send and passed onto the transport layer; transport-layer segment: generated by the transport layer and encapsulates application-layer message with transport layer header; network-layer datagram: encapsulates transport-layer segment with a network-layer header; link-layer frame: encapsulates network-layer datagram with a link-layer header.

25. Routers process network, link and physical layers (layers 1 through 3). (This is a little bit of a white lie, as modern routers sometimes act as firewalls or caching components, and process Transport layer as well.) Link layer switches process link and physical layers (layers 1 through 2). Hosts process all five layers.
26. a) Virus  
Requires some form of human interaction to spread. Classic example: E-mail viruses.
- b) Worms  
No user replication needed. Worm in infected host scans IP addresses and port numbers, looking for vulnerable processes to infect.
27. Creation of a botnet requires an attacker to find vulnerability in some application or system (e.g. exploiting the buffer overflow vulnerability that might exist in an application). After finding the vulnerability, the attacker needs to scan for hosts that are vulnerable. The target is basically to compromise a series of systems by exploiting that particular vulnerability. Any system that is part of the botnet can automatically scan its environment and propagate by exploiting the vulnerability. An important property of such botnets is that the originator of the botnet can remotely control and issue commands to all the nodes in the botnet. Hence, it becomes possible for the attacker to issue a command to all the nodes, that target a single node (for example, all nodes in the botnet might be commanded by the attacker to send a TCP SYN message to the target, which might result in a TCP SYN flood attack at the target).
28. Trudy can pretend to be Bob to Alice (and vice-versa) and partially or completely modify the message(s) being sent from Bob to Alice. For example, she can easily change the phrase "Alice, I owe you \$1000" to "Alice, I owe you \$10,000". Furthermore, Trudy can even drop the packets that are being sent by Bob to Alice (and vice-versa), even if the packets from Bob to Alice are encrypted.

## Chapter 1 Problems

### Problem 1

There is no single right answer to this question. Many protocols would do the trick. Here's a simple answer below:

#### Messages from ATM machine to Server

Msg name	purpose
-----	-----
HELO <userid>	Let server know that there is a card in the ATM machine
	ATM card transmits user ID to Server
PASSWD <passwd>	User enters PIN, which is sent to server
BALANCE	User requests balance
WITHDRAWL <amount>	User asks to withdraw money
BYE	user all done

#### Messages from Server to ATM machine (display)

Msg name	purpose
-----	-----
PASSWD	Ask user for PIN (password)
OK	last requested operation (PASSWD, WITHDRAWL)
	OK
ERR	last requested operation (PASSWD, WITHDRAWL) in ERROR
AMOUNT <amt>	sent in response to BALANCE request
BYE	user done, display welcome screen at ATM

#### Correct operation:

client		server
HELO (userid)	----->	(check if valid userid)
	<-----	PASSWD
PASSWD <passwd>	----->	(check password)
	<-----	OK (password is OK)
BALANCE	----->	
	<-----	AMOUNT <amt>
WITHDRAWL <amt>	----->	check if enough \$ to cover withdrawl
	<-----	OK
ATM dispenses \$		
BYE	----->	
	<-----	BYE

In situation when there's not enough money:

```

HELO (userid) -----> (check if valid userid)
<----- PASSWD
PASSWD <passwd> -----> (check password)
<----- OK (password is OK)
BALANCE ----->
<----- AMOUNT <amt>
WITHDRAWL <amt> -----> check if enough $ to cover
withdrawl
<----- ERR (not enough funds)
error msg displayed
no $ given out
BYE ----->
<----- BYE

```

## Problem 2

At time  $N*(L/R)$  the first packet has reached the destination, the second packet is stored in the last router, the third packet is stored in the next-to-last router, etc. At time  $N*(L/R) + L/R$ , the second packet has reached the destination, the third packet is stored in the last router, etc. Continuing with this logic, we see that at time  $N*(L/R) + (P-1)*(L/R) = (N+P-1)*(L/R)$  all packets have reached the destination.

