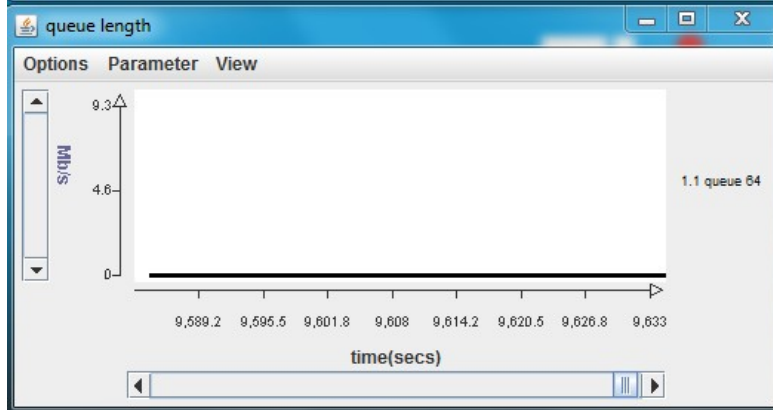
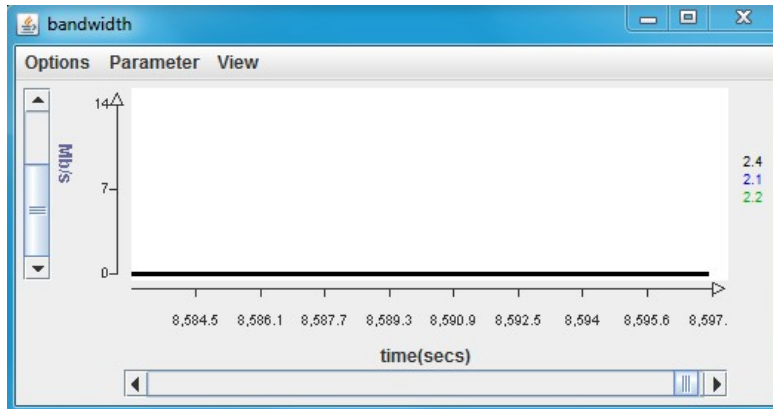
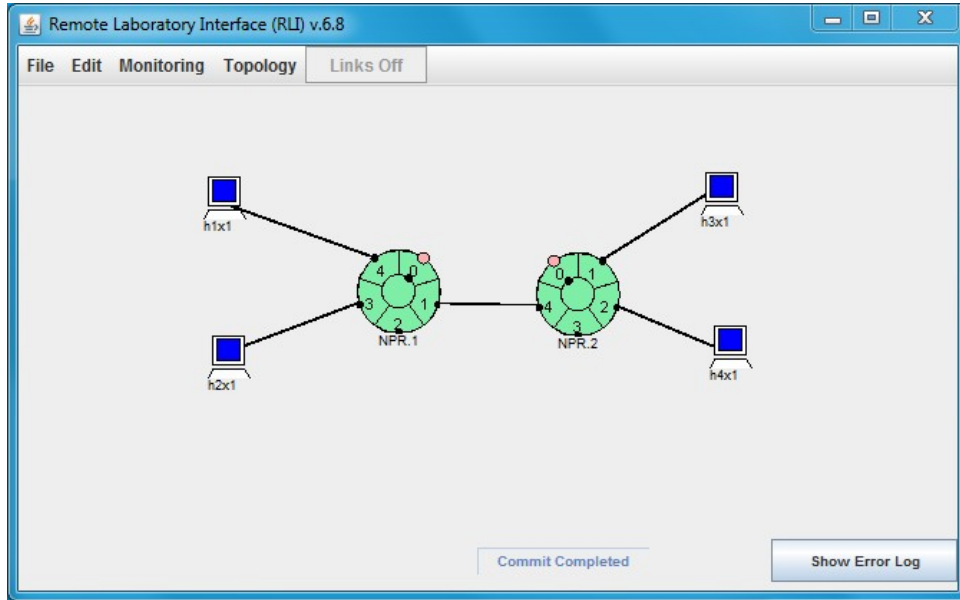


# ECE 671 – Computer Networks

## Fall 2011 – Lab 2-Solution

### Topology



## Part 1: Single TCP Connection

Parameters of TCP connection as follows:

- Window size: **64kB**
- Delay: **50ms**
- Link rate (between two routers): **10Mbps and 100Mbps**
- **Link rate 10Mbps**

