EE482B Advanced Computer Architecture and Organization
Midterm Exam

May 4th, 1999

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

Name: (please print)______________________________

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 open notes open book. You may not, however collaborate in any manner on this exam. You have 1 hour 15 minutes to complete the exam. Before starting, please check to make sure that you have all 6 pages.

1	**35**	
2	**20**	
3	**25**	
4	**20**	
Total	**100**	

Name:_______________

1. Short Answer (35 points, 5 points each)

A. When a new packet arrives at a router in a network using virtual channel flow control, what resource must be allocated to the packet?

A virtual channel. *Later channel bandwidth is allocated to each flit, not to the packet.*

B. To allow a single long packet to use the full bandwidth of the physical channel, what is the minimum size a virtual channel buffer can be?

The buffer must be large enough to hold the data that arrives over the channel during the round-trip delay of the channel, including the 'think' time at both ends.

C. An 8×6 mesh network uses wormhole routing, has only a single virtual channel per physical channel, and allows arbitrary minimal routes. Is this network deadlock free? (explain briefly 8-words or less).

No. All eight 'turns' are allowed, so its easy to construct deadlock using the turn model.

D. An 8×6 mesh network uses wormhole routing and has two virtual channels per physical channel, one in each of two subsets, C1 and C2. At each hop along the route, a packet may use any productive (gets it closer to its destination) C2 channel and may use a productive C1 channel in the x-direction. It may only use a productive C1 channel in the y-direction if it is already at the correct x-coordinate. Is this network deadlock free?

Yes. Subset C1 is routed dimension-order and thus is a deadlock-free channel subset. Subset C2 by routing only in productive directions does not add any indirect dependencies to C1. Hence, by Duato's theorem this network is deadlock free.

E. A router chip has four input ports, four output ports, and 16 virtual channels per physical channel. The router currently uses a 8×4 crossbar switch internally, with two switch inputs per input port. Assuming good switch scheduling in both cases, will router performance be increased if this is increased to a 64×4 crossbar with separate switch inputs for each virtual channel?

Yes, but only by a small amount. All of the advantage occurs when the crossbar is increased to 16x4. There is no further advantage going to 64x4.

F. Suppose you have a router with four inputs, four outputs, and 2 virtual channels per physical channel. The router uses a 4×4 crossbar switch internally. At a given point in time, the 8 input virtual channels all have traffic to forward and are connected to the following output ports (1,2;3,4;1,1;3,3). The semicolons ';' here separate the input ports. What assignment will a 'greedy' allocator make in this case? What is a good assignment?

Greedy: (1;3;-;-). Good:(2;4;1;3)

G. Suppose you have a 4×4 mesh network that employs a CAM-based node routing table. What entries are needed at node (0,0) at the bottom left corner of the mesh?

```
x  y
00 00   this node          X
1X XX   not this column    E
X1 XX   not this column    E
00 1X   this column        N
00 X1   this column        N
```

This does dimension-order routing. X-first, then Y.

2. Topology (20 points)

Consider a network with a mixed-radix alternating 3-cube topology with k_x=4, k_y=8, and k_z=10. This network has three dimensions. Each node has four bi-directional connections to its neighbors. A node at (x,y,z) is connected to its neighbors in the x direction (x+1,y,z) and (x-1,y,z). Nodes at even x coordinates are connected to their neighbors in the y-dimension, and nodes at odd x coordinates are connected to their neighbors in the z-dimension. The network has a different radix in each dimension, spanning k_x=4 nodes in x, k_y=8 nodes in y, and k_z=10 nodes in z.

On uniformly distributed random traffic, what is γ_{max} for this network?

Solution:

The channels along z-dimension are the most heavily loaded. For a normal torus, the load on the z channels is = average number of hops / number of channels = (k_z/4) / 2 = (10/4) /2 = 2.5 / 2 = 1.25.

In this alternating cube, since there is only 1 z-row for every two z-coordinates, γ_{max} becomes twice as loaded = 1.25 * 2 = 2.5.

Another way to look at this is that every node on the z-dimension contributes on average 1 channel along the z-dimension. Hence, the load on the z channels is = average number of hops / 1 = 2.5.

3. Routing (25 points)

Consider the network from problem 2 above. Suppose the following routing algorithm is used on this network:

1. Subtract the coordinates of the present node (x,y,z) from the coordinates of the destination (x_d,y_d,z_d) to get a difference vector (Dx, Dy, Dz), and from this compute a preferred direction vector (d_x,d_y,d_z) with elements +,-, or 0. If both positive and negative directions are minimal (the source node is halfway around from the destination node along a dimension), the positive direction is preferred.
2. If the x-coordinate is even and Dy $\neq$ 0, take one hop in the preferred direction in the y-dimension, then return to step 1.
3. If the x-coordinate is odd and Dz $\neq$ 0, take one hop in the preferred direction in the z-dimension, then return to step 1.
4. If Dx $\neq$ 0, take one hop in the preferred direction in the x-dimension, then return to step 1.
5. If at an even x and Dy = Dx = 0, and Dz $\neq$ 0, take a hop in the positive x-dimension, then go to step 1.
6. If at an odd x and Dz = Dx = 0, and Dy $\neq$ 0, take a hop in the negative x-dimension, then go to step 1.

