CS144
Intro to Computer Networks
Final Exam – Tuesday, December 10, 2019

2 NOTE PAGES, CLOSED BOOK, COMPUTERS OFF

Your Name: Answers

SUNet ID: root @stanford.edu

In accordance with both the letter and the spirit of the Stanford Honor Code, I neither received nor provided any assistance on this exam.

Signature: ____________________

- The exam has 12 questions totaling 86 points.
- You have 120 minutes to complete them.
- Some questions may be much harder than others.
- Keep your answers concise.
- For multiple-choice questions, circle all true answers. You will be credited points for correctly circled answers and well as for answers correctly left blank. You must circle at least one answer to get credit for the question.
- You may express numerical answers as expressions if you don't do the arithmetic, but please explain each term when you do.
- Please box your final answers.

1	/3
2	/2
3	/3
4	/2
5	/2
6	/2
7	/18
8	/14
9	/9
10	/14
11	/12
12	/5
Total	/86

I Multiple Choice Questions

1. [3 points]:

Which of the following statements are true?

A In a network, such as in their B4 network interconnecting their data-center sites, Google can use low and high priorities in the switches to keep the network highly loaded with bulk traffic, while giving low latency to user client traffic.

B When building datacenter networks, cable bundling increases number of fiber runs to decrease number of hops a packet must take.

C Network providers place major content servers like Netflix away from major population areas to even out the load on their servers.

D Some data-center owners use proprietary alternatives to TCP to control congestion for their applications.

E TIMELY is a congestion control algorithm that uses RTT as a signal to decide the rate at which the source host sends data.

Answer:

True - When near full utilization, need a way to give high-priority traffic ways to skip ahead of low-priority to avoid both delays and drops.

False - Cable bundling reduces the number of fiber runs, as the many different fiber cables are bundled together into 1. This is a cost-saving measure often used in datacenters. Depending on whether you count fiber junctions as hops or not, it keeps the number of hops the same or increases it.

False - Internet Providers place servers as close to population centers as possible to avoid extra traffic over longer links (which costs more).

True - Since the network administrators know almost everything about the system, they can have a tuned protocol to use as much of the network as possible.

True - TIMELY uses RTT timings to estimate what the actual throughput is and tries to roughly match that.

2. [2 points]:

Which of the following statements about Manchester coding are true?

A Manchester coding does not negatively affect bandwidth usage.

B Manchester coding was one of the first examples of QAM-16.

C Manchester coding indicates whether a bit is 1 or 0 based on a transition at the end of each period.

D Manchester coding makes clock recovery easier than 4b/5b encoding.

Answer:

Manchester coding indicates whether a bit is a one or a zero in a given bit period by transitioning high-to-low or low-to-high in the middle of the period. Due to the high density of transitions, clock recovery is easier than with 4b/5b encoding, but Manchester coding can double the bandwidth usage (in the worst case).

3. [3 points]:

Which of the following is/are true?

A In multicast routing, routers replicate packets on behalf of the sending host.

B In some circumstances, an ISP might select a route with a longer AS-PATH over a route with a shorter AS-PATH.

C Path-vector routing protocols are susceptible to the "counting to infinity" or "bad news travels slowly" problem.

D Both IP multicast and anycast require special support from routers.

E ICMP messages are delivered unreliably.

4. [2 points]:

Which of the following are true about different network applications?

A DHCP is a network protocol that is used to dynamically share a single public IP address between several devices on a network.

B Because of the hierarchical design of DNS, a single query may touch several servers before returning the final result.

C When accessing `mail.google.com`, the TCP checksum will always guarantee the integrity of the data under all error conditions, while Transport Layer Security (TLS) protocol guarantees confidentiality on top of TCP.

D HTTP/2 addresses the head-of-line blocking issue in HTTP/1.1 by allowing out-of-order responses over a single TCP connection, but the issue remains at the transport layer.

5. [2 points]:

Which of the following statements are true about packet switching (compared to circuit switching)?

A Packet switching uses statistical multiplexing.

B Packet switching provides more precise rate guarantees for flows.

