Problem Set 5 Question 1

a) Consider sending a large file using TCP over a connection with no loss. Assuming we start with a CWND of 6 MSS, how many RTTs would it take to increase to 12 MSS using AIMD?

- After 1 RTT, CWND = 7
- After 2 RTTs, CWND = 8
- After 3 RTTs, CWND = 9
- After 4 RTTs, CWND = 10
- After 5 RTTs, CWND = 11
- After 6 RTTs, CWND = 12

- Increasing from CWND = 6 to CWND = 12 takes **6 RTTs**

Problem Set 5 Question 1

b) Consider sending a large file using TCP over a connection with no loss. Assuming we start with a CWND of 6 MSS, what is the average throughput over the first 6 RTTs?

- Total segments sent over first 6 RTTs:

  6+7+8+9+10+11 = **51**

- 51 MSS / 6 RTT = **8.5 MSS/RTT**

Problem Set 5 Question 2

a) A sender and receiver communicate using TCP. If the MSS is 1500 bytes, draw a diagram showing how long it takes to transfer 15000 bytes.

Use slow start! (after SYN handshake)

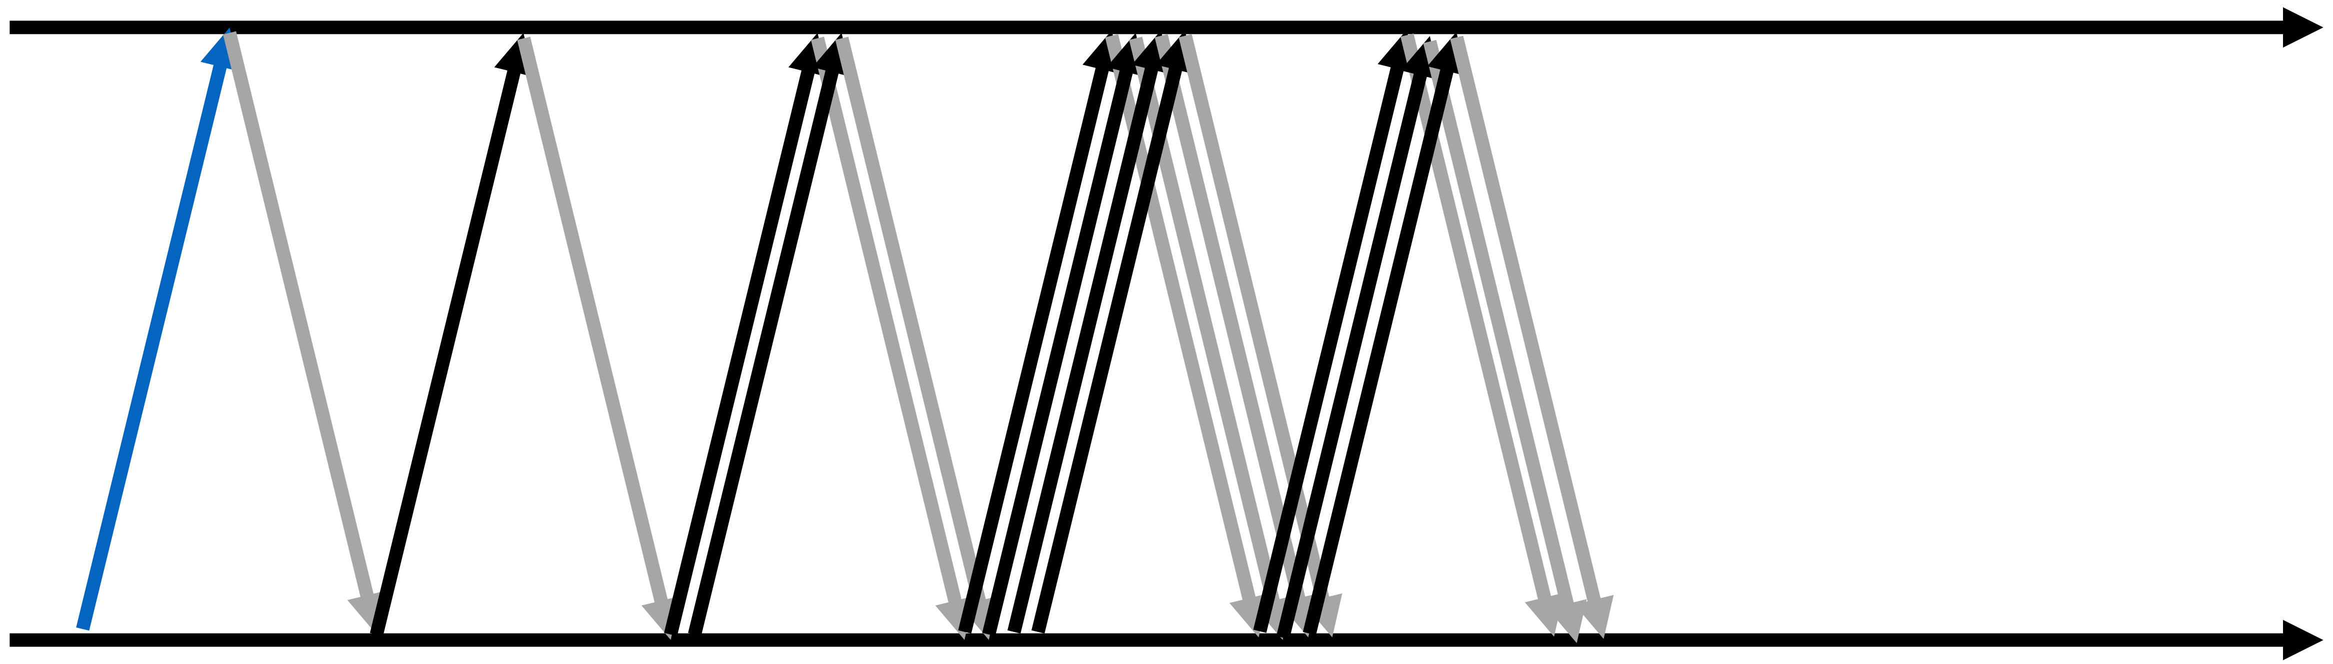

Total time to send: 5 RTTs + 3 transmission delays

Problem Set 5 Question 2

b) Consider a single TCP connection on a 10 Mbps link with no buffering. MSS = 1500 bytes, RTT = 150 ms, ignore slow start.

- **What is the maximum window size (in segments) achieved?**

$$W = \frac{RTT}{TD} = \frac{RTT}{MSS/Bandwidth}$$

$$W = \frac{Bandwidth * RTT}{MSS}$$

$$W = \frac{10\ Mbps\ \ * 150\ ms}{1500\ bytes} = \frac{1572864}{12000} \approx \mathbf{131}\ \boldsymbol{segments}$$

Problem Set 5 Question 2

b) Consider a single TCP connection on a 10 Mbps link with no buffering. MSS = 1500 bytes, RTT = 150 ms, ignore slow start.

- **What is the average window size (in segments) and throughput (in bits per second)?**

The window size fluctuates between $\frac{W}{2}$ and W, for an average window size of $\frac{3W}{4} = \mathbf{98.25\ segments}$

$$98.25 * \frac{MSS}{RTT} = 98.25 * \frac{1500 * 8}{0.150} = 7.49\ Mbps$$

Problem Set 5 Question 2

b) Consider a single TCP connection on a 10 Mbps link with no buffering. MSS = 1500 bytes, RTT = 150 ms, ignore slow start.

- **How long would it take this TCP connection to reach its maximum window after recovering from a loss?**

$$\frac{W}{2} * RTT$$

$$\frac{131}{2} * 0.15s = \mathbf{9.825s}$$

Problem Set 5 Question 3

a) Consider two TCP flows over the following two links. Neither flow ever saturates its link. Both have the same MSS and packet loss probability. What is the ratio of their link utilizations?

- **Link X: 100 Mbps RTT = 10ms**
- **Link Y: 1 Gbps RTT = 1000ms**

$$Throughput = \frac{MSS}{RTT}\sqrt{\frac{3}{2p}} \qquad Utilization = \frac{MSS}{BW*RTT}\sqrt{\frac{3}{2p}}$$

$$\frac{\frac{\cancel{MSS}}{BW_X * RTT_X}\cancel{\sqrt{\frac{3}{2p}}}}{\frac{\cancel{MSS}}{BW_Y * RTT_Y}\cancel{\sqrt{\frac{3}{2p}}}} = \frac{BW_Y * RTT_Y}{BW_X * RTT_X} = \frac{1024 * 1}{100 * .001} = 1024$$

Problem Set 5 Question 3

b) Consider two TCP flows over two wireless links, X and Y. Suppose X has a 1% packet loss rate and Y has a 4% packet loss rate, but the links are otherwise identical. What is the link throughput ratio of X to Y?

$$\frac{\cancel{\frac{MSS}{RTT}}\sqrt{\frac{3}{2p_X}}}{\cancel{\frac{MSS}{RTT}}\sqrt{\frac{3}{2p_Y}}} = \frac{\sqrt{\frac{\cancel{3}}{\cancel{2}p_X}}}{\sqrt{\frac{\cancel{3}}{\cancel{2}p_Y}}} = \sqrt{\frac{p_Y}{p_X}} = \sqrt{\frac{4\%}{1\%}} = \sqrt{4} = \mathbf{2}$$

Problem Set 5 Question 3

c) From the previous two parts, we know TCP does not perform well over links with higher RTT or more packet loss. How could we improve this?

- For links with a high RTT: increase CWND more aggressively.

- For lossy links: Loss is no longer a good indicator of congestion. We must use some other mechanism to detect congestion (but still retransmit lost packets).

Problem Set 5 Question 3

d) Suppose a TCP and a UDP flow are sharing a link. The UDP flow attempts to transmit at a rate equal to the link capacity. How quickly would each flow transmit at equilibrium?

- UDP flows do not necessarily follow TCP congestion control.
- Whenever there is a loss, TCP decreases its congestion window, but UDP continues at full speed.
- At equilibrium: TCP reduces its congestion window to almost zero, so UDP uses almost all of the bandwidth.

Problem Set 5 Question 3

d) Suppose a TCP and a UDP flow are sharing a link. The UDP flow attempts to transmit at a rate equal to the link capacity. How quickly would each flow transmit at equilibrium?

- UDP flows do not necessarily follow TCP congestion control.
- Whenever there is a loss, TCP decreases its congestion window, but UDP continues at full speed.
- At equilibrium: TCP reduces its congestion window to almost zero, so UDP uses almost all of the bandwidth.

Problem Set 5 Question 4

Consider the following plot, where the y-axis denotes time, and the y-axis denotes the TCP window size of a TCP flow.

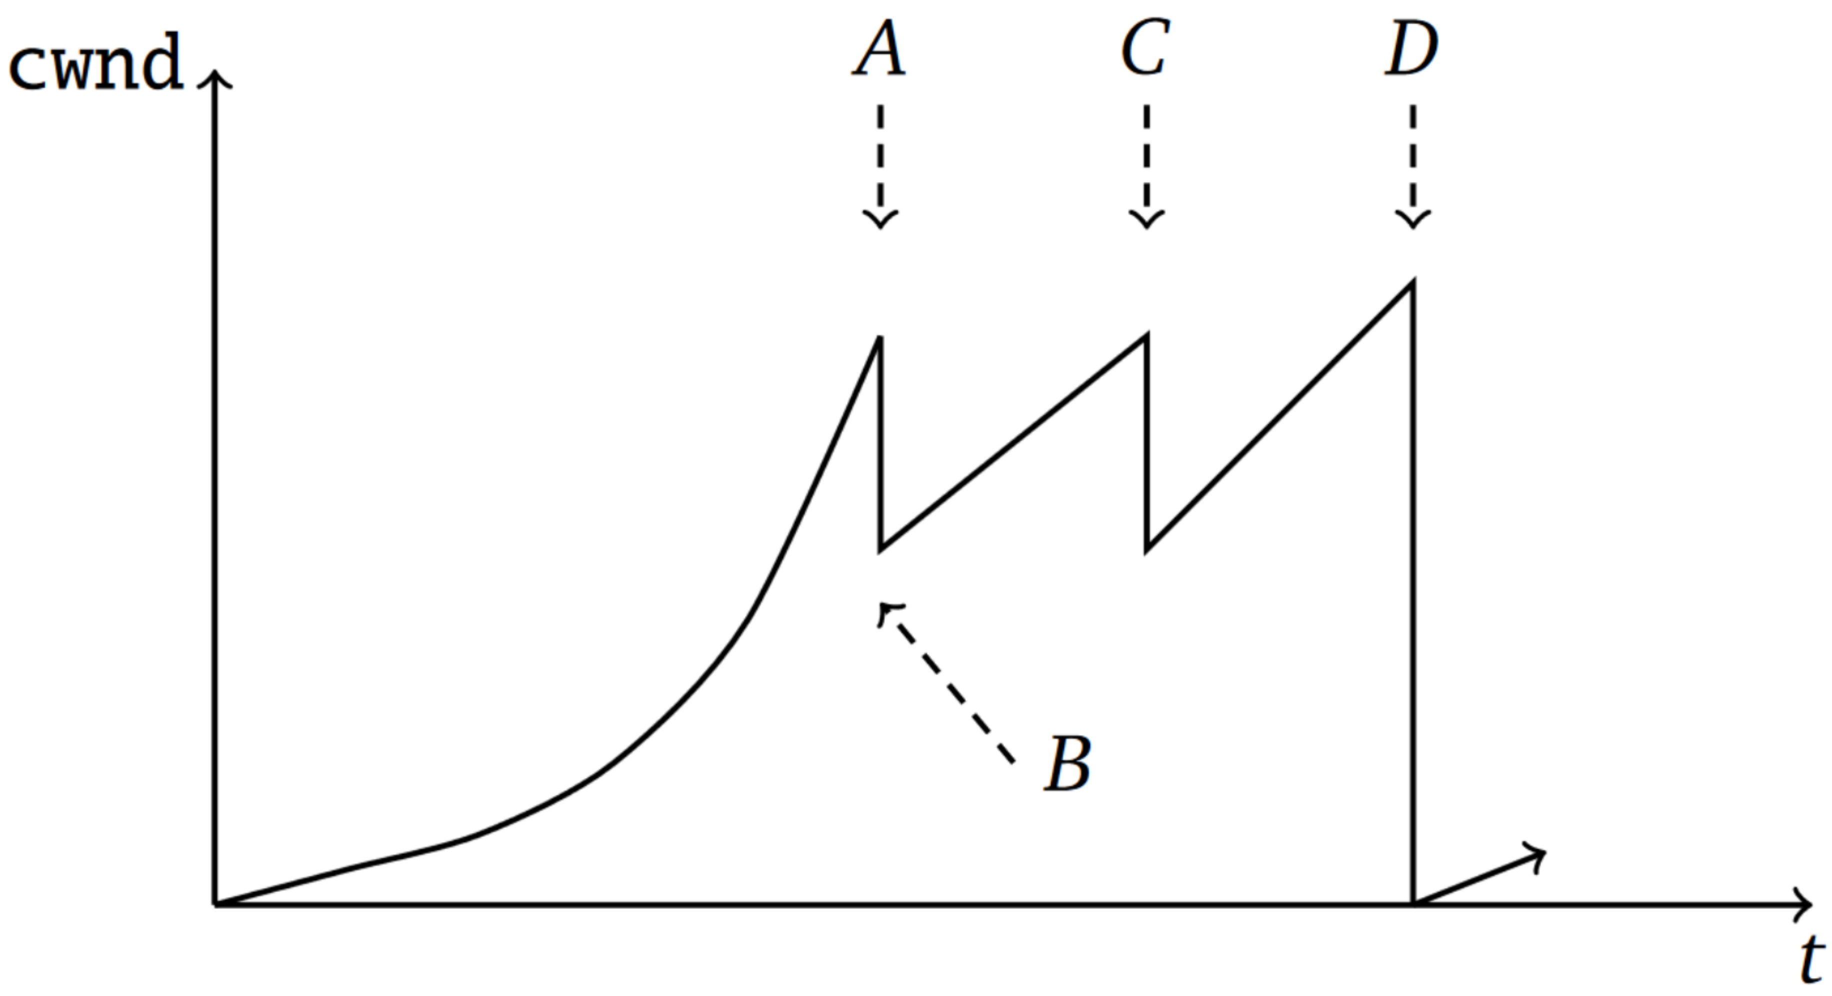

Problem Set 5 Question 4

a) Describe the slow start mechanism (from t=0 to event A) and the AIMD mechanism (between A and C)

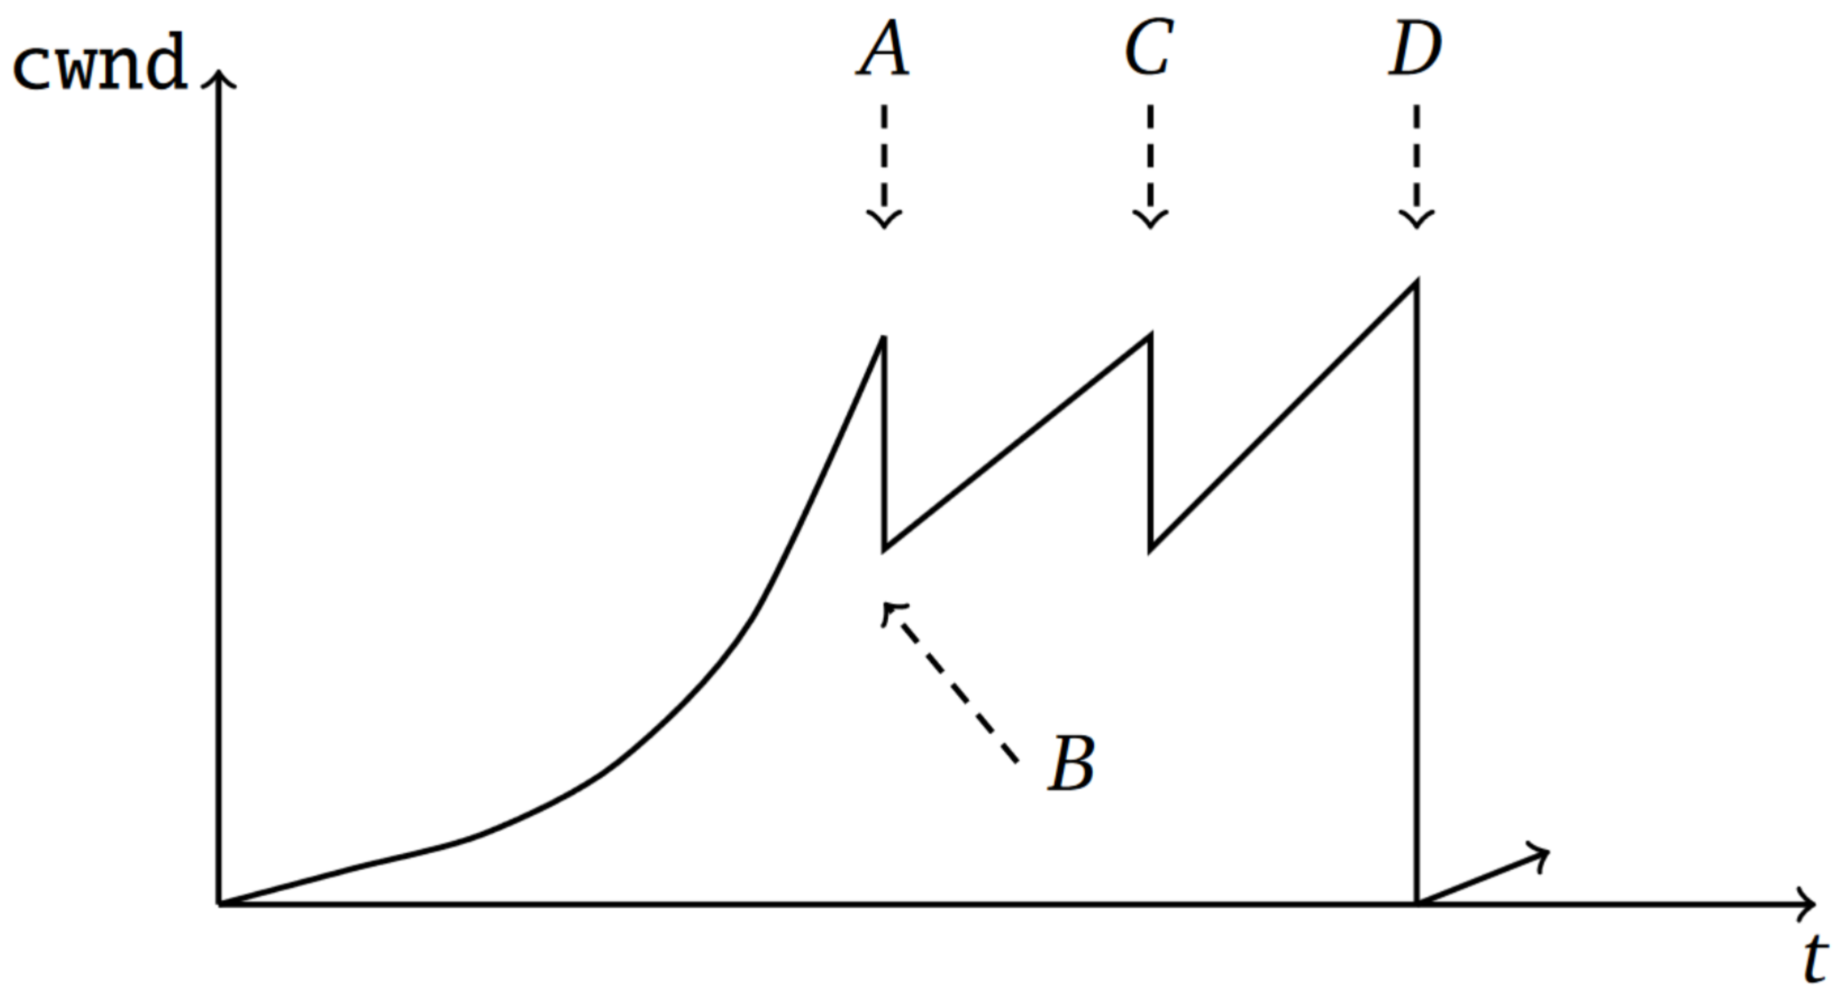

Problem Set 5 Question 4

a) Describe the slow start mechanism (from t=0 to event A) and the AIMD mechanism (between A and C)

Slow Start: The sender increases the congestion window by 1 MSS every time a segment is ACKed, until a dropped packet is detected.

