

Lab8: Flow Fairness and Network Routing - Solutions

COM-208: Computer Networks

Objectives

- Study how competing flows share network resources.
- Experiment with flow characteristics to discover how each of these affects fairness.
- Examine how routing protocols react to changes to network connectivity.

Prerequisites:

- Lab 7 (Introduction to ns2 and Transport-layer protocols).
- Textbook Chapters 3 and 4.

Flow Fairness

Shared bottleneck link

In this part of the Lab, we study how competing TCP flows with similar characteristics behave when they share a single bottleneck link.

Setup

We will use the following files: `tp_fairness.tcl`, `fairness_pps.plot`, and `fairness_pkt.plot`. You can find them on Moodle in the scripts directory.

tp_fairness.tcl generates 5 source-destination pairs which all share a common network link. Each source uses a single TCP flow which transfers FTP traffic to the respective destination. The flows are created one after the other at 5-second intervals (i.e., flow $i \leftrightarrow i+1$ starts 5 seconds after flow i for i in $[1,4]$). You can invoke tp_fairness.tcl as follows:

```
$ ns tp_fairness.tcl
```

The figure below shows the resulting topology; there are 5 sources (2,4,6,8,10), 5 destinations (3,5,7,9,11), and each source is sending a large file to a single destination.




Figure 1: Topology

Output

The tp_fairness.tcl script produces one output file per flow; fairnessMon*[i]*.tr, for each $i \in [1,5]$. Each of these files contains three columns:

```
time | number of packets delivered so far | throughput (packets per second)
```

You can plot the following statistics for all flows using the plotting scripts provided:

- Throughput as a function of time:

```
$ gnuplot fairness_pps.plot
```

- Sum of packets delivered as a function of time:

```
$ gnuplot fairness_pkt.plot
```

If you want to display the NAM window (graphical interface), modify `tp_fairness.tcl` and uncomment the fifth line of the `finish` procedure:

```
proc finish {} {  
    global ns file1 file2  
    $ns flush-trace  
    close $file1  
    close $file2  
    # exec nam out.nam &  
    exit 0  
}
```

Questions

Run the `tp_fairness.tcl` script and observe the output of the `fairness_pps.plot` plotting script (explained above). Answer the following questions:

- Does each flow get an equal share of the capacity of the common link (i.e., is TCP fair)?
 - Explain which observations lead you to this conclusion.
- What happens to the throughput of the pre-existing TCP flows when a new flow is created?
 - Explain the mechanisms of TCP which contribute to this behavior.
 - Argue about whether you consider this behavior to be fair or unfair.