C Packet switching requires less per-flow state.

D Packet switching always provides lower per-packet delay guarantees.

Answer:

Circuit switching requires per-flow state at each intermediate router, whereas packet switching does not. Circuit switching (typically) uses this state to avoid over-allocating the network, preventing statistical multiplexing while providing better throughput and delay guarantees.

6. **[2 points]:**

A router receives an ARP request whose target hardware address is broadcast, and whose target IP does not match any of the router's IPs. The router caches the source IP/hardware address. What else should the router do?

A Send an ICMP error to the frame's sender, in order to inform the sender that it sent the frame to a host other than the intended target.

B Send an ARP reply to the frame's sender, in order to inform the sender of the router's hardware address.

C Forward the frame over all ports except the one on which it was received, in case the intended target of the frame is reachable via one of those ports.

D Nothing.

Answer:

ARP requests are always broadcast, so no error has occurred. A router should only reply to ARP requests corresponding to its own IP. ARP requests should not be forwarded, since ARP is only meant to discover hosts within one hop of the requestor.

II Additive Increase, Additive Decrease

7. [18 points]:

In this question we will explore the dynamics of two flows using the AIAD (additive increase, additive decrease) algorithm instead of AIMD. Consider the topology shown in which two hosts A, B are sending packets to D through router R.

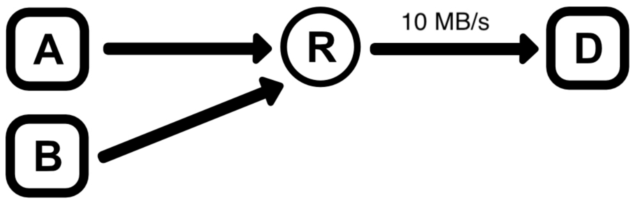

The bottleneck link in the path from R to D runs at $10\,\text{MB/s}$ (i.e. 10 Megabytes per sec ond). The RTTs from A to D and B to D are $100\,\text{ms}$. Assume that both A and B can communicate with R instantaneously and with arbitrarily large rate. Assume that each packet has a length of 1250 bytes and the buffer on router R can hold 200 packets.

a. What is the bandwidth-delay product, in Megabytes (MB)? (2 points.)

Answer: ___*1*___ MB

Answer:

100ms and 10MB/s means BPD is $1\,\text{MB}$.

b. How many packets can the network hold (on the wire plus in R)? (2 points.)

Answer: ___*1000*___ packets

Answer:

The wire contains 1MB or $10^6/(1250) = 800$ packets. i.e. 800 packets on the wire, and 200 in the router.

Suppose that at time $t = 0$, host A has a congestion window of 999 packets, B has a congestion window of 0, and R's buffer is empty.

You should assume that packet drops are always experienced by the flow with more packets in flight. You should assume that each flow updates its congestion window after each RTT, decreasing it if any packet was dropped, and increasing it otherwise.

AIMD: We first consider the evolution of the system when each flow is governed by AIMD, in which the congestion window grows by being incremented and shrinks by being halved (rounding down). In roundtrips 0 through 3, the system evolves as shown in the table below.

t (RTTs)	0	1	2	3
A's cong. window (packets)	999	1000	500	501
B's cong. window (packets)	0	1	2	3

c. At which value of t will the number of outstanding packets next exceed the network's capacity? Give your answer in RTTs, not seconds. (2 points.)

Answer: ___*252*___ RTTs

Answer:

Let us consider the time before the second drop. At t RTTs, A has $498 + t$ packets in flight, and B has t packets in flight. Thus, after $t = 252$ RTTs, there will be 1002 packets in flight, and a drop will occur.

d. What are the congestion windows in the RTT after this? Give your answer in packets. (2 points.)

A: ____*375*____ packets **B:** ____*253*____ packets

Answer:

375 packets for A, 253 for B (at $t = 253$ RTTs).

e. What are the congestion windows of both hosts at $t = 300$ RTTs? Give your answer in packets. (2 points.)

A: ____*422*____ packets **B:** ____*300*____ packets

Answer:

This is 47 RTTs after the time from the last quest, so the windows will be 422 and 300 for A and B respectively.

