ECE 697J
Midterm Exam

Name: Solutions

	Maximum	Achieved
Question 1	12	
Question 2	12	
Question 3	12	
Question 4	30	
Question 5	14	
Question 6	8	
Question 7	12	
Total	100	

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

Time: 75 minutes.

Question 1 (12 points (3+3+3+3)):
Answer the following general questions regarding Active Networks

a) What is an Active Network and how is it different from a traditional network? What problem does Active Networking address?

An active network allows processing of packets on active/programmable routers

An active network addresses the difficulty of deploying new protocols and services. These can be deployed by changing software instead of hardware.

b) There are several approaches to Active Networking. One is the “capsule” approach another is the “programmable switch” approach. Explain each and compare the differences. Discuss the benefits and drawbacks of each.

In the capsule approach, the instruction code is sent with each packet. This code is then executed on the active router.

In the programmable switch approach, code is referenced (either explicitly by the packet or implicitly by a filter installed by an administrator).

Capsules are simpler to use and more flexible, but have a higher overhead because code must be carried in each packet.

c) What additions to packet headers are necessary to support Active Networking? Specifically, how is this handled in an IP network?

There has to be a way of identifying active packets and carrying code (references) in the packet. In IP networks, ANEP (Active Network Encapsulation Protocol) is used. The ANEP header has fields to carry this information.

d) What are the differences between traditional routers and active routers? What functions does an active router have to support?

Active routers have to support packet processing in the data path (traditional routers only allow control path processing). Also, active routers have to support per-flow state.

Question 2 (12 points (6+6)):
Answer the following questions regarding Active Network router architectures:

a) Describe how ANTS, Smart Packets, and the Active Network Node are implemented. Distinguish between software/operating system issues and hardware issues.

The ANTS environment is implemented as a Java Virtual in user space. No special hardware for ANTS was developed.

Smart Packet processing is done in user space on the control processor of the switch.

In ANN, processing is done on special port processors in the kernel of the operating system.

b) Contrast these three systems regarding the following characteristics:
    a. Scalability

ANN: very scalable due to hardware architecture
SP: limited scalability, because processing is done on control processor
ANTS: scalability is not addressed / not scalable

    b. Performance

ANN: high performance due to kernel processing and hardware support
SP: limited performance – uses VM
ANTS: limited performance – uses VM

    c. Ease of programming

ANN: hard to program (uses kernel code)
SP: easier to program
ANTS: easy to program (JAVA)

    d. Security

ANN: only authentication of plugin code
SP: relatively heavy authentication of each packet
ANTS: not discussed

    e. Safety

ANN: zero
SP: sandboxing avoids interference with other components
ANTS: VM isolates packet execution

    f. Interoperability

ANN: limited due to kernel code
SP: limited to sprocket/spanner interpreters
ANTS: limited to JAVA

Question 3 (12 points (4+4+4)):
Answer the following questions regarding Active Reliable Multicast and Overlay Multicast.

a) Describe how Active Reliable Multicast (ARM) works and contrast it to Scalable Reliable Multicast (SRM). What performance aspects are better in ARM?

ARM uses programmable routers to aggregate NACK messages, perform local recovery (if missing packet is cached), and limit retransmissions to local sub-tree. In SRM, recovery request need to go all the way to the source. To avoid NACK implosion in SRM, end-systems wait a randomized interval to see if someone else is requesting the packet. Also, retransmissions are sent to the entire tree.

ARM performs better in terms of recovery latency, bandwidth consumption, and implosion control.

b) Describe how Overlay Multicast (OM) is implemented

Overlay Multicast uses end-systems to implement the multicasting functionality. All participating end-systems are connected via multiple unicast connections that form a mesh. Routing is performed on this mesh to generate a virtual tree topology. Using probe message, an end-system can determine the performance of its connections and chose the best one.

c) Discuss the benefits and drawbacks of OM over IP multicast.