## Problem 3

- a) A circuit-switched network would be well suited to the application, because the application involves long sessions with predictable smooth bandwidth requirements. Since the transmission rate is known and not bursty, bandwidth can be reserved for each application session without significant waste. In addition, the overhead costs of setting up and tearing down connections are amortized over the lengthy duration of a typical application session.
- b) In the worst case, all the applications simultaneously transmit over one or more network links. However, since each link has sufficient bandwidth to handle the sum of all of the applications' data rates, no congestion (very little queuing) will occur. Given such generous link capacities, the network does not need congestion control mechanisms.

## Problem 4

- a) Between the switch in the upper left and the switch in the upper right we can have 4 connections. Similarly we can have four connections between each of the 3 other pairs of adjacent switches. Thus, this network can support up to 16 connections.
- b) We can 4 connections passing through the switch in the upper-right-hand corner and another 4 connections passing through the switch in the lower-left-hand corner, giving a total of 8 connections.

- c) Yes. For the connections between A and C, we route two connections through B and two connections through D. For the connections between B and D, we route two connections through A and two connections through C. In this manner, there are at most 4 connections passing through any link.

### Problem 5

Tollbooths are 75 km apart, and the cars propagate at 100km/hr. A tollbooth services a car at a rate of one car every 12 seconds.

- a) There are ten cars. It takes 120 seconds, or 2 minutes, for the first tollbooth to service the 10 cars. Each of these cars has a propagation delay of 45 minutes (travel 75 km) before arriving at the second tollbooth. Thus, all the cars are lined up before the second tollbooth after 47 minutes. The whole process repeats itself for traveling between the second and third tollbooths. It also takes 2 minutes for the third tollbooth to service the 10 cars. Thus the total delay is 96 minutes.
- b) Delay between tollbooths is  $8 \cdot 12$  seconds plus 45 minutes, i.e., 46 minutes and 36 seconds. The total delay is twice this amount plus  $8 \cdot 12$  seconds, i.e., 94 minutes and 48 seconds.

### Problem 6

- a)  $d_{prop} = m / s$  seconds.
- b)  $d_{trans} = L / R$  seconds.
- c)  $d_{end-to-end} = (m / s + L / R)$  seconds.
- d) The bit is just leaving Host A.
- e) The first bit is in the link and has not reached Host B.
- f) The first bit has reached Host B.
- g) Want

$$m = \frac{L}{R} s = \frac{120}{56 \times 10^3} (2.5 \times 10^8) = 536 \text{ km.}$$

### Problem 7

Consider the first bit in a packet. Before this bit can be transmitted, all of the bits in the packet must be generated. This requires

$$\frac{56 \cdot 8}{64 \times 10^3} \text{ sec} = 7 \text{ msec.}$$

The time required to transmit the packet is

$$\frac{56 \cdot 8}{2 \times 10^6} \text{ sec} = 224 \mu \text{ sec.}$$

Propagation delay = 10 msec.  
The delay until decoding is

$$7 \text{ msec} + 224 \mu \text{ sec} + 10 \text{ msec} = 17.224 \text{ msec}$$

A similar analysis shows that all bits experience a delay of 17.224 msec.

### Problem 8

a) 20 users can be supported.

b)  $p = 0.1$ .

$$\text{c) } \binom{120}{n} p^n (1-p)^{120-n}.$$

$$\text{d) } 1 - \sum_{n=0}^{20} \binom{120}{n} p^n (1-p)^{120-n}.$$

We use the central limit theorem to approximate this probability. Let  $X_j$  be independent random variables such that  $P(X_j = 1) = p$ .

$$P(\text{“21 or more users”}) = 1 - P\left(\sum_{j=1}^{120} X_j \leq 21\right)$$

$$\begin{aligned} P\left(\sum_{j=1}^{120} X_j \leq 21\right) &= P\left(\frac{\sum_{j=1}^{120} X_j - 12}{\sqrt{120 \cdot 0.1 \cdot 0.9}} \leq \frac{9}{\sqrt{120 \cdot 0.1 \cdot 0.9}}\right) \\ &\approx P\left(Z \leq \frac{9}{3.286}\right) = P(Z \leq 2.74) \\ &= 0.997 \end{aligned}$$

when  $Z$  is a standard normal r.v. Thus  $P(\text{“21 or more users”}) \approx 0.003$ .