AIAD: Now, suppose that instead we are using an AIAD protocol. In this case, every drop decreases the congestion window by 5 packets (to a minimum of 1 packet) instead of halving the congestion window.

f. What are the congestion windows of both hosts at $t = 300$ RTTs? Give your answer in packets. (2 points.)

A: ___*300*___ packets **B:** ___*699*___ packets

Answer:

A's window increases by 1 each RTT, so it will be 300. If you make a table for small values of t, you see that for each time value that is a multiple of 3, A's congestion window is $999 - t$.

g. How does the number of dropped packets in AIMD and AIAD compare? (2 points.)

Answer: AIAD results in a **(circle one) larger/smaller/equal** number of dropped packets.

Answer:

In AIAD, more packets are dropped.

h. Does AIAD eventually converge to a fair allocation between flows? (2 points.)

Answer (circle one): yes/no

Answer:

Yes – the smaller flow always grows, and the larger shrinks if necessary, so it does.

i. On the real internet, the smaller flow can have its packets dropped too. Suppose that the flow to experience the drop was chosen at random with 50/50 probability. Explain the effect of AIAD on fairness in this new model. (2 points.)

Answer:

The system would never converge to a fair allocation between flows, since each flow would now evolve in the same way: growth until the network is overloaded, followed by a 50% chance of a flow reduction.

III Elasticity Buffer

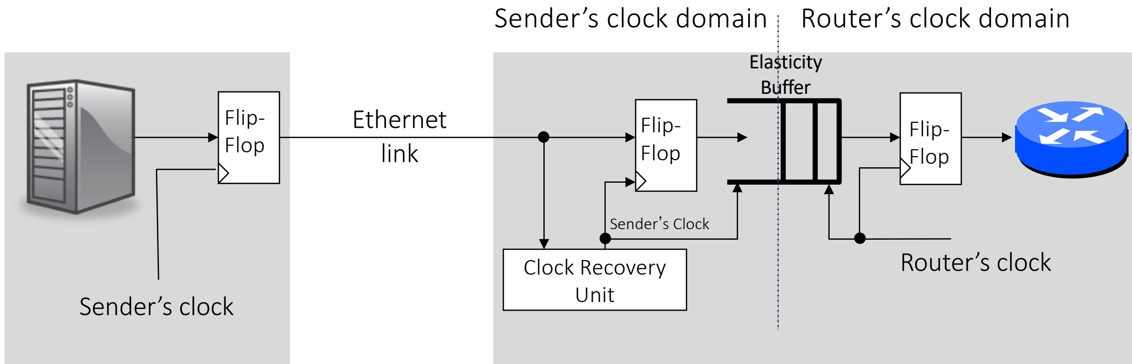

8. [14 points]:

An end host is connected to a router over an Ethernet link. The router maintains an elasticity buffer to move the arriving bits from the end host's clock domain to the router's clock domain. The elasticity buffer is sized so that it doesn't overflow or underflow. As before, the clock rate may vary between R_{max} and R_{min}; the maximum packet size is P_{max}.

Unlike the elasticity buffer we saw in class, when a new packet starts arriving the router waits until the elasticity buffer is only 1/3 full (i.e. it reaches a threshold $B/3$) before it starts reading from the elasticity buffer. Your job is to determine how big the buffer should be (B) in order to prevent the buffer overflowing or underflowing.

(You may find it useful to use the simplifying assumption that $R_{max}/R_{min} \approx 1$.)

a. On the two sets of axes below, sketch the cumulative arrivals and departures from the elasticity buffer on the router for the two cases: (1) When the end host is at the highest clock rate and the router is at the slowest, and (2) When the router is at the highest clock rate and the end host is at the slowest. (2 points.)