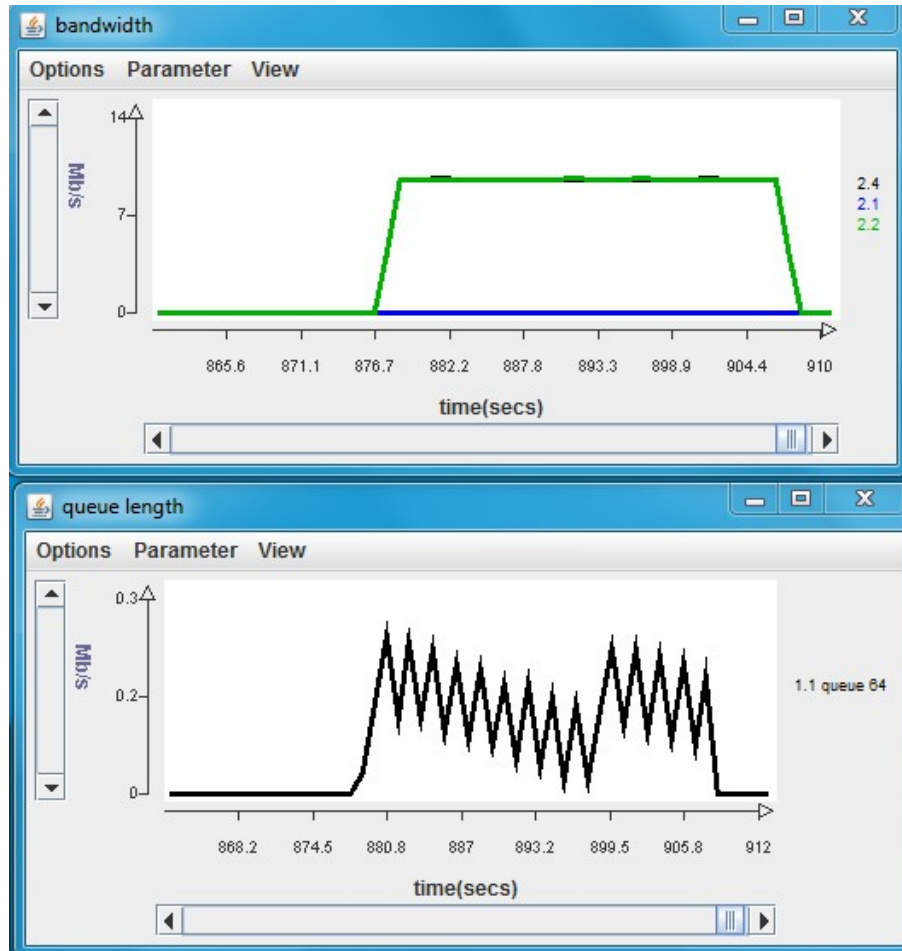


Figure 1. Bandwidth and Queue Length of single TCP, link rate 10Mbps

- Link rate 100Mbps

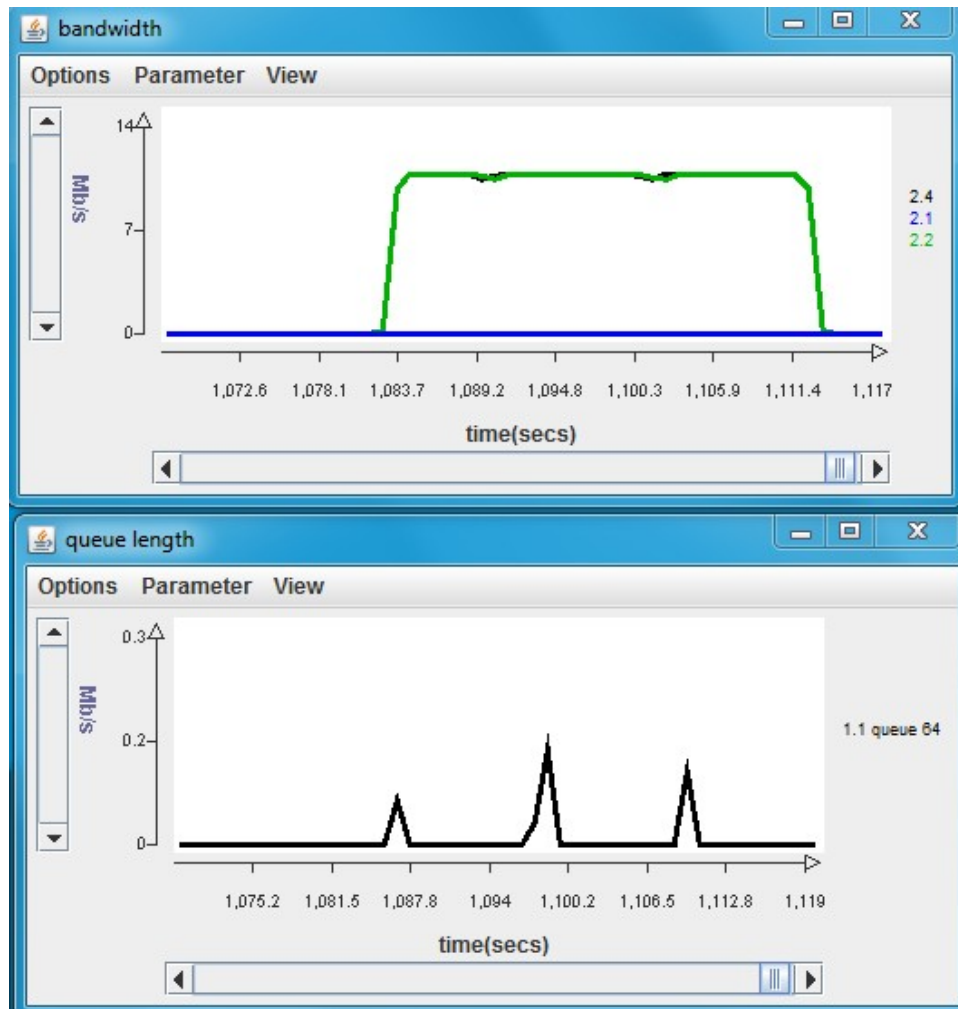


Figure 2. Bandwidth and Queue Length of single TCP, link rate 100Mbps

Key point of explanation:

- In Figure1, the queue chart shows that queue 64 has packets during the congestion period [878, 908] when both flows are active. The maximum value of queue length is approach 0.3MB. The minimum value of queue length is approach 0.03MB. The throughput is around 10.2Mbps.
- There is “Saw tooth” showed in the queue length. The reason of appeared “saw tooth” is due to the TCP congestion-control algorithm. You should illustrate the TCP congestion-control algorithm in detail. The algorithm has three major components: 1) additive-increase. multiplicative-decrease (AIMD), 2) slow start, and 3) reaction to timeout events.