The virtual topology generated by OM requires more overall bandwidth than a multicast tree in the network. OM needs to replicate routing information on the end-systems, which might be a lot of data for large groups. Also, low bandwidth access links might become bottlenecks, because they need to receive and possibly transmit at the same time.

OM is much easier to deploy, though. It does not require any support form the network.

Question 4 (30 points (3+10+2+5+10)):
Answer the following questions regarding the WaveVideo application.

a) What is the idea behind WaveVideo over active networks? Why is particularly useful for video multicast?

WaveVide is a scalable video encoding that can be adapted to the available bandwidth. An active network node can detect congestion and perform the scaling by dropping less important high-frequency packets. For multicast this works well, because the bandwidth to different receivers might differ significantly. The network can adapt to each receiver individually.

b) Assume the following scenario:

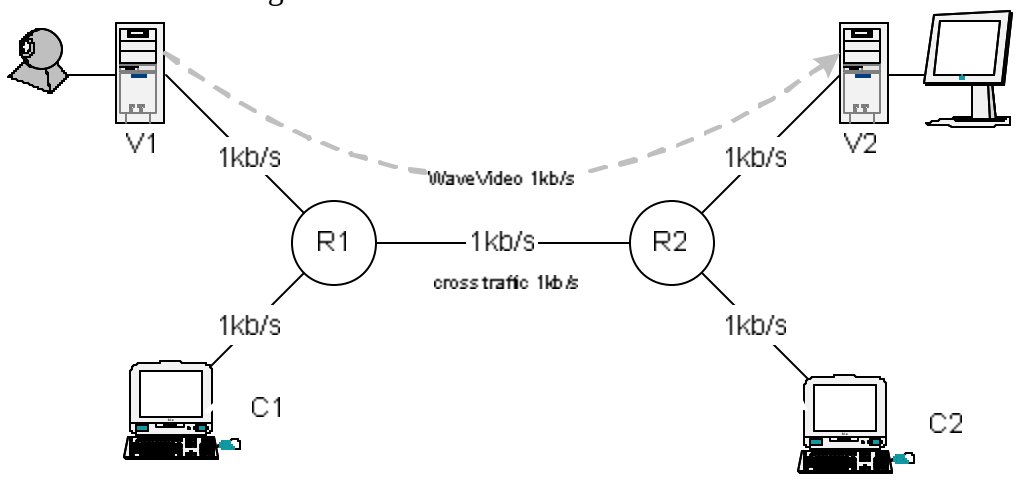

Assume the video transmission starts at time 0 and cross traffic starts at time 3. All sources stop sending at time 11. Then the arrival of packets on router R1 is shown below. Show the state of the packet queue at router R1 (show the packets as boxes). Also, show the WaveVideo packets that arrive at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