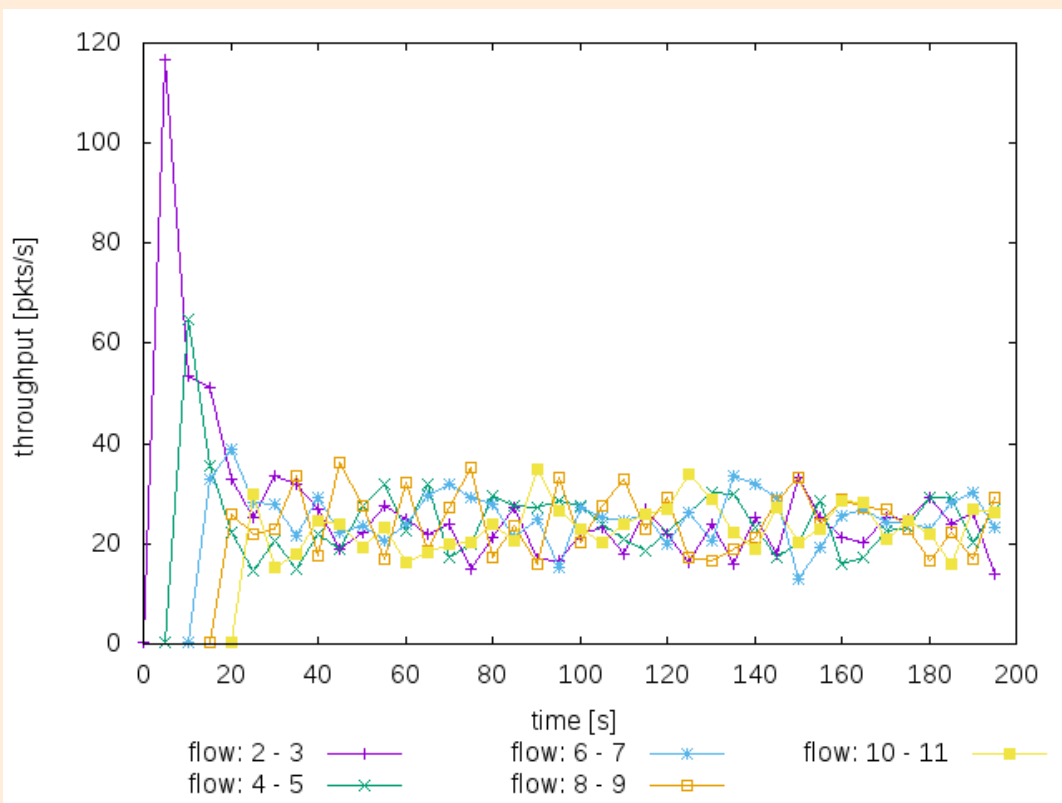


Figure 2

We observe that the throughput of the nodes decreases sharply whenever a new connection starts. It fluctuates a lot, but it is reasonably fair among all connections.

The mechanism responsible for this type of behavior is congestion control. Since all flows experience the same network conditions, they will react in the same manner.

Now, use the `fairness_pkt.plot` plotting script and answer the same set of questions. Does this plot confirm your intuitions?




Figure 3

The graph shows the total number of packets received at each destination as a function of time. The slopes of these curves roughly correspond to the throughput of these connections (number of packet received/time unit).

We can first point out that at any given point in time, all curves corresponding to existing flows have approximately the same slope in the graph. Thus all flows have, more or less, the same throughput and each gets a fair share of the bandwidth on the bottleneck link.

We can also notice that when a new flow enters the network, the throughput of all pre-existing flows is reduced. This happens because all TCP connections detect losses through duplicate ACKs and timeouts, and adapt the size of their congestion window in order to avoid overwhelming the network.

RTT and fairness.

In this part of the Lab, we will examine how two competing flows behave when they experience different round-trip times to their destination. (i.e., the RTT of one flow is longer than the RTT of the other flow).

Setup

We will use the following files: `tp_fairness2.tcl`, `fairness2_pps.plot`, and `fairness2_wnd.plot`. You can find the files on Moodle in scripts file.




Figure 4: Topology

The `tp_fairness2.tcl` script takes 4 arguments and generates the simple 3-node network topology shown in the figure above. The 4 arguments of the `ns` script are:

1. The capacity of link 0-1 (i.e., the link between node 0 and node 1).
2. The capacity of link 1-2.
3. The delay of link 0-1.
4. The delay of link 1-2.

Output

`tp_fairness2.tcl` also generates some TCP flows. You can plot the following statistics for all flows using the plotting scripts provided:

- Throughput as a function of time:
\$ gnuplot fairness2_pps.plot
- Size of TCP congestion window as a function of time:
\$ gnuplot fairness2_wnd.plot

Questions

Execute the `tp_fairness2.tcl` script as follows:

```
$ ns tp_fairness2.tcl 1Mb 1Mb 10ms 50ms
```

Observe the NAM window output and try to answer the following questions:

- Which nodes communicate with which other nodes? Hint: Different colors of NAM flows show different connections.
- Is there any flow that appears to starve another flow?

