CS 168 Fall 2022 — Introduction to the Internet: Architecture and Protocols — Sylvia Ratnasamy

Discussion 2: Packets in Flight

1 True or False

(1) On a fast cross-continental link ($\approx$ 100Gbps), *propagation delay* usually dominates *end-to-end packet delay* (Most messages are smaller than 100MB).

Solution: True. On a 100Gbps link, even a 100MB file download would only take 0.008 seconds to get on the wire, compared to 0.02 seconds propagation delay from New York to London (in the best case). Most communications (web page, emails) don't come close to this size.

(2) On the same cross-continental link ($\approx$ 100Gbps), when transferring a 100GB file, *propagation delay* still dominates end-to-end file delivery.

Solution: False. Sending a 100GB file over a 100Gbps link will have at least 8 seconds transmission delay.

(3) On-demand circuit-switching is adopted by the Internet.

Solution: False. Circuit-switching shares bandwidth through reservation. Packet-switching shares bandwidth on demand. Packet-switching is adopted by the Internet.

(4) The aggregate (i.e., sum) of peaks is usually much larger than peak of aggregates in terms of bandwidth usage.

Solution: True. Statistical multiplexing leverages this to use available scarce resources more effectively.

(5) Bursty traffic (i.e., when packet arrivals are not evenly spaced in time) always leads to queuing delays.

Solution: False. Queuing delay happens when arrival rate is larger than transmission rate (ignoring processing delays). Bursty traffic does not necessarily imply arrival rate is larger than transmission rate. Queuing delay depends on traffic patterns, router internals, and link properties.

2 End-to-End Delay

Consider the diagram below. Link 1 has length L_1 m (where m stands for meters) and allows packets to be propagated at speed $S_1 \frac{\text{m}}{\text{sec}}$, while Link 2 has length L_2 m but it only allows packets to be propagated at speed $S_2 \frac{\text{m}}{\text{sec}}$ (because two links are made of different materials). Link 1 has transmission rate $T_1 \frac{\text{bits}}{\text{sec}}$ and Link 2 has transmission rate $T_2 \frac{\text{bits}}{\text{sec}}$.

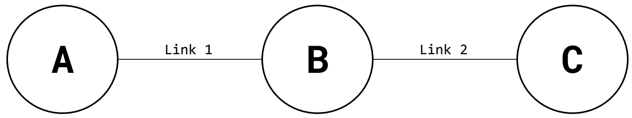

Assuming nodes can send and receive bits at full rate and ignoring processing delay, consider the following scenarios:

(1) How long would it take to send a packet of 500 Bytes from Node A to Node B given $T_1 = 10000$, $L_1 = 100000$, and $S_1 = 2.5 \cdot 10^8$?

Solution: The total time needed is the sum of the transmission delay to push the packet onto Link 1 and the propagation delay for the packet to travel from Node A to Node B.

$$\begin{aligned} t_{\text{total}} &= t_{\text{transmission}} + t_{\text{propagation}} \\ &= \frac{\text{packet size}}{\text{transmission rate of Link 1}} + \frac{\text{distance between A and B}}{\text{propagation speed}} \\ &= \frac{500B \cdot 8\frac{b}{B}}{10000\frac{b}{s}} + \frac{100000m}{2.5 \cdot 10^8 \frac{m}{s}} \\ &= 0.4s + 0.0004s = \boxed{0.4004s} \end{aligned}$$

Notice that transmission delay dominates more than 99.9% in this case.

(2) Compute RTT (round trip time) for a packet of B Bytes sent from Node A to Node C (packet gets transmitted back from Node C immediately after Node C receives it).

Solution: There is only one packet so no need to worry about queuing delays. Consider the diagram below:

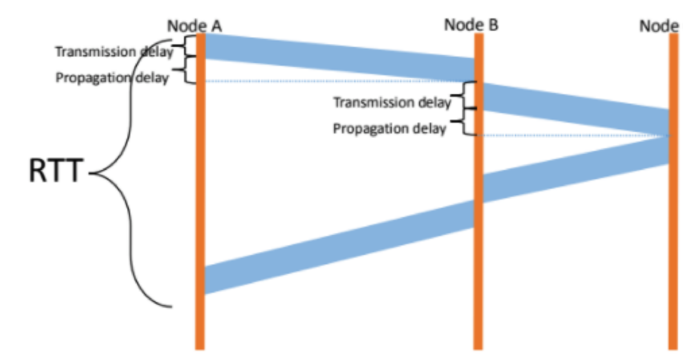

Note the sequence of delays the packet experiences during its route from A to C:

1. Transmission delay to push the packet onto Link 1.
2. Propagation delay as the packet travels from Node A to Node B.
3. Transmission delay to push the packet onto Link 2.
4. Propagation delay as the packet travels from Node B to Node C.
5. Transmission delay to push the packet onto Link 2.
6. Propagation delay as the packet travels from Node C to Node B.

