Dept. Computer Science **Computer Networks**
Vrije Universiteit **26.08.2005**

Always explain your answers concisely and be sure to be to-the-point.

1a Explain the difference between a connectionless and a connection-oriented communicaton service. *5pt*

In both cases, a service is offered that will allow two parties to communicate through message passing. In the connection-oriented case, the both parties are required to first setup a logical connection before they can exchange data. This is analogous to the telephone model. In the connectionless case, a sender can immediately send a data packet to the intended recipient, which is analogous to the post-office model.

1b Give an example of a connection-oriented service that is implemented by means of a connectionless service, and vice versa. *5pt*

TCP is a connection-oriented service that is generally implemented using the IP connectionless service. RPC is arguably a connectionless service that is often implemented on top of TCP. Another example is IP implemented on top of the virtual circuits offered by ADSL.

1c What is the difference between a (virtual) circuit and a connection-oriented service? *5pt*

A (virtual) circuit is an implementation of a route from A to B, such that all packets that are sent from A to B follow the same route. In a connection-oriented service, only the concept of a connection is provided; there is no guarantee that all packets follow the same route.

2a Explain how the transmission speed of a modem having a fixed baud rate can be increased. *5pt*

The baud rate tells us how many signal changes per time unit can take place. By encoding several bits per signal value, it is possible to increase the transmission speed. For example, if a modem can support 16 different signal values, you can encode $4 = \log_2 16$ *bits per signal, effectively increasing the speed by a factor 4 compared to sending 1 bit per signal value.*

2b Telephone modems such as those for ADSL make use of modulation techniques. Why? *5pt*

Sending digital signals through telephone lines is generally not a good idea as the higher frequencies will not make it to the receiving station, effectively distorting the signal beyond recognition at the receiver. The only solution is to encode the digital signal by means of an analog signal, which can be done by varying the frequency or amplitude, or changing phases.

3 The following program is an implementation of a protocol. Explain the principal working of that protocol (i.e., do *not* describe the program). *10pt*

```
void protocol (void) {
  seq_nr next_frame_to_send, frame_expected;
  frame r, s;
  packet buffer;
  event_type event;

  next_frame_to_send = 0; frame_expected = 0;
  from_network_layer(&buffer);
  s.info = buffer;
  s.seq  = next_frame_to_send;
  s.ack  = 1 - frame_expected;
  to_physical_layer(&s); start_timer(s.seq);

  while (true) {
    wait_for_event(&event);
    if (event == frame_arrival) {
      from_physical_layer(&r);
      if (r.seq == frame_expected){
         to_network_layer(&r.info);
         inc(frame_expected);
      }
22       if (r.ack == next_frame_to_send){
23          from_network_layer(&buffer);
24          inc(next_frame_to_send);
25       }
26     }
27     s.info = buffer;
28     s.seq = next_frame_to_send;
29     s.ack = 1 - frame_expected;
30     to_physical_layer(&s); start_timer(s.seq);
31   }
32 }
```

The program implements a 1-bit sliding window protocol. This means that a sender will send a packet and wait for an acknowledgement before sending the next one. If no ack is received after some time, a retransmission of that packet will take place. Likewise, if the wrong ack is sent (i.e., for a different packet), the original packet is sent again. The receiver's job is simply to send acks for the right packets. If it receives a wrong packet, it will return the ack for the packet it had most recently received correctly. ***Important****: A description of the program (like "in line 17, the process waits for a frame to a arrive, after which it checks whether the sequence number is as expected," is* wrong.

4a Explain how fragmentation of frames in a wireless protocol can increase reliability. *5pt*

Let p be the probability of a single-bit transmission error. A frame of length k will then have a probability of $(1-p)^k$ of successful transmission. Obviously, if we can send a frame in parts, then if something goes wrong, the chance that the retransmission *will succeed is higher, as only a part of the original frame needs to be retransmitted.*

4b IEEE 802.11 frames can have a payload of 2312 bytes, whereas an IEEE 802.3 (Ethernet) frame can carry up to 1500 bytes of data. Assume the 802.11 frames are fragmented into parts containing no more than 1500 bytes of data. Explain what happens at a bridge connecting the two types of networks. *5pt*

The bridge will have to assemble the 802.11 frame as this cannot be done by as receiving Ethernet host. At that point, a reassembled 802.11 frame containing more than 2312 bytes will have to be dropped as it is impossible to send this over the 802.3 network.

5a Explain the principal working of distance vector routing. *10pt*

With DVR, each node initially measures the costs to getting to its immediate neighbors, and advertizes these costs to its neighbors. Its neighbors do the same, so that gradually nodes will discover paths to all other nodes. Assume that node S knows a path to node B via its neighbor $N1$. Let $C(S,N1)$ denote the cost of the link $S \rightarrow N1$, and $D(N1,B)$ the distance from $N1$ to B as advertized by $N1$. If neighbor $N2$ of S advertizes a path to B with length $D(N2,B)$, such that $D(N2,B)+C(S,N2) < D(N1,B)+C(S,N1)$, then node S will adjust its routing table by routing packets for B to $N2$.