There are two TCP connections which originate from node 0. The first has node 2 as its destination and starts at time 0.5. The second one has node 1 as its destination and starts at time 1.5.

If the generated traffic exceeds the capacity of link 0-1, the flow from node 0 to node 1 appears to starve the other flow.

Observe the outputs of `fairness2_pps.plot` and `fairness2_wnd.plot` plotting scripts.

- How does the network delay of link 1-2 affect fairness?
- Why does the size of the TCP congestion window differ for the competing flows? Try to come up with a simple theory about why this happens. Hint: the link queue size is 10 packets and the maximum congestion window is 30 packets (see Line 9 and 10 of `tp_fairness2.tcl` script).
- What will happen if the link queue size becomes considerably larger than the maximum congestion window? Do you believe that this will affect fairness?

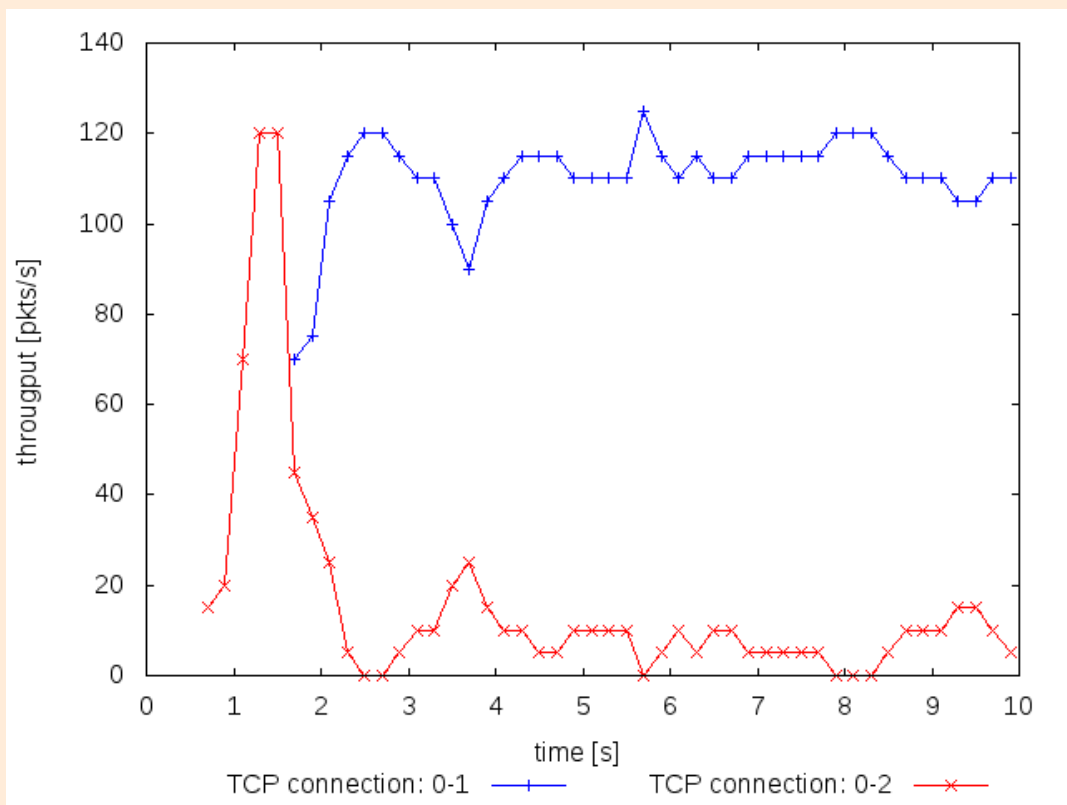


Figure 5

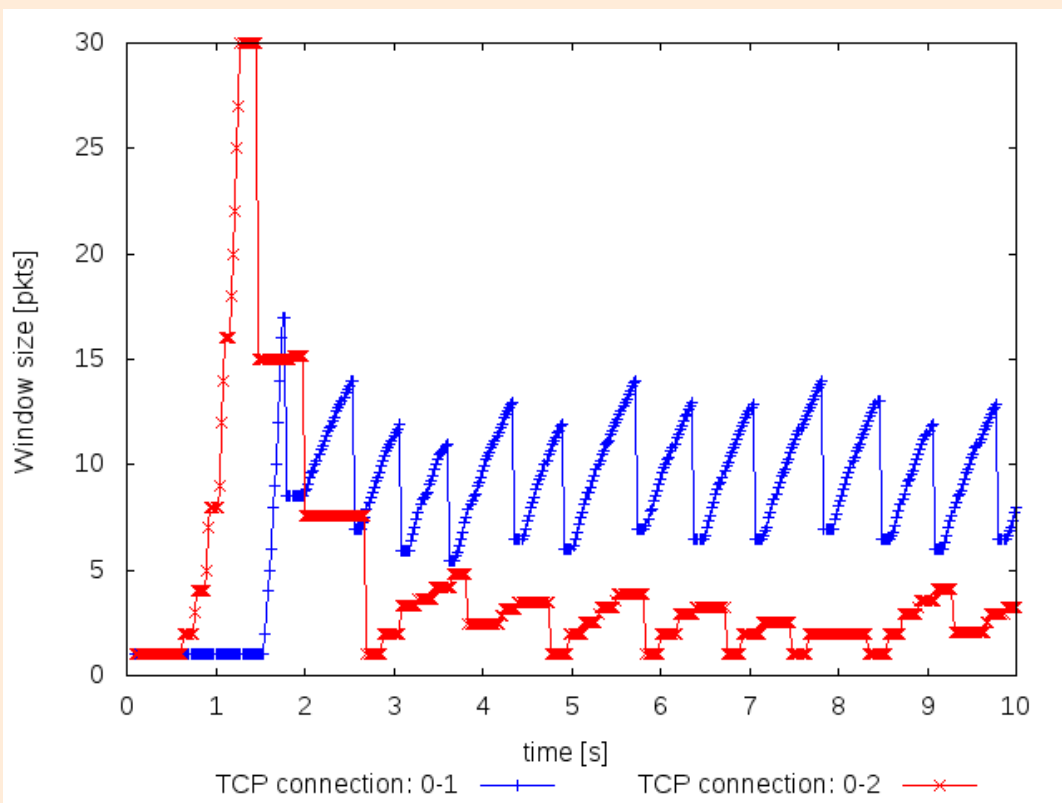


Figure 6

The congestion control algorithm of TCP can be very unfair for connections with significantly larger round-trip times. In our case, the slow-start procedure becomes the limiting factor for the connection with the longest delay. Node 0 receives the ACKs from node 2 at a slower rate due to the additional delay and, as such, the congestion window for the first flow grows at a slower pace.

The queue size is smaller than the sum of the maximum window size of the two flows. Hence, the congestion control algorithm adjusts the window sizes accordingly.

If queue size would be larger than the sum of the maximum window size of the two flows, then flows would share the link more fairly.

Now, it is time for you to verify your hypotheses. Make the following changes to the `tp_fairness2.tcl` file: increase the link queue size from 10 packets to 60 packets, and increase the simulation duration from 10 secs to 100 secs (the latter will give more

time for the flows to stabilize).

- Does this experiment confirm your intuition (i.e., did this modification affect fairness at all)?