When a TCP connection begins, the value of congestion window is typically initialized to 1MSS, resulting in an initial sending rate of roughly  $MSS/RTT$ . TCP sender increases its rate exponentially by doubling its vale of congestion window

every RTT, increases its sending rate exponentially. When the congestion window size achieve to the threshold, the slow start will end and congestion avoidance will begin. After TCP enters the congestion avoidance phase, the congestion window size ramps up linearly. Whenever a loss event occurs, the value of Threshold will decrease to one half of the current value of congestion window size and the new congestion window size decrease to the value of threshold. TCP enters the congestion avoidance phase when packet loss occurs. The queue length counter is read at the polling period, so there is “Saw tooth” showed in the graph.

- In Figure 2, when the link rate is 100Mbps, there is no “Saw tooth ” showed in the queue length graph. The throughput is much smaller than the maximum bandwidth, so there is almost no packet in the buffer queue.
- When the link rate increase to 100Mbps, the throughput could only achieve around 10.2 Mbps. Because the RTT is 50ms, the time need to send 64KB data is

$$t = \frac{64KB}{100Mbps} = \frac{64 * 1000 * 8bits}{100Mbps} = 5.12ms$$

, which is smaller than 50ms, so the throughput is

$$throughput = \frac{window\ size\ 64kB}{RTT} = \frac{64 * 1000 * 8bits}{50ms} = 10.24Mbps$$

## Part 2: Two TCP connections

Parameters of TCP connection as follows:

- Window size: **64kB**
- Delay: **50ms**
- Link rate (between two routers): **1Mbps, 10Mbps and 100Mbps**