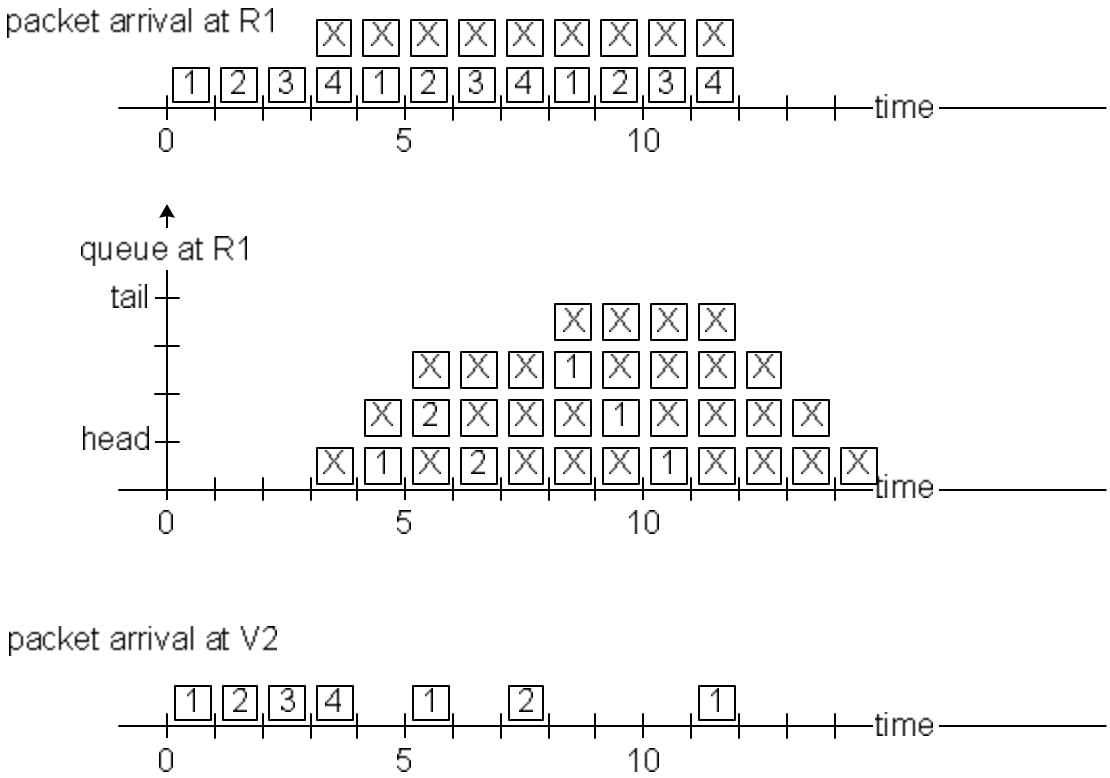

e) Show how your solution works using the above example. You can use the following time lines. Please label clearly what is shown on each time line that you use.