To prevent overflow, we require that (a) – (b) $\leq B$

$$P_{max} - R_{min}\left(\frac{P_{max}}{R_{max}} - \frac{B}{3R_{max}}\right) \leq B$$

$$P_{max}\left(\frac{R_{max} - R_{min}}{R_{max}}\right) + \left(\frac{BR_{min}}{3R_{max}}\right) \leq B$$

$$\boxed{\therefore P_{max}\left(\frac{R_{max} - R_{min}}{R_{max}}\right) \leq \frac{2B}{3}}$$

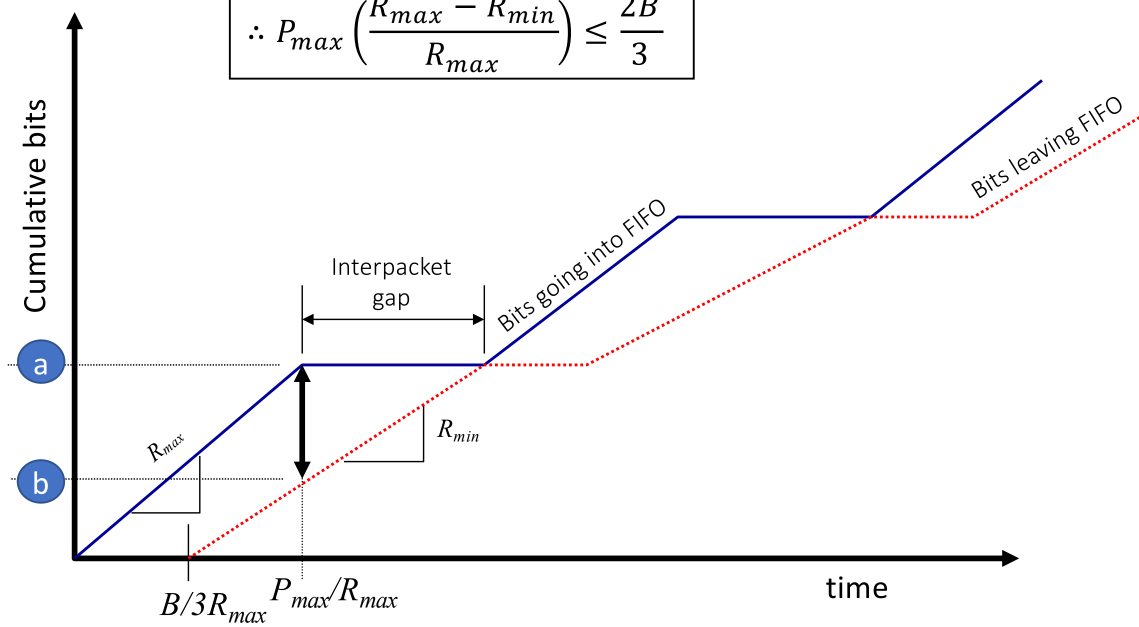

To prevent underflow, we require that

$$\left(\frac{P_{max}}{R_{max}} + \frac{B}{3R_{max}}\right) \geq \frac{P_{max}}{R_{min}}$$

$$P_{max}\left(1 - \frac{R_{max}}{R_{min}}\right) \geq \frac{B}{3}$$

$$\boxed{\therefore P_{max}\left(\frac{R_{max} - R_{min}}{R_{min}}\right) \leq \frac{B}{3}}$$

