Name .... ID .... Section ...... Seat No ......

Sirindhorn International Institute of Technology
Thammasat University

Final Exam Answers: Semester 2, 2010

Course Title: ITS413 Internet Technologies and Applications

Instructor: Steven Gordon

Date/Time: Wednesday 9 March 2011; 13:30–16:30

Instructions:

- This examination paper has 19 pages (including this page).
- Conditions of Examination: Closed book; No dictionary; Non-programmable calculator is allowed
- Students are not allowed to be out of the exam room during examination. Going to the restroom may result in score deduction.
- Students are not allowed to have communication devices (e.g. mobile phone) in their possession.
- Write your name, student ID, section, and seat number clearly on the front page of the exam, and on any separate sheets (if they exist).
- The space on the back of each page can be used if necessary.
- Reference material at the end of the exam may be used.

Internet Technologies and Applications, Semester 2, 2010

Prepared by Steven Gordon on 4 March 2011
ITS413Y10S2E01, Steve/Courses/ITS413/Assessment/Final-Exam.tex, r1721

Question 1 [34 marks]

Consider an IPTV network to be built across Bangkok by the company called *False Corporation*. The network will deliver the "triple-play service" to subscribers: video, voice and Internet across the single network. The video will include standard definition digital TV and video-on-demand, while voice will be using the G.726 codec. You are a consultant to *False*, advising them on the network design.

(a) Explain the difference between ADSL2+, FTTH and FTTN, as options for the service provider access network. State the transmission media they use and the advantages/disadvantages of each technology in your explanation. [3 marks]

Answer. *ADSL2+ uses copper line to the home from a telephone exchange serving a small suburb. FTTN uses optical fibre to a special node that is closer to homes, then uses copper to the homes. It serves 100's of users. FTTH delivers optical fibre direct to a home, avoiding copper lines. The advantage of ADSL2+ is that it uses existing telephone networks. FTTN increases speeds (because the copper link is shorter) compared to ADSL2+ but requires extra cost for the nodes. FTTH delivers the highest speed but at the highest of installing optical fibre all the way to each home.*

(b) *False* will need to sell equipment for the home network. What network technology do you advise to be used inside the home network and why (give two reasons)? [3 marks]

Answer. *IEEE 802.11 Wireless LAN (g or n) or IEEE 802.3 Ethernet. They are relatively cheap, readily available in products and have sufficient capacity (50Mb/s+) for the services.*

(c) Draw a diagram illustrating the network topology of the entire network, showing at least: core network, service provider access network, home network, PSTN, Internet, TV network. Also identify the video headend, and two subscribers on separate access networks. [4 marks]

You estimate that each standard definition video session requires 5Mb/s of core network bandwidth. The core network capacity that can be allocated to video is planned to be 10Gb/s. (In the following three parts, if you do not have enough information to determine an exact number, then explain what will impact on the number).

(d) What is the maximum number of users and/or channels of standard definition TV that can be supported by the core network if *multiple unicast* is used as the delivery mechanism? Explain your answer, stating any assumptions. [2 marks]

