Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Spring Term 2016
Prof. Adrian Perrig

Operating Systems and Networks
Solution 11

Note: these solutions constitute material *supplemental* to the exercise sessions.

1 Congestion 1

RFC 791 `http://tools.ietf.org/html/rfc791` defines a TCP header field known as *Type of Service* (TOS). TOS was originally designed to allow senders of TCP traffic (i.e., applications) to specify whether traffic they create should be treated with priority, preference for high reliability or preference for low delay.

Explain why it is a bad idea to allow routers to read this field and make congestion control decisions based on it.

Answer: Applications cannot be trusted to reliably set that TOS field. All applications would prefer being prioritized over other applications, leading to congestion anyway.

2 Congestion 2

Consider the arrangement of hosts `H1` to `H8` and routers `R1` to `R7` in Figure 1.

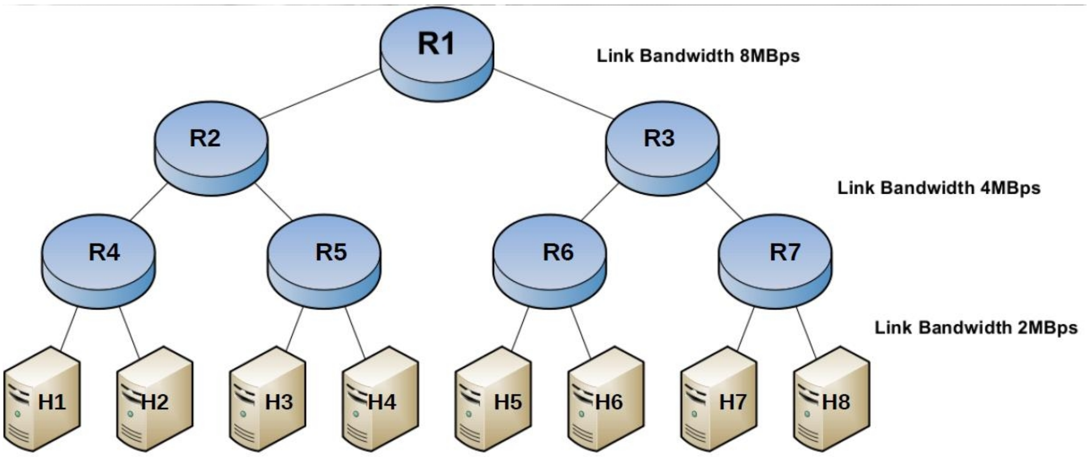

Figure 1: **Network Topology for Exercise**

(a) Show that the links of `R1` cannot become a bottleneck of the network.

Answer: Examining the left link of `R1`:

The worst-case scenario is where all the hosts from the left side of R1 attempt to talk to the respective hosts in the right side of it and vice versa.

- From traffic generated by hosts H1 to H4 towards a host in H5 to H8, the throughput equals to: $\sum_{i=1}^{4} x_i = 2\text{MBps} * 4 = 8\text{MBps}$.
- From traffic generated by hosts H5 to H8 towards a host in H1 H4, the same argument applies.

Consequently, for both directions, the link's utilization reaches its peak (100%), but its capacity is exactly as the demanding throughput, so it does not become a bottleneck.

The network is symmetric, thus, same things apply for the right link of R1.

(b) For all other links, give an example traffic pattern that congests that link. Assume that all the traffic is generated exclusively by messages between two hosts; that is, routers only forward messages and are never the source nor the destination of a message.

Answer:

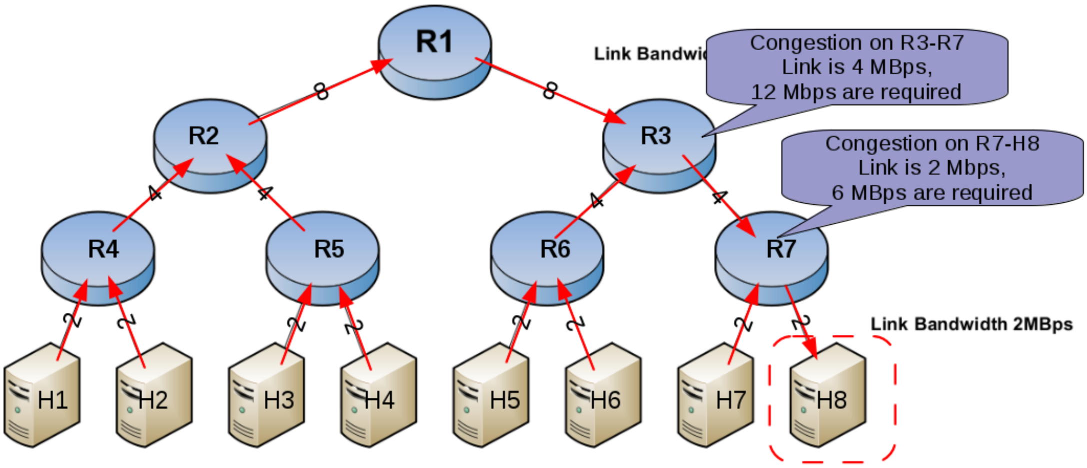

Figure 2: **Two cases where congestion can occur in lower level routers**

(i) We examine the right link l of R3:

Consider a scenario in which hosts H1-H6 generate traffic towards H8. The traffic that passes over l equals to: $\sum_{i=1}^{6} x_i = 2\text{MBps} * 6 = 12\text{MBps}$. In this case, the demanding throughput exceeds the maximum throughput of l, which is 4MB/sec, thus the latter becomes a bottleneck.