7. Transmission delay to push the packet onto Link 1.
8. Propagation delay as the packet travels from Node B to Node A.

Summing these delays yields the total RTT:

$$\boxed{\text{RTT} = \frac{8B}{T_1} + \frac{L_1}{S_1} + \frac{8B}{T_2} + \frac{L_2}{S_2} + \frac{8B}{T_2} + \frac{L_2}{S_2} + \frac{8B}{T_1} + \frac{L_1}{S_1}}$$

(3) At time 0, Node A sends packet P_1 with D_1 Bytes and then it sends another packet P_2 with D_2 Bytes immediately after it pushes all bits of P_1 onto Link 1. When will Node C receive the last bit of P_2?

Solution: There are two packets, and so we might need to consider queueing delays. There will be a queueing delay at Node B if P_2 arrives at B before P_1 is finished being pushed onto Link 2.

Let's start by computing the time at which P_1 finishes being pushed onto Link 2. P_1 takes $\frac{8D_1}{T_1}$ seconds to be pushed onto Link 1, $\frac{L_1}{S_1}$ seconds to propagate from Node A to Node B, and then $\frac{8D_1}{T_2}$ seconds to be pushed onto Link 2. Hence P_1 leaves Node B at time

$$t_1 = \frac{8D_1}{T_1} + \frac{L_1}{S_1} + \frac{8D_1}{T_2}$$

Next, let's figure out the time when P_2 arrives at Node B. It first waits $\frac{8D_1}{T_1}$ seconds for P_1 to be completely pushed onto Link 1, then takes $\frac{8D_2}{T_1}$ seconds of transmission delay to be pushed onto Link 1 itself, before finally needing $\frac{L_1}{S_1}$ seconds of propagation delay to reach Node B. With this, we know that P_2 reaches Node B at time

$$t_2 = \frac{8D_1}{T_1} + \frac{8D_2}{T_1} + \frac{L_1}{S_1}$$

There's queueing delay if $t_1 > t_2$, and the length of the delay can be expressed as

$$t_1 - t_2 = \left(\frac{8D_1}{T_1} + \frac{L_1}{S_1} + \frac{8D_1}{T_2}\right) - \left(\frac{8D_1}{T_1} + \frac{8D_2}{T_1} + \frac{L_1}{S_1}\right) = \frac{8D_1}{T_2} - \frac{8D_2}{T_1}$$

With this analysis in hand, we can express the time at which Node C receives the last bit of P_2 as follows:

$$\boxed{t_{\text{total}} = \frac{8D_1}{T_1} + \frac{8D_2}{T_1} + \frac{L_1}{S_1} + \max\left(\frac{8D_1}{T_2} - \frac{8D_2}{T_1}, 0\right) + \frac{8D_2}{T_2} + \frac{L_2}{S_2}}$$

From left to right, the terms in this sum are:

1. The transmission delay to push P_1 onto Link 1.
2. The transmission delay to push P_2 onto Link 1.
3. The propagation delay as P_2 travels from Node A to Node B.
4. The queueing delay at Node B. Note that the use of the max operator allows us to express the two cases when there is and when there isn't queueing delay compactly.
5. The transmission delay to push P_2 onto Link 2.
6. The propagation delay as P_2 travels from Node B to Node C.

Below is the time-graph of a packet in flight without queuing delay:

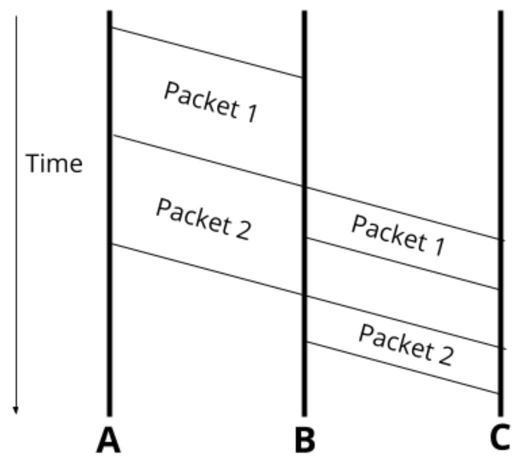

And with queuing delay:

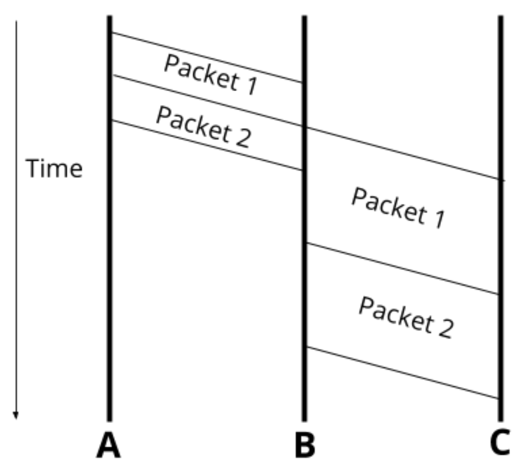

(4) Find the variable relations that need to be satisfied in order to have no queuing delays for part (c).

Solution: From the analysis we conducted in part (c), we know there will be queueing delays if $t_1 > t_2$, or $\frac{8D_1}{T_2} > \frac{8D_2}{T_1}$. Hence, there's no queueing delays if $\frac{8D_1}{T_2} \leq \frac{8D_2}{T_1}$. After simplifying it, we see the relation that must be satisfied is

$$\boxed{\frac{D_1}{T_2} \leq \frac{D_2}{T_1}}$$

3 First one to TikTok wins!

As you may know, the Washington Post has a pretty spicy TikTok account. Execs at the New York Times, the Wall Street Journal, and USA Today have also noticed the Post's success and want to promote their brand on the platform. Each organization has filmed a take on the 9 to 5 challenge video to use as their first upload to TikTok. They are all waiting for the perfect moment to post.

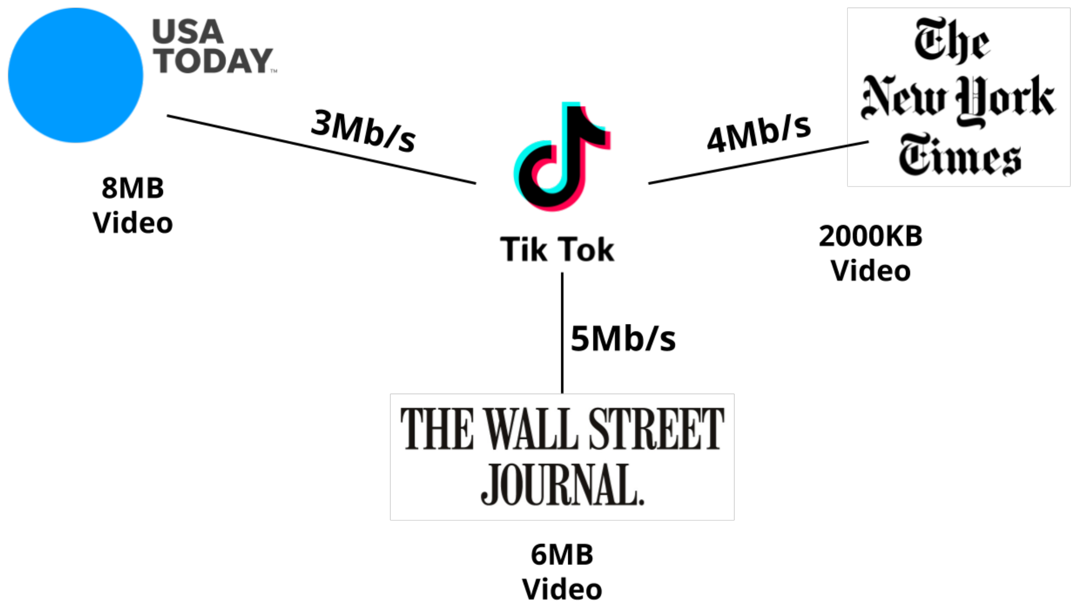

Upon seeing the perfect opportunity, all three organizations begin uploading their video at almost the same time (within 3-seconds of each other). If we assume that propagation time is negligible, which news organization will be the first to publish their video (and ultimately become #1 trending on TikTok)?

Solution: Let's first compute the transmission time for the videos uploaded:

The NYTime's video is 2000KB and takes

$$\frac{2000\text{KB}}{1000\frac{\text{KB}}{MB}} \cdot 8\frac{\text{bits}}{\text{byte}} \div 4\frac{Mb}{s} = \boxed{4.0 \text{ seconds}}$$

to upload the entire video.

