

THE UNIVERSITY OF HONG KONG

FACULTY OF ENGINEERING
DEPARTMENT OF COMPUTER SCIENCE AND INFORMATION SYSTEMS

CSIS0234B Computer and Communication Networks

Date: May 21, 2003

Time: 2:00–5:00 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, 20%) The following is the specification of a type used in an RPC program:

```
struct node_t {  
    string node_name<>;  
    int node_types<10>;  
    struct node_t *next;  
};
```

- a. Write an equivalent C++ structure *node_t*.
b. Suppose you want to avoid using RPC. Write a function, with the prototype

*encode_node(node_t &node, char *message, int &size);*

that, given a *node_t* reference, builds a message with the information in the *node_t*, and fill the size of the message in the variable *size*. Write a function, with prototype

*decode_node(char *message, node_t &node);*

that does the reverse (note that *decode_node()* should not need a *size* argument). Make sure the created message is the same for all architectures. (You may use any standard C and C++ functions, including those for the Unix socket interface, and you may ignore all possibility of runtime errors.)

2. (Datalink, 15%) A data link between 2 computers has the following parameters:

- Amount of time to send a data or ack frame: 1 ms.

- Time between the first byte of a frame is sent by one computer and the first byte of that frame is received by the other computer: 6 ms.
- Flow control protocol: go-back-n with 3-bit window, employing piggybacking with a timeout of 4 ms (started when a frame is fully received).

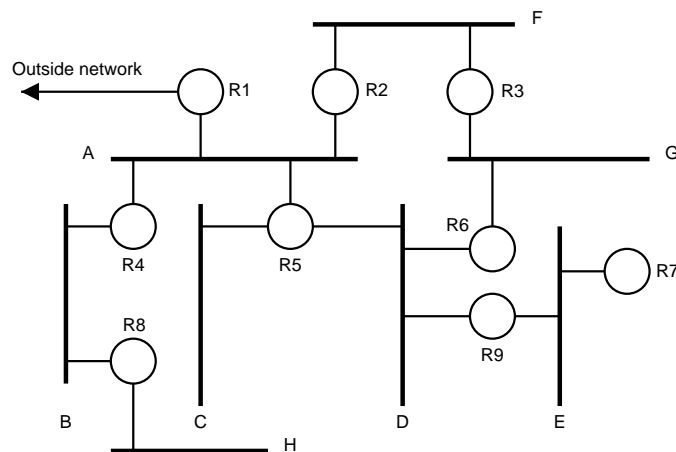
Suppose the network layer of one end has many frames to send to the other end, but there is no reverse traffic at all. Show the detailed sequence of events in the first 32ms, and hence estimate the average frame rate experienced by the network layer.

3. (*Broadcast Networks, 15%*) You have a few Ethernet networks, which are connected to form a larger network. Suppose, at around the same time, each of two computers (perhaps in different physical networks) broadcast a frame to every host in the (combined) network. Let's call the frames A and B.

- If the Ethernet networks are combined with hubs, explain why it is impossible that A is received completely before B in one host, but after B in another host.
- Why such mis-ordering becomes possible if the networks are combined using switches? Explain with the help of a scenario in which the frames can be mis-ordered.
- When the distributed spanning tree algorithm is invoked, it is important that frames do not mis-order. Explain why. Why the algorithm runs correctly even when the networks are combined using switches?

4. (*Network routing, 20%*)

- Give the network address and network mask for the smallest IP subnet that contains both 159.6.13.25 and 159.6.14.175. What is its broadcast address?
- Suppose the subnet is further divided into 8 equal-sized subnets. What are their network addresses and network masks, and how many IP addresses can be allocated to hosts in total? Why the number of hosts that can be accommodated is normally even fewer?
- Some subnets are arranged as follows, and OSPF is used on all routers. Show an optimal routing table at R2, assuming all the networks have equal costs.



- Give an area partitioning in which R1, R2, R6 and R9 are the only backbone routers. Give a minimal set of virtual links that must be configured.

5. (*Connection management, 15%*)

- a. Suppose two TCP connections are established at around the same time between the same pair of computers. What is the relationship between the initial sequence number allocated for these two connections? What about the sequence numbers used later during the communication?
- b. Suppose somehow the two TCP connections use exactly the same initial sequence number. What problem would result? Explain.
- c. Consider the termination of a POP3 connection. For each side of the connection, is it more likely to be performing an active close or a passive close? Why? How can you verify your prediction?

6. (*Congestion Control, 15%*)

- a. Describe a network topology under which the use of choke packets is much better than RED (Random Early Detection) as congestion control scheme.
- b. Explain why the Internet abolished choke packets and adopted RED as congestion control mechanism, by giving the key difference between the network you specify in part (a) and the Internet.
- c. When packets are lost due to RED, what are the normal steps taken by a UDP client and a TCP client to (1) recover the packet, and (2) avoid further congestion? Is it done by the OS kernel or by the application?

END OF PAPER