Prof. Noronha
Spring 2001

EE384B: Multimedia Networking and Communications

Quiz #3 Solutions
Closed Book
Time: 30 minutes

May 31, 2001

Score: ___ / 24

Name:____

E-mail:____

☐ SITN student ☐ On-campus student

I agree to abide by the Stanford Honor Code:____
(signature)

Question 1: RTCP Reports (3 points)

For your reference, the RTCP Sender and Receiver Report headers are reproduced below (from RFC 1889).

V=2	P	RC	PT=SR=200	LENGTH
SSRC OF SENDER				
NTP TIMESTAMP, MOST SIGNIFICANT WORD				
NTP TIMESTAMP , LEAST SIGNIFICANT WORD				
RTP TIMESTAMP				
SENDER'S PACKET COUNT				
SENDER'S OCTET COUNT				
. . . RECEPTION REPORTS . . .				
. . . PROFILE-SPECIFIC EXTENSIONS. . .				

V=2	P	RC	PT=RR=201	LENGTH
SSRC OF SENDER				
SSRC_1 (SSRC OF FIRST SOURCE)				
FRACTION LOST	CUMULATIVE NUMBER OF PACKETS LOST			
EXTENDED HIGHEST SEQUENCE NUMBER RECEIVED				
INTERARRIVAL JITTER				
TIME OF LAST SR (LSR)				
DELAY SINCE LAST SR (DLSR)				
. . . ADDITIONAL RECEPTION REPORTS. . .				
. . . PROFILE-SPECIFIC EXTENSIONS. . .				

a) Based on the Receiver Reports, the sender is capable of computing the round-trip delay between it and the receiver. Explain how the sender does it. (1 point)

The sender does it based on the RR. The fields are the LSR, which is the sender NTP timestamp for the last SR received from that sender, and the DLSR, which is the delay from the reception of the last SR to the transmission of this RR. When the sender receives the RR, it looks at the time; let A denote the time at which the RR was received. The round-trip delay is simply A -LSR - DLSR.

b) RTCP reports can be used by a sender to adapt to network conditions. Identify which information the sender will use to adapt, and explain how the sender adapts to network conditions based on the information you identified. (2 points)

The network conditions are indicated by three fields in the RR: the ***fraction lost****, the* ***cumulative number of packets lost****, and* ***the interarrival jitter****. There is really nothing the sender can do about the interarrival jitter; however, if packets are being lost in the network (as indicated by the fraction lost and the cumulative number of packets lost), the sender may assume that there is network congestion, and will reduce its transmission rate (if it is capable of doing it); for example, if the sender is transmitting an MPEG stream, it may reduce the encoder data rate.*

Question 2: RTP Profiles (4 points)

a) The RTP specification is incomplete, on purpose. It needs a profile for actual interoperation of specific media. Why was RTP designed this way? (1 point)

RTP was designed this way for flexibility. As new types of media are introduced, the protocol can be adapted by defining new profiles, which will specify the usage of some of the fields in the RTP header, and may define additional headers.

b) There are a few fields in the RTP header that are defined by the profile. The timestamp is one of them. In some profiles, the timestamp is monotonically increasing; however, in some other profiles, you can find successive RTP packets with the same timestamp, or even RTP successive RTP packets where the timestamp of the later packet is ***less*** than the timestamp of an earlier packet. Why is that? (1 point)

The timestamp is generally defined as the playback time for the media content in the packet. When it comes to video, a frame may span several packets (as in the case of MPEG and JPEG); in this case, all the packets containing data from the same frame will have the same timestamp. In the case of MPEG video elementary streams, because of the re-ordering of P and B frames for transmission, the timestamp can even go back (because the earlier frame is actually transmitted later).

c) RFC 2733 defines an RTP profile for Forward Error Correction. However, RTP runs on top of UDP (which includes a checksum over the data), and most physical layers also include a CRC over the layer-2 packet. Therefore, the likelihood that an RTP receiver will get a packet with errors in it is infinitesimal; packets with errors will be discarded in the lower protocol layers prior to being delivered to the application. What is the purpose, then, of RFC 2733? Why is it useful? (2 points)

One special case of Forward Error Correction is what is called "Erasure Codes", whereby the data is not wrong, it is missing (and can be recovered from the FEC packets). This is important on networks to recover from packet loss due to congestion or any other problems (or even a packet that is discarded due to errors, although with current technology this is not very common).

Question 3: Session Protocols (SAP/SDP) (5 points)

a) The Session Description Protocol (SDP) is a textual protocol, i.e., the protocol fields use human-readable ASCII text. Other, more traditional protocols, such as IP, RTP and even SAP use binary encoding for their fields. Identify one advantage and one disadvantage of using text encoding over binary encoding. (2 points)

Advantages of using text encoding: easier to implement and debug; data can be processed by standard text-processing tools.

