Prof. Noronha
Spring 2001

EE384B: Multimedia Networking and Communications

Quiz #2 Solutions
Closed Book
Time: 30 minutes

May 10, 2001

Score: ___ / 30

Name:____

E-mail:____

☐ SITN student ☐ On-campus student

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

Question 1: Data Traffic Characteristics (4 points)

Data traffic measurements over the years have found that data traffic is *"self-similar"* and *"heavy-tailed"*, and have proposed some analytical models to fit the observed data.

a) Define *"self-similar"* and *"heavy-tailed"*. (2 points)

Self-Similar: the statistical characteristics of the traffic are the same over several ranges of time scales (this implies in large-scale autocorrelation). Formally, self-similar processes are characterized by hyperbolically decaying autocorrelation functions, which are non-summable.

Heavy-tailed: from a conceptual point of view, a heavy-tailed distribution has a significant "probability mass" on its tail. For example, if the distribution represents waiting time and it is heavy-tailed, the longer you have waited the longer is your expected future waiting time. The formal definition is that $P[X \geq x]$ *tends to* $cx^{-\beta}$ *as* $x \rightarrow \infty$, $\beta \geq 0$.

b) What are the consequences of self-similarity when it comes to mixing data traffic with multimedia traffic? (2 points)

The consequences are:

- Traffic Smoothing is not possible.
- Large buffers and high bandwidth are required to prevent congestion and keep packet loss low; this results in low utilization and high latencies.
- There is a fundamental incompatibility between data traffic and multimedia traffic; in the general case, they cannot be mixed without any kind of differentiated treatment for multimedia traffic.

Question 2: Advances in Network Infrastructure - Bridges (4 points)

The GARP Multicast Registration Protocol (GMRP) is a mechanism to provide multicast forwarding services at layer-2 similar to those available at layer-3 (provided by protocols such as IGMP). Why is this needed? Give an example.

Without a mechanism such as GMRP, bridges have no knowledge of the multicast group membership. Without that knowledge, they cannot intelligently forward multicast packets; they need to treat multicast in the same fashion as broadcast and send the packets over all links, except the one it came from. This will cause additional useless traffic in the network, as multicast packets get forwarded to segments where there are no group members.

One example would be something like IP television, where multiple channels are multicast, and a viewer "tunes in" by joining a multicast group. If we assume 1.5 Mb/s MPEG-1 streams, and 10 channels, the overall required bandwidth is 15 Mb/s. If a certain user is connected through a 10 Mb/s Ethernet, her segment will be overloaded even if she is not watching any channels.

Question 3: Wireless Networking Infrastructures (4 points)

HomeRF and Bluetooth are very similar to IEEE 802.11 for data traffic. However, both of them have a special feature not present in IEEE 802.11 for carrying voice traffic. What is this feature?

This special feature is the presence of an isochronous TDM time-slot for voice, interleaved with packet traffic. It amounts to reserved, deterministic bandwidth for digital voice applications.

Question 4: Resource Reservation Protocols (4 points)

In the older resource reservation protocol, ST-II, the reservation originates from the source, and reserves resources as it proceeds to the destination. In the newer RSVP, the reservation originates from the destination, in response to PATH messages. Point out two advantages of having the reservation come from destination instead of from the source.

- *Scalability in regards to the number of receivers: the source does not need to create individual reservations for each receiver.*
- *Allows for different receivers to make different reservations for the same stream depending on their individual capabilities in terms of bandwidth and local resources.*
- *For multicast traffic, the source does not need to know the identities of the receivers as they are responsible for the reservation.*

Question 5: RSVP Reservation Styles (4 points)

Consider the following situation: a multipoint-to-multipoint audio/video-conference between a certain number of receivers. Audio and video are carried as separate flows using different multicast addresses. All the receivers are capable of receiving and playing all the audio streams simultaneously. Some receivers are also capable of receiving and displaying a single video stream in addition to the audio. ("Capable" here means that the receiver has both the functionality and is connected with suitable bandwidth). The conference is closed, i.e., only available to a pre-selected group. What RSVP reservation styles would you use for the audio and the video streams? Justify your answers.

