ID Number:	Mark:
Last Name:	
First Name:	

CMPUT 313, Winter 2001, Final Examination (B1)

April 19th, 2001

This is an closed book exam. All questions have equal weights and must be answered in whatever space is available on this form. Percentages indicate the relative weight of subquestions. No additional sheets are allowed. To receive full marks present all intermediate steps. If necessary, state any additional assumptions that you made. The exam is marked out of 38 which is the percentage of its contribution to the final mark. All questions carry an equal weight. You have five days to contest your mark from the day the grades are announced.

[1] Consider the data 0010010011011011 and the generator polynomial $G(x) = x^3+x+1$. What is the CRC remainder calculated by applying G(x) on the data? (60%) What will be the transmitted message? (20%) Can a single bit insertion (or deletion) error be always detected using G(x) and why (or why not)? (20%) [Note in bit deletion (or insertion) a bit's value does not change, instead a bit disappears completely (or an additional one is introduced) due to faulty synchronization between sender and receiver.]

[2] Assume a single TCP connection traverses an entire path from source S to destination D. Assume that this is the only connection in the network. Also, let us assume that the path from source to destination consists of ten links. Half of the links are of speed 10 Mbit/sec and the rest are of speed 100 Mbit/sec. The round-trip propagation delay over the entire path is 500 msec. If TCP had the ability to guess the best window size to use, what would this window be? (60%) Draw how TCP's congestion window will likely evolve in reality? (in the same illustration point out the ideal window size, if you have been able to derive it from the first question) (40%). [Assume that the packets are of relatively small size.]

[3] In the token ring protocol, why are the A and C bits not covered by the CRC of the frame? (50%) Why is it that in token ring protocols, in general, the incoming frames are delayed for at least a bit time before relayed downstream? (50%)

[4] When Class-Less Interdomain Routing (CIDR) was introduced, not all routers in the Internet were capable of understanding the new (prefix/mask) format for routing entries, and some routers kept using the Class-Full routing by interpreting the first few bits of a network ID as defining its class (0=Class A, 10=Class B, 110=Class C). Describe two problems that such mixture of routers had to cope with.

[5] Somebody claims that they have used Go-Back-N over the same link on which previously they were using Selective Repeat and it turns out that Go-Back-N improved the throughput! Can you take this claim seriously or not and why? [Be precise. Note that the link (being the same) has the

same characteristics in both cases, i.e., frame error rate, bitrate and propagation delay.]

[6] Describe why the count-to-infinity problem cannot be solved even when the split horizon technique is used. (50%) Give an example. (50%)

[7] A group of N stations share a 56 kbits/sec ALOHA channel. Each station outputs a 1000-bit frame on an average of once every 100 sec, even if the previous one has not yet been sent (i.e., the stations are buffered). What is the maximum value of N? (ALOHA throughput: $S=Ge^{-2G}$.)

[8] The following figure represents a network of bridges and LAN segments. For convenience we assume that the forwarding tables at all the bridges are initially empty. We also assume for convenience that on each LAN segment there is only one host attached and we use the host's MAC address to name the segments as well (that is, the host MAC addresses are A, B, C, D, E, F, G, H, I, J). Show the contents of the forwarding tables of all the bridges (B1, B2, B3, B4, B5, B6, B7) after all of the following transmissions have taken place: (a) a frame from A to C, (b) a frame from C to A, (c) a frame from D to C.