Suggest a permutation (permutation of nodes, not necessarily bits of the node address) traffic pattern that gives the highest possible γ_{max} given this topology and routing algorithm. Describe your permutation and compute the resulting γ_{max}.

Solution:

The permutation traffic which gives the highest possible γ_{max} is :-

(x, y, z) -> (x', y, (z+5) mod 10), where x': {0->3, 1->1, 2->2, 3->0}

This routing algorithm is essentially a twist on dimension-ordered routing to cater for this alternating torus. The algorithm first goes all the way in y (or z), and then all the way in z (or y), and then all the way in x, moving one hop in x to reach y or z channels to connect the dimensions.

Take the (3,0,z) channels as an example:-

(0,0,z) -> (3,0,z+5): Since there is no y-offset, and (0,0,z) does not have z-channels, the routing algorithm will take step 4 and go to (3,0,z). It will then traverse the z-dimension before reaching its destination. Within the z-dimension (3,0,z), it loads the channels to 5.

(2,0,z) -> (2,0,z+5): Since there is no y-offset, and (2,0,z) does not have z-channels, and x-offset is 0, the routing algorithm will take step 5 and go one hop along the x-dimension and arrive at (3,0,z). It will then traverse the z-dimension and then take step 4 of the routing algorithm to arrive at its destination (2,0,z+5). Within the z-dimension, it loads the channels to 5.

(3,0,z) -> (0,0,z+5): There is no y-offset, but the current node (3,0,z) has z channels, so the routing algorithm will take step 3 and traverse the z-dimension before taking step 4 and arriving at its destination (0,0,z+5). It again contributes a load of 5 to the z-channels.

Hence, the highest possible γ_{max} given this topology and routing algorithm is 15.

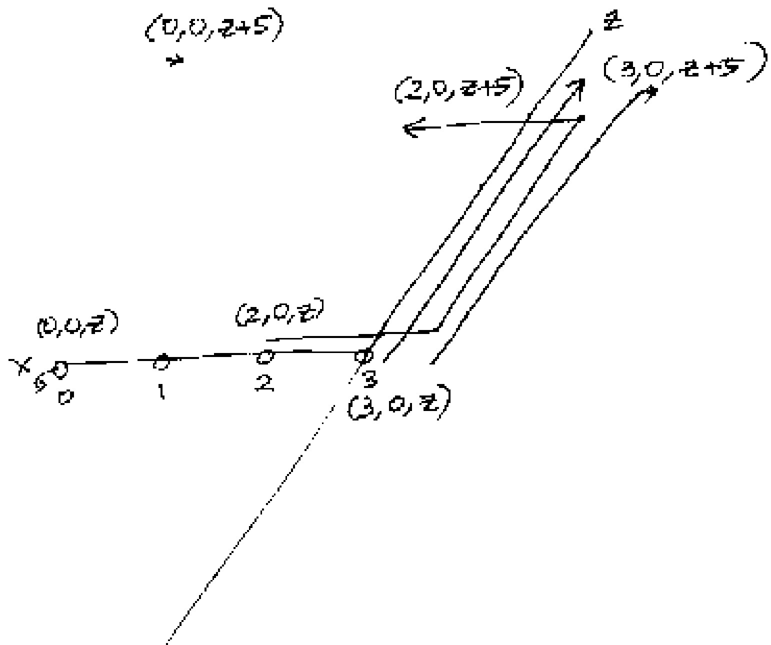

4. Flow Control (20 points)

Draw a time-space diagram for the following flow control method (method X). With method X, packets may be from 2 to 1024 flits in length. The first flit of the packet contains the destination. Packet transmission begins by sending just the first flit of the packet to the destination to establish the path (as with circuit switching). As the head flit passes each node it allocates a virtual channel and a buffer with a capacity of four flits to the packet. When the head reaches the destination, an acknowledgement with a credit to send four flits is returned to the source along the same path. When this acknowledgement reaches the source, the source sends the next four flits of the packet to the destination. When the fourth flit arrives at the destination, another acknowledgement with four credits is again returned to the source. This process continues until the tail flit reaches the destination. At this point a terminating acknowledgement is returned to the source. As this terminating acknowledgment passes each node, it frees the resources associated with the packet. Sketch the time-space diagram for sending a 9-flit packet using method X and compute the latency in terms of message length, L, and hop count, H. Clearly define the terms used in the computation and detail all assumptions made.

Assumptions:

A message consists of 1 header flit + L-1 data flits
tr = routing latency
ts = switching time
tp = propagation delay

There is no think time when the header arrives before acknowledgements are sent, or when acknowledgements arrive before flits are sent.

Latency = H*(tr+ts+tp) {time for header to arrive at destination}
+ H*(ts+tp) {time for acknowledgement to arrive at source}
+ ceil((L-1)/4)*(H*(ts+tp) + {time for 1st of 4 flits to arrive at destination}
H*tp) {time for ack to arrive at source}
+ L*tp {serialization delay}