THE UNIVERSITY OF HONG KONG

FACULTY OF ENGINEERING
DEPARTMENT OF COMPUTER SCIENCE AND INFORMATION SYSTEMS

CSIS0234B Computer and Communication Networks

Date: May 27, 2004 Time: 9:30am–12:30 pm

Candidates may use any calculator which fulfils the following criteria: (a) it should be self-contained, silent, battery-operated and pocket-sized; (b) it should have numeral-display facilities only and should be used only for the purpose of calculation; (c) it should not have any printing device, alphanumeric keyboard, or graphic display; and (d) it should not contain any recorded data or program. It is the candidate's responsibility to ensure that the calculator operates satisfactorily and the candidate must record the name and type of the calculator on the front page of the examination scripts. Lists of permitted/prohibited calculators will not be made available to candidates for reference, and the onus will be on the candidate to ensure that the calculator used will not be in violation of the criteria listed above.

This is an "open-book" examination. Be reminded that, in your answer, you may refer to materials of the course including (1) the textbook and reference book of the course, (2) sample solutions of any assignment and quiz, and (3) reading materials, worksheet and solutions of tutorials. If you do so, please clearly list the source of your reference together with the section, page and line number. You don't need to copy materials verbatim to the answer book.

Answer all questions in the answer book provided.

1. *(Network programming, 15%)* A TCP server usually has code like this (error handling code is hidden to save space):

```
void network_loop() {
    int sockfd = socket(...);
    bind(sockfd, ...);
    listen(sockfd, 5);
    for (;;) {
        int connfd = accept(sockfd, 0, 0);
        switch (fork()) {
        case 0: /* child */
            if (fork() != 0) exit();
            run_service(connfd);
            exit(0);
        default: /* parent */
            wait(0);
            close(connfd);
        }
    }
}
```

   a. What is the use of *listen*()? What does its argument, 5, means?

   b. The code above puts no limit on the number of clients connected. Modify it to allow at most 5 clients to be served at the same time. (Hint: use the *wait*() call to wait for the completion of connections when too many clients are connecting.)

2. *(Domain name system, 10%)* Many programs (e.g., `ssh`) that use the Domain Name Service (DNS) perform reverse lookups: after translating a domain name (e.g., `abc.com`) provided by the user into an IP address (say, `123.45.67.89`), it tries to "go back" and see if the IP address matches the domain name. If not, it gives a warning.

   a. What is the domain name and resource type in each of the DNS queries?

   b. Suppose that the program queries a DNS server which does not have the resources in the cache. So full DNS searches are initiated by the DNS server, starting from the root domain. Give two factors which determine how many DNS servers will need to handle the query, and show example to illustrate how each affects the path of search.

3. *(Physical links, 15%)* Many ISPs use the UTPs of the telephone infrastructure to allow customers to obtain Internet access. A commonly used protocol is ADSL, which usually supports 1.5Mbps communication.

   a. Given that modems utilizing the phone line can only get 56Kbps for Internet access, why ADSL can find capacity to carry 1.5Mbps traffic?

   b. To avoid noisy frequencies, the bandwidth of the UTP is divided into many bands, and a separate modulation scheme is used in each of them. Why is it impossible for Pulse Code Modulation to be used in these bands?

   c. Suppose we want a modulation scheme that allows the channel bandwidth to be adjusted based on the noise level. Suggest a commonly used scheme for that, and explain how to adjust the bandwidth in the scheme.

4. *(LAN, 15%)* A network is built with Ethernet switches. All the switches in the network are shown in the following diagram, together with three of the hosts.

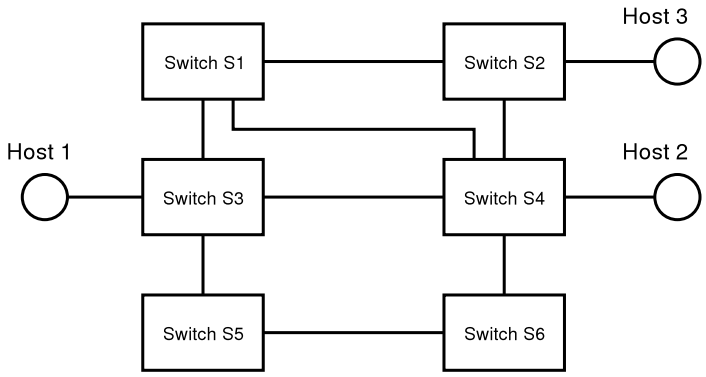

   a. Suppose the IDs of the switches are such that ID of Switch 1 is smallest, Switch 2 is the next, etc., and Switch 6 is the largest, and all link-costs are the same. Suppose Host 1 is sending a stream of packets destinated to Host 2. Which switches will see the packets? Show all your workings.

   b. Two hosts A and B are using the path A - Sa - Sb - B to communicate. A switch other than Sa and Sb is removed from the network. Is it possible that the path changes? If not, explain why. If so, give an example to illustrate.

   c. Suppose Host 1 sends a sequence of multicast packets to a group that host 2 and host 3 are listening. Which switches will receive the packet?

5. *(Network, 15%)* OSPF is proposed to replace RIP as the Internet Interior Gateway Protocol, since the distance vector approach used by RIP suffers from the count-to-infinity problem.

   a. Somebody claims that slow convergence shouldn't be a problem: if a host disappears, it cannot be reached anyway, so we shouldn't care whether the routing tables are correct for such routes. With an example, explain why he is wrong.

   b. Split horizon is used by RIP to avoid the problem. Why OSPF is still used instead? Give an example to show the problem.

   c. Both OSPF and BGP use the distance vector idea. Explain where they are used, and how the count-to-infinity problem are solved by these protocols.

6. *(Transport, 15%)* In modern TCP, timestamps are attached to all packets.

   a. State the two primary uses of them. For each use, explain the problem being addressed, explain how properties of modern networking equipment causes the problem, and how the timestamps solve the problem.

   b. The TCP standard proposes that if a host does not store extra states, during boot the computer should be quiet (don't make TCP connections) for 2 minutes. Suppose a host has a clock, and uses TCP initial sequence numbers that increase by 64000 every half second. What is the minimum quiet time that will make connections safe? Why?

   c. Suppose the host has TCP timestamps which increments every 10 milliseconds. Why this should not be used to determine the quiet time? Does this affect the above guarantee in any other way?

7. *(Congestion Control, 15%)* Suppose that, in a TCP connection which the sender always has data to send, the MSS (Maximum Segment Size) is 1024 bytes (1 KiB). The receiver read the data fast enough that the buffer is never full.

   a. Suppose the connection has round trip time (RTT) being 200ms, and the initial TCP threshold is set to 16 MSS. When will the transfer rate experienced by the application first exceeds 100 KiB per second, assuming the network can handle that?

   b. The slowest connection in the network currently has 512KiB per second bandwidth (counted in frames in the datalink layer) usable by this connection; above this Explicit Congestion Notification (ECN) is used to trigger speed reduction in the sender. Estimate the average rate of transfer in the long run, given that the TCP header has 12 bytes in options fields, and no IP level option is used.

   c. In the ECN scheme, two bits in the TOS field of the IP header is used. In an alternative scheme that have been suggested, only one bit is used: if it is 1, it means the packet supports ECN, and congestion is not experienced; routers that notice congestion will set this bit to 0. Why this scheme is inferier to the currently adopted scheme?

END OF PAPER