Disadvantage: efficiency - text encoded messages use more bits than binary-encoded messages, and this leads to higher bandwidth utilization for what is essentially overhead.

b) SAP announcements are multicast periodically on a well-known address and port. However, RFC 2974 requires that the actual interval between two successive announcements be set to **2(1+*a*)*T*/3**, where ***T*** is the announcement period and ***a*** is a random variable, uniformly distributed between 0 and 1. Why is that? (1 point)

This is done to randomize transmission times and try to "spread" announcements from multiple sources. If they were all to use the exact same interval, there is a possibility that announcements get synchronized and cause instantaneous network congestion.

c) RFC 2974 requires that the announcement period ***T*** in seconds be set to **max(300; (8*n*S)/B)**, where ***n*** is the number of announcements in the system, ***S*** is the announcement size in bytes, and ***B*** is the configured bandwidth. Explain how ***n*** is determined and what is the purpose of setting the announcement period according to this formula. (2 points)

*The number of announcements in the system (**n**) is determined just by listening to the announcement channel, and counting the announcements. The purpose of setting the announcement as indicated above is to keep the overall announcement traffic load under **B**; each announcer will throttle its messages so that it only takes **1/n** of **B**.*

Question 4: RTSP (2 points)

In the "universe" of IP multimedia protocols, we have RTP to carry the stream; SDP to describe the session; SAP to carry the SDP to the receiver; and RSVP to reserve the resources. Given all that, what exactly is the purpose of RTSP? (In other words, why do we need it? What function does it perform that is not included in the protocols above?)

RTSP is a control protocol, used by a client to control a video server. In a live multicast (or unicast) stream, SAP/SDP are enough to make the "connection" between the sender and the receiver. However, if a client wants to request content from a video server (or wants a video server to record a stream), there is no provision in any of the protocols above to make that happen. RTSP is the means to do that.

Question 5: H.323 (4 points)

Opening an H.323 Version 1 connection is a four-step process. Each of the four steps is indicated below. For each step, identify which information is typically exchanged. (1 point each)

a) H.225.0 Terminal to Gateway Signaling (RAS)

- *Registration (done only once; not needed for every call)*
- *Endpoint location*
- *Admission Request*

b) H.225.0 Call Signaling

- *Actual call setup and establishment between the two parties*
- *Call acceptance/rejection*
- *User information*

c) H.245 Control Channel

- *Capabilities Exchange*
- *Determination of the Selected Communication Mode*
- *If this is a conference, negotiation of the conference mode*
- *Opening of the logical channels*

d) H.245 Logical Channel(s)

- *Actual media (audio, video, etc.)*

Question 6: H.323 Version 2 Fast Connect (3 points)

Use the signaling exchange in question 5 as a reference. One of the new features introduced in H.323 Version 2 is called "Fast Connect". Regarding "Fast Connect":

a) What is the issue that "Fast Connect" is solving? (1 point)

The issue is that the standard H.323 negotiation is complex and takes several round-trip delays to actually open a channel. Fast Connect is used to speed-up this process and decrease setup latency.

b) How does "Fast Connect" solve this issue? (2 points)

Fast Connect solves the latency issue by skipping the third phase of the negotiation - the setup of the H.245 control channel. The originator requests the use of Fast Connect during the H.225 Call Signaling phase; if the other side agrees, it can start sending as soon as it receives the request (and the originator must be prepared to accept the stream as soon as the request is made).

Question 7: VoIP Issues (3 points)

When designing a VoIP transport, one of the variables is the IP packet size to use. There are reasons for using large packets, and reasons for using small packets. Make the argument for both cases, i.e., indicate what are the reasons for using large packets, and what are the reasons for using small packets. Conclude by making a recommendation on which way to go - large or small.

Reason for using large packets: efficiency. The IP/UDP/RTP headers represent a fixed overhead in bytes; therefore, the larger the packet, the more efficient the transfer is. If the stream rate is ***R*** *bits/second, the data payload size is* ***P*** *bytes, and the header is* ***H*** *bytes, the bandwidth wasted on headers is* ***RH/(H+P)****; the larger* ***P*** *is, the lower the overhead.*

Reasons for using small packets: (1) latency; the first bit in the packet has to wait until the packet is completely full before it is transmitted; (2) susceptibility to packet loss; if a small packet is lost, the impact in the data is small; if a large packet is lost, the impact is large.

Recommendation: in VoIP, latency is a big issue due to the fact that the communication is interactive; in other words, if latency is high, the system just does not work. Therefore, small packets are indicated, and one just has to accept the overhead.

Bonus Question: T.120 MCS Topology (4 additional points)

Consider the following scenario: a number of T.120 terminals connected by an Ethernet segment. Logically, the T.122/125 MCS providers are arranged in a tree, but physically they are connected to a bus. Now, assume that one node in this conference wants to transmit a binary file to a set of other nodes, using T.127. The standards require that this file be transmitted using the logical tree, from the source node to the top MCS provider and then back to the destinations. In this Ethernet, this is really inefficient; the source could directly multicast the data to all the destinations simultaneously and be done with it. Why does T.120 force the communication to be this way? In other words, what are we gaining by sending the data over this logical tree?

T.120 requires ordered, reliable multipoint communication to follow the tree. What is being gained is the guarantee that the packets are all delivered in order, and that packet loss is detected and dealt with. This is made possible by decomposing the multipoint communications problem into multiple point-to-point communications, independently managed.

In the case of an Ethernet, this is clearly sub-optimal. However, the same abstraction is used independently of the actual physical topology of the network, and this is the price to pay for that logical separation. Could a better protocol be devised, specifically for Ethernet? Yes.