5b Explain what the count-to-infinity problem is. *5pt*

This problem occurs when a router S advertizes a route to A (through its neighbor $N1$), but discovers that the link from S to $N1$ fails. At that point, it may notice that another neighbor, $N2$, also advertizes a route to A, but S is unaware of the fact that this route is via the link $S \rightarrow N1$. At that point, S will increase its advertized distance to A, after which $N2$ will do the same, ad infinitum.

5c BGP uses distance vector routing, but does not have the count-to-infinity problem. Why not? *5pt*

It's very simple: BGP advertizes the complete path to a destination, so that a source can discover that an alternative route actually crosses a broken link.

6 Transactional TCP is a refinement of TCP for handling request-reply behavior. Its principal working is sketched below. All T/TCP processes can also follow TCP.

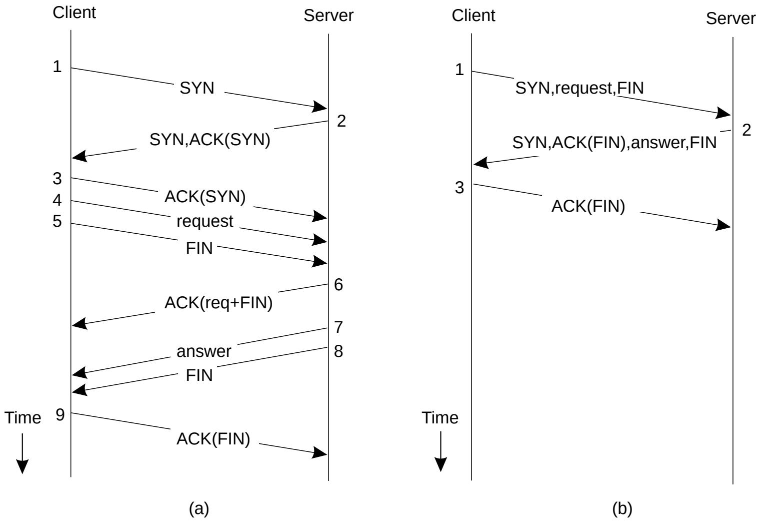

(a) Request-reply behavior for normal TCP; (b) the same for transactional TCP.

6a Which messages from TCP correspond to message #1 in transacational TCP? How would you interpret this message #1? *5pt*

Message 1, 4, and 5 from TCP correspond to message #1 from T/TCP. Its interpretation is "I want to send this single request, after which the connection should be torn down again."

6b Assume the server does not implement T/TCP, but only TCP. How it will it react after receiving message #1 from the (T/TCP) client? What will the client then do? *5pt*

It will simply react by sending message #2 from the original TCP protocol, as part of the 3-way handshake. After that, the client simply switches to the normal TCP operations, by completing the 3-way handshake and following the message sequence from Figure (a).

7a Explain why it is useless to provide authentication without checking message integrity and vice versa. *5pt*

If you do authentication without integrity, Bob may be able receive messages that are certain to come from Alice, but which have been modified by Trudy. Likewise, messages that can be proven not to have been tampered with are useless if you can verify who actually sent them.

7b Give a simple authentication protocol that is based on public-key cryptography. *5pt*

Variations of the following figure will do, but be sure to explain your answer:

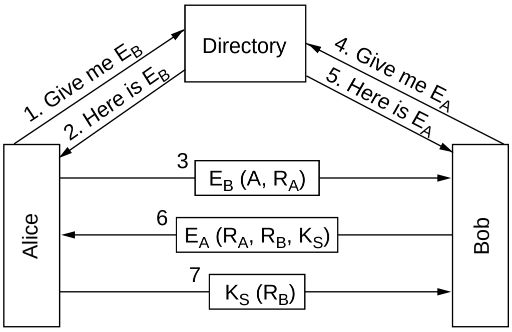

7c In many security protocols, after Alice sends a message $K(N)$ to Bob, where she encrypts the nonce N with a key K, Bob responds with $K^*(N-1)$. Why is Bob forced to change N in his response? *5pt*

The keyword here is replay attack. By forcing Bob to modify N and sendign ity back, he proves that (1) he could decrypt the message from Alice, and (2) sends a response back to that original message.

In this way, we have a tight connection between a request and a reply, that will make it much harder for Trudy to later simply replay Bob's response to a message sent by Alice: the response will not match Alice's request.

Grading: *The final grade is calculated by accumulating the scores per question (maximum: 90 points), and adding 10 bonus points. The maximum total is therefore 100 points.*