### Problem 9

a) 10,000

$$\text{b) } \sum_{n=N+1}^M \binom{M}{n} p^n (1-p)^{M-n}$$

## Problem 10

The first end system requires  $L/R_1$  to transmit the packet onto the first link; the packet propagates over the first link in  $d_1/s_1$ ; the packet switch adds a processing delay of  $d_{proc}$ ; after receiving the entire packet, the packet switch connecting the first and the second link requires  $L/R_2$  to transmit the packet onto the second link; the packet propagates over the second link in  $d_2/s_2$ . Similarly, we can find the delay caused by the second switch and the third link:  $L/R_3$ ,  $d_{proc}$ , and  $d_3/s_3$ .

Adding these five delays gives

$$d_{end-end} = L/R_1 + L/R_2 + L/R_3 + d_1/s_1 + d_2/s_2 + d_3/s_3 + d_{proc} + d_{proc}$$

To answer the second question, we simply plug the values into the equation to get  $6 + 6 + 6 + 20 + 16 + 4 + 3 + 3 = 64$  msec.

## Problem 11

Because bits are immediately transmitted, the packet switch does not introduce any delay; in particular, it does not introduce a transmission delay. Thus,

$$d_{end-end} = L/R + d_1/s_1 + d_2/s_2 + d_3/s_3$$

For the values in Problem 10, we get  $6 + 20 + 16 + 4 = 46$  msec.

## Problem 12

The arriving packet must first wait for the link to transmit  $4.5 * 1,500$  bytes = 6,750 bytes or 54,000 bits. Since these bits are transmitted at 2 Mbps, the queuing delay is 27 msec. Generally, the queuing delay is  $(nL + (L - x))/R$ .

## Problem 13

- a) The queuing delay is 0 for the first transmitted packet,  $L/R$  for the second transmitted packet, and generally,  $(n-1)L/R$  for the  $n^{th}$  transmitted packet. Thus, the average delay for the  $N$  packets is:

$$\begin{aligned} & (L/R + 2L/R + \dots + (N-1)L/R)/N \\ &= L/(RN) * (1 + 2 + \dots + (N-1)) \\ &= L/(RN) * N(N-1)/2 \\ &= LN(N-1)/(2RN) \\ &= (N-1)L/(2R) \end{aligned}$$

Note that here we used the well-known fact:

$$1 + 2 + \dots + N = N(N+1)/2$$

- b) It takes  $LN/R$  seconds to transmit the  $N$  packets. Thus, the buffer is empty when a each batch of  $N$  packets arrive. Thus, the average delay of a packet across all batches is the average delay within one batch, i.e.,  $(N-1)L/2R$ .

### Problem 14

- a) The transmission delay is  $L/R$ . The total delay is

$$\frac{IL}{R(1-I)} + \frac{L}{R} = \frac{L/R}{1-I}$$

- b) Let  $x = L/R$ .

$$\text{Total delay} = \frac{x}{1-ax}$$

For  $x=0$ , the total delay =0; as we increase  $x$ , total delay increases, approaching infinity as  $x$  approaches  $1/a$ .

### Problem 15

$$\text{Total delay} = \frac{L/R}{1-I} = \frac{L/R}{1-aL/R} = \frac{1/\mu}{1-a/\mu} = \frac{1}{\mu-a}$$

### Problem 16

The total number of packets in the system includes those in the buffer and the packet that is being transmitted. So,  $N=10+1$ .

Because  $N = a \cdot d$ , so  $(10+1)=a \cdot (\text{queuing delay} + \text{transmission delay})$ . That is,  $11=a \cdot (0.01+1/100)=a \cdot (0.01+0.01)$ . Thus,  $a=550$  packets/sec.