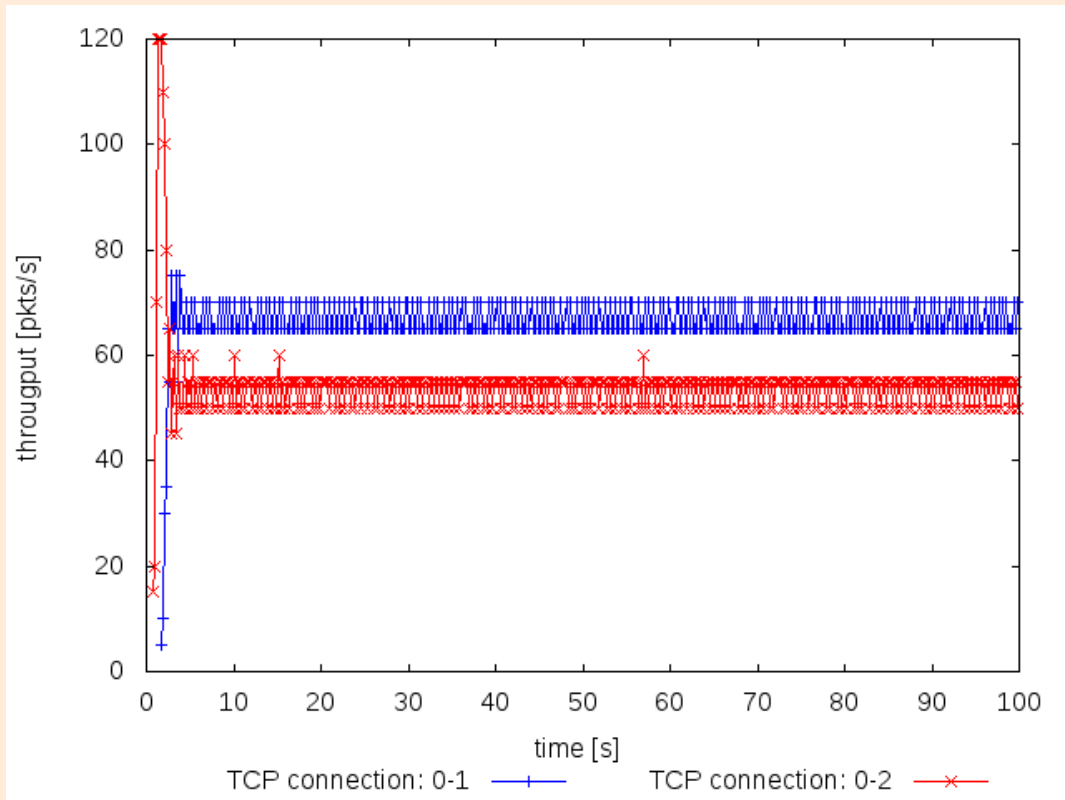


Figure 7




Figure 8

The queue size is bigger than the sum of the maximum possible values for the two congestion windows. Both connections can reach their maximum value after a few seconds, and no packets will be dropped. On the other hand, this should increase the round-trip time delay for both flows, since the queuing delay increases.

The difference in throughput between the two flows diminishes significantly. The reason why the throughput is not the same is that both flows are forced to use the same congestion window size. Thus, the flow with higher RTT value will achieve lower utilization than the one with lower RTT value.

TCP competing with UDP

In this part of the Lab, we are going to see how a TCP flow reacts when it has to share a bottleneck link that is also used by a UDP flow.

Setup

We will use the following files: `tp_TCPUDP.tcl` and `TCPUDP_pps.plot`. You can find them on Moodle in scripts file.

The `tp_TCPUDP.tcl` script takes a link capacity value as an argument. It creates a link with the given capacity and creates two flows which traverse that link, one UDP flow and one TCP flow. A traffic generator creates new data for each of these flows at a rate of 4Mbps.

Output

When you are done running the simulation, you can plot the throughput of the two flows as follows:

```
$ gnuplot TCPUDP_pps.plot
```

Questions

- How do you expect the TCP flow and the UDP flow to behave if the capacity of the link is 5Mbps?

Now, you can use simulation to test your hypothesis. Execute the `tp_TCPUDP.tcl` script as follows:

```
$ ns tp_TCPUDP.tcl 5Mb
```

Use `TCPUDP_pps.plot` to plot the throughput of the two flows.

