ECE 697J
Final Exam

Name:_____

	Maximum	Achieved
Question 1	16	
Question 2	12	
Question 3	6	
Question 4	20	
Question 5	10	
Question 6	12	
Question 7	14	
Question 8	10	
Total	100	

Please write legibly! Be concise. Unreadable answers will not be graded.

Time: 90 minutes.

Question 1 (16 points (4+4+8)):
Answer the following general questions regarding Network Processors:

a) What is the motivation behind developing network processors instead of simply using workstation processors?

b) What are the main characteristics of a network processor?

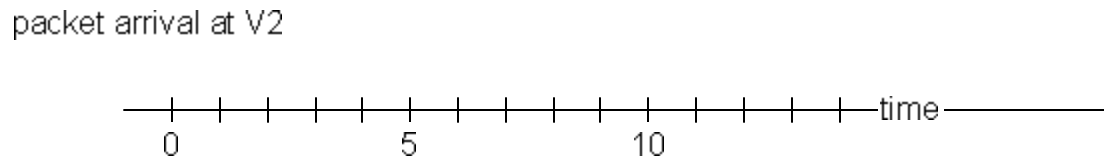c) Show a system outline of a network processor (does not have to be a realistic commercial processor).

Question 2 (12 points (4+4+4)):
Answer the following questions regarding Benchmarking and Network Processor Design:

a) What are the differences in the workload characteristics of network processors as compared to workstation processors?

b) Suggest three benchmark applications that cover a broad range of network processing and argue why these are good choices.

c) In the context of CommBench, "computational complexity" is defined as a benchmark metric. What is the exact definition and why is it relevant?

Question 3 (6 points (3+1+2)):
Answer the following questions regarding Parallelism in Networks

a) In class we have discussed that there are several levels of parallelism in networking. Describe where this parallelism occurs.

b) How can these levels of parallelism be exploited by a network processor architecture?

c) What are the limitations to exploiting parallelism on a router and on a network processor?

Question 4 (20 points (2+6+1+5+6)):
Answer the following questions regarding the WaveVideo application. This question is similar to question 4 in the midterm.

a) What is the idea behind WaveVideo scaling on active routers? For what type of applications is this mechanism useful?
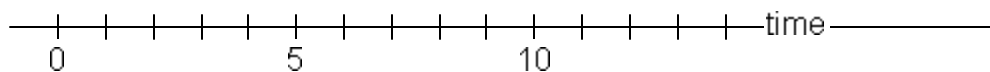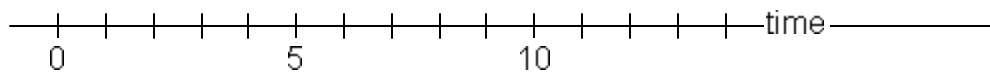
b) Assume the following scenario:

V1
1kb/s
WaveVideo 1kb/s
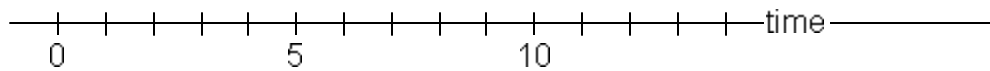R1 —1kb/s— R2
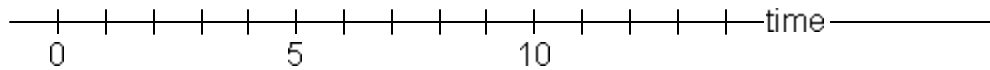cross traffic 1kb/s
1kb/s
1kb/s
V2
1kb/s
C1
C2
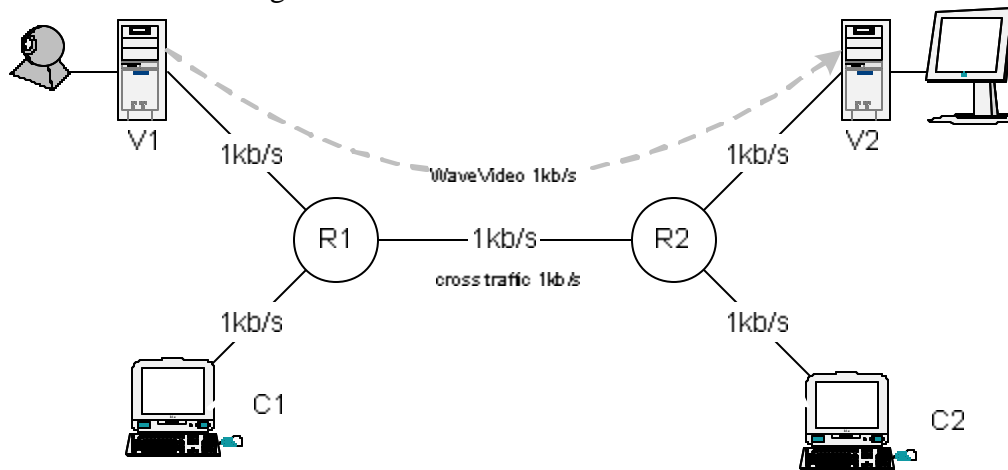
Assume the video transmission starts at time 0 and cross traffic starts at time 3. All sources stop sending at time 11. Then the arrival of packets on router R1 is shown below. Show the state of the packet queue at router R1 (show the packets as boxes). Also, show the WaveVideo packets that arrive at a given time at V2:

packet arrival at R1

X X X X X X X X X

1 2 3 4 1 2 3 4 1 2 3 4

time

0   5   10

queue at R1

tail

head

time

0   5   10

packet arrival at V2

time

0   5   10
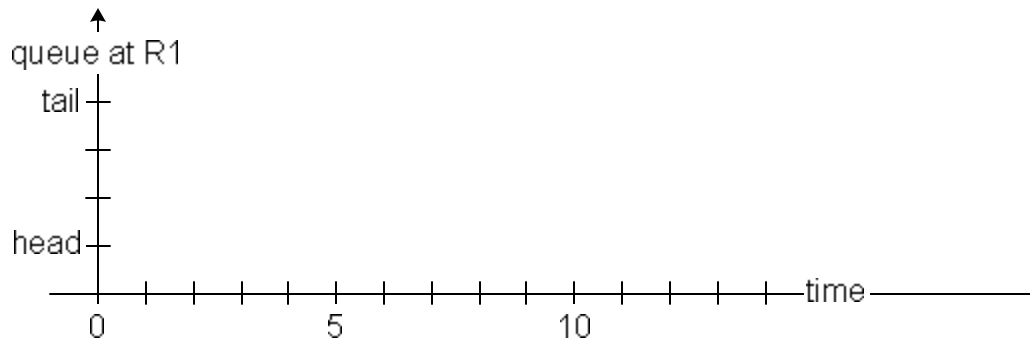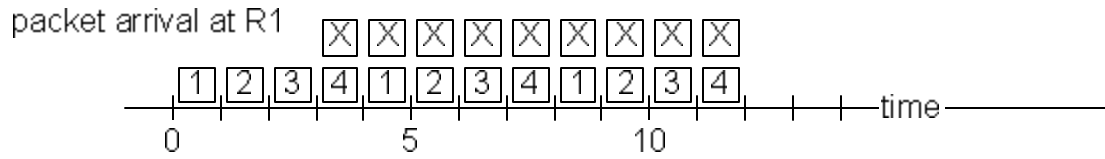
e) Show how your solution works and how it solves the problems you have observed in c) using the above example. You can use the following time lines. Please label clearly what is shown on each time line that you use.

packet arrival at R1

| X | X | X | X | X | X | X | X | X |
| 1 | 2 | 3 | 4 | 1 | 2 | 3 | 4 | 1 | 2 | 3 | 4 |

0   5   10   ──time──────

0   5   10   ──time──────→

0   5   10   ──time──────

0   5   10   ──time──────

0   5   10   ──time──────

0   5   10   ──time──────
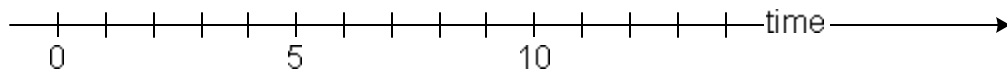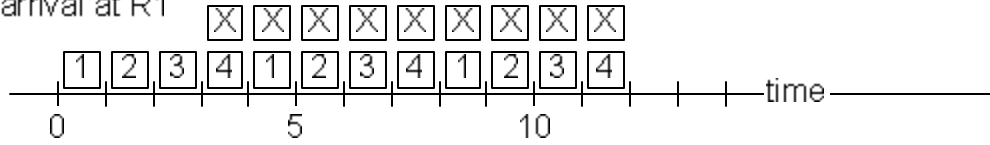
Question 5 (10 points):
Active networks and network processors are both approaches to providing flexibility in how network traffic is handled on a router. Discuss and contrast the goals of both research areas. Can the systems developed in these areas be combined easily? Where do you see differences between active networking and network processors?

Question 6 (12 points (3+4+4+1)):
In class, we discussed a paper on "Design Space Exploration of Network Processor Architectures." Answer the following questions relating to this research area.

a) What is the "design space" in this context?

b) What process is used to explore the design space?

c) What performance criteria can be used to determine the quality of a point in the design space?

d) What alternatives are available to b)?

Question 7 (14 points (5+4+5)):
Assume a network processor implements the typical functions of a port on a programmable router.

a) What scheduling problems appear in the system (i.e., what resource needs to be scheduled for and who/what is competing for the resource and what is the scheduler trying to achieve)? Can they all be solved by the same algorithms? Explain your answer.

b) Why are OS scheduling algorithms only of limited use for network processor scheduling problems?

c) In class, we briefly discussed a scheduling algorithm that is "locality-aware? What are the ideas behind this scheduling algorithm and why is it useful for network processors?

Question 8 (10 points):
Consider all research ideas and systems discussed in this course. Which do you think will have the biggest impact on commercial networks (i.e., the Internet)? Support your argument. Which do you think will have least impact or is most unrealistic and why?