packet arrival at R1

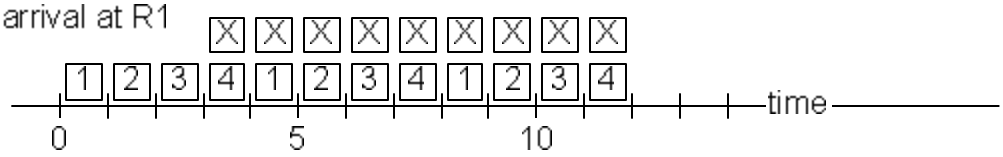

cross traffic queue

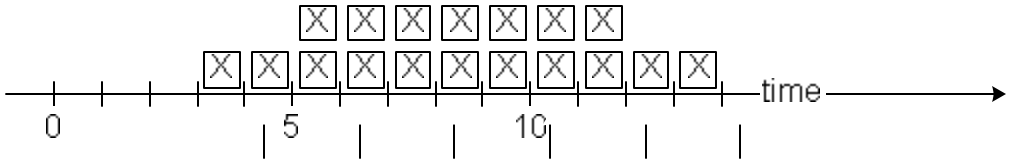

WaveVideo queue

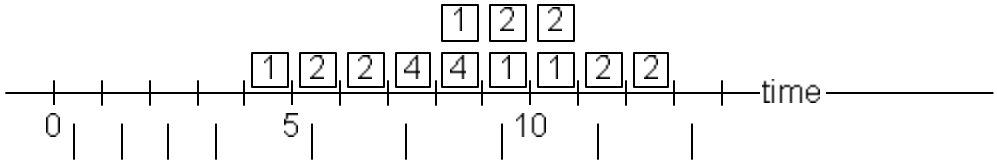

packet arrival at V2

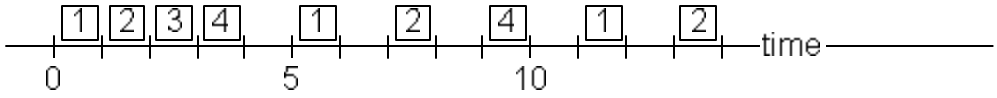

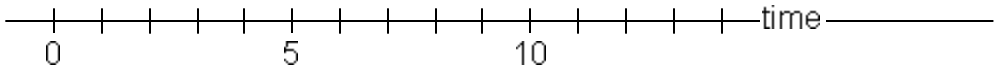

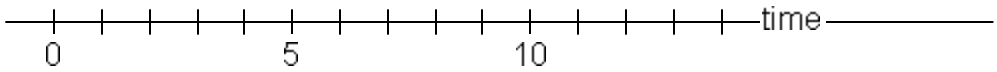

Question 5 (14 points (2+2+2+2+6)):
Answer the following questions regarding safety and security in active networks.

a) SNAP achieves safety in terms of resource usage. Which resources are addressed by SNAP?

Processor usage, memory usages, bandwidth usage.

b) How does SNAP achieve bounds on usage for each of these resources?

By forcing the resource consumption to be proportional to the code size (which is proportional to the packet size and the transmission time). For processing, only forward references are allowed. For memory, a stack is used. For bandwidth, the packet size can only grow by a limited amont.

c) What are the challenges of using cryptographic approaches in an active network?

Once a cryptographic method has been applied to data, that data cannot be modified (either because it's encrypted or because the signature wouldn't match anymore). However, active networking is based on the ability to modify packets.

d) How does the packet format proposed by Murphy et al. address these challenges?

They propose to split the packet into static and varying payload. Cryptographic methods are only applied to the static payload.

(over)

e) Which of the following functions could or could not be implemented with SNAP? Explain why or why not.
    a. Compute checksum over packet header.

Can be implemented, because header is of fixed size.

    b. Compute checksum over packet payload.

Can be implemented if the program code is not part of the payload.

    c. Compute checksum over entire packet.

Cannot be implemented, because it includes the program. The code would have to push itself onto the stack and do computations.

    d. Implement traceroute by adding each hop's address to the packet payload.

Could be implemented up to a certain number of addresses.

    e. Implement multicast by duplicating packet.

Cannot be implemented for all cases. The resource bound needs to be split among duplicated packets and that can be done only so many times. If the multicast tree is larger, some nodes cannot be reached

    f. Encrypt packet payload.

Could be done, but it would not be truly encrypted, because the key would have to be carried in the packet (no flow state on routers allowed).

Question 6 (8 points (4+4)):
Answer the following questions regarding resource management in Darwin:

a) Darwin uses Xena, delegates, and hierarchical schedulers for resource management. Discuss these components and contrast their levels of operation.

Xena performs application- and flow-level allocation of resources. Xena operates on the (sub-)network level. Delegates operate on node levels where they monitor resource usage and might make local decisions. Schedulers operate on a per-packet basis and enforce the resource allocations done by Xena and the delegates.

b) In Darwin, a strong emphasis is put on hierarchical structuring. Why do the authors think this is important?

Because with a hierarchical allocation of resources, these resources can be sub-allocated locally. I.e., if a service provider sells some bandwidth to a company, the company can allocate parts to different departments. This makes the allocation process much simpler and more manageable (because the service provider does not have to get involved for each sub-allocation).

Question 7 (12 points (3+3+3+3)):
If you were working for a network service provider and your manager asked you to implement the new services below, how would you go about it (i.e., would you use an active network or a non-active solution)? Give some supporting arguments for your decision. (There is no right or wrong answer.)

a) Multicast.

Use end-system multicast, because it can be deployed easily.

Or: use IP multicast, because overlay multicast uses more bandwidth (especially on access links).

b) Web caching.

Has been implemented successfully as proxies on hosts (non-active).

Or: use active network to make it transparent for end-system.

c) Transcoding of web documents for mobile, wireless devices.

Use proxy (non-active) for simplicity or use active network for transparency.

d) Blocking mechanism for peer-to-peer networks.

Use simple firewall, but it's limited to know patterns.

Or: use active network to be able to scan connection content and decide to block or not.