## 1. Link rate 1Mbps

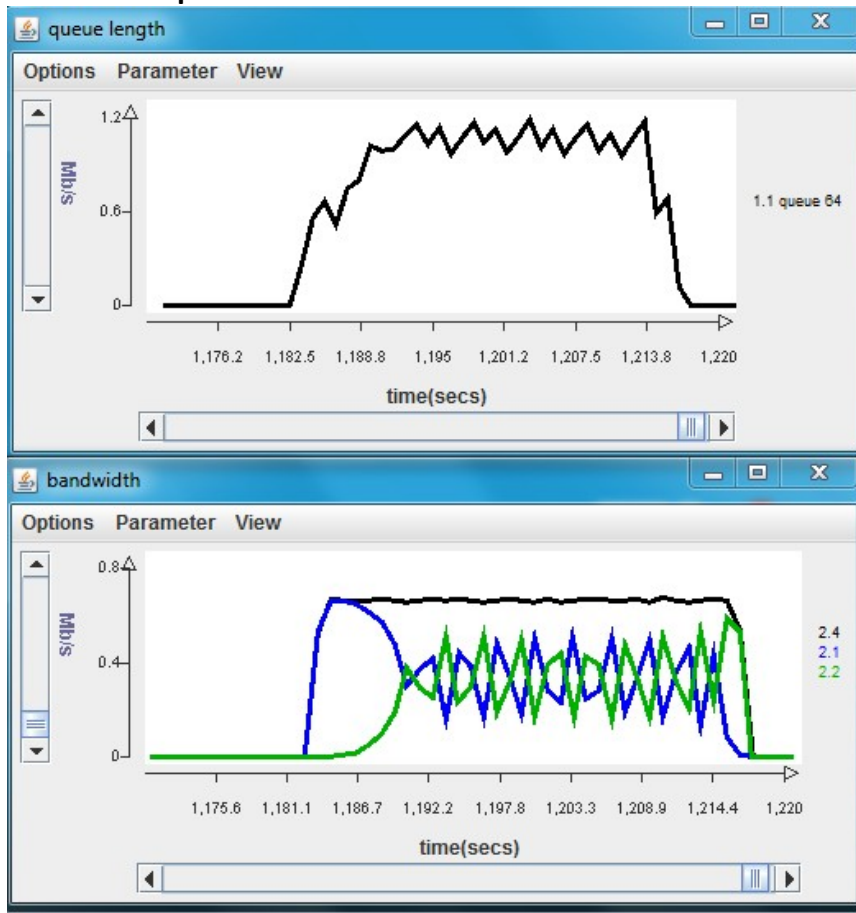


Figure 3. Bandwidth and Queue Length of two TCP connections, link rate 1Mbps

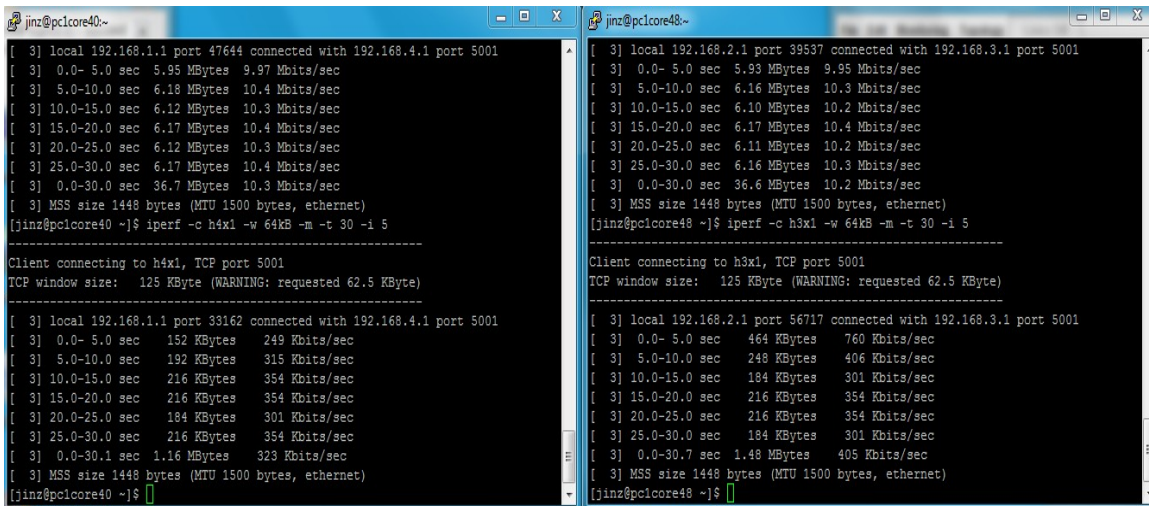


Figure 4. Bandwidth of two separate clients in TCP connection, link rate 1Mbps

## 2. Link rate 10Mbps

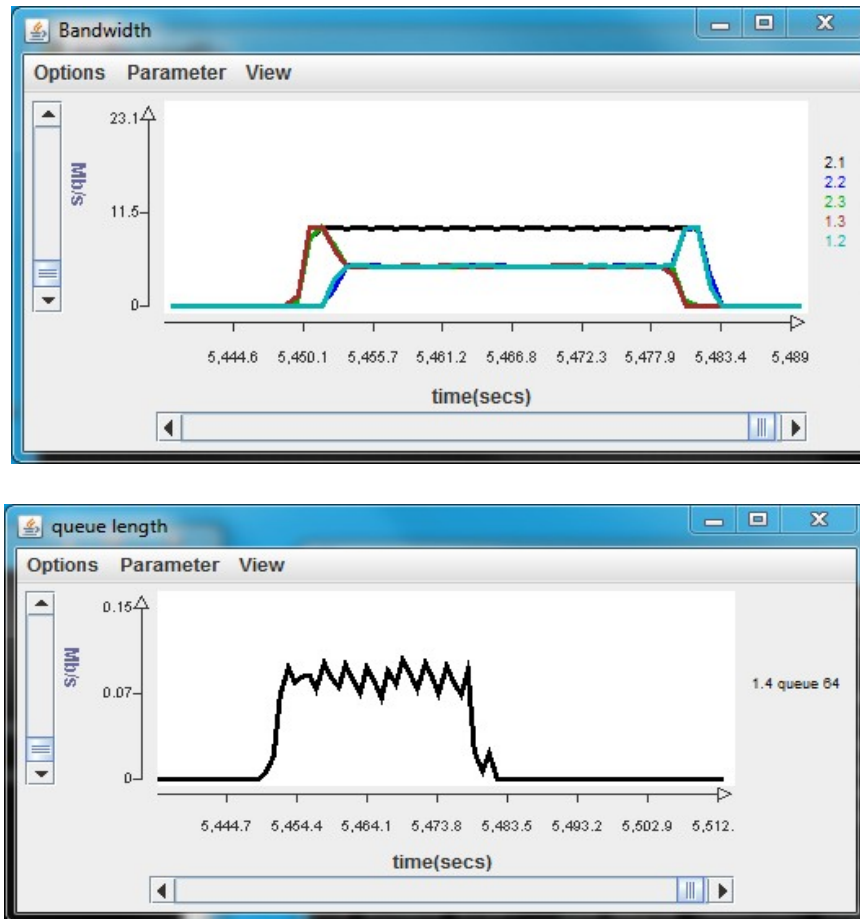


Figure 5. Bandwidth and Queue Length of two TCP connections, link rate 10Mbps

```
jinz@pclcore61:~$ iperf -c h5x2 -w 64kB -t 30 -i 5
-----
TCP window size: 125 KByte (WARNING: requested 62.5 KByte)
-----
[ 3] local 192.168.1.1 port 41254 connected with 192.168.5.2 port 5001
[ 3] 0.0- 5.0 sec 4.05 MBytes 6.79 Mbits/sec
[ 3] 5.0-10.0 sec 2.76 MBytes 4.63 Mbits/sec
[ 3] 10.0-15.0 sec 2.75 MBytes 4.61 Mbits/sec
[ 3] 15.0-20.0 sec 2.76 MBytes 4.63 Mbits/sec
[ 3] 20.0-25.0 sec 2.76 MBytes 4.63 Mbits/sec
[ 3] 25.0-30.0 sec 2.75 MBytes 4.61 Mbits/sec
[ 3] 0.0-30.0 sec 17.8 MBytes 4.98 Mbits/sec
[jinz@pclcore61 ~]$ iperf -c h5x2 -w 64kB -t 30 -i 5
-----
Client connecting to h5x2, TCP port 5001
TCP window size: 125 KByte (WARNING: requested 62.5 KByte)
-----
[ 3] local 192.168.1.1 port 41255 connected with 192.168.5.2 port 5001
[ 3] 0.0- 5.0 sec 4.05 MBytes 6.79 Mbits/sec
[ 3] 5.0-10.0 sec 2.73 MBytes 4.57 Mbits/sec
[ 3] 10.0-15.0 sec 2.75 MBytes 4.61 Mbits/sec