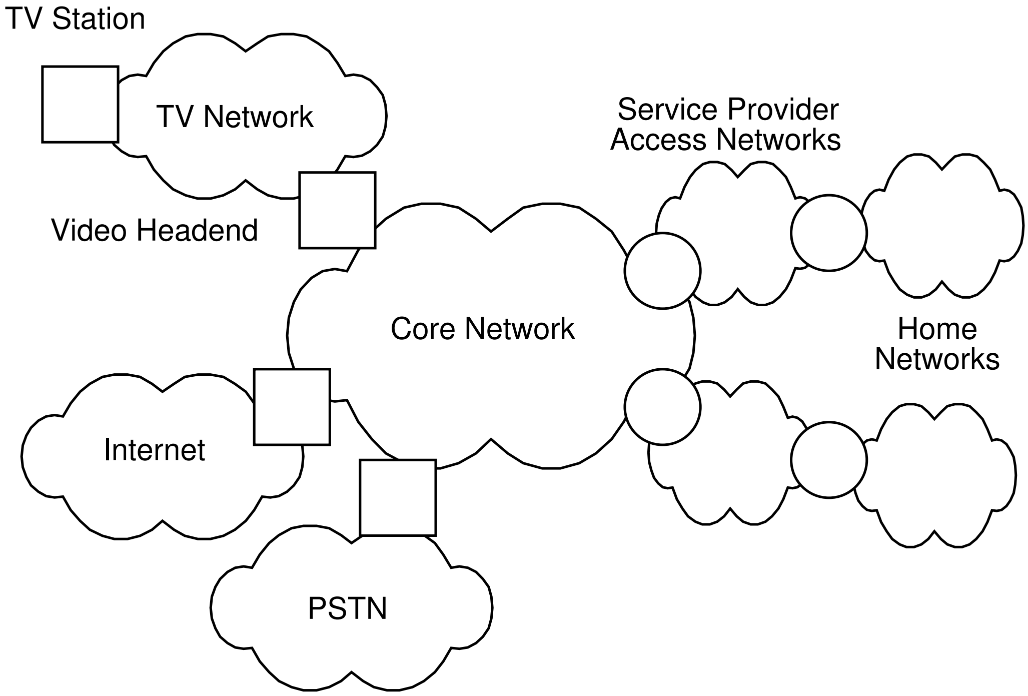

Answer. *If each user is watching 1 channel, then 2000 users can be supported.*

(e) What is the maximum number of users and/or channels of standard definition TV that can be supported by the core network if *multicast* is used as the delivery mechanism? Explain your answer, stating any assumptions. [2 marks]

Answer. *2000 channels (unlimited number of users in theory, limited by multicast)*

(f) What is the maximum number of users and/or channels of standard definition video-on-demand that can be supported by the core network? Explain your answer, stating any assumptions. [2 marks]

Answer. *If each user is watching one video-on-demand, then 2000 users can be supported.*

(g) If the core network supported multicast, then explain what IGMP is used for. [1 mark]

Answer. *The Internet Group Management Protocol is used for users to subscribe to a multicast group. For example, if each group is a channel, then switching channels can be performed by subscribing to the appropriate multicast group with IGMP.*

Now consider the voice traffic to be sent across the network. The G.726 codec produces a sample size of 20 Bytes and has a sample interval of 5ms. It generates a RTP packet every 20ms.

(h) What protocol should be used for signalling in the IPTV network? [1 mark]

Answer. *SIP or H.323*

(i) If 1Gb/s of the core network capacity is allocated to voice calls, what is the maximum number of voice calls supported? (You should assume the data link layer and physical layer contribute 5 Bytes of header in total to each IP datagram) [5 marks]

Answer. *Four samples per RTP packet means 80 Bytes of payload plus 12 Bytes RTP header, 8 Bytes UDP header, 20 Bytes IP header and 5 Bytes DLL/PHY header. 125 Bytes sent every 20ms, means 6250 Bytes per second or 50kb/s. That allows 20,000 simultaneous calls.*

(j) An alternative to the G.726 codec is the G.711 codec which has 80 Byte samples, 10ms sample intervals and generates a RTP packet every 20ms. What is an advantage of using G.711 instead of G.726 in the network? [2 marks]

Answer. *G.711 produces higher quality voice calls, at the expense of using more bandwidth.*

You need to recommend QoS mechanisms for the network. Two options available are service differentiation or soft QoS using DiffServ, and hard QoS using IntServ.