AIMD: At event A, a packet drop is detected and the congestion window is decreased to half (event B). Afterwards, the sender increases the congestion window by 1 MSS every RTT until the next packet drop (event C).

Problem Set 5 Question 4

b) Name the event that triggers the window size reduction at event A.

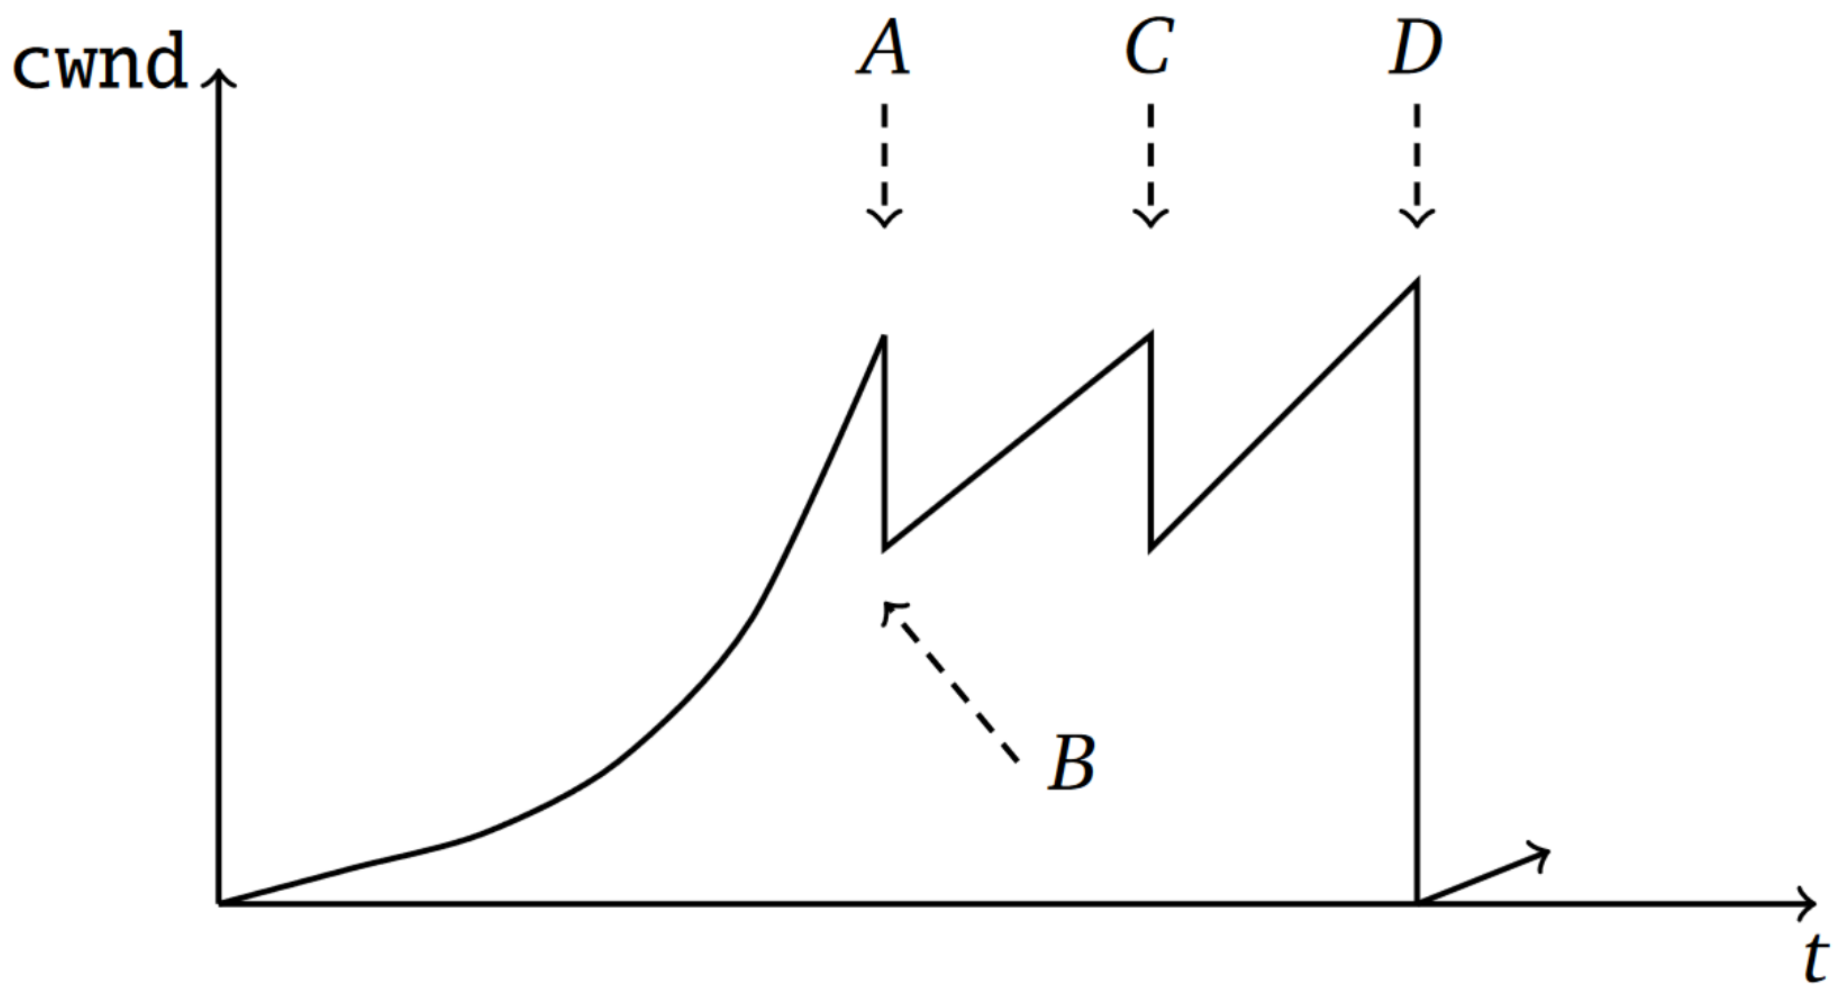

Answer: 3 duplicate ACKs

Problem Set 5 Question 4

c) Do three duplicate ACKs mean that a packet was definitely dropped?

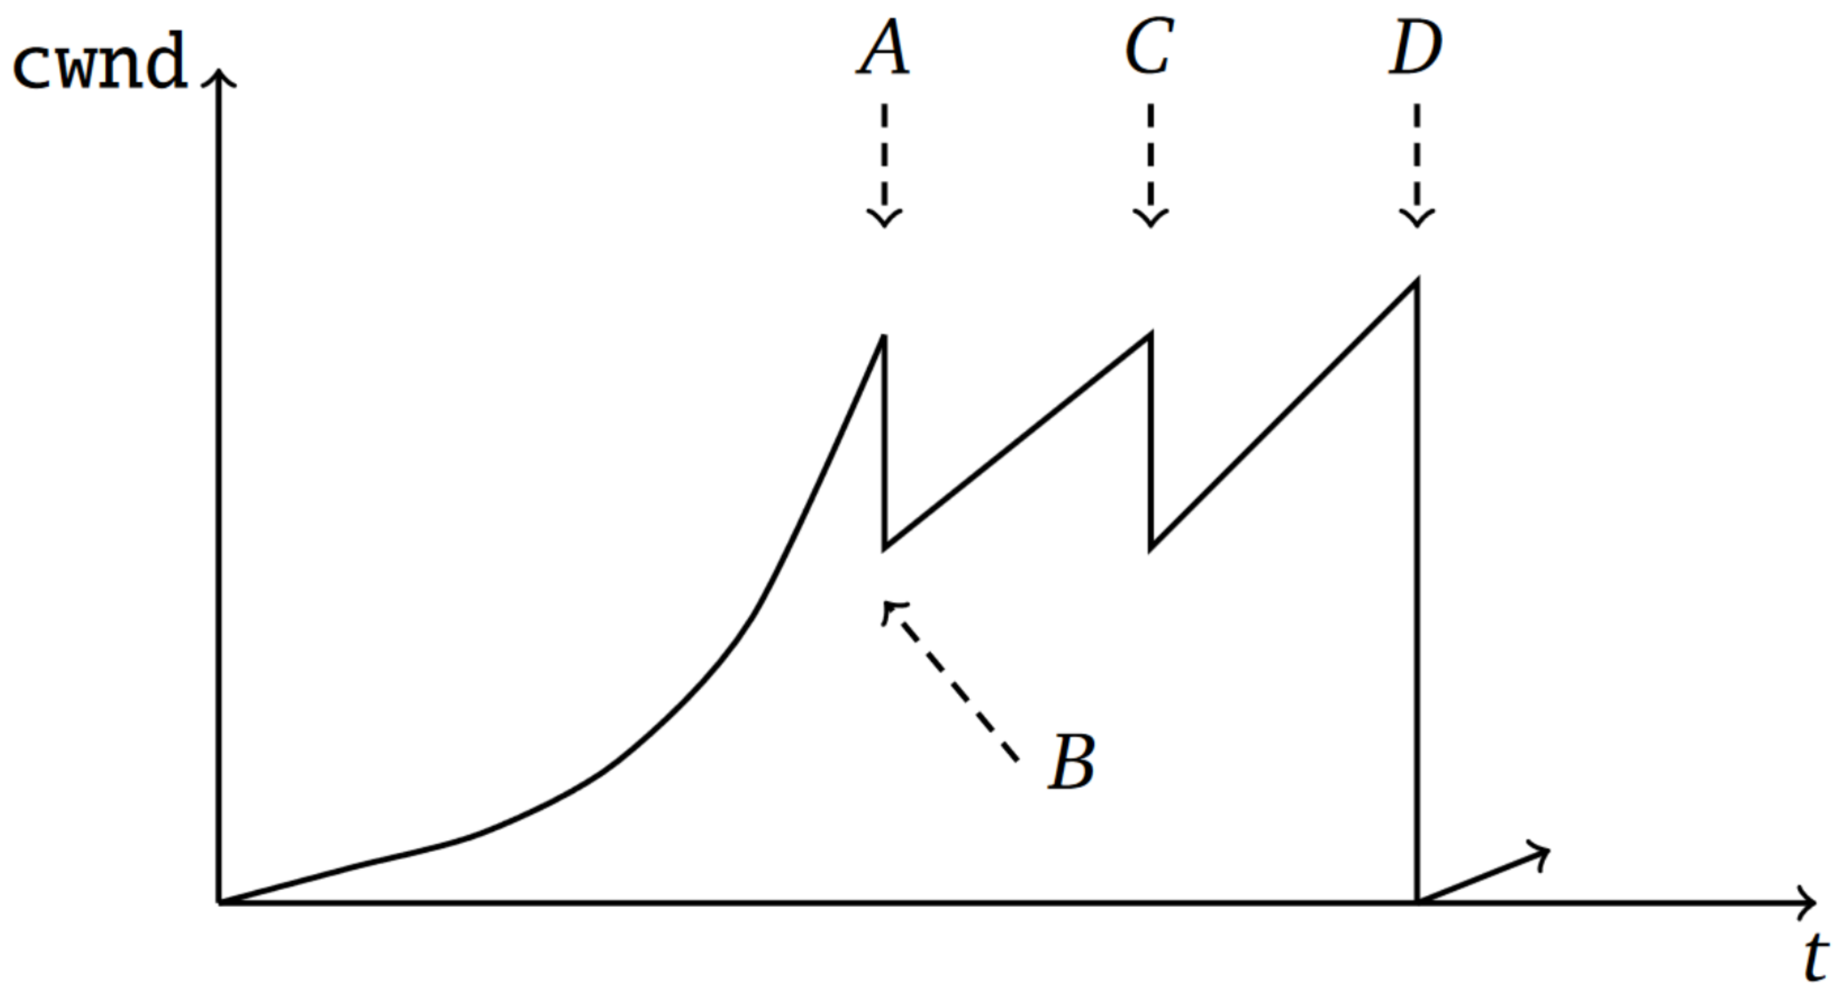

Answer: No. It might mean that packets were reordered.

Problem Set 5 Question 4

d) Name the event that triggers the window size reduction at event D.