[ 3] 15.0-20.0 sec 2.76 MBytes 4.63 Mbits/sec
[ 3] 20.0-25.0 sec 2.76 MBytes 4.63 Mbits/sec
[ 3] 25.0-30.0 sec 2.75 MBytes 4.61 Mbits/sec
[ 3] 0.0-30.0 sec 17.8 MBytes 4.97 Mbits/sec
[jinz@pclcore61 ~]$
```

```

jinz@pc1core51:~$ iperf -c h4x2 -w 64kB -t 30 -i 5
TCP window size: 125 KByte (WARNING: requested 62.5 KByte)
-----
[ 3] local 192.168.2.1 port 59694 connected with 192.168.4.2 port 5001
[ 3] 0.0- 5.0 sec 2.70 MBytes 4.52 Mbits/sec
[ 3] 5.0-10.0 sec 2.76 MBytes 4.63 Mbits/sec
[ 3] 10.0-15.0 sec 2.76 MBytes 4.63 Mbits/sec
[ 3] 15.0-20.0 sec 2.75 MBytes 4.61 Mbits/sec
[ 3] 20.0-25.0 sec 2.76 MBytes 4.63 Mbits/sec
[ 3] 25.0-30.0 sec 4.02 MBytes 6.74 Mbits/sec
[ 3] 0.0-30.0 sec 17.7 MBytes 4.96 Mbits/sec
[jinz@pc1core51 ~]$ iperf -c h4x2 -w 64kB -t 30 -i 5
-----
Client connecting to h4x2, TCP port 5001
TCP window size: 125 KByte (WARNING: requested 62.5 KByte)
-----
[ 3] local 192.168.2.1 port 59695 connected with 192.168.4.2 port 5001
[ 3] 0.0- 5.0 sec 2.70 MBytes 4.52 Mbits/sec
[ 3] 5.0-10.0 sec 2.76 MBytes 4.63 Mbits/sec
[ 3] 10.0-15.0 sec 2.76 MBytes 4.63 Mbits/sec
[ 3] 15.0-20.0 sec 2.75 MBytes 4.61 Mbits/sec
[ 3] 20.0-25.0 sec 2.76 MBytes 4.63 Mbits/sec
[ 3] 25.0-30.0 sec 3.98 MBytes 6.68 Mbits/sec
[ 3] 0.0-30.0 sec 17.7 MBytes 4.95 Mbits/sec
[jinz@pc1core51 ~]$

