Final Test (April 24)

1. (25%) Given a system with three processes: P_1 (with events a_1, a_2, and a_3 in order), P_2 (b_1 and b_2), and P_3 (c_1, c_2, and c_3), there are four inter-process communication events: (a_1, b_1), (b_2, c_1), (a_2, c_3), and (c_3, a_3).
   a. Draw a time-space diagram for the system.
   b. Show all the pairs of events that are related.
   c. Provide logical time for all the events using (1) linear time, (2) vector time, and (3) matrix time. Assume that each logical clock in initialized zero, and d for P_1, P_2, and P_3 is 1, 2, and 3, respectively.
2. (25%) Given a 6-node directed graph with the directed edge set {(1, 2), (1, 3), (2, 5), (3, 5), (3, 6), (4, 6)}. Express this DAG using (1) fork/join, (2) parbegin/parend together with semaphore, (3) parbegin/parend only, after **adding** a minimum number of dependence links to the original graph, and (4) parbegin/parend only, after **deleting** a minimum number of dependence links from the original graph.
3. Show the **resource allocation time** for each of the three processes P_1, P_2, P_3 competing on three resources A, B, C in the table below:

Process ID	Priority	1st request time	Length	Retry interval	Resource & locking order
P_1	3	1	4	1	A,B
P_2	4	1	3	1	C
P_3	1	2	1	1	A,C
P_4	2	2	2	2	C,A

   a. When the **wait-die** scheme is used.
   b. When the **wound-wait** scheme is used, assuming each locking process takes a short, but none-zero, time to execute.

   Explain the **scheduling policy** for each scheme in the following scenario: When a process releases a resource R at time T, there is another process P_i in the waiting queue of R and a new process P_j comes for R exactly at time T. Show how this policy works when P_1's length is reduced to 3 in the above example.
4. (25%) Using Glass and Ni's **Turn model**, we can divide the routing space of a 2-D mesh into four regions (relative to the location of the source): NE, SE, NW, SW. For a particular Turn model (such as West-last), it is fully adaptive in some region(s) and non-adaptive in the regions region(s). Two Turn models A and B are said to be **complement** to each other, if in any region, either A or B (or both) is fully adaptive.
   a. Determine a complement of the West-last (and Negative-first) Turn model.
   b. Demonstrate the correctness of your results by showing routing adaptivity of these two Turn models in each region.
   c. Use two **virtual networks** to implement A and B. Show how routing should be constructed to avoid deadlock, reach balanced adaptivity in each region, and achieve efficiency in channel usage.
5. (25%) Refer to the 4th question in Homework 4, consider a 3-cube with one faulty node 111 and one faculty link (000, 001).
   a. How would you propose to handle the faculty link? (**Hint**: treat both end nodes of a faculty link faulty.)

b. Show safe and unsafe status using (1) Lee/Hayes' safety node: a nonfaulty node is unsafe if and only there are at least two unsafe or faulty neighbors, (2) Wu/Fernandez' safety node, (3) Wu's safety level, and (4) Wu's safety vector. Show each case by drawing a different 3-cube.
c. Give a **simple prove** that safety vector is no worse than safety level, which in turn is no worse than Wu/Fernandez' and Lee/Hayes' safety node models, in terms of accurately depicting optimal routing capability of each node in a cube.

6. (**Bonus**: 10 points) The way to treat a faulty link as faulty nodes (i.e., two end nodes of the faulty link) is still too conservative. For example, node 101's safety level is 1 in Question 5. In reality, however, node 101 can reach every node within 2 hops through a Hamming distance path. Find a more efficient **vector coding** method to better represent faulty links. (**Hint**: Similar to safety vector, but use exact coding for routing capability of nodes within 2-hop and apply safety vector coding for nodes beyond 2-hop.)