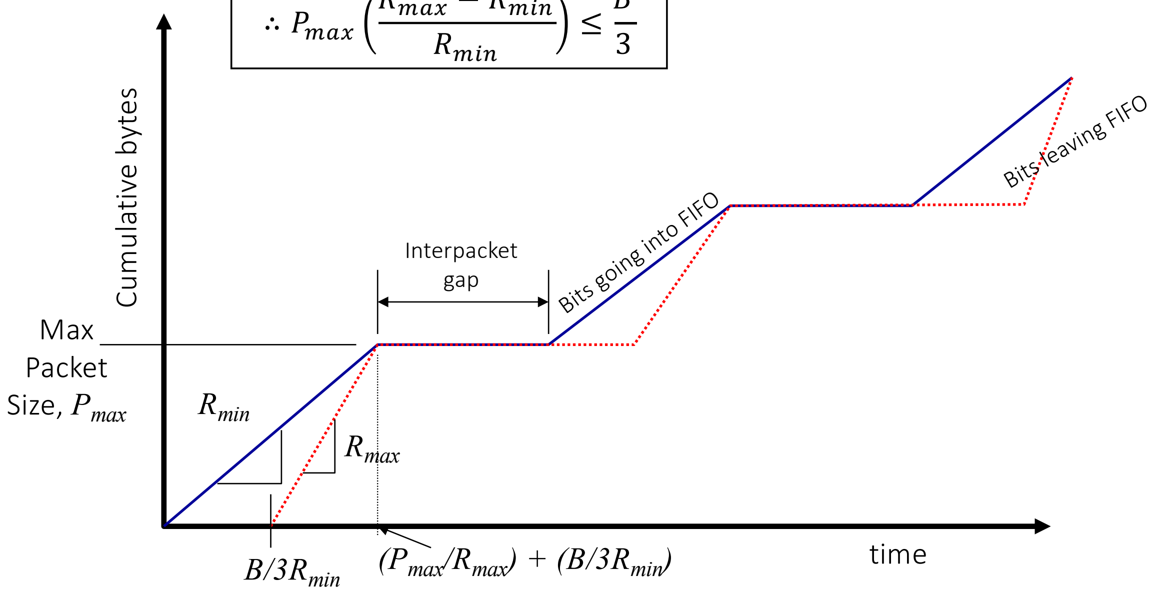

b. Write down an expression for the conditions required so that the buffer does not overflow. Your answer should be an inequality for the buffer size, B, in terms of R_{max}, R_{min}, and P_{max}. (2 points.)

Answer:

See figure above.

c. Simplify your expression above, to prove that $(\frac{R_{max}-R_{min}}{R_{max}}) \times P_{max} \leq \frac{2B}{3}$. (2 points.)

Answer:

See figure above.

d. Write down an expression for the conditions required so that the buffer does not underflow. Your answer should be an inequality for the buffer size, B, in terms of R_{max}, R_{min}, and P_{max}. (2 points.)

Answer:

See figure above.

e. Simplify your expression above, to prove that $(\frac{R_{max}-R_{min}}{R_{min}}) \times P_{max} \leq \frac{B}{3}$. (2 points.)

Answer:

See figure above.

f. Write down an expression for the minimum size of the gap we need between packets (i.e. the *interpacket gap*) to prevent the elasticity buffer from overflowing. Explain briefly why we need an interpacket gap. (2 points.)

Answer:

Interpacket gap must be large enough for buffer to drain before the next packet starts. In the worst case, this is the time for the whole EB to drain from full; i.e. B/R_{min}.

g. The end host and the router each have an internal clock that runs at nominally 1GHz, but can vary from 999.9 MHz to 1000.2 MHz. i.e. in the worst case, the clock might be 0.1 MHz slower (100 parts per million) or 0.2 MHz faster (200 parts per million). Note that unlike the examples we used in class the clock range is asymmetric around the center frequency of 1GHz. Explain why it might make sense to use the expressions above to design an elasticity buffer for this system. (2 points.)

Answer:

In this asymmetric case, the clock is more likely to be running fast than slow, and so it likely makes sense to use more of the buffer (in fact, twice as much) for the overflow case than the underflow case. Hence, the equations we derived above make sense in this case.

IV Are you still watching?

9. [9 points]:

Consider an *adaptive bitrate* (ABR) video streaming system, where the video is delivered in 2-second chunks (the playback cannot start until the whole chunk is downloaded). At $T = 0$, a client that has a link speed of 4 Mb/s starts streaming a movie with an encoded bitrate of 2 Mb/s.

a. How long does it take the video to start playing (startup time)? (3 points.)

Answer: ___*1*___ s

Answer:

$$\begin{aligned} \textit{chunk-size} &= 2\,\text{Mb/s} \times 2\,\text{s} \\ &= 4\,\text{Mb} \\ T_{\textit{start}} &= \frac{4\,\text{Mb}}{4\,\text{Mb/s}} = 1\,\text{s} \end{aligned}$$

b. Let's say that at $T = 10\,\text{s}$, the link speed suddenly drops to 0. At what time T will the video stop playing? (Hint: See (a) for when video starts playing) (3 points.)