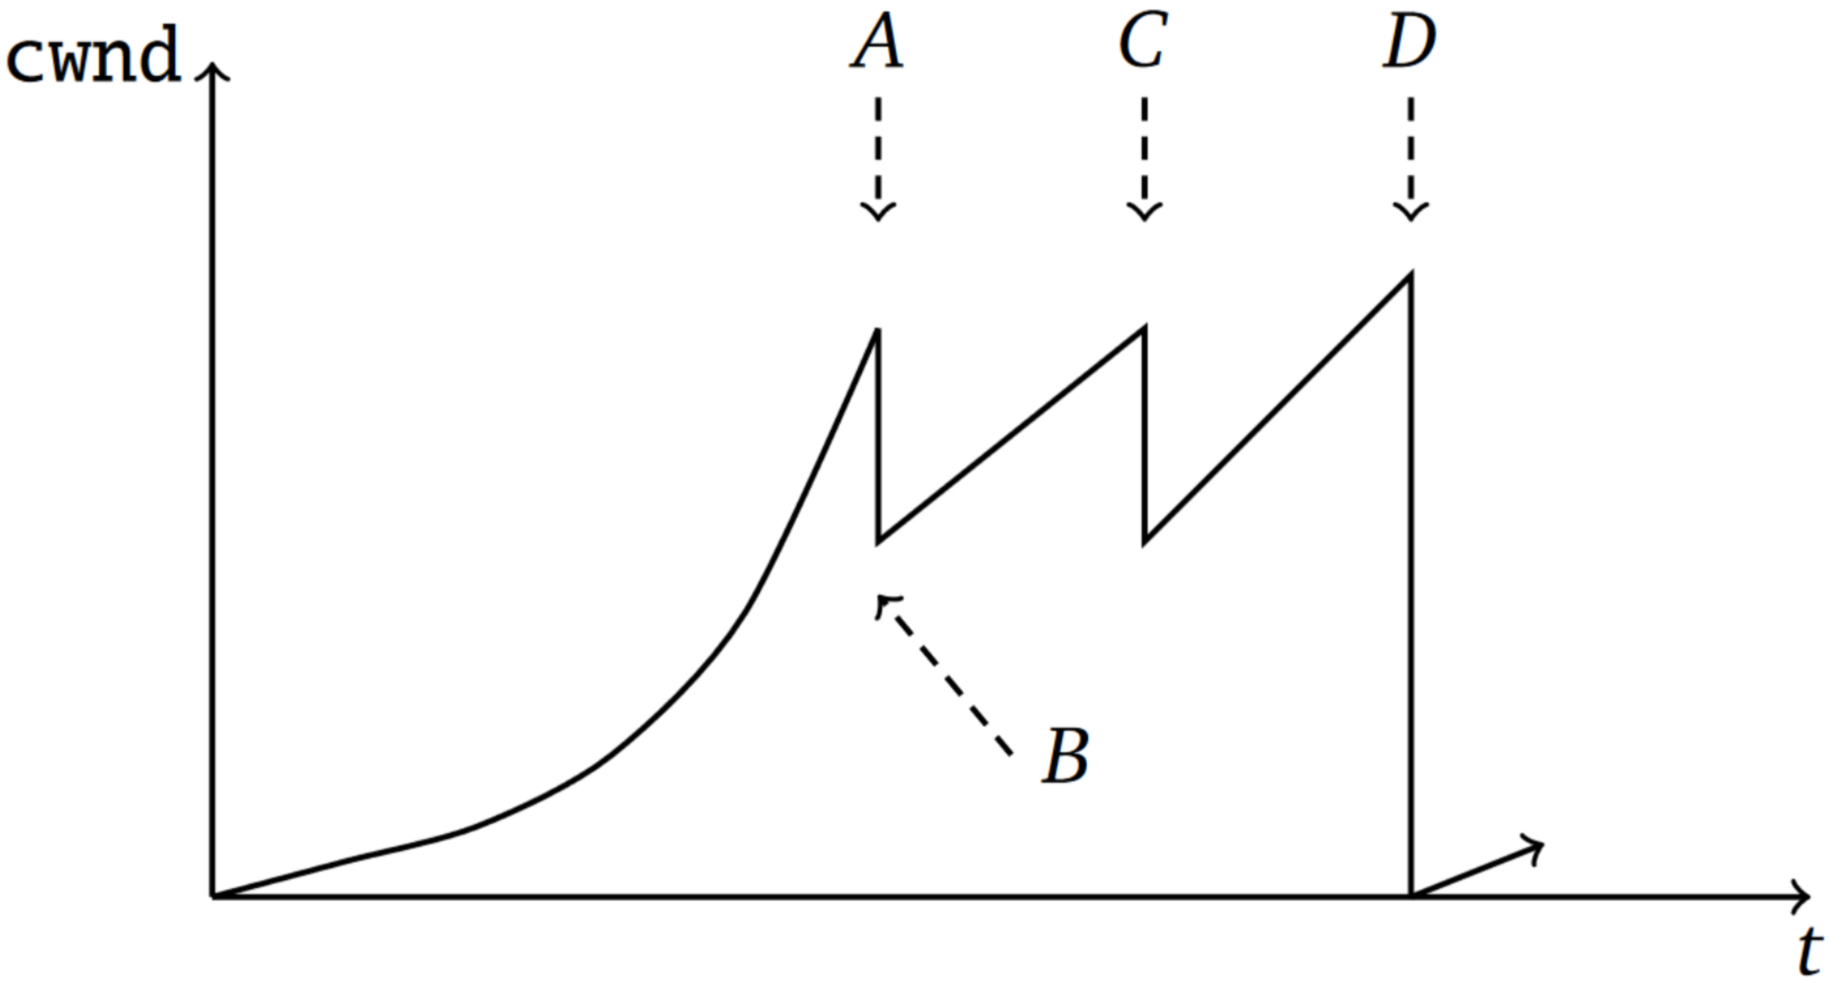

Answer: Timeout

Problem Set 5 Question 4

e) Does a timeout mean that a packet was definitely dropped

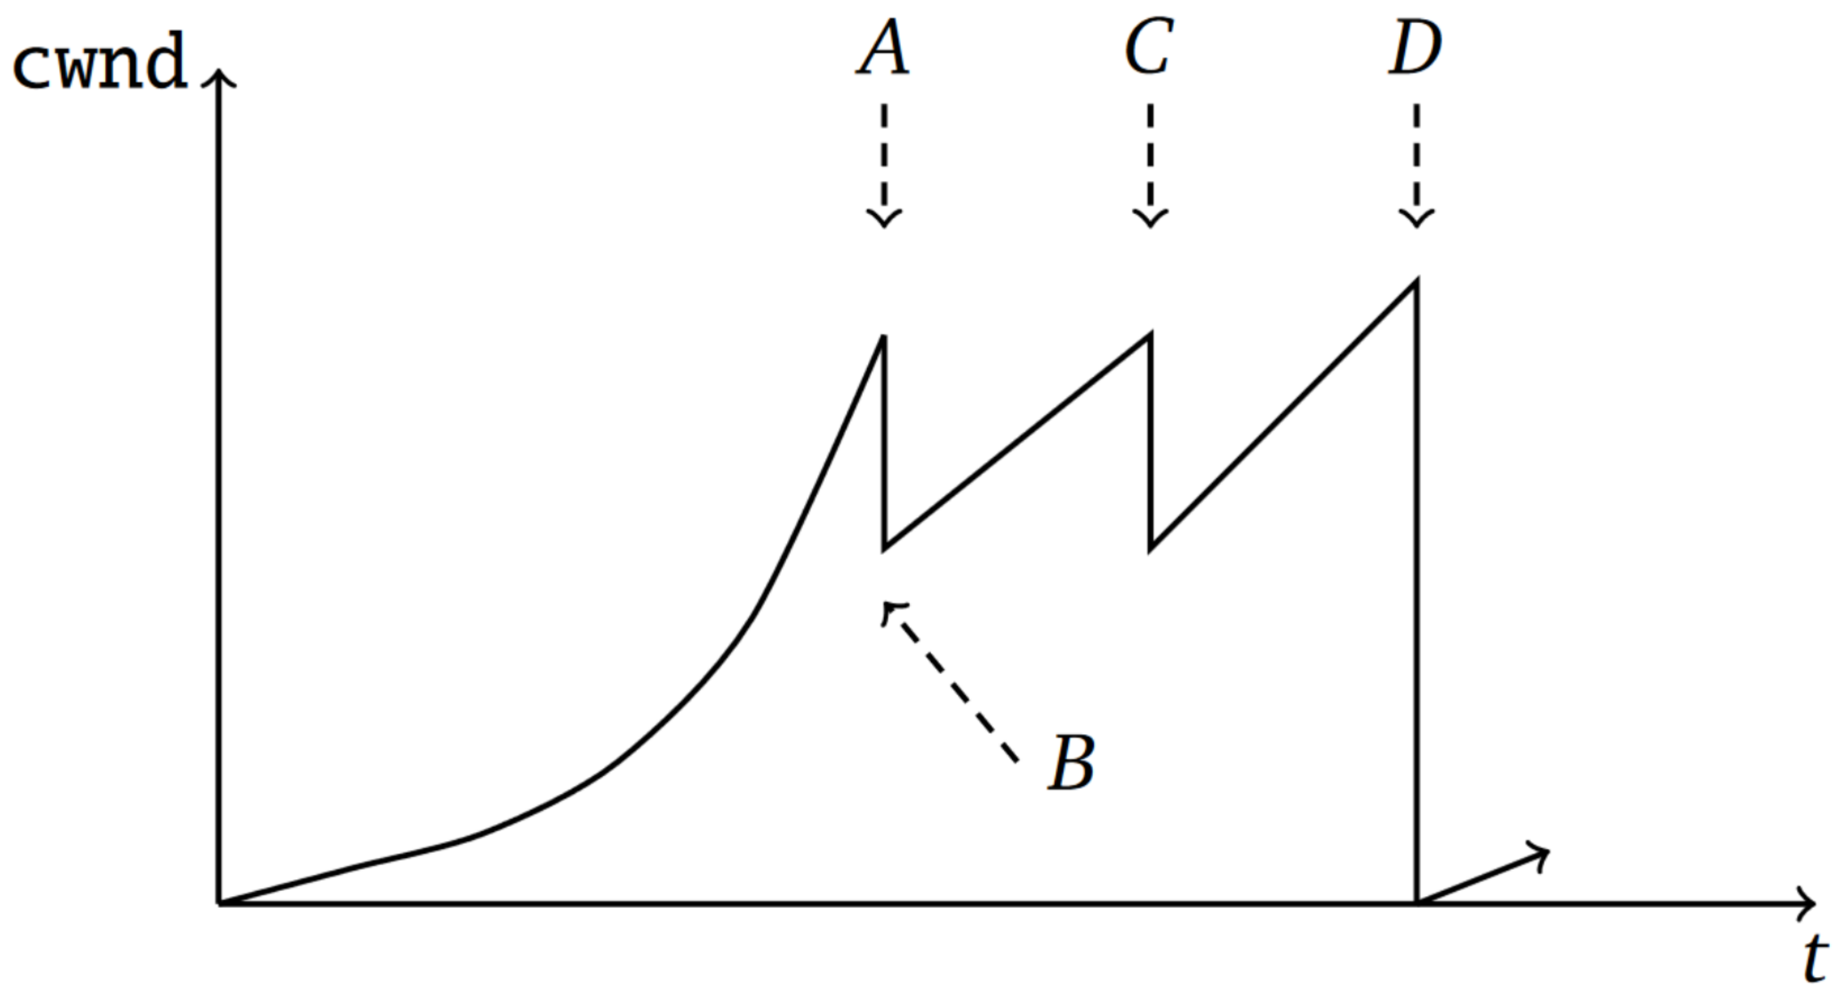

Answer: No. It might mean that the delay on the link increased suddenly, and the packets are still in flight.

Problem Set 5 Question 4

f) Suppose the window size before the drop at event D is 16. Could the window size ever increase beyond 16 after D?

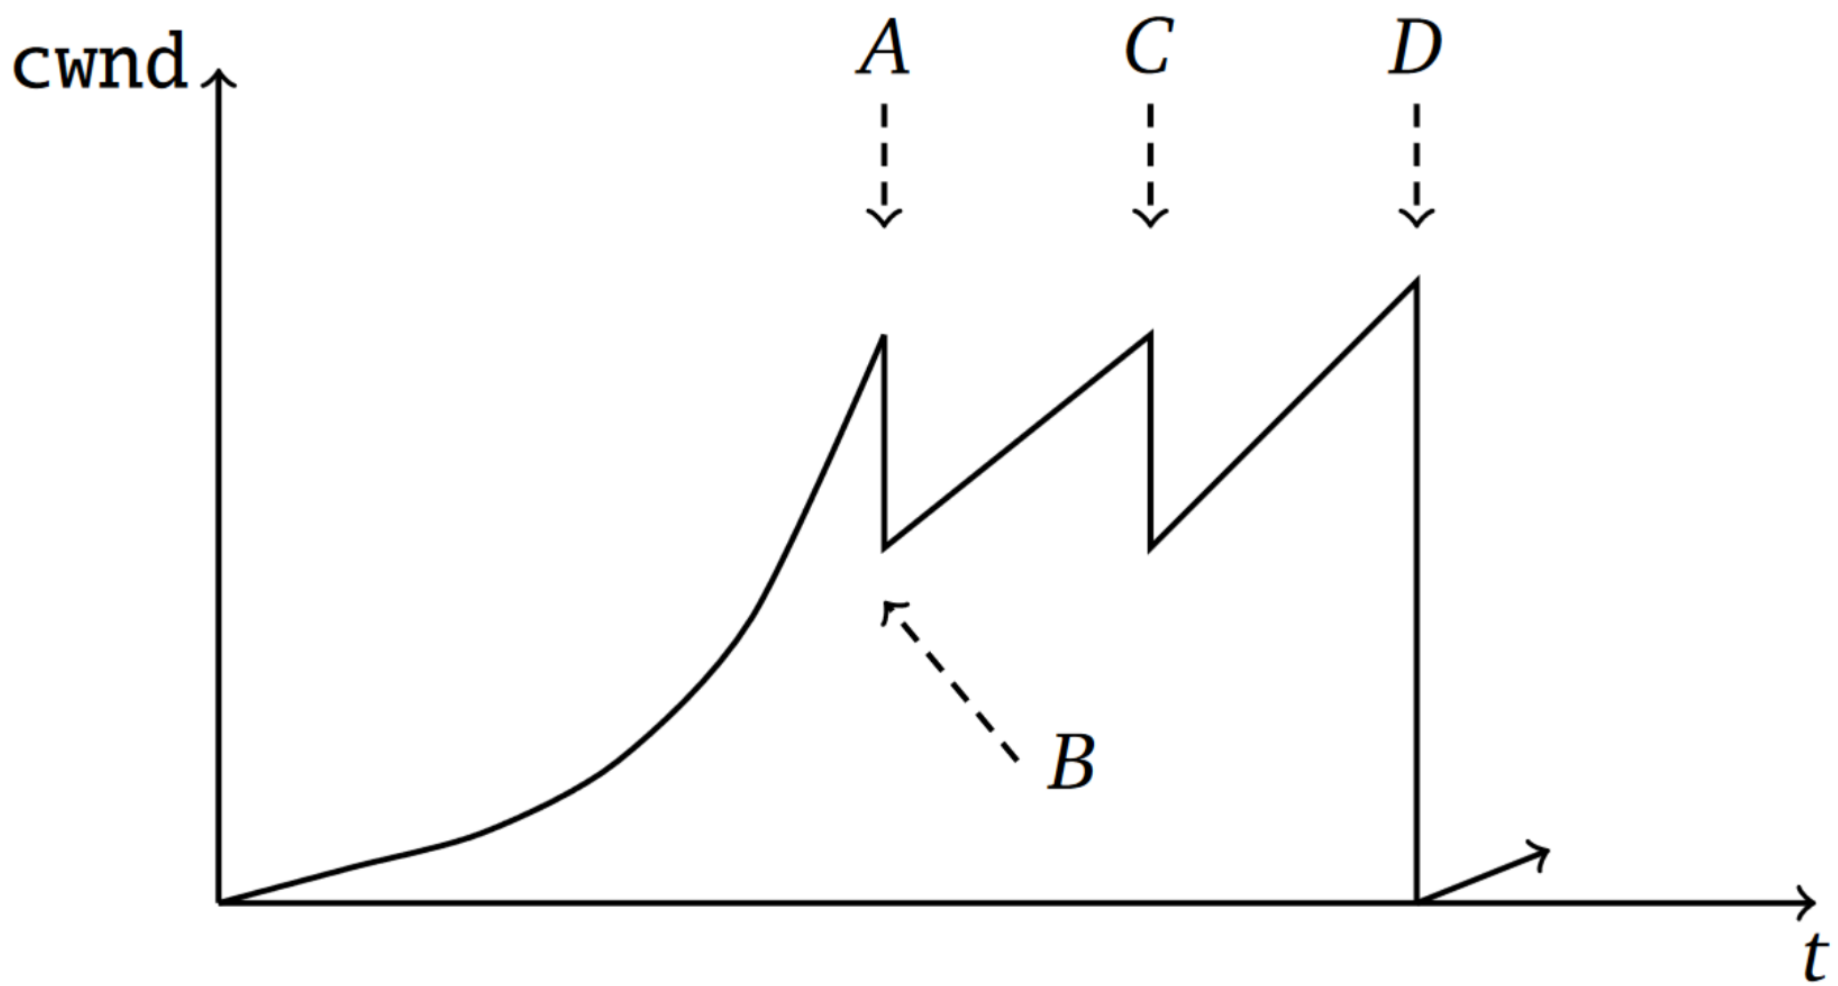

Answer: Yes. As long as there is no congestion, the window size will increase indefinitely.

Problem Set 5 Question 5

Describe what is happening in this TCP trace at various times:

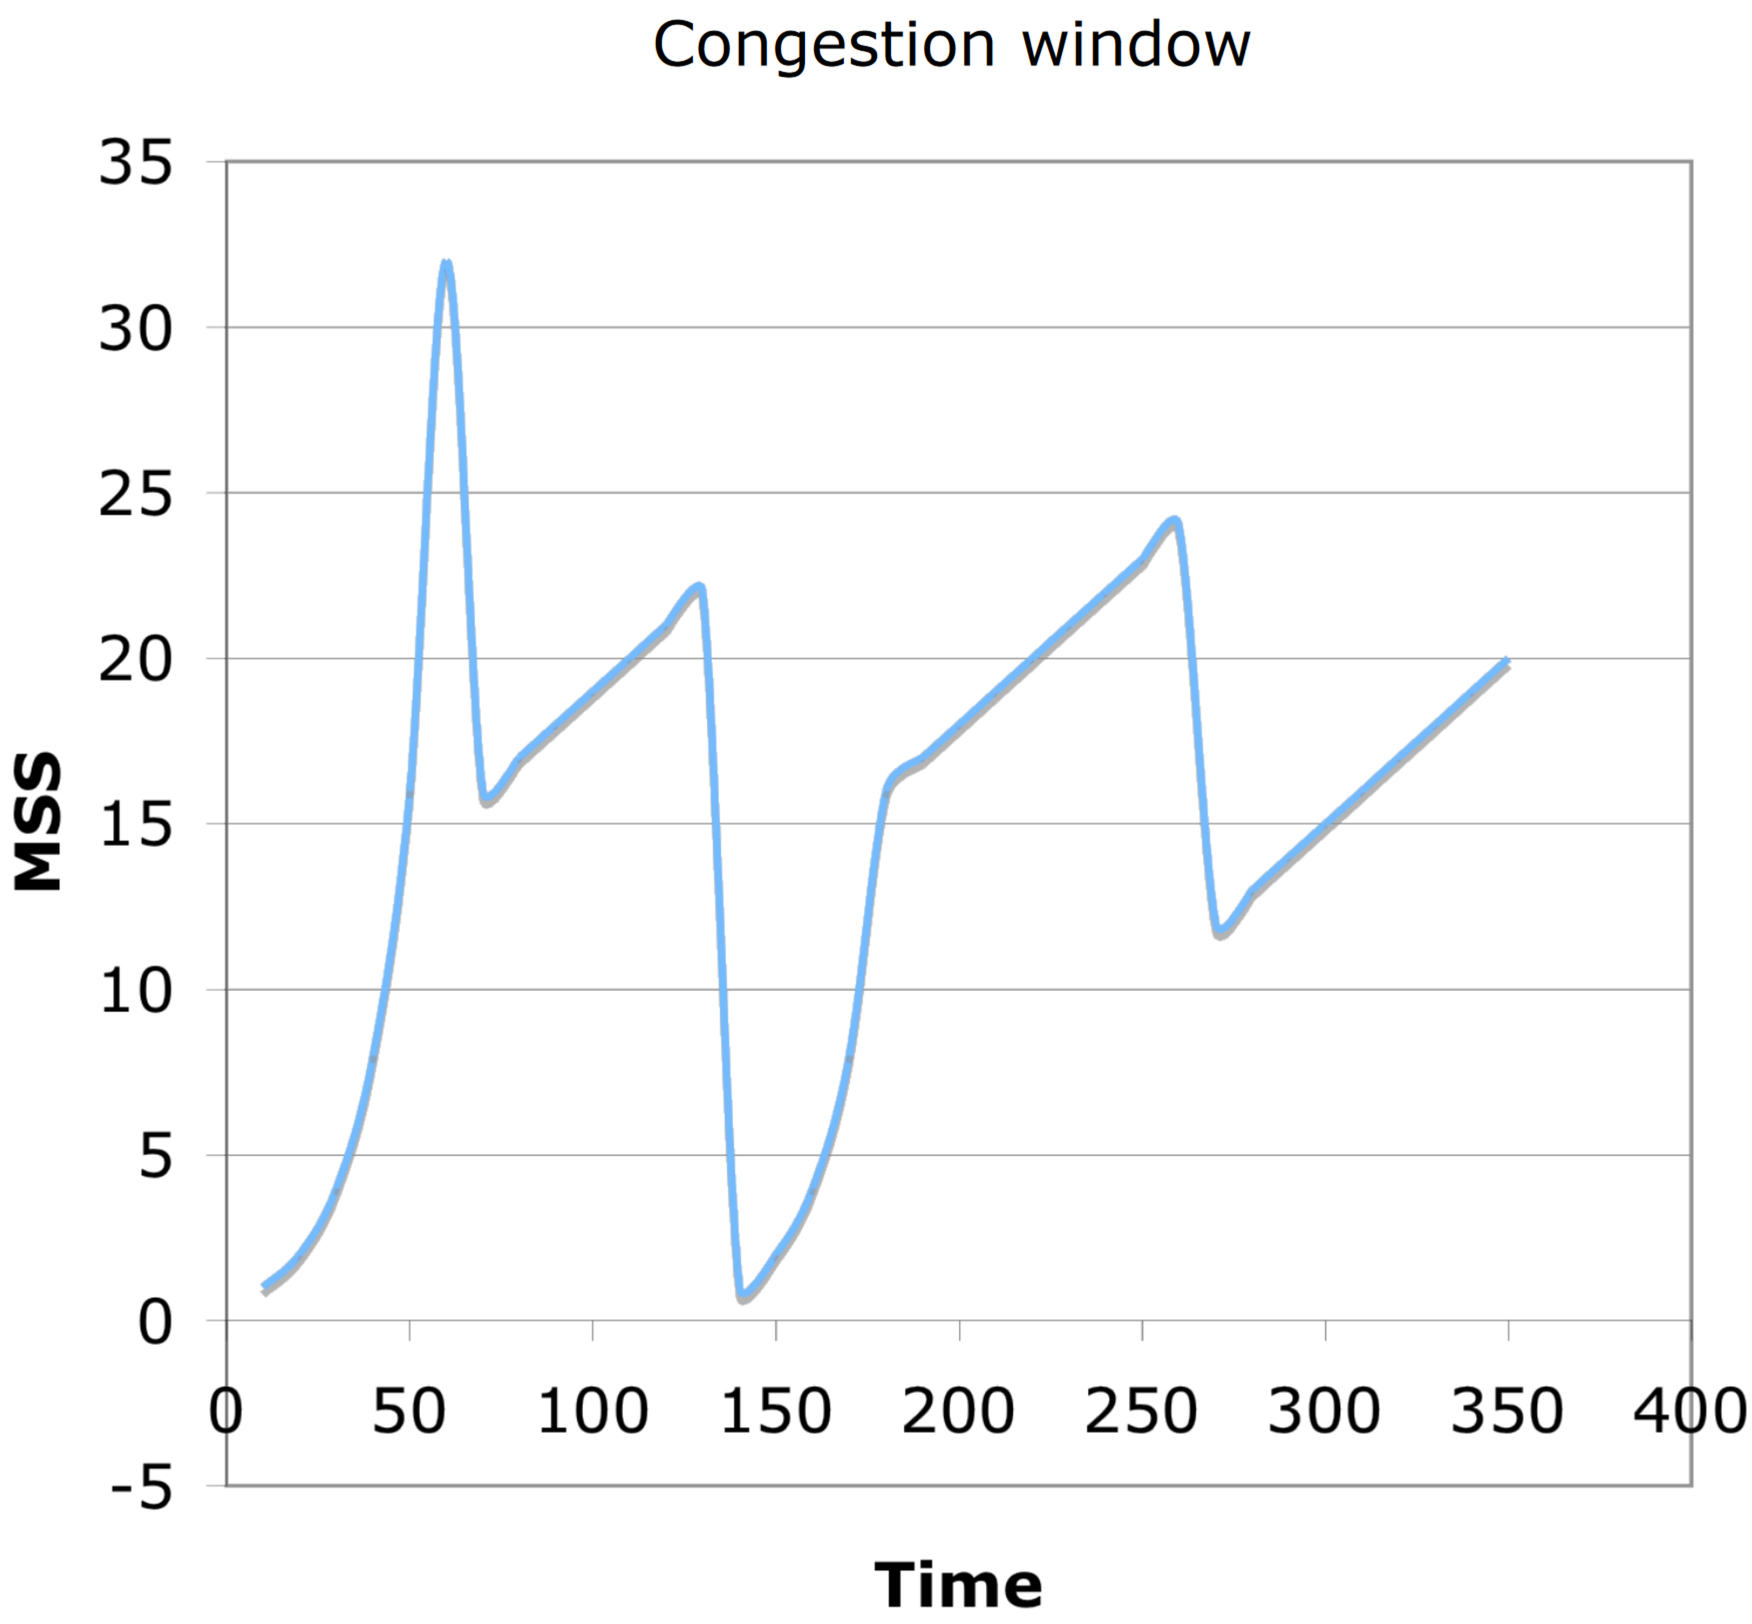

Problem Set 5 Question 5

Describe what is happening in this TCP trace at various times:

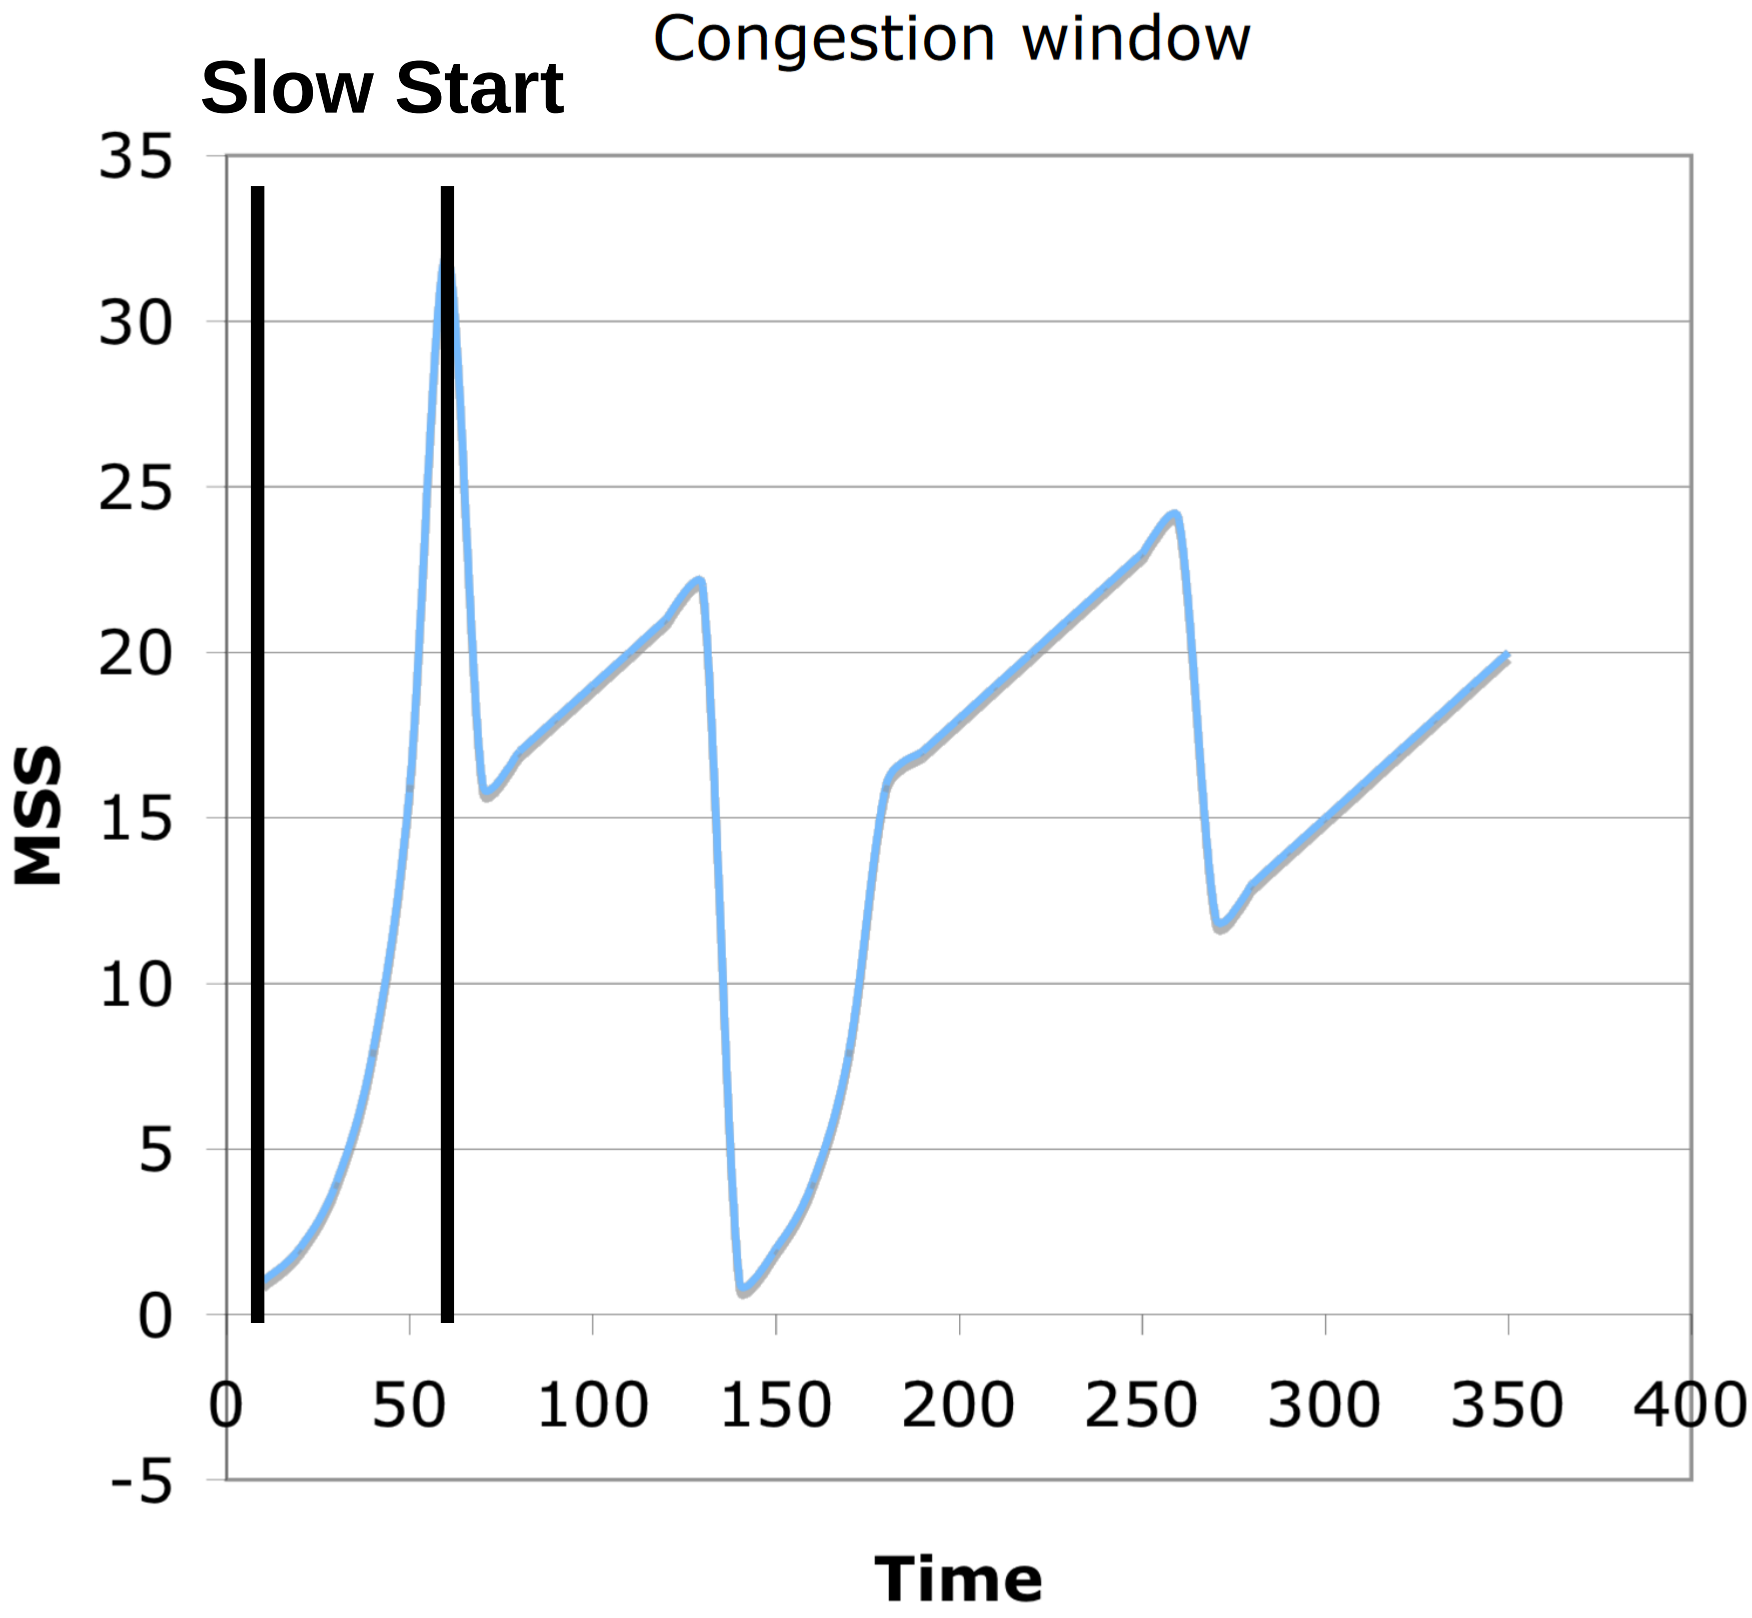

Problem Set 5 Question 5

Describe what is happening in this TCP trace at various times:

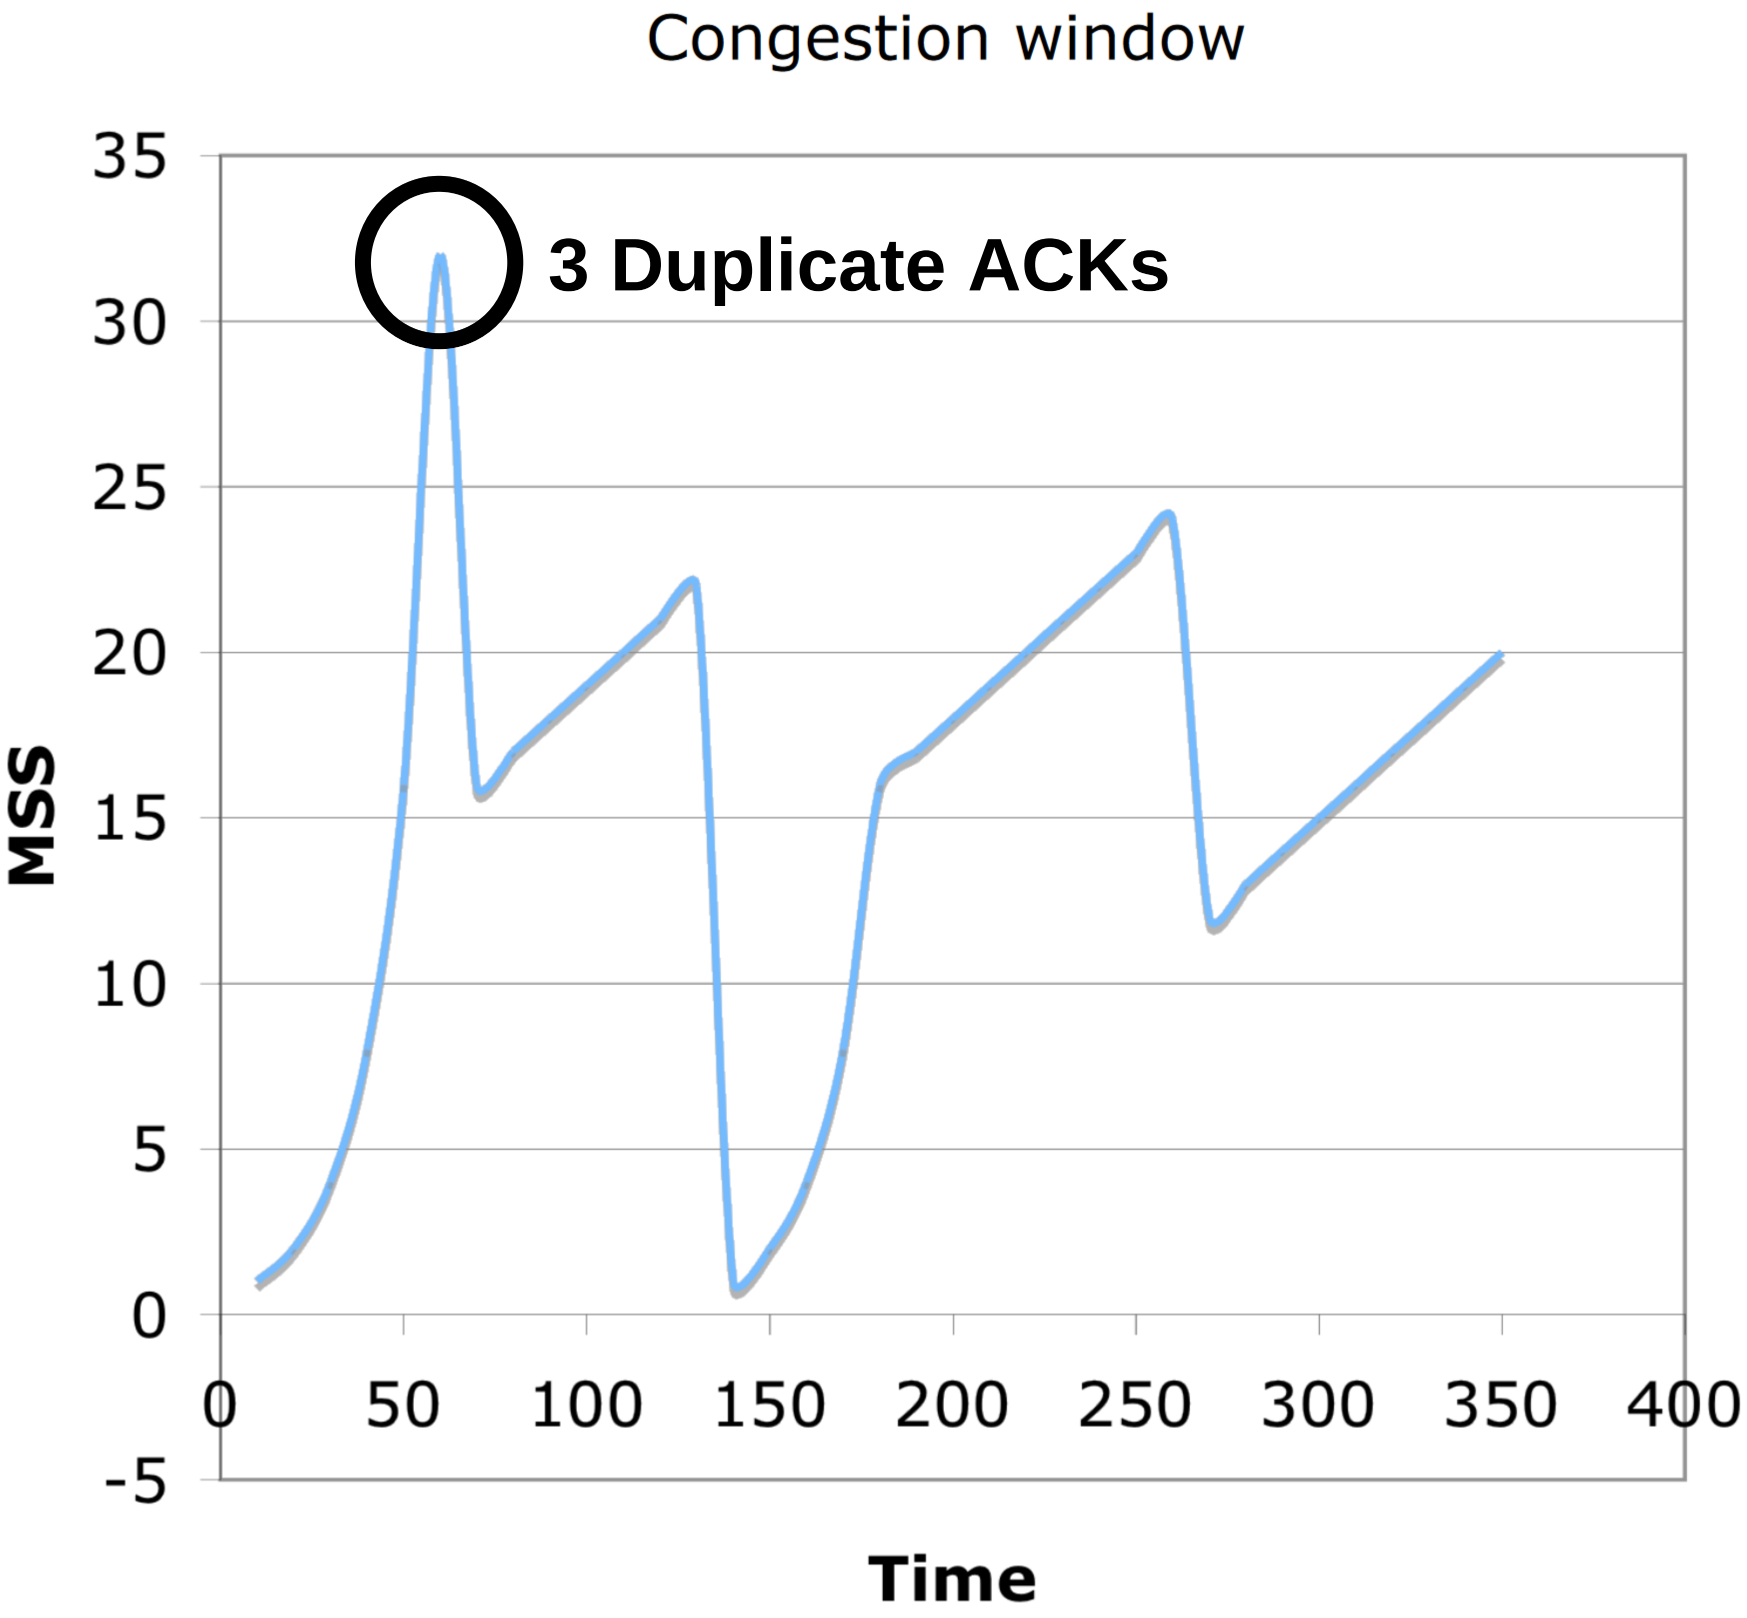

Problem Set 5 Question 5

Describe what is happening in this TCP trace at various times:

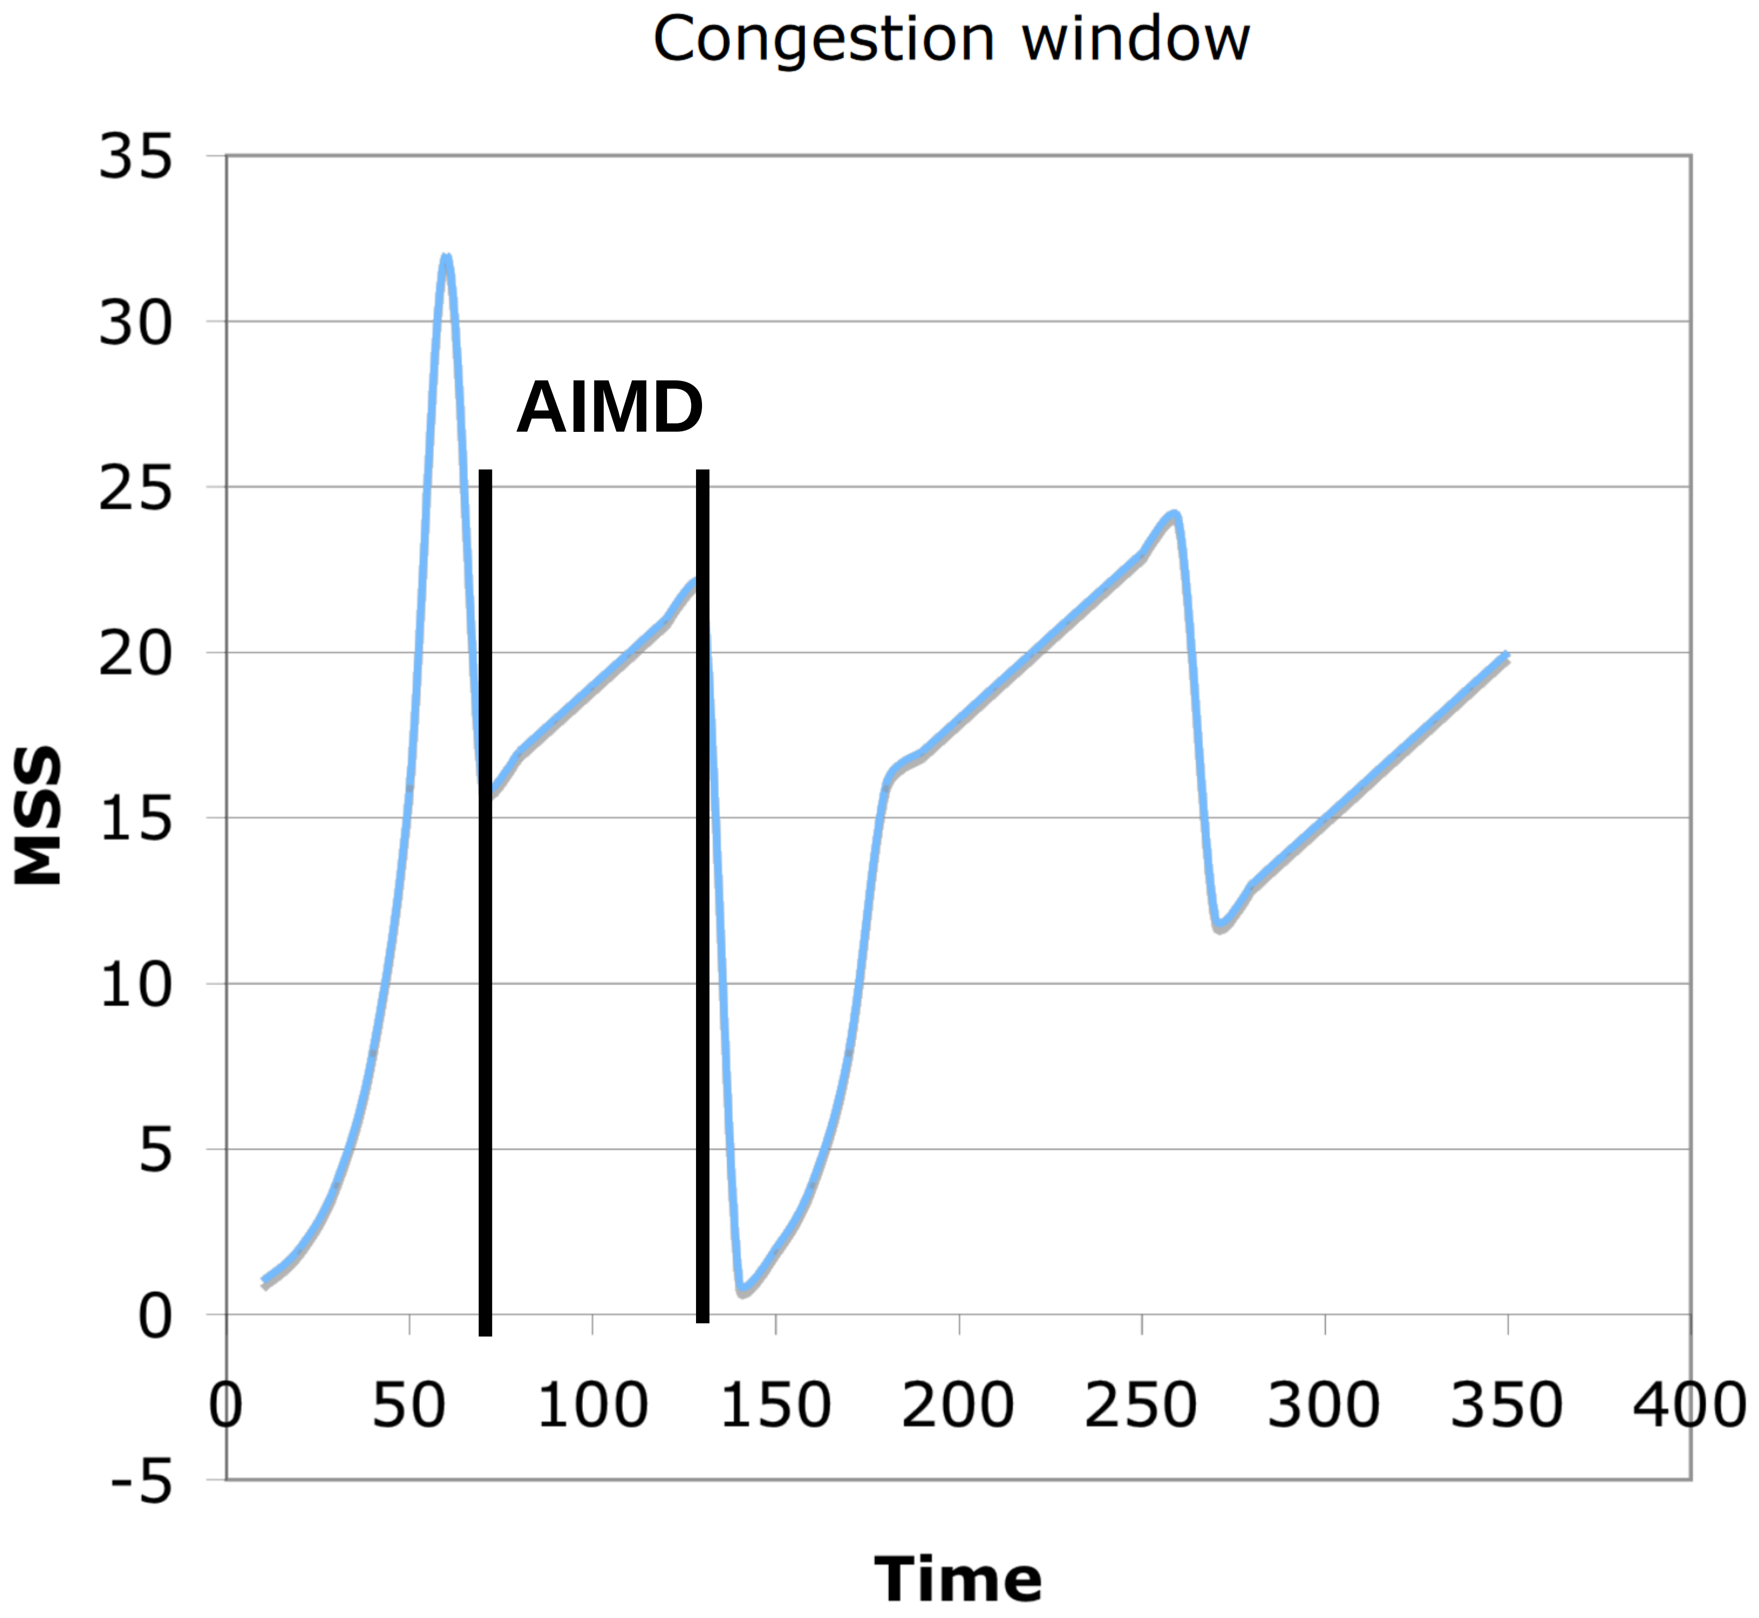

Problem Set 5 Question 5

Describe what is happening in this TCP trace at various times:

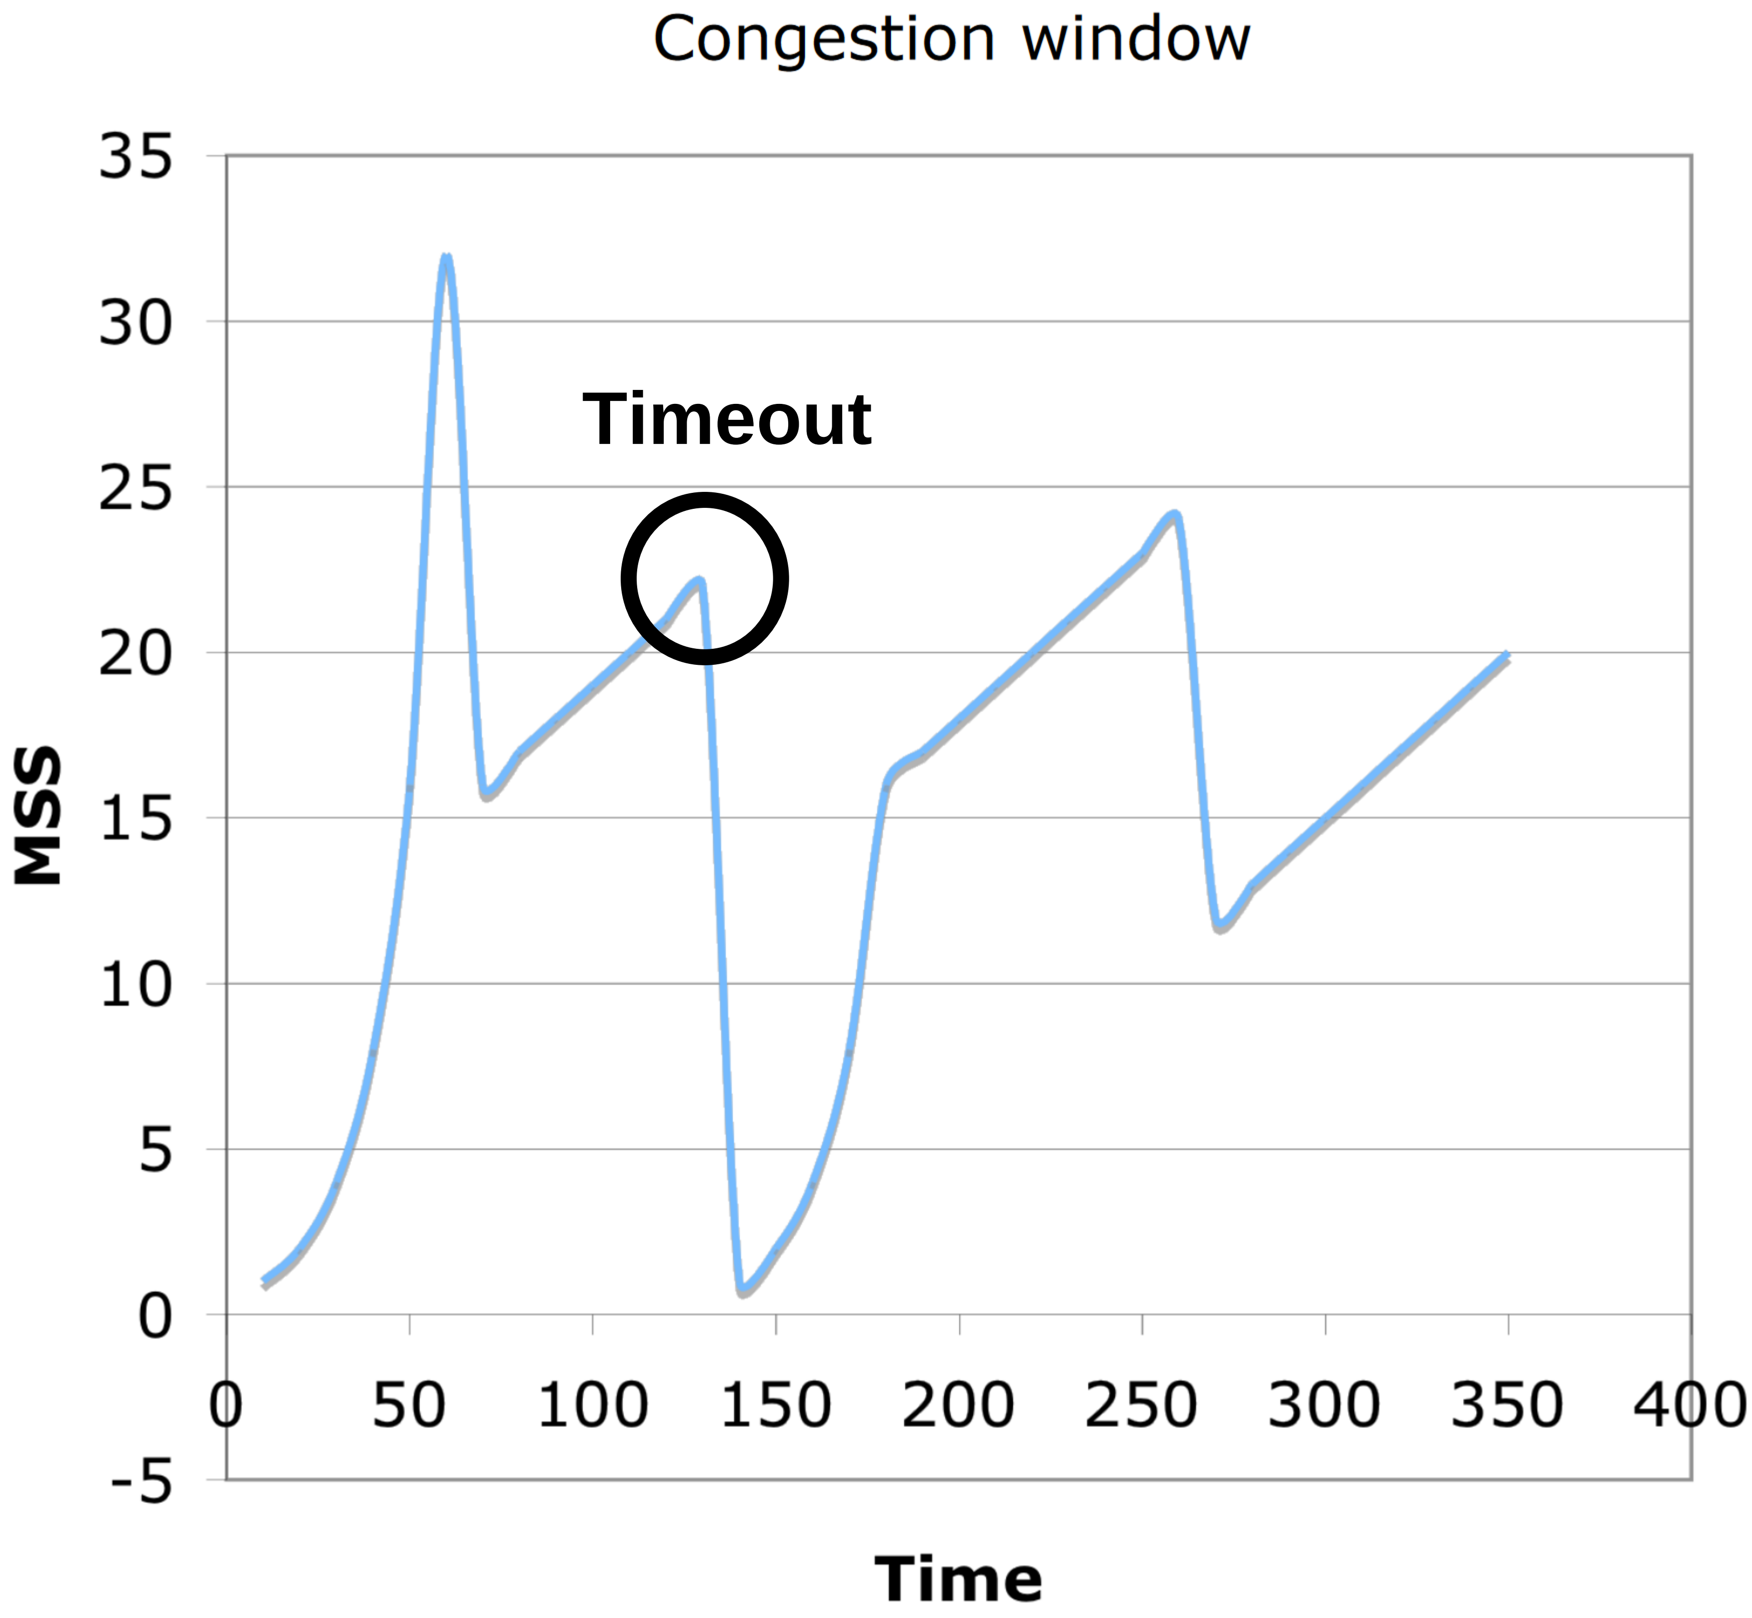

Problem Set 5 Question 5

Describe what is happening in this TCP trace at various times:

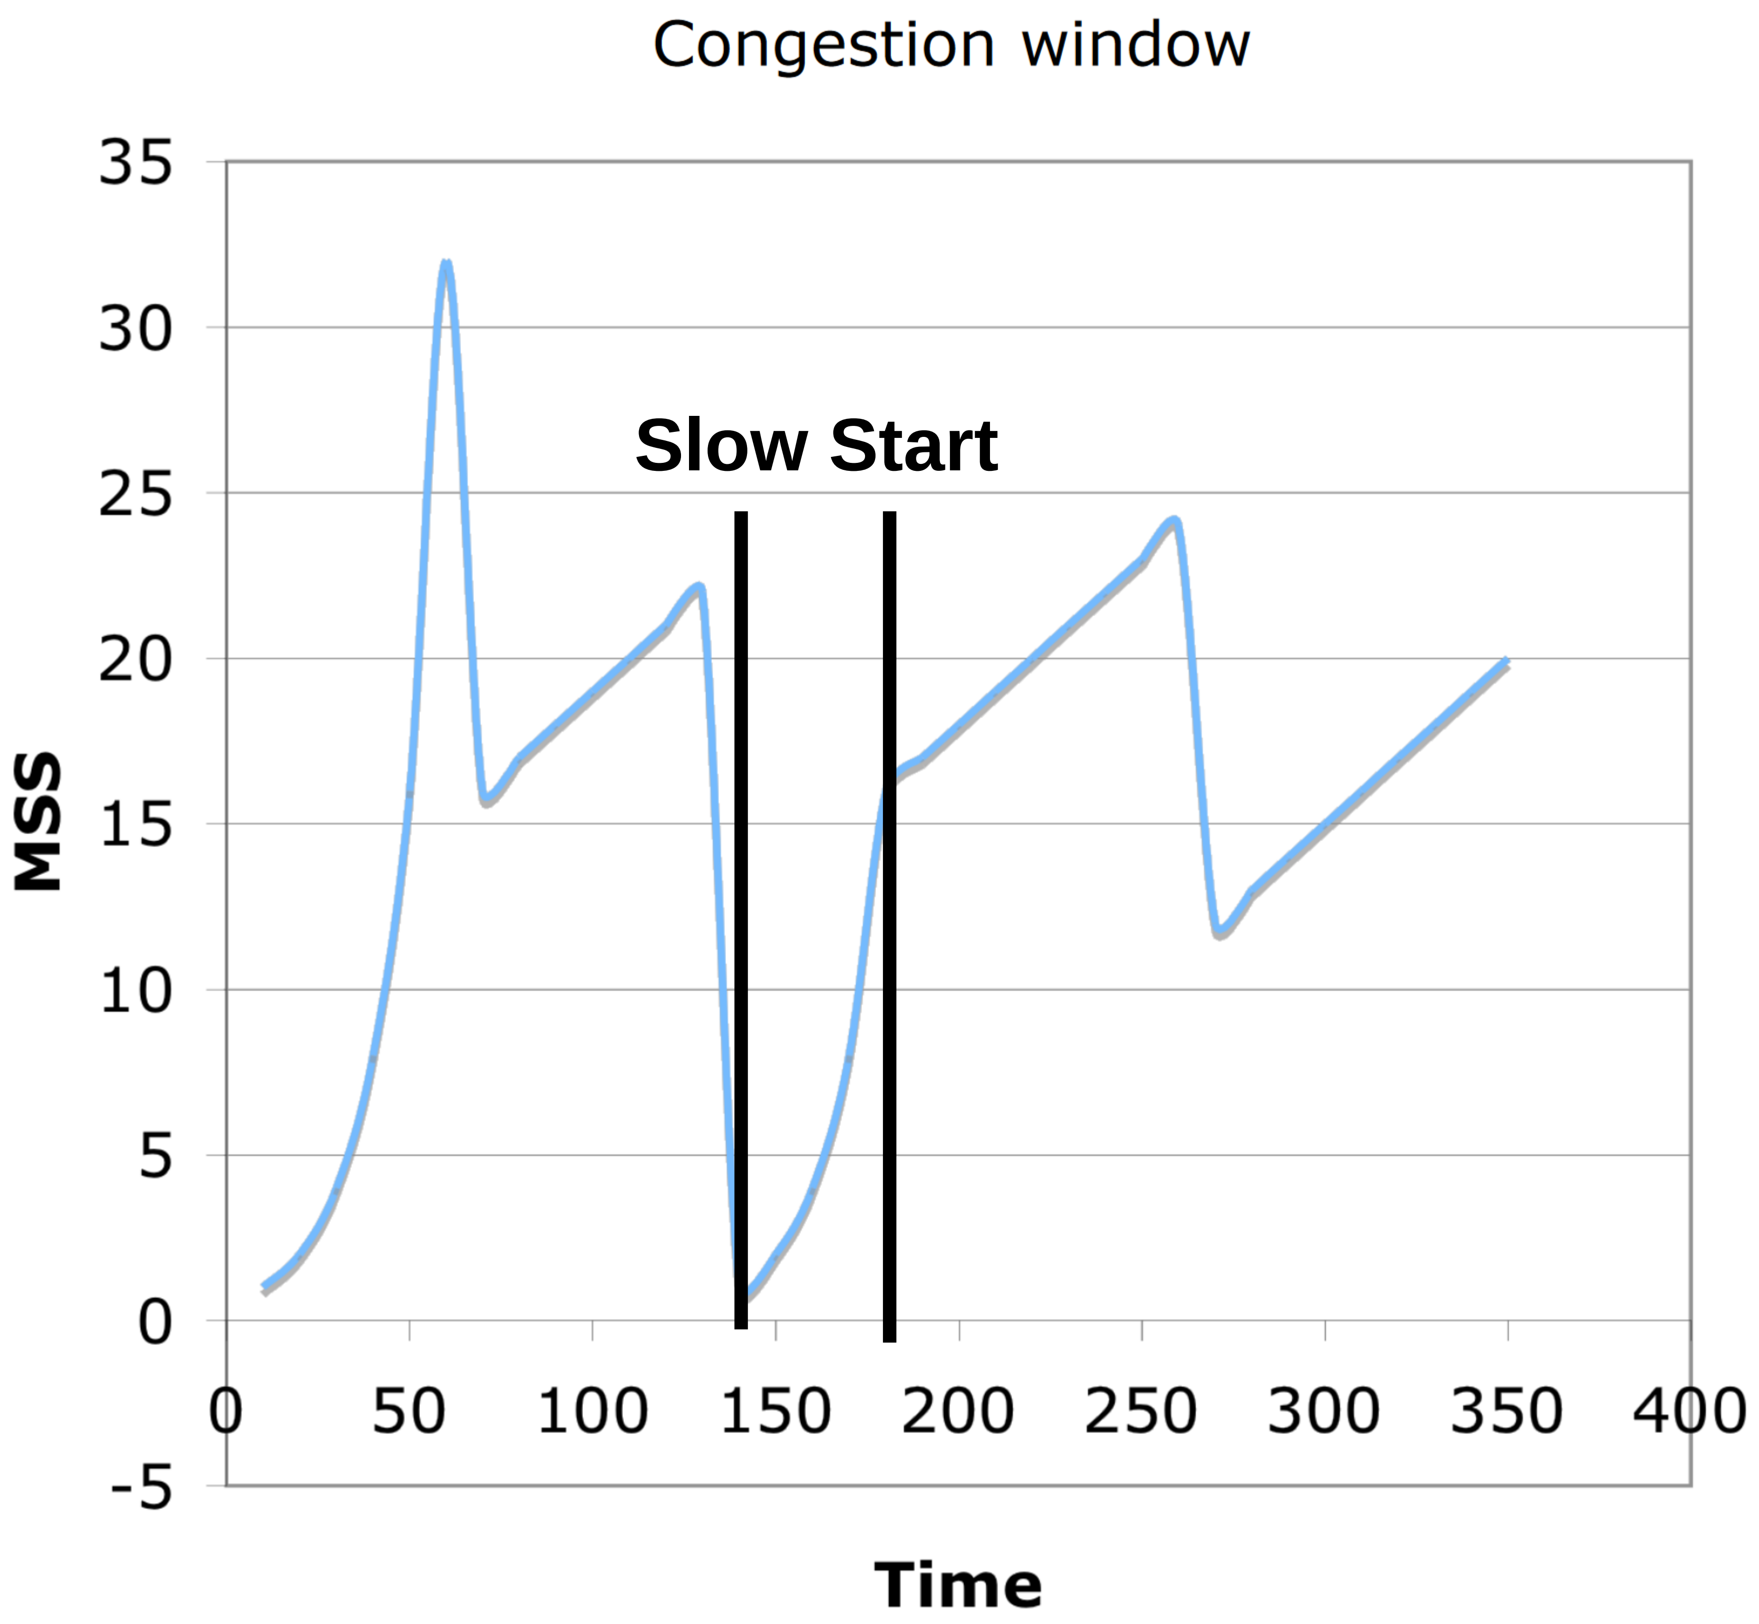

Problem Set 5 Question 5

Describe what is happening in this TCP trace at various times:

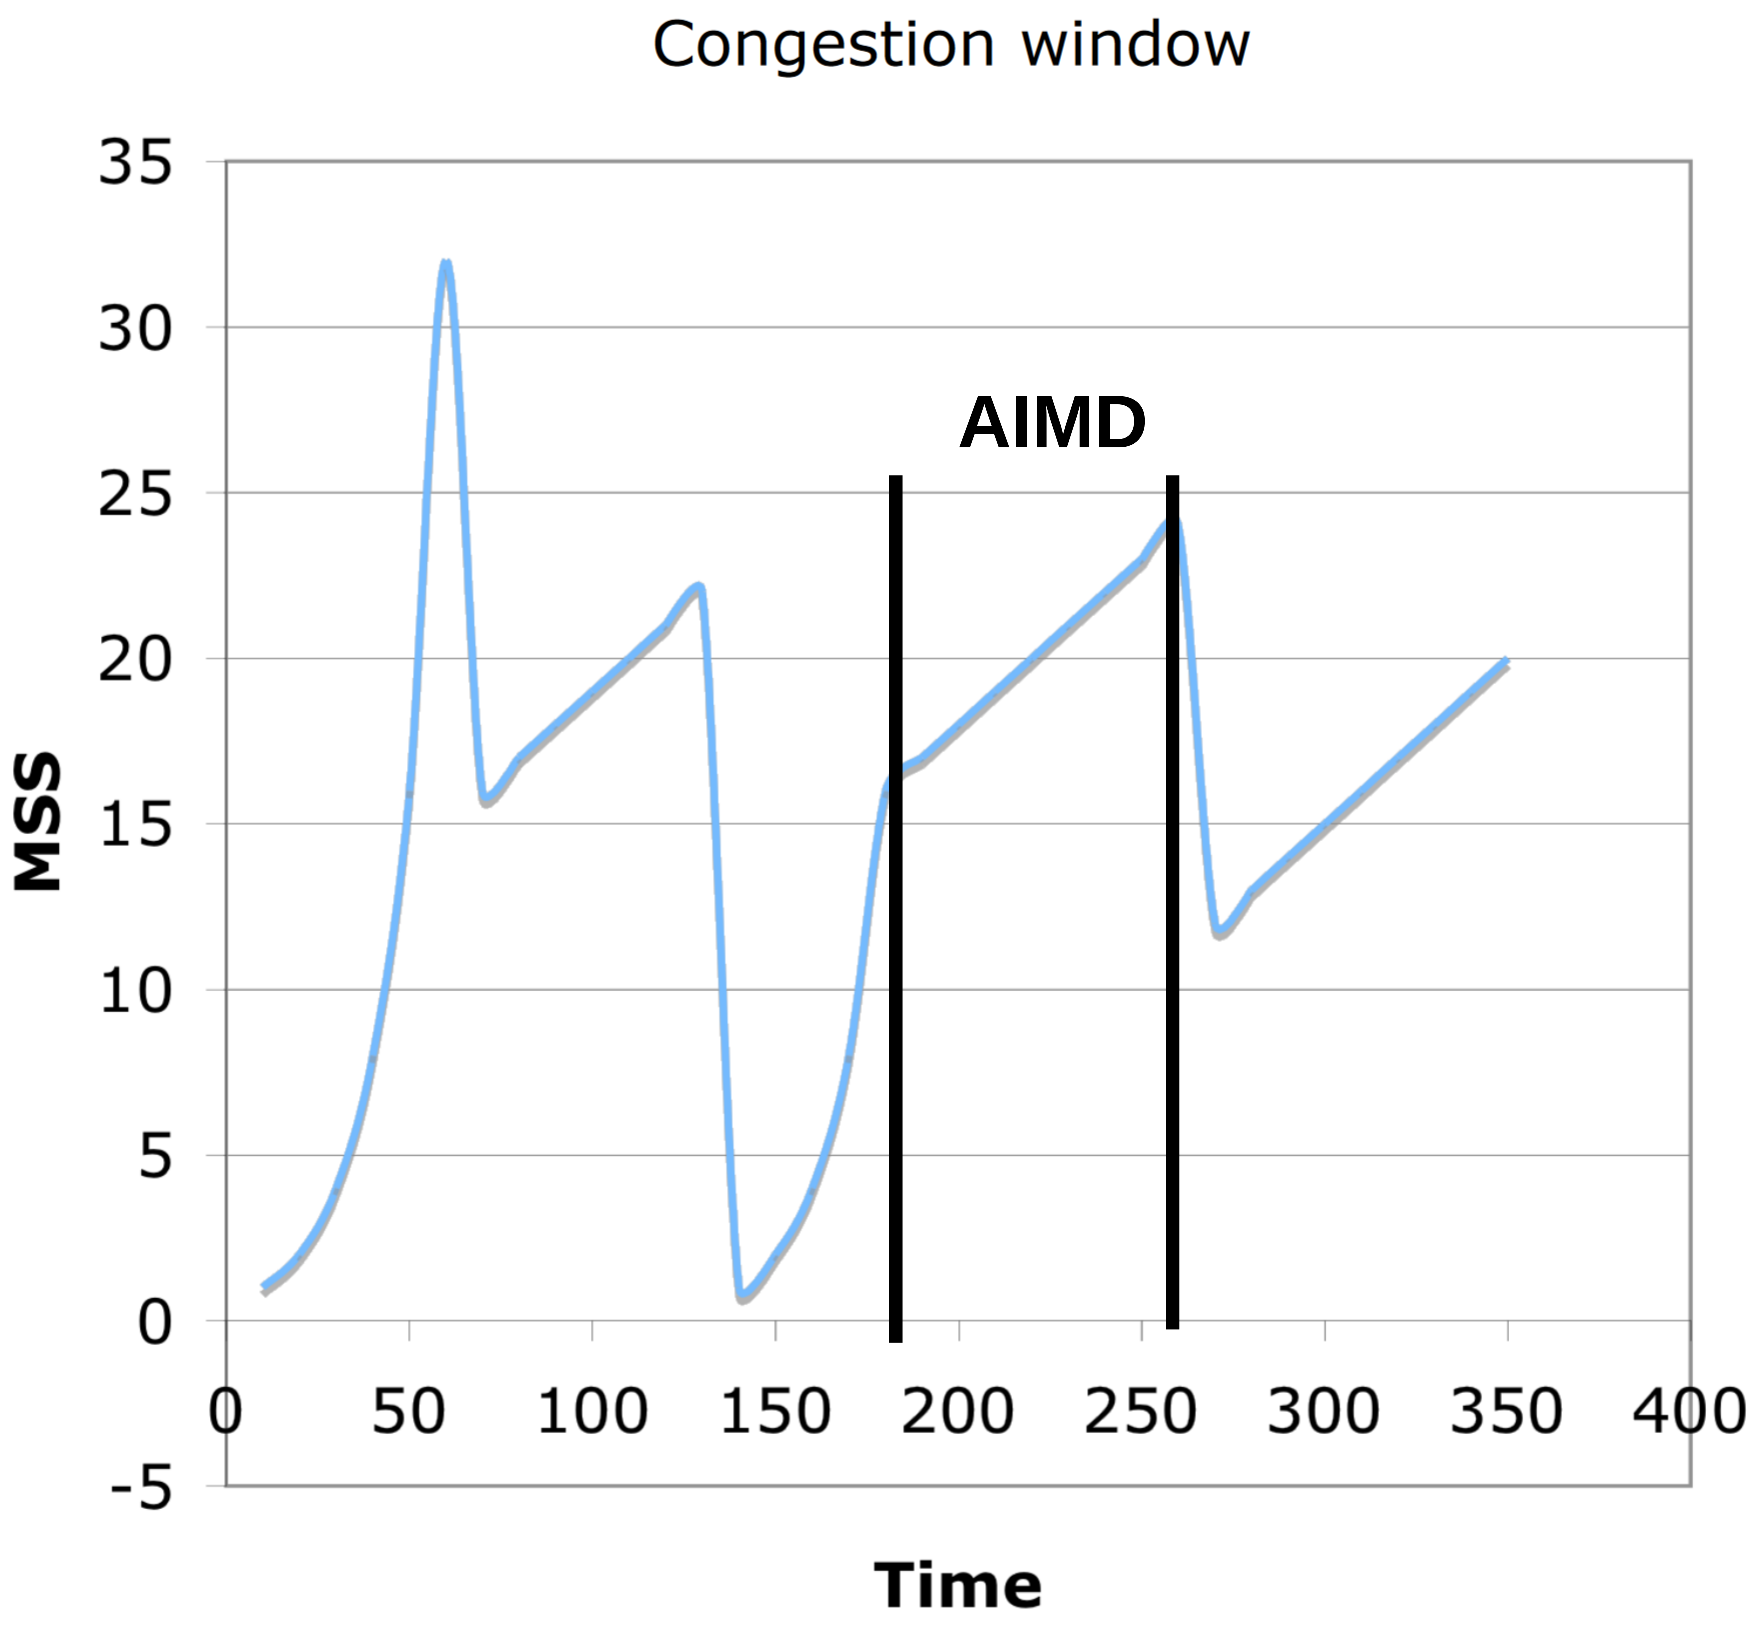

Problem Set 5 Question 5

Describe what is happening in this TCP trace at various times:

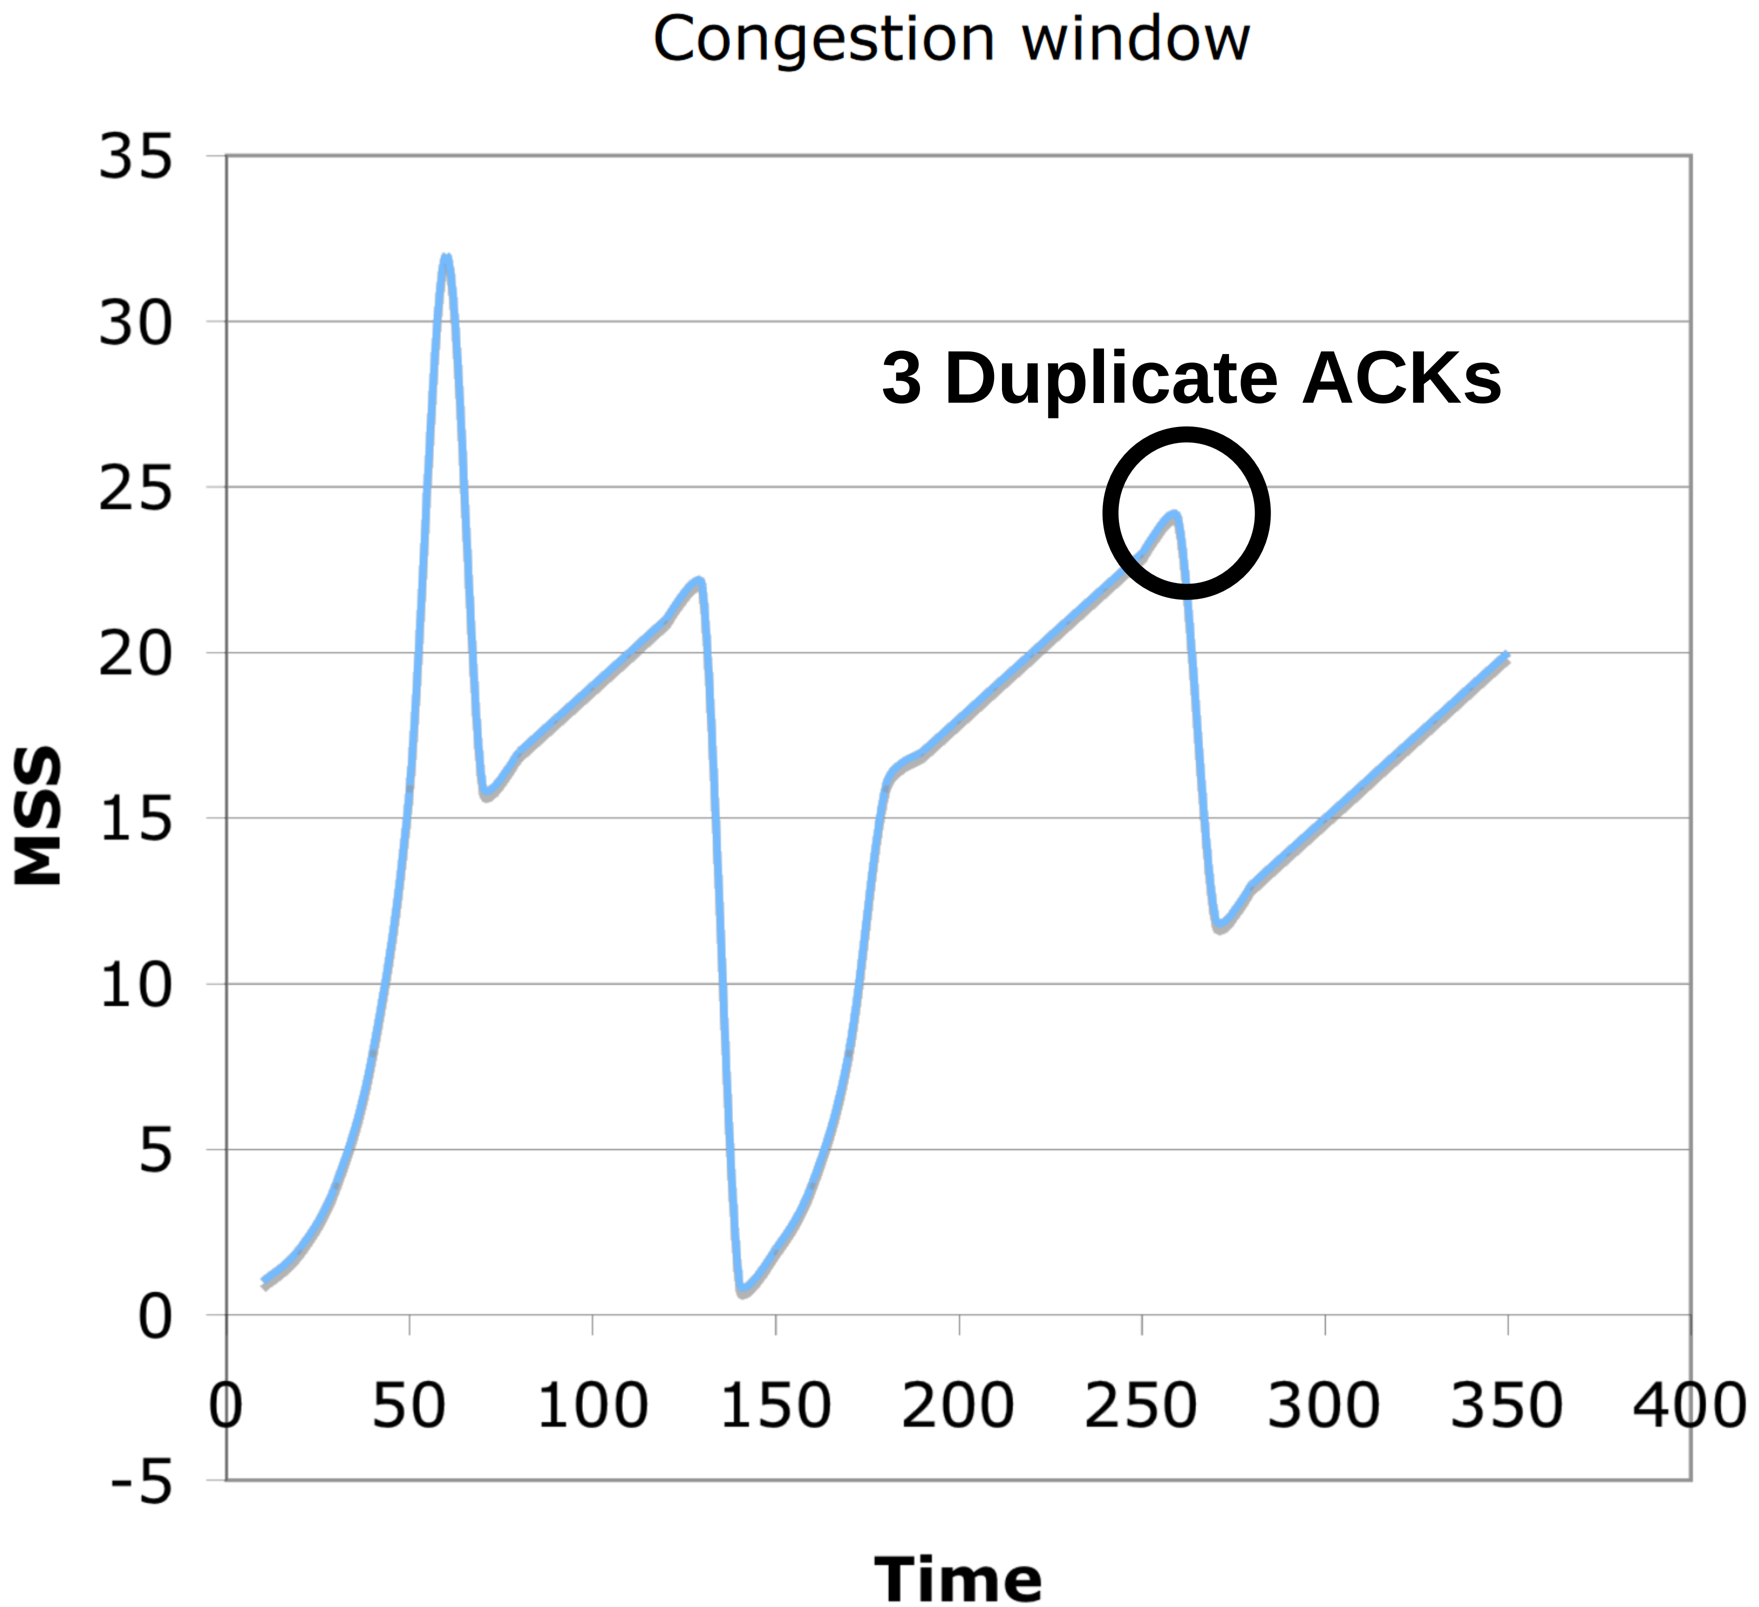

Problem Set 5 Question 5

Describe what is happening in this TCP trace at various times:

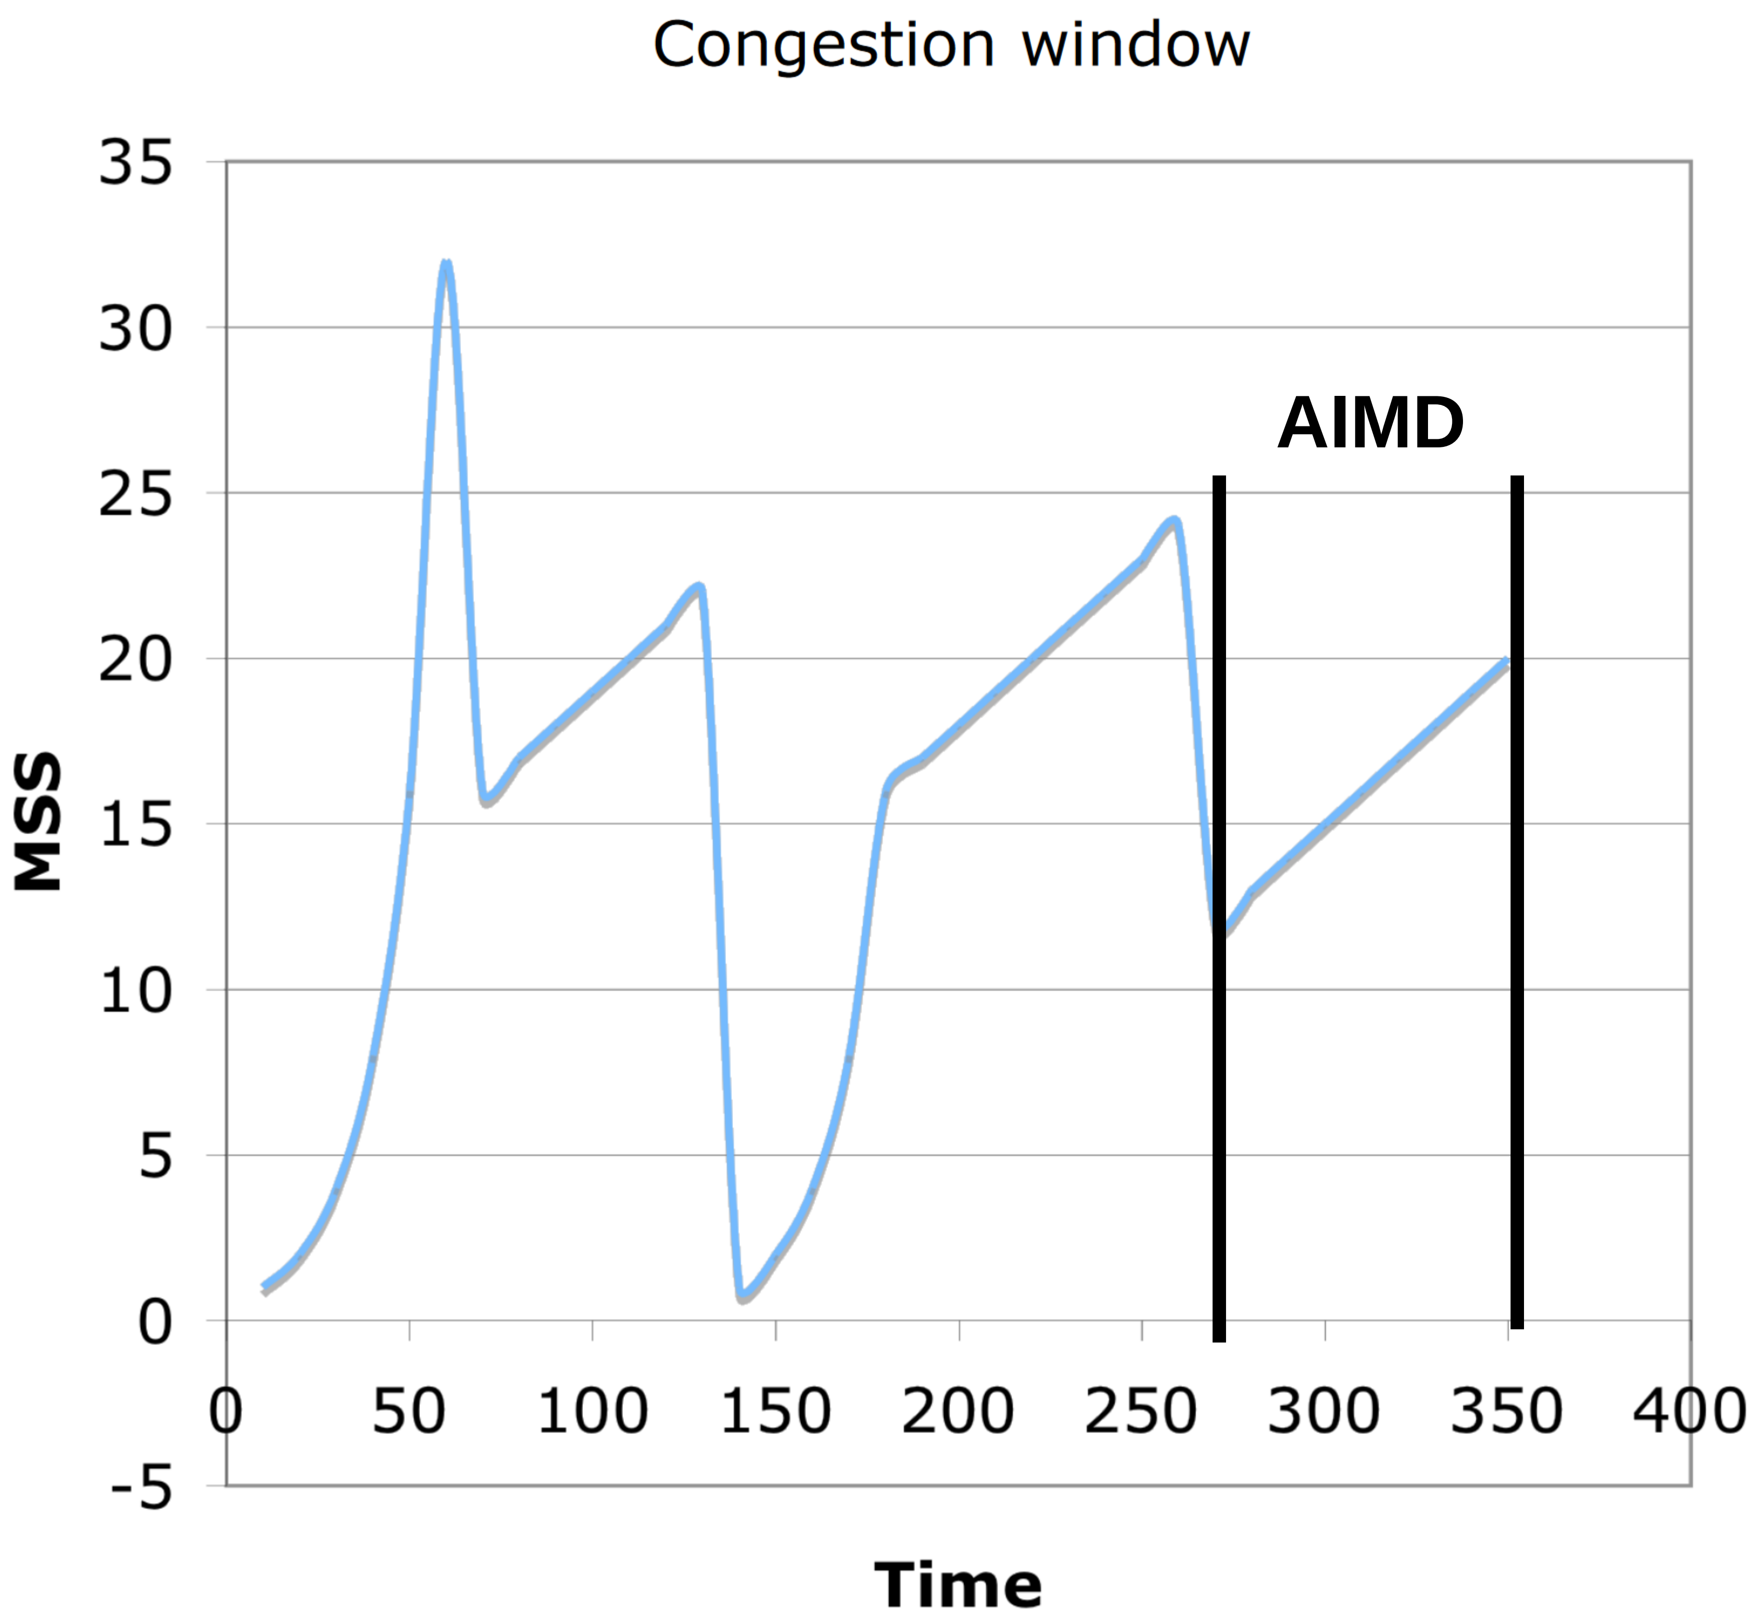

Problem Set 5 Question 5

b) One reason that congestion leads to poor performance is large queuing delays. Would eliminating queues improve the performance of the Internet?

No. Eliminating queues would eliminate queuing delay, but we would have high packet loss instead.

Problem Set 5 Question 5

c) In TCP, each endpoint picks a random ISN (initial sequence number). Why is this done instead of starting at 0 every time?

- Avoids confusion from heavily delayed packets from the previous connection
- Makes it more difficult for an attacker to guess the current sequence number
- Any other ideas?

Problem Set 5 Question 5

d) RFC 7323 defines a "TCP Window Scale Option" that allows window sizes up to 2^{30} instead of the usual 2^{16}. When would this option be useful?

- Links with high bandwidth and high latency.
- On such a link, it may take less than one RTT to transmit 2^{16} packets. We don't want to limit CWND.
- Example: If BW = 1 Gbps, RTT = 1s, and MSS = 1500, more than 2^{16} packets can be sent per RTT.

Problem Set 5 Question 6

a) Find the "sawtooth time" (STT) of TCP in terms of the RTT and W, the maximum congestion window.

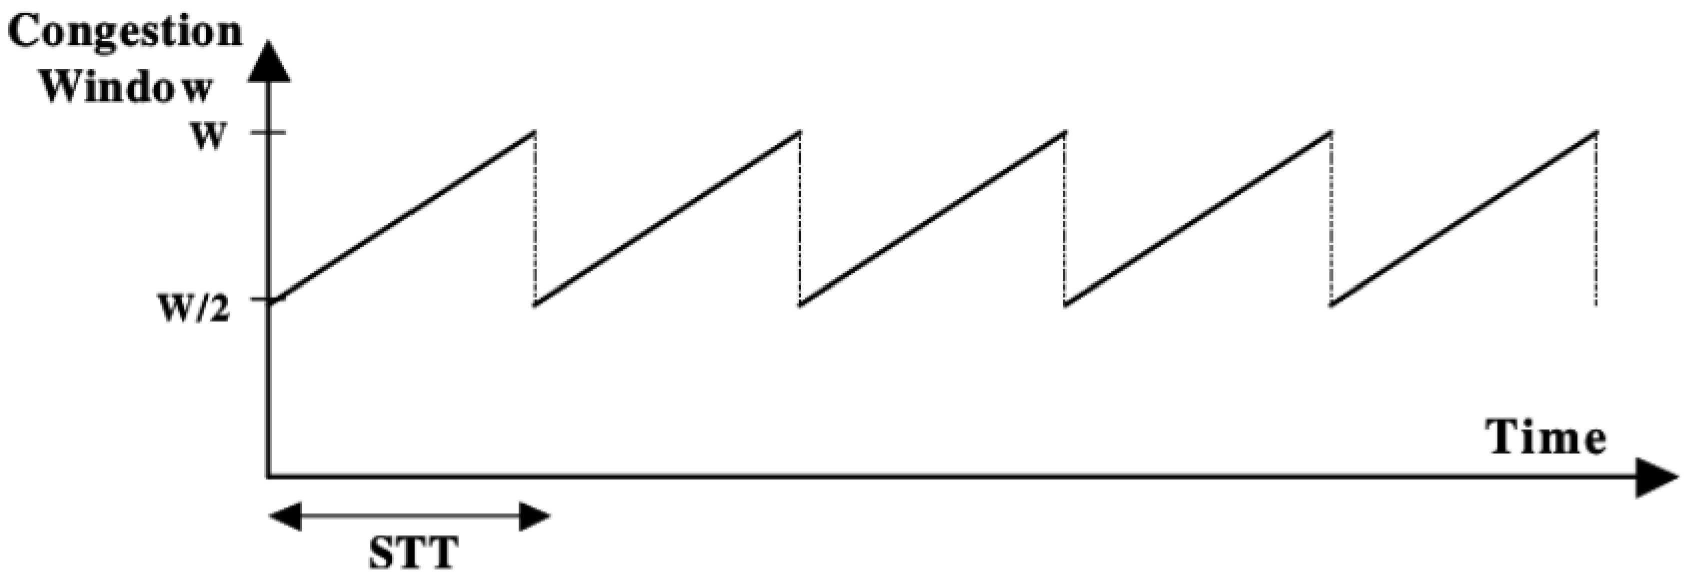

During AIMD, it takes 1 RTT to increase the congestion window by 1. It must increase by $\frac{W}{2}$ during each sawtooth.

$$\boldsymbol{STT = RTT * \frac{W}{2}}$$

Problem Set 5 Question 6

b) How many segments can be sent during each sawtooth?

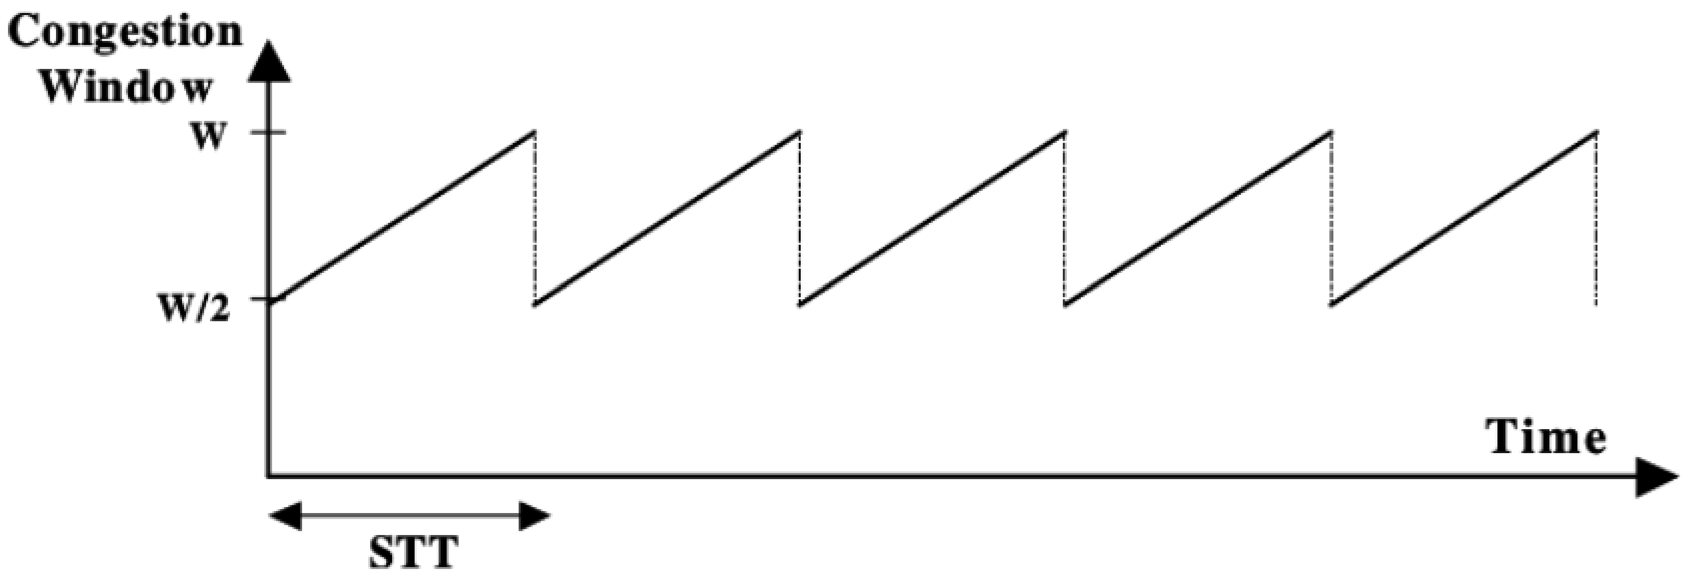

$$\frac{W}{2} + \left(\frac{W}{2} + 1\right) + \left(\frac{W}{2} + 2\right) + \left(\frac{W}{2} + 3\right) + \cdots + \left(\frac{W}{2} + \frac{W}{2}\right)$$

$$= \left(\frac{W}{2} + 1\right) * \frac{W}{2} + \left(1 + 2 + 3 + 4 + \cdots + \frac{W}{2}\right)$$

$$\approx \frac{W}{2} * \frac{W}{2} + \frac{\frac{W}{2} * \left(\frac{W}{2} + 1\right)}{2} \approx \frac{W^2}{4} + \frac{W^2}{8} = \frac{3}{8} W^2$$

Problem Set 5 Question 6

c) Exactly one packet loss occurs per sawtooth. Using this fact, express p in terms of W.

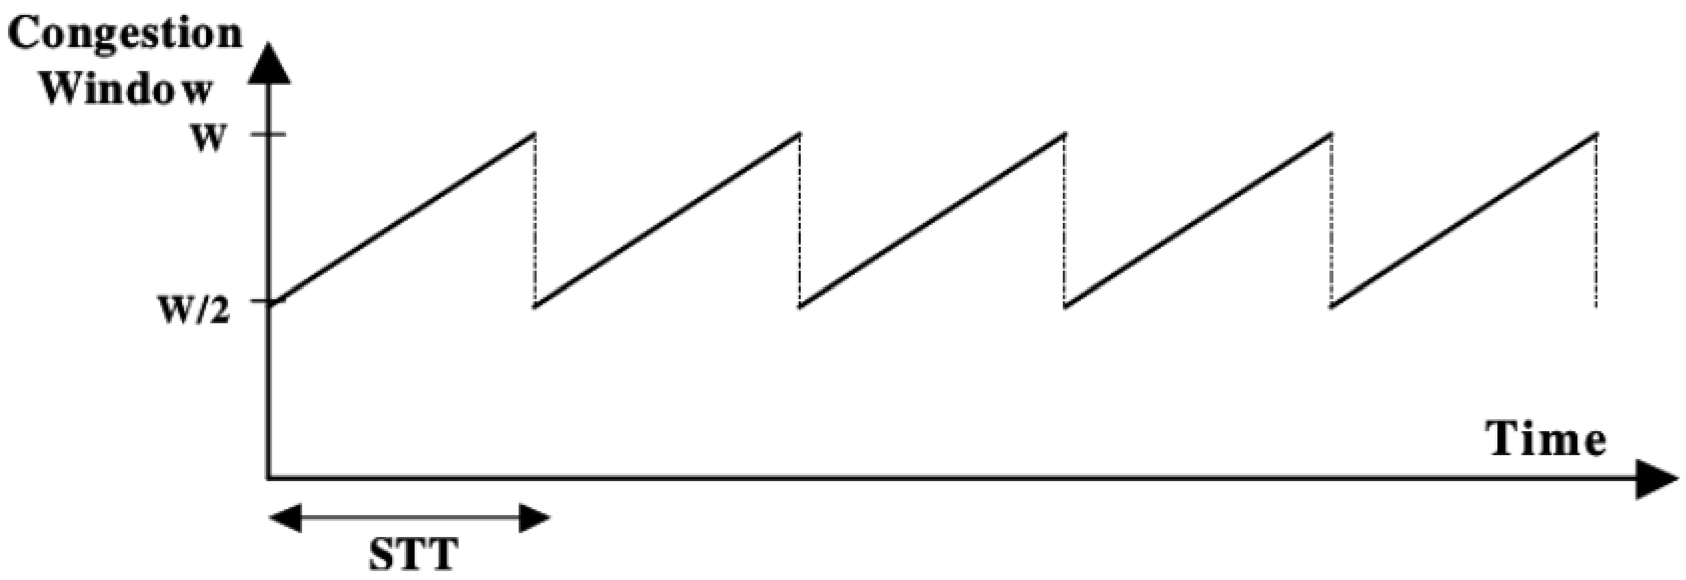

$$p = \frac{1}{\frac{3}{8}W^2} = \frac{8}{3W^2}$$

Problem Set 5 Question 6

d) Calculate the average throughput of this perfect sawtooth graph as a function of RTT, MSS, and p.

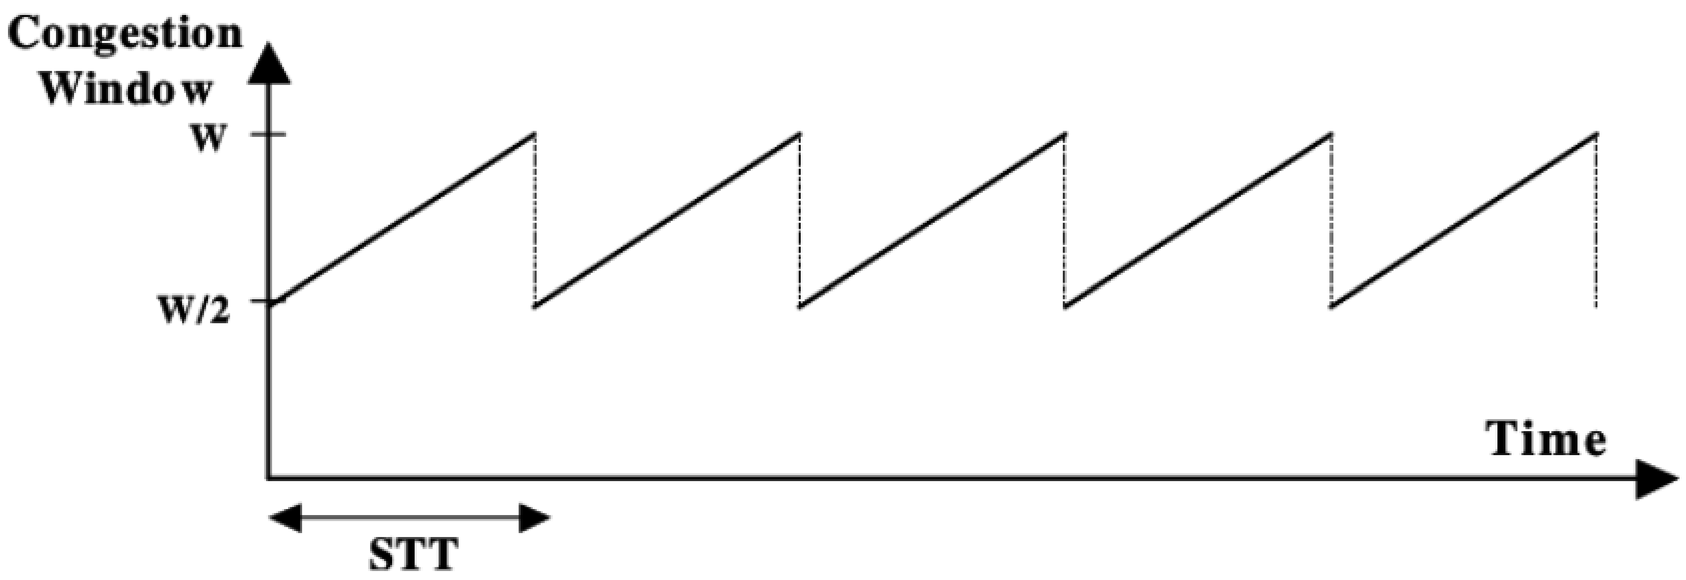

$$T = \frac{Bytes\ Transferred\ per\ Sawtooth}{STT}$$

$$T = \frac{\frac{3}{8}W^2 * MSS}{RTT * \frac{W}{2}} = \frac{\frac{3}{4}W * MSS}{RTT} = \sqrt{\frac{3}{2}}\frac{\sqrt{\frac{3W^2}{8}} * MSS}{RTT} = \sqrt{\frac{3}{2}}\frac{MSS}{\sqrt{p} * RTT}$$

Problem Set 6 Question 1

a) Starting at the root server, show the list of queries used to find www.cs.cornell.edu.

```
> dig +norecurse @b.edu-servers.net www.cs.cornell.edu

...
;; AUTHORITY SECTION:
edu.                        172800 IN NS      b.edu-servers.net.
edu.                        172800 IN NS      f.edu-servers.net.
edu.                        172800 IN NS      i.edu-servers.net.
...

;; ADDITIONAL SECTION:
b.edu-servers.net.          172800 IN A       192.33.14.30 b.edu-
servers.net.          172800 IN AAAA    2001:503:231d::2:30
f.edu-servers.net.          172800 IN A       192.35.51.30
...
```

Problem Set 6 Question 1

a) Starting at the root server, show the list of queries used to find www.cs.cornell.edu.

```
> dig +norecurse @b.edu-servers.net www.cs.cornell.edu


...
;; AUTHORITY SECTION:
cornell.edu.            172800  IN  NS    bigred.cit.cornell.edu.
cornell.edu.            172800  IN  NS    dns.cit.cornell.edu.
cornell.edu.            172800  IN  NS    drdns.cit.cornell.edu.
cornell.edu.            172800  IN  NS    drdns2.cit.cornell.edu.

;; ADDITIONAL SECTION:
bigred.cit.cornell.edu. 172800  IN  A     128.253.180.35
dns.cit.cornell.edu.    172800  IN  A     192.35.82.53
drdns.cit.cornell.edu.  172800  IN  A     192.195.74.252
drdns2.cit.cornell.edu. 172800  IN  A     52.45.47.12
```

Problem Set 6 Question 1

a) Starting at the root server, show the list of queries used to find www.cs.cornell.edu.

```
> dig +norecurse @bigred.cit.cornell.edu www.cs.cornell.edu


...
;; ANSWER SECTION:
www.cs.cornell.edu.      86400   IN  CNAME   web1.cs.cornell.edu.
web1.cs.cornell.edu.     86400   IN  A       132.236.207.20


...
```

Problem Set 6 Question 1

b) Suppose we can access the caches of the local DNS servers of the CS department. How could we determine which websites are most popular among the users in the CS department?