The USA Today's video would need around

$$8\text{MB} \cdot 8\frac{\text{bits}}{\text{byte}} \div 3\frac{Mb}{s} = \boxed{21.3 \text{ seconds}}$$

to be put onto the link.

The WSJ's video would need around

$$6\text{MB} \cdot 8\frac{\text{bits}}{\text{byte}} \div 5\frac{Mb}{s} = \boxed{9.6 \text{ seconds}}$$

to be put onto the link.

The NYTimes video is the first uploaded because it's time to upload is 5.6 seconds faster than the next fastest upload from the WSJ (which is greater than the 3-second range of uploading).

4 Statistical Multi-What?

Consider three flows (F_1, F_2, F_3) sending packets over a single link. The sending pattern of each flow is described by how many packets it sends within each one-second interval; the table below shows these numbers for the first ten intervals. A perfectly smooth (i.e., non-bursty) flow would send the same number of packets in each interval, but our three flows are very bursty, with highly varying numbers of packets in each interval:

Time (s)	1	2	3	4	5	6	7	8	9	10
F_1	1	8	3	15	2	1	1	34	3	4
F_2	6	2	5	5	7	40	21	3	34	5
F_3	45	34	15	5	7	9	21	5	3	34

(Questions on next page)

(1) What is the peak rate of F_1? F_2? F_3? What is the sum of the peak rates?

Solution: The peak rate is the highest the flow gets throughout the whole period. The peak rate of F_1 34, the peak rate of F_2 is 40, and the peak rate of F_3 is 45.

The sum of their peaks is 34 + 40 + 45 = 119.

(2) Now consider all packets to be in the same aggregate flow. What is the peak rate of this aggregate flow?

Solution: Summing the flows together, we get the following values for an aggregate flow:

Time (s)	1	2	3	4	5	6	7	8	9	10
Aggregate Flow	52	44	23	25	16	50	43	42	40	43

The peak of the aggregate flow happens at 1s, where it is 52.

(3) Which is higher - the sum of the peaks, or the peak of the aggregate?

Solution: The sum of the peaks is 119, whereas the peak of the aggregate is 52, so the sum of the peaks is much higher. This is the insight from Statistical Multiplexing! The peak of the aggregate can only be at most the sum of the peaks, but that only happens in the case that all of the peaks happen at the same time. This is very unlikely, so usually, the peak of the aggregate is much lower than the sum of the peaks.

5 Plenty of Packets

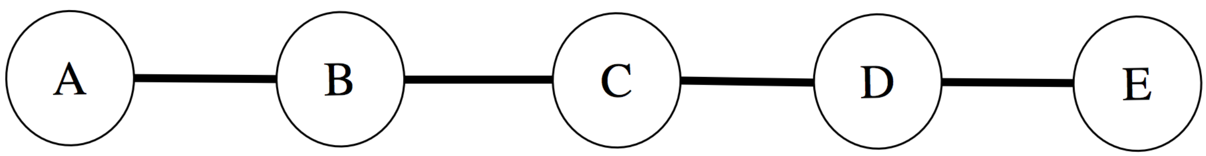

Let's suppose we have three packets (P_1, P_2, P_3) of size x, y, z bytes respectively. We want to send packet P_1 from A to C and packets P_2 and P_3 from C to E.

Use the following values for all calculations:

x	2000 bytes
y	500 bytes
z	1000 bytes
L_{AB}	100 km
L_{BC}	50 km
L_{CD}	75 km
L_{DE}	100 km
S_{AB}	$1.5 * 10^8 \frac{m}{s}$
S_{BC}	$2 * 10^8 \frac{m}{s}$
S_{CD}	$2 * 10^8 \frac{m}{s}$
S_{DE}	$1.75 * 10^8$
T_{AB}	$5000 \frac{b}{s}$
T_{BC}	$8000 \frac{b}{s}$
T_{CD}	$10000 \frac{b}{s}$
T_{DE}	$7000 \frac{b}{s}$

(1) Will P_1 reach its destination first or will P_2 and P_3 arrive before?

Solution: Let's first calculate the time it takes for P_1 to reach its destination. Since there's only one packet, we can disregard queuing delay. Then, we just need to add the transmission and propagation delays from each segment, A to B and B to C. Transmission delay is defined as the time it takes to push a packet onto the wire, which is $\frac{\text{packet size (bits)}}{\text{transmission rate}}$. Propagation delay is defined as the time it takes for the packet to be sent to its destination along the wire, which is $\frac{\text{distance between A and B}}{\text{propagation speed}}$.