Answer: ___*21*___ s

Answer:

$$\begin{aligned} \textit{downloaded} &= \frac{10\,\text{s} \times 4\,\text{Mb/s}}{\textit{chunk-size}} = 10\ \textit{chunks} \\ T_{\textit{stop}} &= T_{\textit{start}} + \textit{downloaded} \times 2\,\text{s} \\ &= 21\,\text{s} \end{aligned}$$

c. New scenario. The link runs at its original speed of 4 Mb/s between $T = 0$ s until $T = 15$ s. Then, at $T = 15$ s, the link speed drops to 400 kb/s, but the client continues streaming video that was encoded at a bitrate of 2 Mb/s. The video will eventually stall. At what time will it resume playing (even if only for a short time)? (3 points.)

Answer: ___*40*___ s

Answer:

Stall when downloaded = played.

$$
\begin{aligned}
\textit{downloaded} &= \frac{(15\,\mathrm{s} \times 4\,\mathrm{Mb/s}) + (x \times 400\,\mathrm{Kb/s})}{\textit{chunk-size}} \\
\textit{played} &= (14 + x)\textit{s} \times 2\,\mathrm{Mb/s} \\
x &= 20\,\mathrm{s} \\
T_{\textit{stall}} &= 35\,\mathrm{s} \\
\Delta T_{\textit{stalled}} &= \frac{\textit{chunk-size}}{400\,\mathrm{Kb/s}} \\
&= 5\,\mathrm{s} \\
T_{\textit{resume}} &= T_{\textit{stall}} + \Delta T_{\textit{stalled}} = 35 + 5 = 40\,\mathrm{s}
\end{aligned}
$$

V Web page on a cold start

10. **[14 points]:**

After talking to your friend, you decide to look up a fact they told you on Wikipedia. You pull out your laptop, power it on, and connect to Wikipedia. Of the following messages sent as a result, which are sent by your laptop, which are sent by the first hop router, and which are sent by the Wikipedia server? We define ”sent” as originating a message, not forwarding it. Assume all caches are empty initally, and the ARP cache is populated thereafter as in the Network Interface lab (i.e. cache mappings from both requests and replies for 30 seconds – assume the entire sequence of events takes less than 30 seconds). Assume the first hop router is also the DHCP Server on the network, and that the DNS server is neither the router nor the Wikipedia server. Assume the Wikipedia server is on the same network as the router, and both are already connected to the Internet. Each row and column may have any number of boxes selected.

Message	Laptop	First Hop Router	Wikipedia Server
DHCP Request	■	□	□
DHCP Response	□	■	□
ARP Request	■	■	□
ARP Response	□	■	■
DNS Request	■	□	□
DNS Response	□	□	□
TCP Segment(s)	■	□	■

VI Ethernet and CSMA/CD

11. [12 points]:

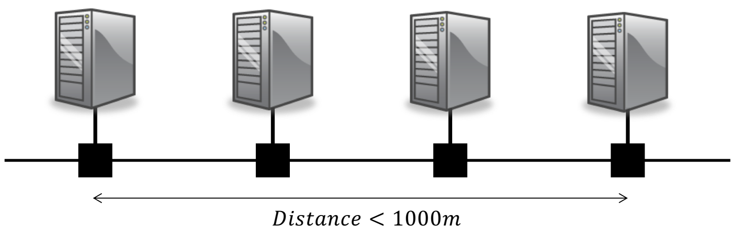

a. A 1000m long 1Gb/s Ethernet network uses CSMA/CD to control access to a shared copper cable as shown in the figure. The Ethernet specification requires that if a collision occurs, it must detect the collision before it finishes transmitting a packet. What is the size of the minimum packet in this network? Express your answer in bits or bytes. (Assume the speed of propagation is 2×10^8m/s.) (3 points.)