Because of the topology symmetry, relevant use cases can be constructed for the left link of R3 and the two links of R2

(ii) We examine the left link l of R7 (connected to H8):

Consider a scenario where hosts H5-H7 generate traffic towards H8. The generated throughput equals to $\sum_{i=5}^{7} x_i = 2\text{MBps} * 3 = 6\text{MBps}$. Again, this exceeds the maximum capacity of the link l (2MBps) which becomes a bottleneck.

3 Congestion 3

In the lectures, we saw two types of congestion control techniques: *Additive Increase/Multiplicative Decrease* and *Slow Start*. This question focuses on Slow Start. A TCP connection uses a threshold of 8KB for congestion control. The maximum segment size should be 1KB and the receiver's window is

16KB. After the 8th, the 11th, and the 17th transmission, timeouts are occurring, which are interpreted as network overload.

Sketch the size of congestion window and the threshold into the following diagram.

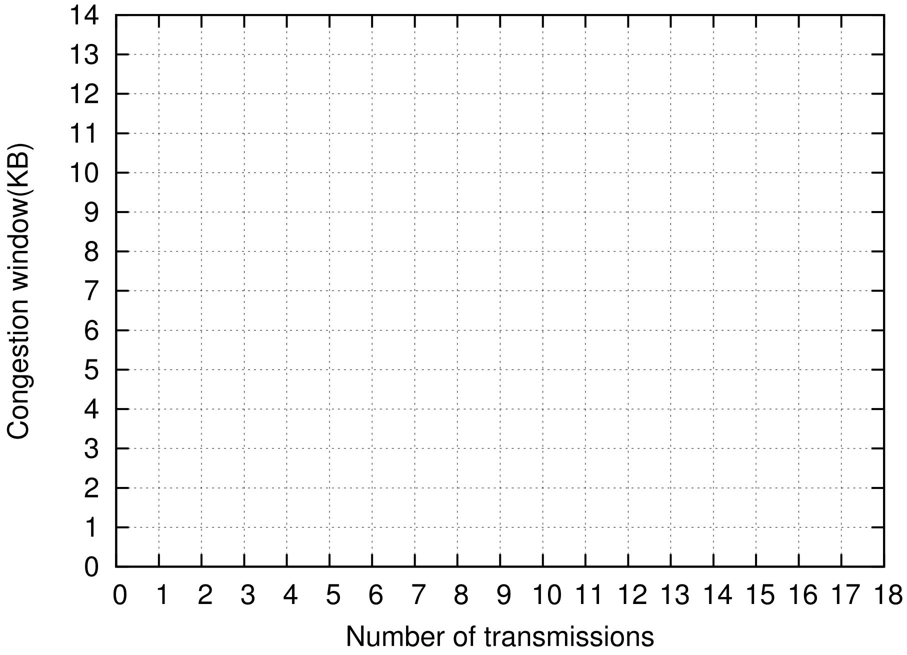

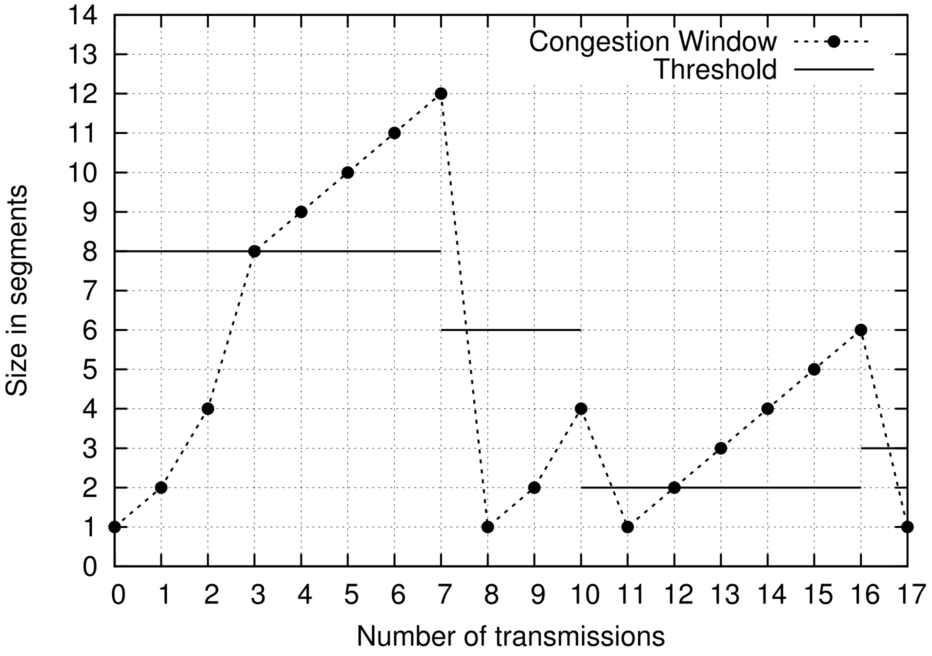

4 Plain DNS Hands On

In this section, you will use the Linux command line tool `dig` to query DNS servers for information.

Using the following command: `$dig www.ethz.ch`

(1) What is the TTL of the A record for `www.ethz.ch`

Answer: The ANSWER section shows:

```
  www.ethz.ch.  292 IN A 129.132.19.216
```

The second column shows the TTL of the record. This number should refresh periodically, and the maximum number displayed is the TTL. In this case it is 300.

(2) Looking at the output, a list of authoritative name servers, A records, and AAAA records is returned by a DNS server. Which DNS server is returning this information?

Answer: The third line from the bottom says: `SERVER: X.X.X.X`, where `X.X.X.X` is the DNS server that is returning that information. It is likely a device inside your local network.

(3) Notice the `flags` line in the output. What do the flags `qr rd ra` mean? Hint: look up the DNS header format at `http://www.networksorcery.com/enp/protocol/dns.htm`

Answer: These are flags set in the DNS header. `qr` set to 1 means it is a reply, as opposed to a query; `rd` means recursion is desired (likely set by the client, and echoed by the server); `rd` means recursion was available.

(4) Can the server in answer 2 replace the information, for example, returning `1.2.3.4` as the A record for www.ethz.ch?

Answer: Yes. This query was not protected by any security mechanism. The query and the reply are sent in plain text. Thus, the server (or any device in the middle) can replace the query or the answer.

5 DNSSEC Hands On

The `dig` command supports the `+dnssec` parameter, which enables the retrieval of DNSSEC data if available.

Issue the following command `$dig www.google.com +dnssec`