```

Figure 6. Bandwidth of two separate clients in TCP connection, link rate 10Mbps

### 3. Link rate 100Mbps

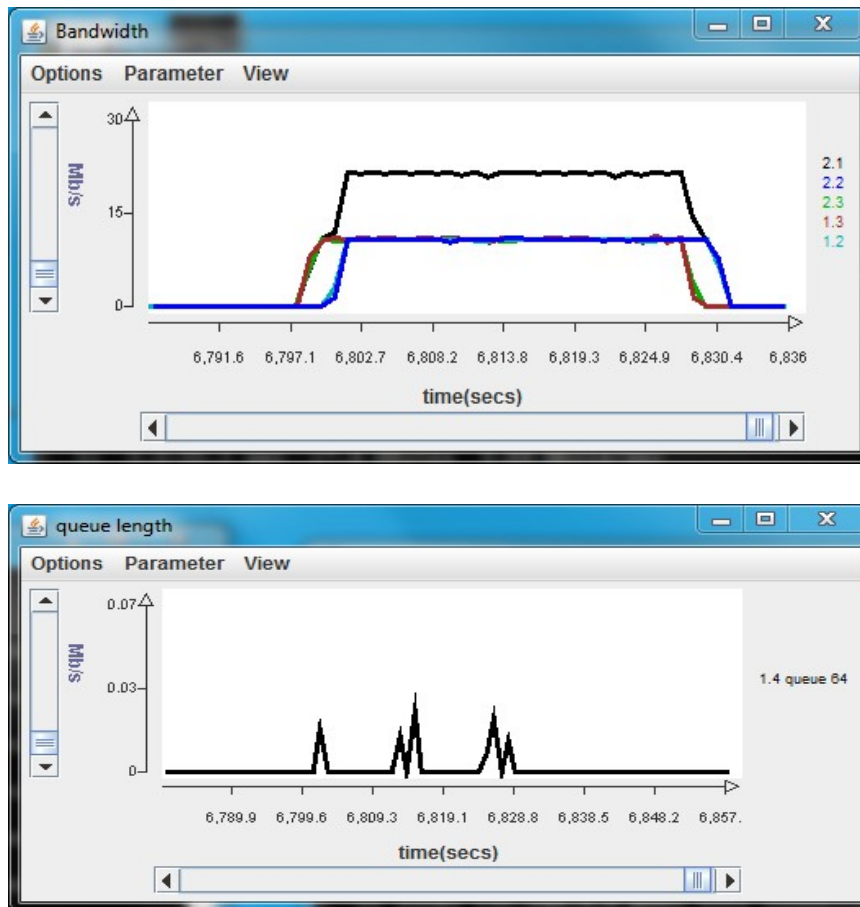


Figure 7. Bandwidth and Queue Length of two TCP connections, link rate 100Mbps

```

jinz@pclcore61:~
TCP window size: 125 KByte (WARNING: requested 62.5 KByte)
-----
[ 3] local 192.168.1.1 port 37654 connected with 192.168.5.2 port 5001
[ 3] 0.0- 5.0 sec 5.93 MBytes 9.95 Mbits/sec
[ 3] 5.0-10.0 sec 6.16 MBytes 10.3 Mbits/sec
[ 3] 10.0-15.0 sec 6.09 MBytes 10.2 Mbits/sec
[ 3] 15.0-20.0 sec 6.16 MBytes 10.3 Mbits/sec
[ 3] 20.0-25.0 sec 6.13 MBytes 10.3 Mbits/sec
[ 3] 25.0-30.0 sec 6.12 MBytes 10.3 Mbits/sec
[ 3] 0.0-30.0 sec 36.6 MBytes 10.2 Mbits/sec
[jinz@pclcore61 ~]$ iperf -c h5x2 -w 64kB -t 30 -i 5
-----
Client connecting to h5x2, TCP port 5001
TCP window size: 125 KByte (WARNING: requested 62.5 KByte)
-----
[ 3] local 192.168.1.1 port 37655 connected with 192.168.5.2 port 5001
[ 3] 0.0- 5.0 sec 5.93 MBytes 9.95 Mbits/sec
[ 3] 5.0-10.0 sec 6.16 MBytes 10.3 Mbits/sec
[ 3] 10.0-15.0 sec 6.10 MBytes 10.2 Mbits/sec
[ 3] 15.0-20.0 sec 6.16 MBytes 10.3 Mbits/sec
[ 3] 20.0-25.0 sec 6.15 MBytes 10.3 Mbits/sec
[ 3] 25.0-30.0 sec 6.09 MBytes 10.2 Mbits/sec
[ 3] 0.0-30.0 sec 36.6 MBytes 10.2 Mbits/sec
[jinz@pclcore61 ~]$

```

```

jinz@pclcore51:~
TCP window size: 125 KByte (WARNING: requested 62.5 KByte)
-----
[ 3] local 192.168.2.1 port 38475 connected with 192.168.4.2 port 5001
[ 3] 0.0- 5.0 sec 5.93 MBytes 9.95 Mbits/sec
[ 3] 5.0-10.0 sec 6.18 MBytes 10.4 Mbits/sec
[ 3] 10.0-15.0 sec 6.11 MBytes 10.2 Mbits/sec
[ 3] 15.0-20.0 sec 6.16 MBytes 10.3 Mbits/sec
[ 3] 20.0-25.0 sec 6.17 MBytes 10.4 Mbits/sec
[ 3] 25.0-30.0 sec 6.11 MBytes 10.2 Mbits/sec
[ 3] 0.0-30.0 sec 36.7 MBytes 10.2 Mbits/sec
[jinz@pclcore51 ~]$ iperf -c h4x2 -w 64kB -t 30 -i 5
-----
Client connecting to h4x2, TCP port 5001
TCP window size: 125 KByte (WARNING: requested 62.5 KByte)
-----
[ 3] local 192.168.2.1 port 38476 connected with 192.168.4.2 port 5001
[ 3] 0.0- 5.0 sec 5.93 MBytes 9.95 Mbits/sec
[ 3] 5.0-10.0 sec 6.15 MBytes 10.3 Mbits/sec
[ 3] 10.0-15.0 sec 6.11 MBytes 10.2 Mbits/sec
[ 3] 15.0-20.0 sec 6.16 MBytes 10.3 Mbits/sec
[ 3] 20.0-25.0 sec 6.14 MBytes 10.3 Mbits/sec
[ 3] 25.0-30.0 sec 6.13 MBytes 10.3 Mbits/sec
[ 3] 0.0-30.0 sec 36.6 MBytes 10.2 Mbits/sec
[jinz@pclcore51 ~]$

```

Figure 8. Bandwidth of two separate clients in TCP connection, link rate 100Mbps

Link Rate		1Mbps	10Mbps	100Mbps
Bandwidth	h1x1	323Kbps	4.97Mbps	10.2Mbps
	h2x1	405Kbps	4.96Mbps	10.2Mbps

The key point of explanation:

- There is bandwidth oscillation of two TCP connections when the link rate is 1Mbps and 10Mbps. When the link rate is 1Mbps, the bandwidth oscillation is around 0.3 Mbps. When the link rate is 10Mbps, the bandwidth oscillation is around 4Mbps. The oscillation synchronizes over time.
- The bandwidth oscillation due to the competition of two TCP connections. When the link rate is 1 Mbps or 10Mbps, one TCP connection (A) increase the bandwidth, the other one (B) will suffer from the packet loss, due to the TCP



congestion control algorithm, when there is loss event occurs, the value of threshold is set to one half of the current value of congestion window size, and the new congestion window size decrease to the value of threshold. Then the TCP enters the congestion avoidance phase, and the congestion window size ramps up linearly in the next step. This procedure loops again and again and makes the bandwidth oscillation.

- When the link rate set to 100Mbps, there is no competition of two TCP connections. But two TCP connections still could not achieve the maximum link rate, because the restriction of RTT.