- Why does one flow achieve higher throughput than the other? Try to explain what mechanisms force the two flows to stabilize to the observed throughput
- List the advantages and the disadvantages of using UDP instead of TCP for a file transfer, when our connection has to compete with other flows for the same link. What would happen if everybody started using UDP instead of TCP for that same reason?




Figure 9

We expect UDP having higher throughput as it is not restrained by congestion control.

As expected, UDP achieves higher throughput than TCP, because it is not restrained by congestion control; it transmits packets at a constant rate, regardless of whether any of them get dropped. This is unfair to the TCP flow, which decreases its transmission rate when it detects congestion in the network. Thus, the more aggressive UDP flows can starve TCP flows which traverse the same bottleneck links.

If network end-hosts stopped using any form of congestion control in their transmissions, this could incur a congestive collapse to the network. Read this wikipedia article [here](#) for a bit of historical background on what was taking place on the Internet in the 1980s, before congestion control was adopted.

Network Routing

In this part of the Lab, we are going to see how routing protocols react when network conditions change (e.g., a network link fails).

Setup

We will use the following file: `tp_routing.tcl`. You can find the file on Moodle in `scripts.zip` file.




Figure 10: Topology

The `tp_routing.tcl` script takes no arguments and generates the network topology shown in the figure above. You can run the simulation with the following command:

```
$ ns tp_routing.tcl
```

Questions

Run the `tp_routing.tcl` script and observe the NAM window output. You may have to refer to the source code of `tp_routing.tcl` to answer the following questions:

- Which nodes communicate with which other nodes?
- Which is the transport protocol used?
- Which route do the packets follow? Does it change over time?

Node 0 uses UDP to send packets to node 5. The transmitted packets follow route 0-1-4-5, and the route does not change over time.

Modify `tp_routing.tcl` and add the following lines just before command `ns at 0.5 "↔ $cbr0 start"`:

```
$ns rtmodel-at 1.0 down $n1 $n4
$ns rtmodel-at 1.2 up $n1 $n4
```

Rerun the simulation and observe the NAM window output.

- What happens at time 1.0 and at time 1.2?
- Does the route between the communicating nodes change as a result of that?

At time $t=1.0$, link 1-4 goes down, but the route between node 0 and node 5 does not change. This makes node 0 unable to reach node 5.

At time $t=1.2$, link 1-4 goes up again, and the packets that were waiting at node 1 are now forwarded to node 4 and, eventually node 5.

The nodes in the topology use a static routing protocol (i.e., preferred routes do not change over time). We are going to change that, so that they use a Distance-Vector routing protocol. Modify `tp_routing.tcl` and add the following line before the definition of the finish procedure:

```
$ns rtproto DV
```

Rerun the simulation.

- How does the network react to the changes that take place at time 1.0 and time 1.2?

Now, when link 1-4 goes down, the routing protocol discovers a different route (0-1-2-3-5) and uses it. Once link 1-4 becomes available again, the routing protocol reverts to the original route (0-1-4-5), since it has a lower cost (in number of hops to destination).

Remove these two lines you previously added (`$ns rtmodel-at 1.0 down $n1 $n4` and `$ns rtmodel-at 1.2 up $n1 $n4`) and add the following the following line, instead:

```
$ns cost $n1 $n4 3
```

- How does this change affect routing? Explain why.

This changes the cost of link 1-4 to 3 (the cost of the rest of the links is assumed to be 1, implicitly). We can notice that the flow now uses route 0-1-2-3-5, because the cost of this route becomes lower than the cost of 0-1-4-5.

Now, replace the line you just added with:

```
$ns cost $n1 $n4 2
```

and uncomment the following line, which is located right after the finish procedure definition:

```
Node set multiPath_ 1
```

- Describe what happens and deduce the effect of the line you just uncommented.

Both routes now have equal cost to the destination. Since we are using multipath, node 1 will split traffic equally on the shortest paths.