### Problem 17

- a) There are  $Q$  nodes (the source host and the  $Q-1$  routers). Let  $d_{proc}^q$  denote the processing delay at the  $q$ th node. Let  $R^q$  be the transmission rate of the  $q$ th link and let

$d_{trans}^q = L/R^q$ . Let  $d_{prop}^q$  be the propagation delay across the  $q$ th link. Then

$$d_{end-to-end} = \sum_{q=1}^Q [d_{proc}^q + d_{trans}^q + d_{prop}^q]$$

b) Let  $d_{queue}^q$  denote the average queuing delay at node  $q$ . Then

$$d_{end-to-end} = \sum_{q=1}^Q [d_{proc}^q + d_{trans}^q + d_{prop}^q + d_{queue}^q].$$

## Problem 18

On linux you can use the command

```
traceroute www.targethost.com
```

and in the Windows command prompt you can use

```
tracert www.targethost.com
```

In either case, you will get three delay measurements. For those three measurements you can calculate the mean and standard deviation. Repeat the experiment at different times of the day and comment on any changes.

Here is an example solution:

```
traceroute to www.poly.edu (128.238.24.40), 30 hops max, 40 byte packets
 1 thunder.sdsc.edu (132.249.20.5)  2.802 ms  0.645 ms  0.484 ms
 2 dolphin.sdsc.edu (132.249.31.17)  0.227 ms  0.248 ms  0.239 ms
 3 dc-sdg-aggr1--sdsc-1.cenic.net (137.164.23.129)  0.360 ms  0.260 ms  0.240 ms
 4 dc-riv-core1--sdg-aggr1-10ge-2.cenic.net (137.164.47.14)  8.847 ms  8.497 ms  8.230 ms
 5 dc-lax-core1--lax-core2-10ge-2.cenic.net (137.164.46.64)  9.969 ms  9.920 ms  9.846 ms
 6 dc-lax-px1--lax-core1-10ge-2.cenic.net (137.164.46.151)  9.845 ms  9.729 ms  9.724 ms
 7 hurricane--lax-px1-ge.cenic.net (198.32.251.86)  9.971 ms  16.981 ms  9.850 ms
 8 10gigabitethernet4-3.core1.nyc4.he.net (72.52.92.225)  72.796 ms  80.278 ms  72.346 ms
 9 10gigabitethernet3-4.core1.nyc5.he.net (184.105.213.218)  71.126 ms  71.442 ms  73.623 ms
10 lighttower-fiber-networks.10gigabitethernet3-2.core1.nyc5.he.net (216.66.50.106)  70.924 ms  70.959 ms  71.072 ms
11 ae0.nycmnyzrj91.lighttower.net (72.22.160.156)  70.870 ms  71.089 ms  70.957 ms
12 72.22.188.102 (72.22.188.102)  71.242 ms  71.228 ms  71.102 ms
```

```
traceroute to www.poly.edu (128.238.24.40), 30 hops max, 40 byte packets
 1 thunder.sdsc.edu (132.249.20.5)  0.478 ms  0.353 ms  0.308 ms
 2 dolphin.sdsc.edu (132.249.31.17)  0.212 ms  0.251 ms  0.238 ms
 3 dc-sdg-aggr1--sdsc-1.cenic.net (137.164.23.129)  0.237 ms  0.246 ms  0.240 ms
 4 dc-riv-core1--sdg-aggr1-10ge-2.cenic.net (137.164.47.14)  8.628 ms  8.348 ms  8.357 ms
 5 dc-lax-core1--lax-core2-10ge-2.cenic.net (137.164.46.64)  9.934 ms  9.963 ms  9.852 ms
 6 dc-lax-px1--lax-core1-10ge-2.cenic.net (137.164.46.151)  9.831 ms  9.814 ms  9.676 ms
 7 hurricane--lax-px1-ge.cenic.net (198.32.251.86)  10.194 ms  10.012 ms  16.722 ms
 8 10gigabitethernet4-3.core1.nyc4.he.net (72.52.92.225)  73.856 ms  73.196 ms  73.979 ms
 9 10gigabitethernet3-4.core1.nyc5.he.net (184.105.213.218)  71.247 ms  71.199 ms  71.646 ms
10 lighttower-fiber-networks.10gigabitethernet3-2.core1.nyc5.he.net (216.66.50.106)  70.987 ms  71.073 ms  70.985 ms
11 ae0.nycmnyzrj91.lighttower.net (72.22.160.156)  71.075 ms  71.042 ms  71.328 ms
12 72.22.188.102 (72.22.188.102)  71.626 ms  71.299 ms  72.236 ms
```

```

1 thunder.sdsc.edu (132.249.20.5) 0.403 ms 0.347 ms 0.358 ms
2 dolphin.sdsc.edu (132.249.31.17) 0.225 ms 0.244 ms 0.237 ms
3 dc-sdg-aggl--sdsc-1.cenic.net (137.164.23.129) 0.362 ms 0.256 ms 0.239 ms
4 dc-riv-core1--sdg-aggl-10ge-2.cenic.net (137.164.47.14) 8.850 ms 8.358 ms 8.227 ms
5 dc-lax-core1--lax-core2-10ge-2.cenic.net (137.164.46.64) 10.096 ms 9.869 ms 10.351 ms
6 dc-lax-px1--lax-core1-10ge-2.cenic.net (137.164.46.151) 9.721 ms 9.621 ms 9.725 ms
7 hurricane--lax-px1-ge.cenic.net (198.32.251.86) 11.345 ms 10.048 ms 13.844 ms
8 10gigabitethernet4-3.core1.nyc4.he.net (72.52.92.225) 71.920 ms 72.977 ms 77.264 ms
9 10gigabitethernet3-4.core1.nyc5.he.net (184.105.213.218) 71.273 ms 71.247 ms 71.291 ms
10 lighttower-fiber-networks.10gigabitethernet3-2.core1.nyc5.he.net (216.66.50.106) 71.114 ms 82.516 ms 71.136 ms
11 ae0.nycmnyzrj91.lighttower.net (72.22.160.156) 71.232 ms 71.071 ms 71.039 ms
12 72.22.188.102 (72.22.188.102) 71.585 ms 71.608 ms 71.493 ms