(1) Look at the output. Does the domain support dnssec?

Answer: No. No DNSSEC data is returned.

Now compare the output of `$dig .` and `$dig .  +dnssec`

(2) Compared the the output of the command in 1, which additional fields indicate support for DNSSEC?

Answer: The `ad` (Authenticated Data) flag indicates support DNSSEC. When the +dnssec flag is passed, the NSEC and RRSIG records also show support for the DNSSEC signed records.

6 Swiss DNS names

RFC 1912 (`https://www.ietf.org/rfc/rfc1912.txt`) states the following.

> DNS domain names consist of ‘labels’ separated by single dots. The DNS is very liberal in its rules for the allowable characters in a domain name. However, if a domain name is used to name a host, it should follow rules restricting host names. Further if a name is used for mail, it must follow the naming rules for names in mail addresses.
>
> Allowable characters in a label for a host name are only ASCII letters, digits, and the ‘-’ character. Labels may not be all numbers, but may have a leading digit (e.g., 3com.com). Labels must end and begin only with a letter or digit. See [RFC 1035] and [RFC 1123]. (Labels were initially restricted in [RFC 1035] to start with a letter, and some older hosts still reportedly have problems with the relaxation in [RFC 1123].) Note there are some Internet hostnames which violate this rule (411.org, 1776.com). The presence of underscores

> in a label is allowed in [RFC 1033], except [RFC 1033] is informational only and was not defining a standard. There is at least one popular TCP/IP implementation which currently refuses to talk to hosts named with underscores in them.

Explain how names such as `www.grüezi.ch` can be resolved. What does a DNS request for non-ASCII domains look like?

Answer: We refer to `https://www.nic.ch/reg/cm/wcm-page/faqs/idn.jsp?lid=en` for a good overview and the technical details of international domain names.

7 DNSSEC

Cache poisoning or DNS spoofing attacks are used by attackers to attract traffic.

(a) Explain the intention behind traffic attraction attacks.

Answer: Attackers try to attract traffic to be able to respond to legitimate application requests (HTTP, MAIL, DNS). The responses could be riddled with malware or redirections to other malicious destinations.

(b) Explain how attackers mount DNS spoofing attacks.

Answer: As we have seen in previous exercises, answers from DNS servers are unprotected. An attacker can thus alter DNS answers on the fly.

(c) Explain countermeasures.

Answer: The Domain Name System Security Extensions (DNSSEC) constitutes a number of mechanisms to guarantee the authenticity and the integrity of DNS answers. More information can be found in `http://tools.ietf.org/html/rfc3833`.