*Audio: since all receivers are capable of receiving all audio streams simultaneously, and this is a conference, the most appropriate reservation style for the audio is a **fixed-filter** reservation. This allows all the participants to hear all other participants all the time; also, since the reservation is specific to the conference members, no outsiders can use it.*

*Video: since each receiver can only display one video stream, and since the conference is closed, the most appropriate reservation style for the video is **shared-explicit**, with the conference members being named in the reservation. Note that, from a bandwidth point of view, a wildcard filter reservation would work too, but it is not appropriate since the conference is closed.*

Question 6: DiffServ and MPLS (4 points)

One of the main driving forces behind DiffServ and MPLS is to reduce the processing complexity in the core routers and simplify the design of very large, very fast routers. Each of these techniques primarily addresses one area of the processing required in a modern router. For each technique, identify the area of processing it simplifies and explain how it simplifies that area.

<u>DiffServ</u>: *it addresses the QoS portion, i.e., providing different levels of service to different traffic flows. DiffServ moves the classification portion (i.e., figuring out which packet belongs to which flow) to the edge of the network. Packets are "stamped" in the DS field based on the flow they belong to, and the core routers just consult that field, which is a very fast operation. In summary, DiffServ simplifies the flow classification step for different levels of service.*

<u>MPLS:</u> *it addresses the routing part, i.e., figuring out where each packet needs to be sent. This is again done at the ingress, and the core routers can forward the packets based only on label. More specifically, MPLS simplifies the routing step by replacing the longest prefix match operation with a simple table lookup. Note that the MPLS classification can also be used for QoS in addition to routing.*

Question 7: RTP (6 points)

The RTP header is depicted below:

<table>
<tr><td>0</td><td>1</td><td>2</td><td>3</td><td>4</td><td>5</td><td>6</td><td>7</td><td>8</td><td>9</td><td>10</td><td>11</td><td>12</td><td>13</td><td>14</td><td>15</td><td>16</td><td>17</td><td>18</td><td>19</td><td>20</td><td>21</td><td>22</td><td>23</td><td>24</td><td>25</td><td>26</td><td>27</td><td>28</td><td>29</td><td>30</td><td>31</td></tr>
<tr><td colspan="2">V=2</td><td>P</td><td>X</td><td colspan="4">CC</td><td>M</td><td colspan="7">Payload Type (PT)</td><td colspan="16">Sequence Number</td></tr>
<tr><td colspan="32">Timestamp</td></tr>
<tr><td colspan="32">Synchronization Source (SSRC) Identifier</td></tr>
<tr><td colspan="32">Contributing Source (CSRC) Identifiers (between zero and 15)</td></tr>
</table>

Explain how a receiver would use the following fields in the header:

a) Payload Type (2 points)

The payload type identifies the type of media the packet is carrying. The receiver will look into this field and use it to select the appropriate media decoder for the encoding used in this packet (if available; if not, it will drop the packet).

b) Sequence Number (2 points)

The sequence number allows the receiver to perform two functions: (i) detection of lost packets (discontinuities in the sequence number), and (ii) detection of out-of-order packets (non-monotonic sequence numbers).

c) Timestamp (2 points)

The timestamp field allows the receiver to reconstruct the original media timing, removing network jitter. It can also be used for synchronization between multiple media streams.

Bonus Question (4 additional points):

One of the available techniques for router queue management is Random Early Detection (RED). Its purpose is to avoid the undesirable synchronization and traffic oscillation that is introduced by the interaction of tail drop and TCP, and it works fine for that purpose. However, in some situations, the overall performance of RED is lower than that of tail drop. Construct an application scenario where this is the case (i.e., a scenario were we would be better off with plain tail drop than using RED).

The problem with RED is that it drops packets before the buffer is completely full. In other words, it drops packets that it had resources to transmit. One scenario where this would be undesirable is when there is streaming involved. More specifically, assume a fixed-bandwidth WAN link, and a number of VBR video sources being transmitted over this link. The sum of the average rates of these VBR sources is controlled to be less than the link capacity. However, since they are VBR sources, the instantaneous traffic may exceed the link bandwidth for short periods of time and cause the buffer to fill; RED may drop packets in this situation, causing the video to glitch unnecessarily; tail drop would work fine as long as the burst sizes do not exceed the available buffering. TCP can be flow-controlled to a certain extent by dropping packets, but not continuous media.