```

### Traceroutes between San Diego Super Computer Center and www.poly.edu

- The average (mean) of the round-trip delays at each of the three hours is 71.18 ms, 71.38 ms and 71.55 ms, respectively. The standard deviations are 0.075 ms, 0.21 ms, 0.05 ms, respectively.
- In this example, the traceroutes have 12 routers in the path at each of the three hours. No, the paths didn't change during any of the hours.
- Traceroute packets passed through four ISP networks from source to destination. Yes, in this experiment the largest delays occurred at peering interfaces between adjacent ISPs.

```

traceroute to www.poly.edu (128.238.24.40), 30 hops max, 60 byte packets
1 62-193-36-1.stella-net.net (62.193.36.1) 0.500 ms 0.415 ms 0.440 ms
2 62.193.33.29 (62.193.33.29) 0.910 ms 1.065 ms 1.026 ms
3 bg1.stella-net.net (62.193.32.254) 0.972 ms 1.026 ms 1.078 ms
4 62.193.32.66 (62.193.32.66) 1.021 ms 0.988 ms 0.947 ms
5 10gigabitethernet-2-2.par2.he.net (195.42.144.104) 1.537 ms 1.752 ms 1.714 ms
6 10gigabitethernet7-1.core1.ash1.he.net (184.105.213.93) 80.273 ms 80.103 ms 79.971 ms
7 10gigabitethernet1-2.core1.nyc4.he.net (72.52.92.85) 86.494 ms 85.872 ms 86.223 ms
8 10gigabitethernet3-4.core1.nyc5.he.net (184.105.213.218) 85.248 ms 85.424 ms 85.388 ms
9 lighttower-fiber-networks.10gigabitethernet3-2.core1.nyc5.he.net (216.66.50.106) 86.194 ms 85.864 ms 86.116 ms
10 ae0.nycmnyzrj91.lighttower.net (72.22.160.156) 85.796 ms 85.823 ms 85.766 ms
11 72.22.188.102 (72.22.188.102) 87.717 ms 86.817 ms 86.774 ms