Answer: ___10^4___ bits **or** ___*1250*___ bytes

Answer:

The RTT on the cable is $2000m/2 \times 10^8 = 10^{-5}s$. *Therefore* $P/10^9 \geq 10^{-5}$ *which means the packet size must be greater than* 10^4 *bits or 1250bytes.*

b. If the Ethernet network is upgraded to 10Gb/s, how does this affect the minimum size packet we can use? Explain why this might be a problem. (3 points.)

Answer:

If we multiply R by 10 then we also need to make the minimum packet size 10 times bigger, or 10^5 bits, which is very big. If the packets we want to send are smaller than 10^5 bits, then we need to pad them (e.g. with zeros), which might be very inefficient.

c. If we replace the long shared copper cable with a "broadcast star" in which no two hosts are now more than 500m apart, what is the new minimum packet size for the 1Gb/s Ethernet network? (3 points.)

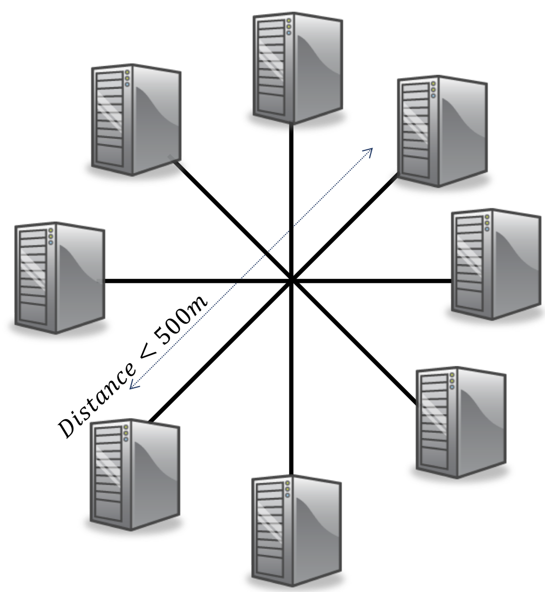

Answer: 5^4 bits **or** *625* bytes

Answer:

The maximum RTT has now halved, so the minimum size packet halves too.

d. Today, CSMA/CD is not commonly used in wired Ethernet networks. Intead, end hosts are typically connected to their nearest Ethernet switch using a full duplex cable (which means there is a separate, independent channel for transmit and receive). Explain why this kind of network does not require CSMA/CD. (3 points.)

Answer:

Because there are only two senders on each link (a switch and end host), and each has its own dedicated channel, there can never be any collisions on the link. There is only one sender. If two packets arrive at the switch destined to the same link, one packet is sent while the other is buffered in a FIFO queue.

VII Manchester Encoding

12. [5 points]:

Encode the bitstream "011001" using Manchester encoding. The first bit has been encoded for you.

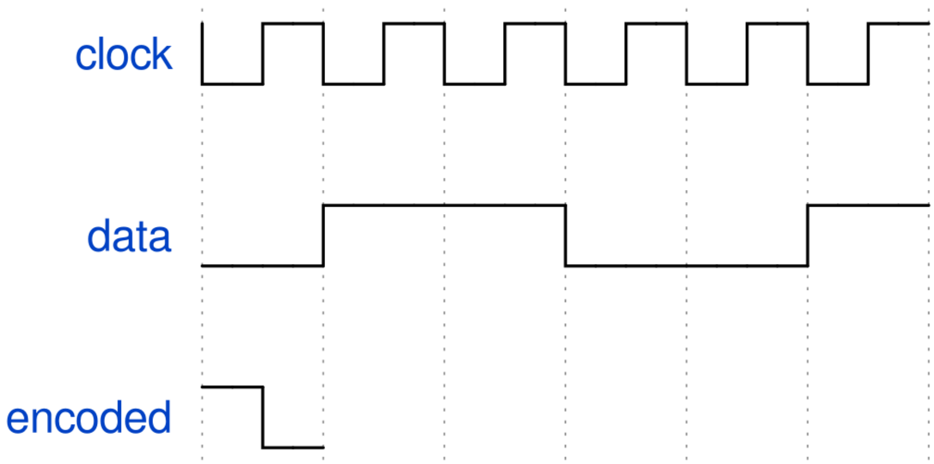

Answer: