EE382C Advanced Computer Architecture and Organization
Midterm Exam Spring 2005

May 4th, 2005
Version γ 0.3

(Total time = 75 minutes, Total Points = 100)

Name: (please print)___________SOLUTION___________________

In recognition of and in the spirit of the Stanford University Honor Code, I certify that I will neither give nor receive unpermitted aid on this exam.

Signature:_______________________________________

This examination is closed notes closed book except that you may refer to at most 1 page of notes that you have prepared yourself. You may not collaborate in any manner on this exam. You have 1 hour and 15 minutes to complete the exam. Before starting, please check to make sure that you have all 8 pages.

1	40	
2	20	
3	20	
4	20	
Total	100	

1. Short Answer (40 points, 5 points each)

A. True or False: In order to implement an adaptive routing algorithm, node-table routing is more preferable than source routing. Give a one sentence rationale for your answer.

True – node-table allows for adaptivity to implemented in each node

B. In a 3-ary 4-fly network, there is only a single path from each source to each destination. If two extra stages are added, what is path diversity of the new topology?

*path diversity is 3*3 = 9*

C. True or False: Credit-based flow control requires more buffering than On/Off flow control. Give a one sentence rationale for your answer.

False – On/Off requires more buffering

D. True or False: Since minimal routing algorithms route only in the minimal direction, they always provide lower *latency* compared to non-minimal routing algorithms. Give a one sentence rationale for your answer.

False – because of contention, minimal can have longer latency

E. In a router, the total bandwidth (B) is divided into k input and k outputs such that each port has a bandwidth of $b = B/2k$. As the bandwidth(B) of a router increases, the bandwidth can be taken advantage of in two ways (i) increase the bandwidth per port(b) and keep the number of ports(k) constant or (ii) increase the number of ports but keep the bandwidth per port constant. For very long packets (1MB packets), which method provides a better latency? Give a one sentence rationale for your answer.

Method (i) – with longer packets, serialization latency becomes critical, thus more bandwidth reduce L/b latency

F. True or False: Channel slicing and bit slicing reduce the network diameter at the expense of higher serialization latency. Give a one sentence rationale for your answer.

False – Channel slicing does reduce the diameter but bit slicing does NOT reduce the diameter.

G. Consider a 64 node 8-ary 2-cube mesh network (no end around connections) in which each network link has the bandwidth of 4 Gb/s. Assuming the routing and flow control are perfect – so that the network achieves 100% of its capacity – how much continuous random traffic can the network support at each node?

$b = \gamma\theta$

$\theta = b / \gamma = b / (k/4) = 4Gbs/2 = 2Gb/s$

2. Flow Control (20 points)

Consider the following flow control method – optimistic circuit switching. With this protocol, a node optimistically sends an entire packet along with a header request to set up a circuit. If the header becomes blocked, the data is dropped and a NACK is sent back the partially reserved circuit, which terminates the transmission of the packet. The source only resends the packet once the circuit has been completely reserved.

A. How does this flow control improve on the circuit switching? Is any bandwidth wasted compared to circuit switching which would reduce the throughput?

The method reduces the zero-load latency of the circuit switching flow control. No bandwidth is wasted since the bandwidth used belongs to channels that are already reserved for circuit switching.

B. Draw a time-space diagram for this flow control method. Assume that the packet must traverse four hops and that the header is blocked in the network for 2 cycles after the second hop. Assume that the packet is 4 flits long.

The space diagram is similar to circuit switching except a NACK is sent back when the header request blocks.

Many people got points deducted since this "header request" is NOT the head flit but only a request for the circuit setup, same as circuit switching. Thus, after the circuit is setup, 4 flits need to be sent (1head, 2 body, and 1 tail).

The termination of the flow control was not properly specified so the termination could occur with the tail flit releasing the channels or an explicit ACK being sent back from the destination.

C. When a NACK is received at the source for a dropped packet, instead of waiting for the entire circuit to be established, if the packet is speculatively resent, how does this effect the latency?

This can still reduce the latency since the packet is sent earlier, compared to previous methods where the packet is sent when the circuit is fully established.

3. Concentrators (20 points)

You are given the task of creating a 16-ary 2-cube network with 1GByte/s channels. Each port has an average bandwidth of 225 MByte/s. When a node presents a request to the network, it does so at a peak bandwidth of 1GByte/s for 200ns.

A. Estimate the cost of this network in terms of total pin bandwidth of the routers.

Pin BW/node $= 4 \times 2 \times 1$ *GB/s = 8 GB/s*

Total pin BW $= 16^2 \times 8 = 2048$ *GB/s*

B. You decide to use concentrators. What degree of concentration should you use? Your goal is to provide sustained average bandwidth at minimum cost.

4, since 4×225 *MB/s < 1GB/s <* 5×225 *GB/s*

C. What are the average and peak bandwidths of the concentrated node?

Average $= 4 \times 225$ *MB/s = 900 MB/s*

Peak $= 4 \times 1$ *GB/s = 4 GB/s*

D. What will be the topology of your new network?

We are combining 4 nodes – do this by combining 2 nodes in each dimension. The new network will have 8 concentrated nodes in each dimension. i.e. an 8-ary-2cube

E. Estimate the cost of your new network with concentrators, still assuming 1GByte/s channels.

Number of nodes = 82 = 64

Since each node still has 8 channels:
Total pin BW = 64 $\times$ 8 GB/s = 512 GB/s

F. What is the serialization latency for an access, still assuming 1GByte/s channels?

With 1 Gb/s channels, serialization latency = 200ns

G. Instead of 1Gb/s, if you size the channels to handle average bandwidth, what is the serialization latency?

By sizing channels 900MB/s,
Serialization latency = 200 ns $\times$ 1000/900 = 222.22 ns

4. Clos Networks (20 points)

Consider a 16-node folded Clos topology as shown in the figure. All channels shown are bidirectional with a bandwidth taper. All channels between the terminals and stage 0 are of bandwidth, b=1 in each direction; channels between stage 0 and stage 1 are b=7/8 and so on as marked in the figure below. Three different routing algorithms can be used for this topology. Over all the channels in the network, determine the maximum overload (O_{max}) in each case. Assume packets are injected at rate 1 by each terminal node. For any channel *c*, the overload, $O_c = \gamma_c / b_c$

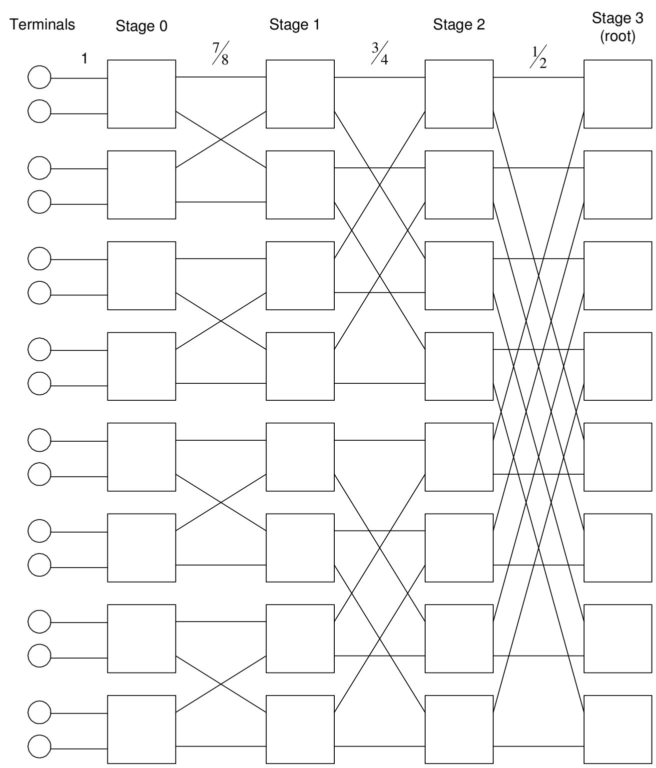

A. *Routing to root* - Every packet has to be routed to a "root" before it can reach the destination. A root node is one of the nodes in the right most stage. Root nodes are chosen randomly. What is

- O_{max} for worst case traffic

All traffic is sent all the way to a root node, thus all channels have γ=1 (since root nodes are chosen randomly, traffic is distributed uniformly).
So maximum overload is at the channels with b_c=1/2, and O_{max} =2

- O_{max} for uniform random traffic

Same as above, as all traffic is routed through randomly chosen root node
O_{max} =2

B. *Closest common ancestor routing* – A packet only needs to be routed to the closest ancestor common to source and destination. If more than one node can be chosen at any step, the choice is made randomly. Determine

- O_{max} for worst case traffic

For worst case, all traffic will go via a root node, so as in part A, $O_{max} = 2$

- O_{max} for uniform random traffic

For UR traffic, 1/8 of the traffic turns around at stage 0, and only 7/8 is seen on the channels from stage 1 to stage 1, so for those channels, $O = (7/8)/(7/8) = 1$.
Similarly, the load on channels between stage 1 and stage 2 is ¾,and $O = (3/4)/(3/4) = 1$
Etc.

$O_{max} = 1$

C. *Valiant* – A packet is first routed to a randomly chosen terminal node first, and then routed to its final destination. Routing to and from the randomly chosen terminal node is done via the closest ancestor. Once again, determine

- O_{max} for worst case traffic

This is equivalent to 2 stages of B with uniform random traffic, so $O_{max} = 2$

- O_{max} for uniform random traffic

Since Valiant routing does not take advantage of any locality, this is the same as above. $O_{max} = 2$

D. Which of these algorithms would you suggest should never be used? Why not?

Valiant or Routing to Root, since load is high for both uniform random and worst case traffic. They do not take advantage of locality in the traffic.

E. Which of these do you consider to be the best routing algorithm? What are the pros and cons of using this routing scheme?

Closest Common ancestor.