```

```

traceroute to www.poly.edu (128.238.24.40), 30 hops max, 60 byte packets
 1 62-193-36-1.stella-net.net (62.193.36.1) 0.375 ms 0.397 ms 0.355 ms
 2 62.193.33.29 (62.193.33.29) 0.810 ms 0.877 ms 0.836 ms
 3 bg1.stella-net.net (62.193.32.254) 1.098 ms 0.991 ms 1.055 ms
 4 62.193.32.66 (62.193.32.66) 0.994 ms 0.960 ms 1.157 ms
 5 10gigabitethernet-2-2.par2.he.net (195.42.144.104) 1.679 ms 1.816 ms 1.768 ms
 6 10gigabitethernet7-1.core1.ash1.he.net (184.105.213.93) 80.416 ms 90.573 ms 90.659 ms
 7 10gigabitethernet1-2.core1.nyc4.he.net (72.52.92.85) 85.933 ms 95.987 ms 96.087 ms
 8 10gigabitethernet3-4.core1.nyc5.he.net (184.105.213.218) 90.268 ms 90.229 ms 90.030 ms
 9 lighttower-fiber-networks.10gigabitethernet3-2.core1.nyc5.he.net (216.66.50.106) 85.833 ms 85.448 ms 85.418 ms
10 ae0.nycmnyzrj91.lighttower.net (72.22.160.156) 87.067 ms 86.025 ms 85.962 ms
11 72.22.188.102 (72.22.188.102) 86.542 ms 86.369 ms 86.170 ms

```

```

traceroute to 128.238.24.40 (128.238.24.40), 30 hops max, 60 byte packets
 1 62-193-36-1.stella-net.net (62.193.36.1) 0.396 ms 0.284 ms 0.239 ms
 2 62.193.33.29 (62.193.33.29) 0.817 ms 0.786 ms 0.848 ms
 3 bg1.stella-net.net (62.193.32.254) 1.150 ms 1.216 ms 1.265 ms
 4 62.193.32.66 (62.193.32.66) 1.002 ms 0.963 ms 0.923 ms
 5 10gigabitethernet-2-2.par2.he.net (195.42.144.104) 1.573 ms 1.534 ms 1.643 ms
 6 10gigabitethernet7-1.core1.ash1.he.net (184.105.213.93) 88.738 ms 82.866 ms 82.783 ms
 7 10gigabitethernet1-2.core1.nyc4.he.net (72.52.92.85) 94.888 ms 90.936 ms 90.877 ms
 8 10gigabitethernet3-4.core1.nyc5.he.net (184.105.213.218) 90.498 ms 90.543 ms 90.482 ms
 9 lighttower-fiber-networks.10gigabitethernet3-2.core1.nyc5.he.net (216.66.50.106) 85.716 ms 85.408 ms 85.637 ms
10 ae0.nycmnyzrj91.lighttower.net (72.22.160.156) 85.779 ms 85.290 ms 85.252 ms
11 72.22.188.102 (72.22.188.102) 86.217 ms 86.652 ms 86.588 ms

```

### Traceroutes from www.stella-net.net (France) to www.poly.edu (USA).

- d) The average round-trip delays at each of the three hours are 87.09 ms, 86.35 ms and 86.48 ms, respectively. The standard deviations are 0.53 ms, 0.18 ms, 0.23 ms, respectively. In this example, there are 11 routers in the path at each of the three hours. No, the paths didn't change during any of the hours. Traceroute packets passed three ISP networks from source to destination. Yes, in this experiment the largest delays occurred at peering interfaces between adjacent ISPs.

