Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich
Spring Term 2014

Operating Systems and Networks
Sample Solution 1

- 1 byte = 8 bits
- 1 kilobyte = 1024 bytes ≈ 10^3bytes

1 Network Performance

1.1 Delays

Given a 1Gbps point to point copper wire (propagation speed 200m/µs) with a length of 3km.

a) *What is the transmission time for a 16 MB file?*

$t_{transmission} = \frac{L}{R} = \frac{16MB}{1Gbps} = \frac{16*10^6 byte}{10^9 bit*s^{-1}} = \frac{1.6*10^7*8bit}{10^9 bit*s^{-1}} = 0.128s$

b) *What is the propagation delay?*

$t_{propagation} = \frac{d}{s} = \frac{3km}{200m/\mu s} = \frac{3000m}{200m/10^{-6}s} = 15\mu s$

1.2 Throughput

What is the throughput when retrieving a 4MB file across a 1Gbps network with a round trip time of 150msec, again ignoring ACKs?

Assuming that the file transfer has to be initiated by a request:

$$R = \frac{4MB}{150msec + \frac{4MB}{1Gbps}} = \frac{4MB}{150msec + 32msec} = \frac{4*8Mbit}{0.182s} = 175.82Mbps$$

Otherwise:

$$R = \frac{4MB}{75msec + \frac{4MB}{1Gbps}} = \frac{4MB}{75msec + 32msec} = \frac{4*8Mbit}{0.107s} = 299.07Mbps$$

1.3 Total Time

Calculate the total time required to transfer a 1'000 KB file in the following cases, assuming an RTT of 100 ms, a packet size of 1-KB data, and an initial 2 x RTT of "handshaking" before data is sent.

The following solutions assume that the sender does not wait for a reply from the receiver after the last packet(s) have been sent, thus requiring only a wait time of $\frac{1}{2}RTT$ instead of a full RTT. However, the exercise could also be interpreted differently. In that case, the total transfer time would be increased by $\frac{1}{2}RTT$.

a) The total transfer time is the sum of

- the time required for the handshaking: $2 \cdot RTT$
- the propagation delay: $\frac{1}{2} \cdot RTT$
- the transmission time: $1000 \cdot \frac{\text{packet size}}{\text{bandwidth}}$

total transfer time $= 0.2s + 0.05s + 1000 \cdot \frac{8*10^3 bit}{1.5 \cdot 10^6 \frac{bit}{s}} \approx 5.583s$

b) The total transfer time is the sum of

- the time required for the handshaking: $2 \cdot RTT$
- the transfer time for 999 packets: $999 \cdot (RTT + \frac{\text{packet size}}{\text{bandwidth}})$
- the transfer time for the last packet: $\frac{1}{2}RTT + \frac{\text{packet size}}{\text{bandwidth}}$

To put it another way, the total transfer time is the sum of

- the time required for the handshaking: $2 \cdot RTT$
- the propagation delay and the time we have to wait due to the protocol: $999\frac{1}{2} \cdot RTT$
- the transmission time for the 1000 packets: $1000 \cdot (\frac{\text{packet size}}{\text{bandwidth}})$

total transfer time $= 0.2s + 999\frac{1}{2} \cdot 0.1s + 1000 \cdot \left(\frac{8*10^3 bit}{1.5 \cdot 10^6 \frac{bit}{s}}\right) \approx 105.483s$

c) The total transfer time is the sum of

- the time required for the handshaking: $2 \cdot RTT$
- the propagation delay and the time we have to wait due to the protocol: $49.5 \cdot RTT$
- the transmission time for the 1000 packets: $1000 \cdot (\frac{\text{packet size}}{\text{bandwidth}}) = 0$

total transfer time $= 0.2s + 4.95s + 0s = 5.15s$

d) We need 10 RTTs to transmit the 1000 packets. The total transfer time is the sum of

- the time required for the handshaking: $2 \cdot RTT$
- the propagation delay and the time we have to wait due to the protocol: $9.5 \cdot RTT$
- the transmission time for the 1000 packets: $1000 \cdot (\frac{\text{packet size}}{\text{bandwidth}}) = 0$

total transfer time $= 0.2s + 9.5 \cdot 0.1s + 0s = 1.15s$

1.4 Transmission Time

How long does it take to transmit x KB over a y Mb/s link? Give your answer as a ratio of x and y.

$t = \frac{8192x}{1'000'000y}$ or with the rule of thumb $t = \frac{8*1'000x}{1'000'000y} = \frac{x}{125*y}$

2 Network Characteristics

2.1 *Width* of a Bit

Determine the width of a bit on a 10 Gbps link. Assume a copper wire, where the speed of propagation is $2.3 * 10^8$ m/s.

$\frac{2.3*10^8*m*s^{-1}}{10*10^9 bit*s^{-1}} = 0.023m/bit$

A single bit has a width of 2,3 cm on the wire.

2.2 Transport media

You have the task to deliver a large data set from Zurich to Chicago (Distance: 6500km) and a choice of different transport media. The first option is a fiber-optical transatlantic link with a bandwidth of 100 Gbps. The signal propagation in the optical fiber is 200'000 km/s. The second option is a Boeing 747 plane with a capacity for delivering 120 tonnes of hard disks with a speed of 800 km/h from Zurich to Chicago. The weight of a hard disk is 600 grams and the capacity is 2TB.

a) *If your task was to broadcast a soccer match from Zurich to Chicago in real time, which of the properties of the two transport links would be dominating? Which transport would be your choice? Explain your arguments with numbers.*

We assume that the bandwidth is sufficient for the live broadcast in either of the two options. Hence, a minimal transmission delay is the determining factor. For the optical fiber V_O or the air plane V_P:

$V_O = \frac{d}{s} = \frac{6500km}{200000km*s^{-1}} = 32.5ms$ $\quad$ $V_P = \frac{d}{s} = \frac{6500km}{800\frac{km}{h}} = 8.125h(= 29250s)$

As one could assume, the optical fiber has the lower transmission delay than the air plane. Furthermore, the plane would have to wait until the end of the match before it could take off to transmit the data.

b) *If your task was to transfer a full backup of all ETH server disks (total capacity ∼2 PB) to Chicago, which parameters would then be most important and what is the numerical size? Where is the crossover point so that choosing the Boeing plane is preferable over the transatlantic link?*

Here, the bandwidth is the dominating criteria. For the optical fiber, it is $R_O = 100Gbps$. For the air plane, the bandwidth R_P is:

$$\#discs = \frac{120t}{600g} = 200000$$
$$capacity\_plane = 200000 * 2TB = 400'000TB$$
$$transmission\_delay = V_P = 29250s$$
$$R_P = \frac{capacity\_plane}{transmission\_delay} = 109.40Tbps$$

The Boeing can transmit 400'000 TB of data in 8.12 h. In the same time, the optical fiber can transmit:

$\frac{100Gbps*29250s}{8} = 365.63TB$

If the amount of data is larger than 292.5 TB, then the plane is faster.

c) *Consider a third location on a ship 2000 km away from Zurich. The link to the computer on the ship is through a geostationary satellite half way between the ship and Zurich and in an altitude above ground of 36'000 km. Let the signal propagation be the speed of light (∼300'000 km/s). If you try to send a ping to the computer in Chicago connected through the fiber-optical link and to the computer on the ship at the same time, which one will be the first to respond? Which of the parameter has the main impact?*

Here, the dominating parameter is the round-trip-time. For the optical fibre, we know that the transmission delay is $V_O = 32.5ms$. As a simplification, we take the RTT as twice the transmission delay, so $RTT_O = 65ms$. If the satellite is in the optimal position, we can approximate the distance with the altitude above ground. $V_S = 2 \cdot \frac{36000km}{300000*km*s^{-1}} = 240ms$ This is again the delay. The time until we receive the reply is therefore $RTT_S = 480ms$. The Chicago-based computer with 65ms via optical fibre will respond before the geographically closer computer on the ship connected through the satellite link with $480ms$.

3 Network tools

3.1 `ping`

a) ping and RTT

(a) —

(b) The RTT can be determined by measuring the time where the request packet has been sent and the time when the response has been received. This measurement uses the same clock since it happens on the same machine.

(c) The latency, in contrast, requires a measurement on two different machines that connect each other with the same path. Otherwise, it would require the two independent clocks in these machines to be perfectly time-synchronized. In practice, this is hard to achieve.

b) Calculate the bandwidth using `ping`

(a) Considering a connection between hosts A and B, we measure the RTT using `ping`. We assume that this RTT Z of a connection having bandwidth R for a given package size L consists of a transmission delay $)\frac{L}{R}$ and a further delay V. Thereby, we ignore any processing delays on the machine B.

$$Z = 2\left(\frac{L}{R} + V\right) \tag{1}$$

The factor 2 comes from the fact that the `ping` response package has the same size as the request package.

This formula is unable to determine the bandwidth having only L and Z since V is unknown. However, taking two measurements for different package sizes L_1 and Z_1 respectively L_2 and Z_2, assuming: $L_1 > L_2$ und $Z_1 > Z_2$, one gets two versions of the equation and is able to solve both equations for V. For L_1 and Z_1, this is:

$$V = \frac{Z_1}{2} - \frac{L_1}{R} \tag{2}$$

Under the assumption that the further delays V do not change much for different package sizes, one can identify the two equations for V. Now solving the resulting equation for R, the formula for the bandwidth is:

$$R = 2\frac{L_1 - L_2}{Z_1 - Z_2} \tag{3}$$

The exercise can also be solved by empirically measuring the RTT for different package sizes, drawing a diagram, and approximating the measurements through a line. Then, the bandwidth corresponds to two times the gradient of the line.

(b) The `ping` package size can be changed on Linux using the `-s` switch (e.g., `ping -s 100 www.sbb.ch`, for sending 100 Bytes of data). On UNIX machines, the `-s` uses a slightly different format (`ping -s www.sbb.ch 100`). On Windows machines, the option to use is `-l` (ping -l 100 www.sbb.ch).

(c) We calculate the mean value of the RTTs of the measurements for different package sizes. With the formula determined before, we calculate R. The value can vary significantly due to variations in network utilization and the relatively small package sizes. The package sizes, however, should not exceed the size limitations of and Ethernet frame (1518 Bytes, 1500 payload plus IP and ICMP headers); larger packages get fragmented which affects the measurement. A typical value for R is in the Mbps range.

3.2 `traceroute`

a) Hosts in Switzerland typically have domain names ending with `.ch`. Our path to `www.joke.com.au` contains these hosts:

```
1  rou-cx-1-service-inf-isg-cx-server-1.ethz.ch (129.132.216.1)
2  rou-ref-hci-service-inf.ethz.ch (10.1.18.38)
3  rou-fw-hci-service-inf-isg.ethz.ch (10.1.18.34)
4  rou-fw-rz-fw-cla.ethz.ch (192.33.92.185)
5  rou-rz-gw-fwrz-gwrz-core.ethz.ch (192.33.92.170)
6  swiez2.ethz.ch (192.33.92.11)
7  swiIX2-10GE-3-1.switch.ch (130.59.36.250)
```

b) One can try to find a machine in a geographically remote location or in a country with a low level of industrialization. For instance, the webserver of the University of Tokio (`www.u-tokyo.ac.jp`) is 30 hops away from ETHZ.

c) The `-t` option (resp. `-i` for Windows) changes the time to live (TTL) of the packets that `ping` generates and the value of TTL corresponds to the maximum number of hops taken. Using this option, one can iteratively determine the maximum number of hops that a packet can traverse.

d) We make use of the following observation: if the TTL expires for a `ping` request, then the hop on which the TTL has expired will respond and tell us. This way, we can get a response from each machine on the path to the target machine when beginning with a TTL of 1 and iteratively increment the TTL until we reach the target host.

3.3 `curl`

a) *discussed in the recitation session*

b) `http://google.com/` sends you a `301 Moved Permanently` and tells you that the new URL is `http://www.google.com/`. This page, in turn is very likely sending you a `302 Found` with a rederict to the swiss google homepage `http://www.google.ch/`.

3.4 `Wireshark`

a) *discussed in the recitation session*

b) *We can use `Wireshark` to capture the traffic and inverstigate the packet payload.*