### Part 3: UDP and TCP Competition

The parameter as follows:

Delay: **50ms**

Outgoing Link rate (link between two routers): **10Mbps, 100Mbps**

TCP connection: window size **64KB**

UDP connection: Bandwidth: **90Mbps**

1. UDP Bandwidth: 90Mbps, Link rate: 10Mbps

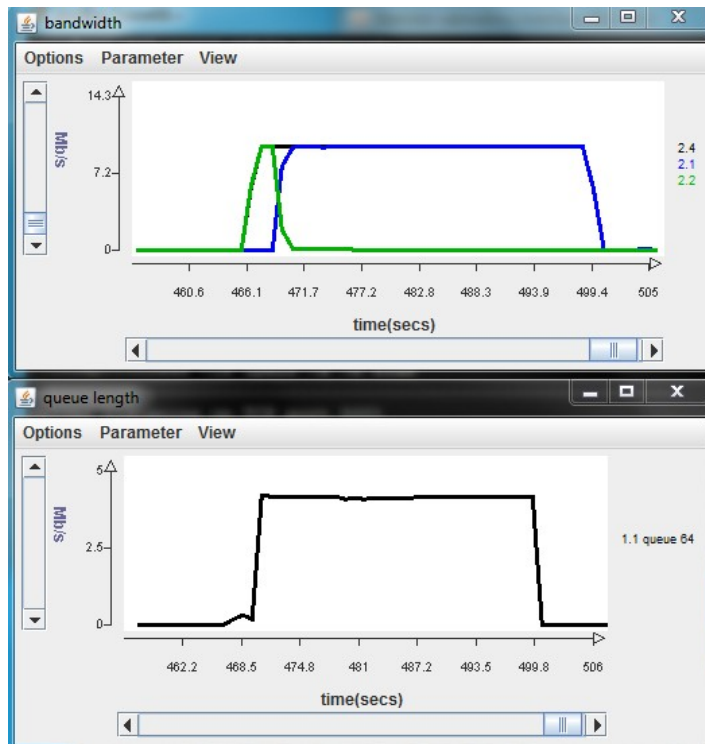


Figure 9 Bandwidth and Queue Length of TCP and UDP connections, link rate 10Mbps

```

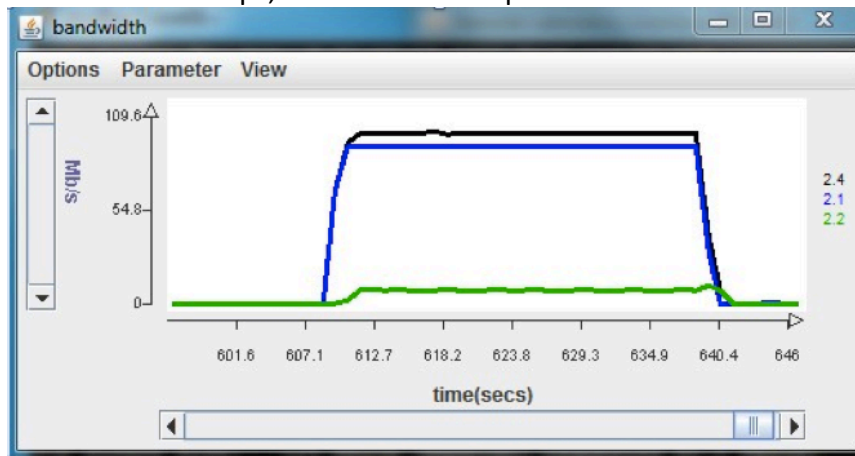
jinz@pclcore67:~$ iperf -c h4x1 -w 64kB -t 30 -i 5
-----
Client connecting to h4x1, TCP port 5001
TCP window size: 125 KByte (WARNING: requested 62.5 KByte)
-----
[ 3] local 192.168.1.1 port 39842 connected with 192.168.4.1 port 5001
[ 3] 0.0- 5.0 sec  96.0 KBytes  157 Kbits/sec
[ 3] 5.0-10.0 sec  0.00 Bytes  0.00 bits/sec
[ 3] 10.0-15.0 sec  0.00 Bytes  0.00 bits/sec
[ 3] 15.0-20.0 sec  0.00 Bytes  0.00 bits/sec
[ 3] 20.0-25.0 sec  0.00 Bytes  0.00 bits/sec
[ 3] 25.0-30.0 sec  0.00 Bytes  0.00 bits/sec
[ 3] 30.0-35.0 sec  0.00 Bytes  0.00 bits/sec
[ 3] 0.0-36.8 sec  104 KBytes  23.1 Kbits/sec
jinz@pclcore67 ~]$

jinz@pclcore64:~$ iperf -c h3x1 -u -b 90m -t 30 -i 5
-----
Client connecting to h3x1, UDP port 5001
Sending 1470 byte datagrams
UDP buffer size: 116 KByte (default)
-----
[ 3] local 192.168.2.1 port 56658 connected with 192.168.3.1 port 5001
[ 3] 0.0- 5.0 sec  53.9 MBytes  90.5 Mbits/sec
[ 3] 5.0-10.0 sec  53.9 MBytes  90.5 Mbits/sec
[ 3] 10.0-15.0 sec  53.9 MBytes  90.5 Mbits/sec
[ 3] 15.0-20.0 sec  53.9 MBytes  90.5 Mbits/sec
[ 3] 20.0-25.0 sec  53.9 MBytes  90.5 Mbits/sec
[ 3] 25.0-30.0 sec  53.9 MBytes  90.5 Mbits/sec
[ 3] 0.0-30.0 sec  324 MBytes  90.5 Mbits/sec
[ 3] Sent 230766 datagrams
[ 3] Server Report:
[ 3] 0.0-30.4 sec  34.1 MBytes  9.39 Mbits/sec  1.075 ms 206445/230757 (89%)
[ 3] 0.0-30.4 sec  1 datagrams received out-of-order
jinz@pclcore64 ~]$