### Problem 19

An example solution:

```

traceroute to www.poly.edu (128.238.24.30), 30 hops max, 60 byte packets
 1 62-193-36-1.stella-net.net (62.193.36.1) 0.426 ms 0.329 ms 0.284 ms
 2 62.193.33.25 (62.193.33.25) 0.810 ms 0.771 ms 0.878 ms
 3 62.193.32.66 (62.193.32.66) 0.815 ms 0.840 ms 0.801 ms
 4 10gigabitethernet-2-2.par2.he.net (195.42.144.104) 1.387 ms 1.506 ms 1.467 ms
 5 10gigabitethernet7-1.core1.ash1.he.net (184.105.213.93) 85.402 ms 85.553 ms 85.353 ms
 6 10gigabitethernet1-2.core1.nyc4.he.net (72.52.92.85) 94.360 ms 96.220 ms 96.355 ms
 7 10gigabitethernet3-4.core1.nyc5.he.net (184.105.213.218) 90.279 ms 87.459 ms 87.709 ms
 8 lighttower-fiber-networks.10gigabitethernet3-2.core1.nyc5.he.net (216.66.50.106) 85.474 ms 85.450 ms 85.983 ms
 9 ae0.nycmnyzrj91.lighttower.net (72.22.160.156) 86.160 ms 85.768 ms 86.016 ms
10 72.22.188.102 (72.22.188.102) 124.111 ms 89.340 ms 89.556 ms

```

```

 1 vl200.hs01.mar01.jaguar-network.net (85.31.192.253) 0.552 ms 0.414 ms
 2 ae1.cr01.mar01.jaguar-network.net (85.31.194.9) 0.340 ms 0.213 ms
 3 xe2-0-0.cr01.par02.jaguar-network.net (78.153.231.201) 9.933 ms 9.841 ms
 4 te1-3.er01.par02.jaguar-network.net (85.31.194.14) 9.828 ms 9.962 ms
 5 10gigabitethernet-2-2.par2.he.net (195.42.144.104) 10.456 ms 10.332 ms
 6 10gigabitethernet7-1.core1.ash1.he.net (184.105.213.93) 88.793 ms 96.781 ms
 7 10gigabitethernet1-2.core1.nyc4.he.net (72.52.92.85) 94.651 ms 99.654 ms
 8 10gigabitethernet3-4.core1.nyc5.he.net (184.105.213.218) 94.786 ms 94.755 ms
 9 lighttower-fiber-networks.10gigabitethernet3-2.core1.nyc5.he.net (216.66.50.106) 91.935 ms 91.776 ms
10 ae0.nycmnyzrj91.lighttower.net (72.22.160.156) 91.909 ms 91.784 ms
11 72.22.188.102 (72.22.188.102) 93.791 ms 93.515 ms

```

## Traceroutes from two different cities in France to New York City in United States

- a) In these traceroutes from two different cities in France to the same destination host in United States, seven links are in common including the transatlantic link.

```

traceroute to www.poly.edu (128.238.24.30), 30 hops max, 60 byte packets
 1
 2 hos-tr3.juniper2.rz10.hetzner.de 213.239.224.65 de 0.224 ms * *
   hos-tr2.juniper1.rz10.hetzner.de 213.239.224.33 de 0.174 ms 0.176 ms
 3 hos-bb1.juniper1.ffm.hetzner.de 213.239.240.224 de 4.746 ms 4.780 ms
   hos-bb1.juniper4.ffm.hetzner.de 213.239.240.230 de 4.823 ms
 4 20gigabitethernet4-3.core1.fra1.he.net 80.81.192.172 de 5.462 ms 5.461 ms 5.456 ms
 5 10gigabitethernet1-4.core1.ams1.he.net 72.52.92.94 us 12.899 ms
   10gigabitethernet5-3.core1.ams1.he.net 72.52.92.77 us 13.197 ms
   10gigabitethernet5-3.core1.lon1.he.net 184.105.213.145 us 26.110 ms
 6 10gigabitethernet1-4.core1.lon1.he.net 72.52.92.81 us 18.720 ms 18.871 ms 18.862 ms
 7 10gigabitethernet7-4.core1.nyc4.he.net 72.52.92.241 us 86.677 ms 85.580 ms 86.560 ms
 8 lighttower-fiber-networks.10gigabitethernet3-2.core1.nyc5.he.net 216.66.50.106 us 118.500 ms
   10gigabitethernet3-4.core1.nyc5.he.net 184.105.213.218 us 90.346 ms
   lighttower-fiber-networks.10gigabitethernet3-2.core1.nyc5.he.net 216.66.50.106 us 118.500 ms
 9 ae0.nycmnyzrj91.lighttower.net 72.22.160.156 us 85.289 ms 85.552 ms 85.283 ms

```