If we periodically take snapshots of the DNS caches, each time we will see which websites were recently accessed.

If we take many snapshots, the websites that appear in the most snapshots will be the most popular, since they are frequently accessed.

Problem Set 6 Question 1

c) Suppose the CS department has a local DNS server for all of the computers in the department. Can we tell whether or not anyone in the department recently accessed a given website?

Yes. If we use dig to query the department DNS server for that website, the query time will be low if it was recently accessed. Otherwise, it will be high since it is not cached.

Problem Set 6 Question 2

a) Which servers are queried by the following dig command?

```
> dig einstein.physics.cornell.edu
```

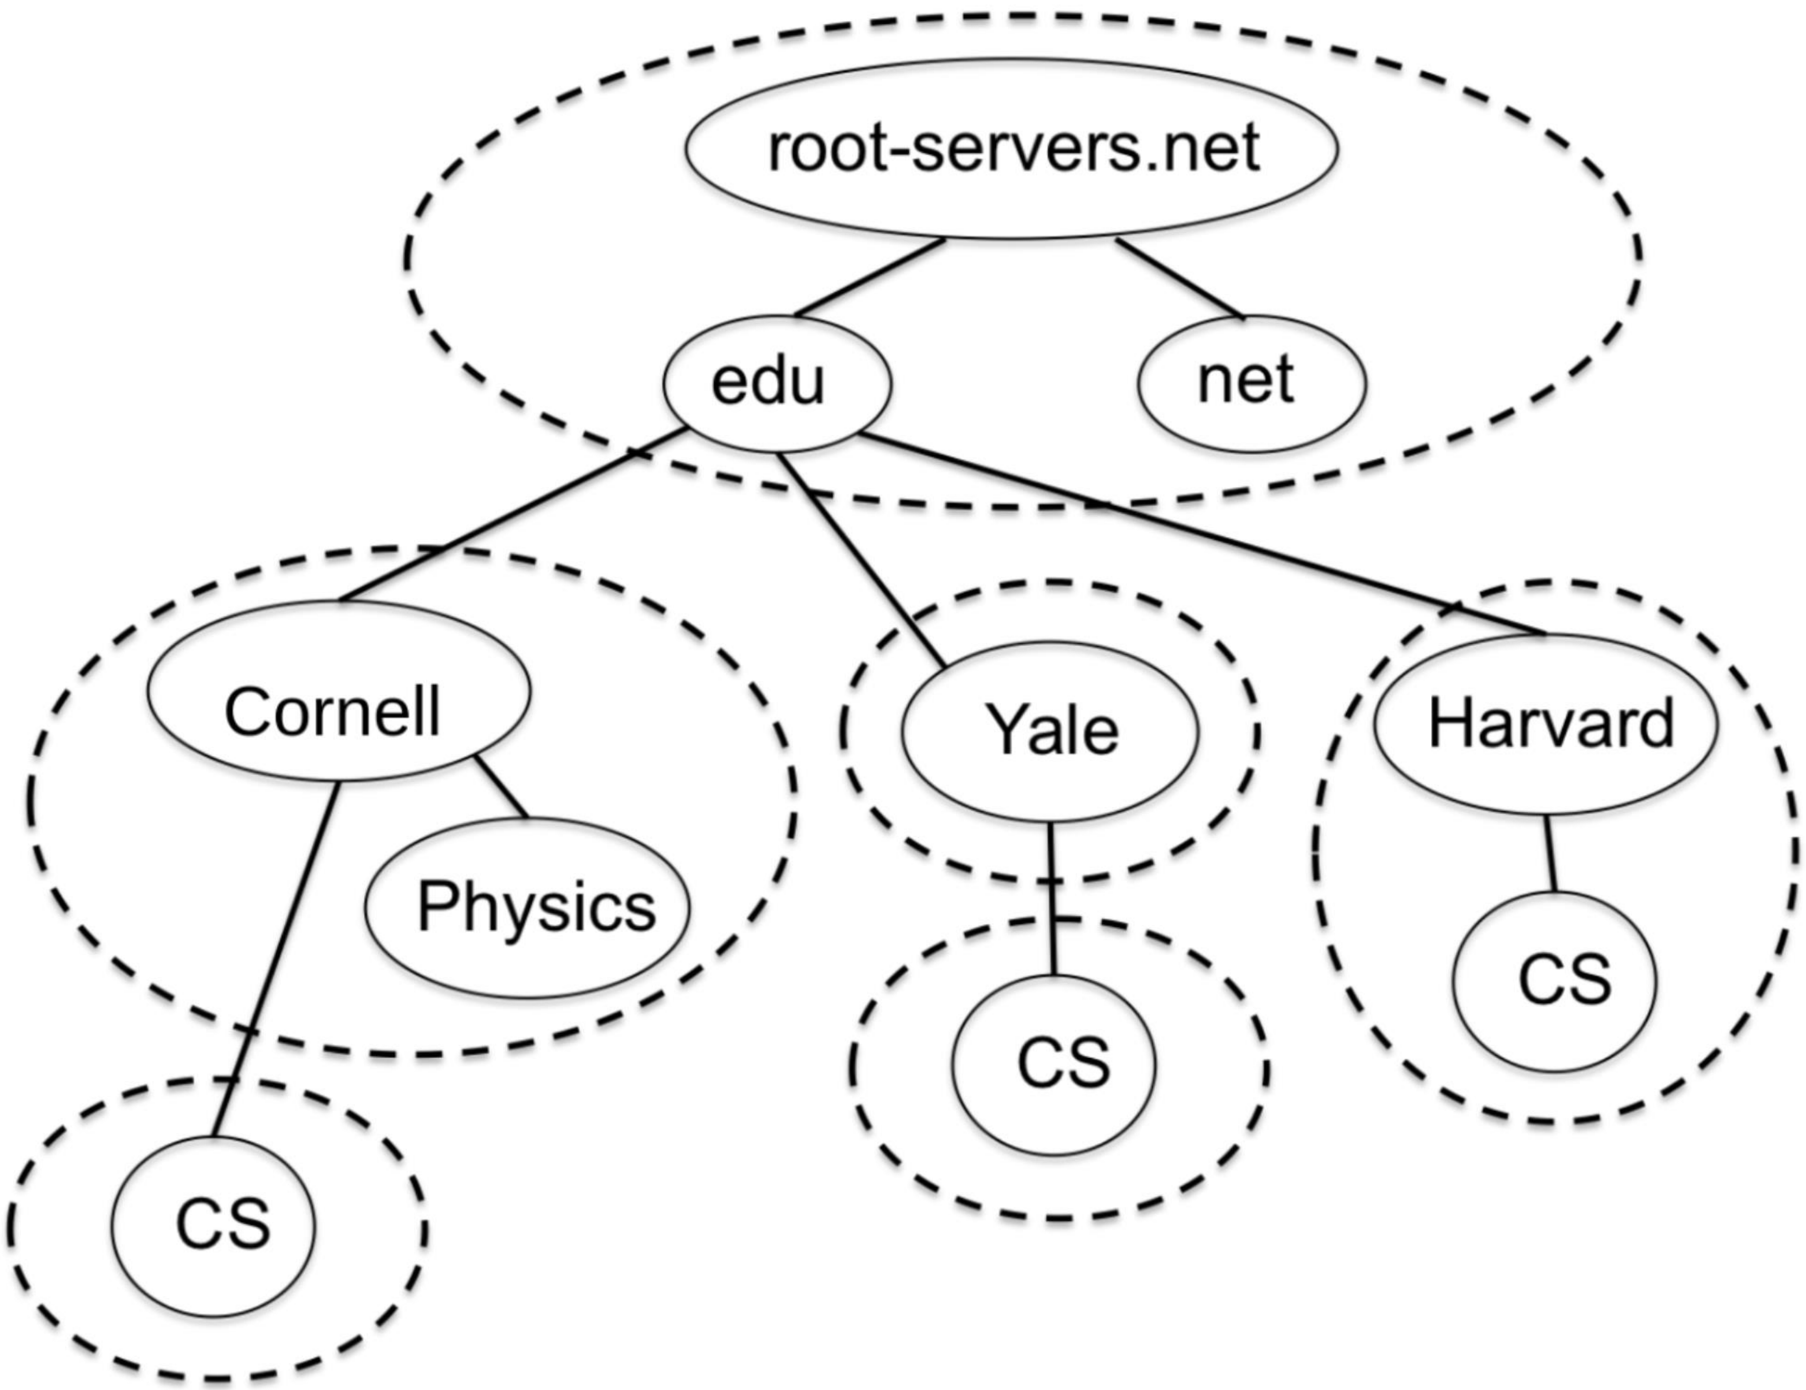

Problem Set 6 Question 2

b) Which servers are queried by the following dig command?

```
> dig dan.cs.yale.edu
```

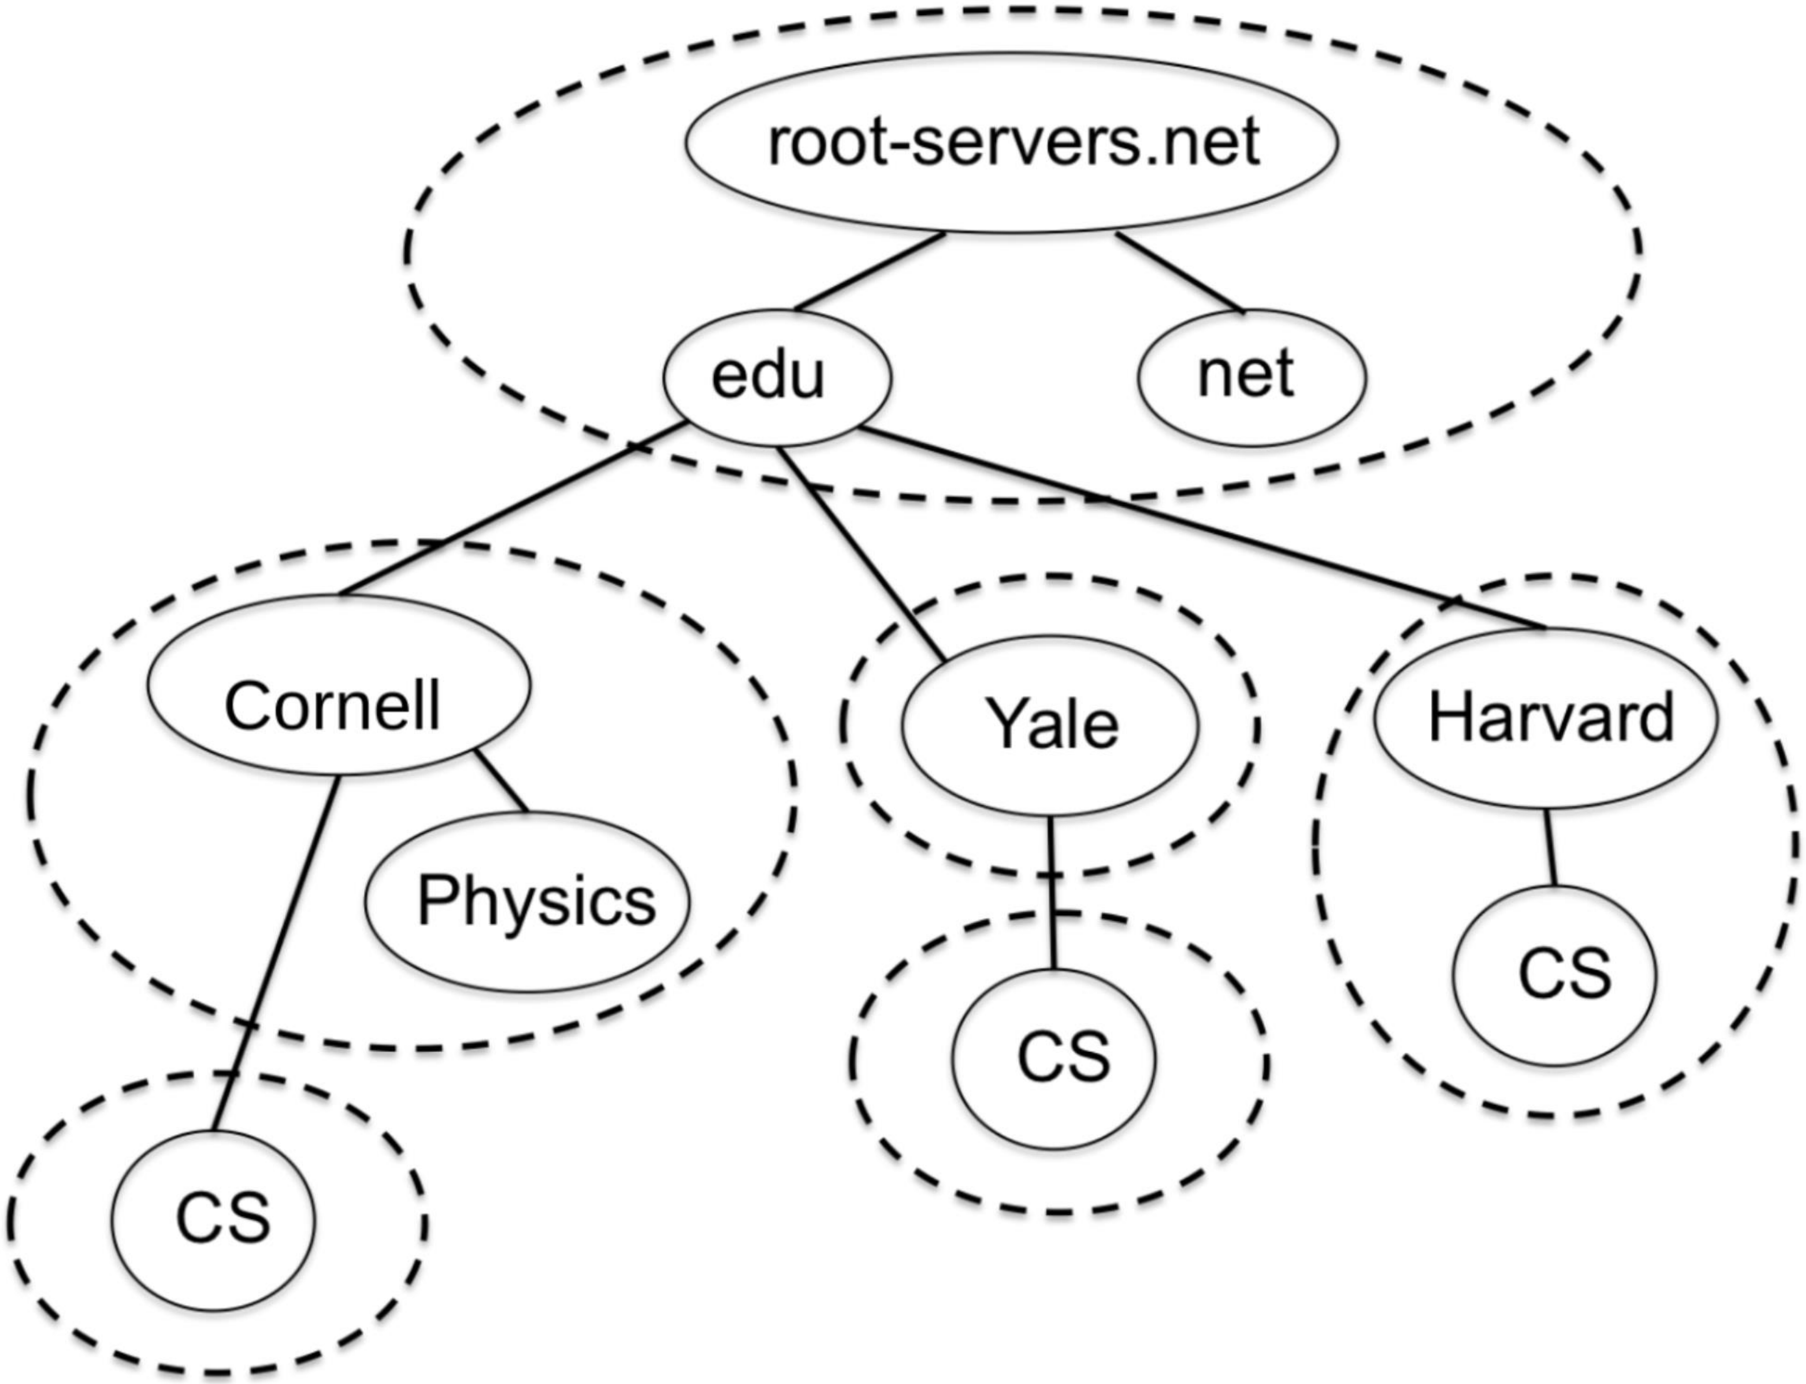

Problem Set 6 Question 1

a) List the IP addresses of the name servers of harvarde.edu

```
;; ANSWER SECTION:
harvard.edu. 1306 IN MX 100 b-00171101.pphosted.com.
harvard.edu. 1630 IN MX 100 a-00171101.pphosted.com.

;; AUTHORITY SECTION:
harvard.edu. 172100 IN NS ext-1.harvard.edu.

;; ADDITIONAL SECTION:
a-00171101.pphosted.com.        1313    IN A 67.231.148.27
b-00171101.pphosted.com.        1797    IN A 67.231.156.27
ext-1.harvard.edu.            172756    IN A 128.103.200.35
```

Problem Set 6 Question 1

a) List the IP addresses of the name servers of harvarde.edu

```
;; ANSWER SECTION:
harvard.edu. 1306 IN MX 100 b-00171101.pphosted.com.
harvard.edu. 1630 IN MX 100 a-00171101.pphosted.com.

;; AUTHORITY SECTION:
harvard.edu. 172100 IN NS ext-1.harvard.edu.

;; ADDITIONAL SECTION:
a-00171101.pphosted.com.        1313    IN A 67.231.148.27
b-00171101.pphosted.com.        1797    IN A 67.231.156.27
ext-1.harvard.edu.            172756    IN A 128.103.200.35
```

Problem Set 6 Question 1

b) List the IP addresses of the mail servers of harvarde.edu

```
;; ANSWER SECTION:
harvard.edu. 1306 IN MX 100 b-00171101.pphosted.com.
harvard.edu. 1630 IN MX 100 a-00171101.pphosted.com.

;; AUTHORITY SECTION:
harvard.edu. 172100 IN NS ext-1.harvard.edu.

;; ADDITIONAL SECTION:
a-00171101.pphosted.com.        1313    IN A 67.231.148.27
b-00171101.pphosted.com.        1797    IN A 67.231.156.27
ext-1.harvard.edu.            172756    IN A 128.103.200.35
```

Problem Set 6 Question 1

b) List the IP addresses of the mail servers of harvarde.edu

```
;; ANSWER SECTION:
harvard.edu. 1306 IN MX 100 b-00171101.pphosted.com.
harvard.edu. 1630 IN MX 100 a-00171101.pphosted.com.

;; AUTHORITY SECTION:
harvard.edu. 172100 IN NS ext-1.harvard.edu.

;; ADDITIONAL SECTION:
a-00171101.pphosted.com.        1313    IN A 67.231.148.27
b-00171101.pphosted.com.        1797    IN A 67.231.156.27
ext-1.harvard.edu.            172756    IN A 128.103.200.35
```

Problem Set 6 Question 1

c) For how many seconds are the entries for the address records of the mail and name servers valid?

```
;; ANSWER SECTION:
harvard.edu. 1306 IN MX 100 b-00171101.pphosted.com.
harvard.edu. 1630 IN MX 100 a-00171101.pphosted.com.

;; AUTHORITY SECTION:
harvard.edu. 172100 IN NS ext-1.harvard.edu.

;; ADDITIONAL SECTION:
a-00171101.pphosted.com.        1313    IN A 67.231.148.27
b-00171101.pphosted.com.        1797    IN A 67.231.156.27
ext-1.harvard.edu.            172756    IN A 128.103.200.35
```

Problem Set 6 Question 1

c) For how many seconds are the entries for the <u>address records</u> of the mail and name servers valid?

```
;; ANSWER SECTION:
harvard.edu. 1306 IN MX 100 b-00171101.pphosted.com.
harvard.edu. 1630 IN MX 100 a-00171101.pphosted.com.

;; AUTHORITY SECTION:
harvard.edu. 172100 IN NS ext-1.harvard.edu.

;; ADDITIONAL SECTION:
a-00171101.pphosted.com.          1313    IN A 67.231.148.27
b-00171101.pphosted.com.          1797    IN A 67.231.156.27
ext-1.harvard.edu.              172756    IN A 128.103.200.35
```