First, P_1 takes $\frac{8x}{T_{AB}}$ to be pushed onto the AB link. Then, it takes $\frac{L_{AB}}{S_{AB}}$ to propagate from node A to node B. Similarly, it will take $\frac{8x}{T_{BC}}$ to be pushed onto the BC link. Then, it will take $\frac{L_{BC}}{S_{BC}}$ to be propagate from B to C. So, the time for P_1 to reach its destination is

$$\frac{8x}{T_{AB}} + \frac{L_{AB}}{S_{AB}} * \frac{1000m}{1km} + \frac{8x}{T_{BC}} + \frac{L_{BC}}{S_{BC}} * \frac{1000m}{1km} = 5.20s.$$

Now, we will calculate the time it takes for packets P_2 and P_3 to reach their destination. Because there are two packets, there may be a queuing delay. There will be a queuing delay if the time it takes for P_3 to arrive at node D, t_3, is shorter than the time it takes for P_2 to be pushed onto the link between D and E, t_2.

Let's start with finding the amount of time it takes for P_2 to be pushed onto the link between D and E. First, we need to push P_2 onto the wire, which takes $\frac{8y}{T_{CD}}$. Then, it takes $\frac{L_{CD}}{S_{CD}}$ for the packet to arrive at D. It takes an additional $\frac{8y}{T_{DE}}$ for P_2 to be pushed onto the wire between D and E. So, $t_2 = \frac{8y}{T_{CD}} + \frac{L_{CD}}{S_{CD}} + \frac{8y}{T_{DE}}$.

Next, we will calculate the time it takes for P_3 to reach D. First, it has to wait for P_2 to be pushed onto the link before pushing P_3, which will take $\frac{8y}{T_{CD}}$. Then, we can push P_3 onto the wire in $\frac{8z}{T_{CD}}$. It will take $\frac{L_{CD}}{S_{CD}}$ for the packet to arrive at D. So, $t_3 = \frac{8y}{T_{CD}} + \frac{8z}{T_{CD}} + \frac{L_{CD}}{S_{CD}}$.

Now, we can compare t_2 and t_3. We will have a queuing delay if the following is true:

$$\begin{gathered} t_2 > t_3 \\ \frac{8y}{T_{CD}} + \frac{L_{CD}}{S_{CD}} * \frac{1000m}{1km} + \frac{8y}{T_{DE}} > \frac{8y}{T_{CD}} + \frac{8z}{T_{CD}} + \frac{L_{CD}}{S_{CD}} * \frac{1000m}{1km} \\ \frac{8y}{T_{DE}} > \frac{8z}{T_{CD}} \\ 0.571428571 > 0.8 \end{gathered}$$

Plugging in the values from the problem statement, we see that $t_2 < t_3$, so there is no queuing delay. We can now calculate the time it takes for P_2 and P_3 to reach their destination.

We will push P_2 and P_3 onto the link between C and D in $\frac{8y}{T_{CD}}$ and $\frac{8z}{T_{CD}}$ time respectively. It will take $\frac{L_{CD}}{S_{CD}}$ time for P_3 to propagate to D. Then, the transmission delay for P_3 to be pushed onto the link between D and E is $\frac{8z}{T_{DE}}$. The propagation delay for P_3 is $\frac{L_{DE}}{S_{DE}}$. So, the total amount of time it takes to transmit both packets will be (it is also in the diagram at the end):

$$\frac{8y}{T_{CD}} + \frac{8z}{T_{CD}} + \frac{L_{CD}}{S_{CD}} * \frac{1000m}{1km} + \frac{8z}{T_{DE}} + \frac{L_{DE}}{S_{DE}} * \frac{1000m}{1km} = 2.34s.$$

We see that P_2 and P_3 will arrive at their destination before P_1.

(2) How small would P_1 have to be for it to take longer to send the other two packets?

Solution: To solve this, we will set the time it takes for P_1 to reach its destination to be less than the time we calculated for the other two packets in the first part.

$$\frac{8x}{T_{AB}} + \frac{L_{AB}}{S_{AB}} * \frac{1000m}{1km} + \frac{8x}{T_{BC}} + \frac{L_{BC}}{S_{BC}} * \frac{1000m}{1km} < 2.34$$

$$x < 901.11$$

So, P_1 can be at most 901 bytes for it to arrive at its destination faster than the other two packets.