```

Figure 10 Bandwidth of TCP and UDP connections, link rate 10Mbps

2. UDP Bandwidth: 90Mbps, Link rate: 100Mbps



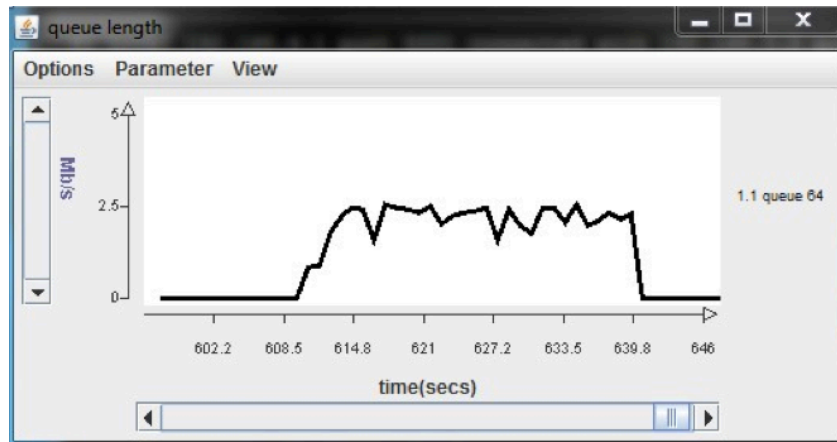


Figure 11 Bandwidth and Queue Length of TCP and UDP connections, link rate 100Mbps

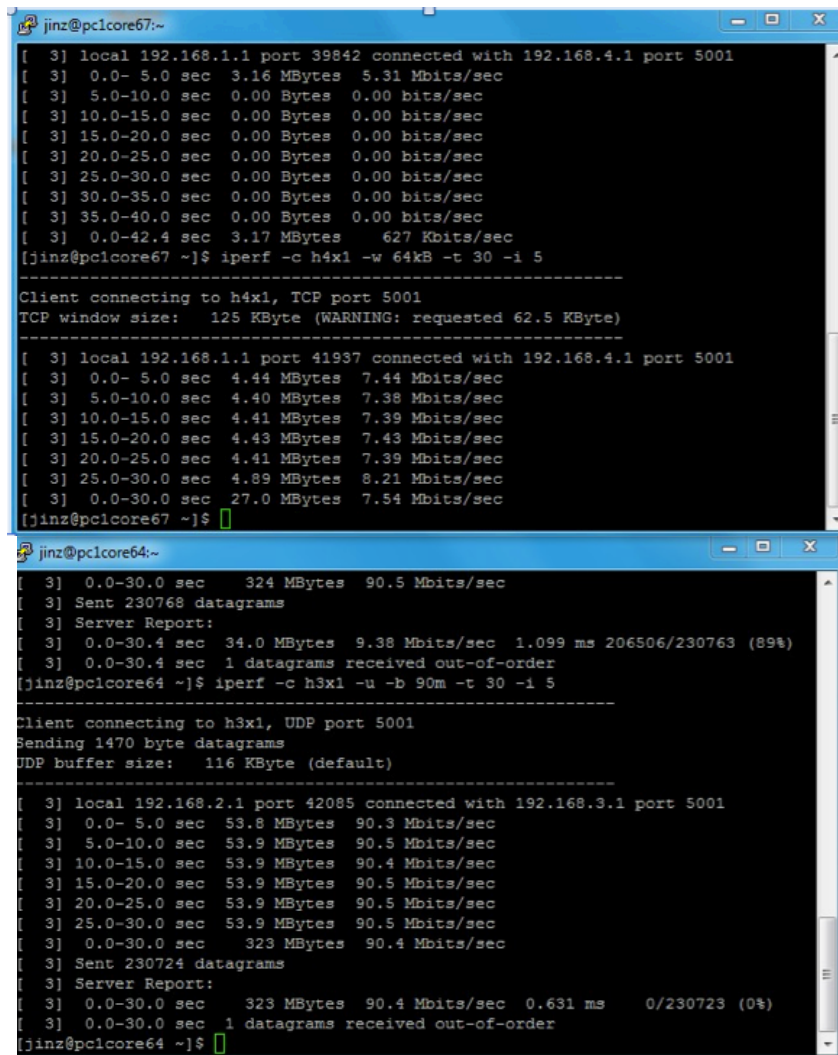


Figure 12 Bandwidth of TCP and UDP connections, link rate 100Mbps

Link rate	TCP	UDP (BW=90Mbps)
10Mbps	627kbps	9.38Mbps
100Mbps	7.39Mbps	90.4Mbps

The key point of explanation:

- UDP provides an unreliable, connectionless service to the invoking application without the congestion control and congestion detection. UDP will send the packets as much as possible, but could also result in high loss rates between a UDP sender and receiver. (here, the lose rate of packets is 89%)
- When TCP and UDP flows share a link, there is no fairness with TCP. In the UDP connection, there is no fall back of the loss packets, but for the TCP connection, the loss packets would occur the congestion control. So when the link rate is 10Mbps, the throughput of TCP connection almost tends to 0. When the link rate is 100Mbps, UDP occupy the most part of bandwidth resource and left 10Mbps to TCP connection.