(h) Explain the difference between soft QoS and hard QoS. [2 marks]

Answer. *Soft QoS allows priority to be given to specific packets, but cannot guarantee applications/hosts/packets will receive an absolute performance guarantee, e.g. never less than 1Mb/s. Hard QoS provides performance guarantees, e.g. an application is guaranteed to receive at least 1Mb/s.*

(i) Explain how hard QoS is provided in IntServ. [2 marks]

Answer. *Before applications that require hard QoS start, a resource reservation request (using RSVP) is sent. If sufficient resources are available to support the application, the session may be admitted into the network (otherwise it is rejected). Once admitted, the network (routers) use queuing mechanism to guarantee the packets for that session are given the allocated resources.*

Consider soft QoS being used to provide priority to video and voice packets from subscribers that have paid for a premium service. Inside the core network, routers check the DiffServ field (also called Type of Service (ToS) or DiffServ Code Point (DSCP)) in the IP header to determine the priority to give packets. There are two values of the DiffServ field available:

000000 Best effort traffic

101110 Voice and video of premium users

For packets sent by subscribers' computers, the DiffServ field is initially ignored. Instead a router on the edge of the core network classifies packets and marks the DiffServ field.

(j) Explain how this edge router can classify packets, giving two examples that illustrate classifying based on type of user (premium or normal) and type of application (voice/video or other). [3 marks]

Answer. *The router can classify based on fields in the packet headers. For example, the source IP address indicates which subscriber sent the packet; a router can determine if this is a premium subscriber. The source/destination port numbers indicate the application: voice, video, web, email.*

(k) Explain how other routers in the core network can provide priority to packets marked with DiffServ field '101110'. [2 marks]

Answer. *Once the router classifies the packet to receive a certain QoS, it may use different queueing schemes to determine when to send the packet. E.g. video packets go to the head of the queue, data packets to the end of a queue.*

Question 2 [9 marks]

Consider the times at which six packets were transmitted by a source and received by the destination (all times are relative to an initial clock value and measured in milliseconds (ms); the clocks at source and destination are synchronised):

Packet 1 Transmit time: 20; Receive time: 42

Packet 2 Transmit time: 40; Receive time: 62

Packet 3 Transmit time: 60; Receive time: 80

Packet 4 Transmit time: 80; Receive time: 103

Packet 5 Transmit time: 100; Receive time: 119

Packet 6 Transmit time: 120; Receive time: 140

(a) What is the average delay experienced in the network? [1 mark]

Answer. *Average of: 22, 22, 20, 23, 19, 20 = 21ms*

(b) What is the jitter experienced in the network? [2 marks]

Answer. *Average of: 0, 2, 3, 4, 1 = 2ms*

(c) Playback buffers are often used to compensate for jitter. Explain how a playback buffer can be used in this case, and how it reduces the effect of jitter. [2 marks]

Answer. *When a packet is received, it is buffered before played at the receiver. The time of buffering is such that the playback occurs at a regular interval.*

(d) Using a playback buffer, what is the preferred playback time of each of the 6 packets? [2 marks]

Answer. *43, 63, 83, 103, 123, 143.*

(e) One disadvantage of playback buffers is the additional complexity/memory needed. What is another disadvantage of using a playback buffer? Use the example six packets to explain. [2 marks]

Answer. *An additional delay is introduced before playback starts. In this example, playback starts at time 43 (delay of 23ms), as opposed 40 (delay of 20ms), i.e. an extra 3ms delay.*

Question 3 [14 marks]

(a) Explain the role of an indexer in a Bittorrent network. [1 mark]

Answer. *An indexer maintains a list of .torrent files and associated descriptive information. Used for searching for .torrent files.*

(b) Explain the role of a tracker in a Bittorrent network. [1 mark]

Answer. *A tracker manages the set of peers in a swarm accessing a torrent. It maintains a list of peers in the swarm and statistics about the swarm.*

(c) What application protocol does a Bittorrent client use to communicate with a tracker? [1 mark]

Answer. *HTTP*

(d) Once a peer obtains a .torrent file, can Bittorrent be considered as a fully distributed (or de–centralised) system? Explain your answer. [2 marks]

Answer. *No. The swarm relies a the centralised tracker which keeps track of the peers and swarm statistics. If the tracker fails, then the system will (eventually) fail. This can be overcome by using multiple trackers and DHT techniques.*

Assume a peer, N_1, has joined a swarm and established two connections to peers N_2 and N_3. Peer N_1 wants to download a torrent with 100 pieces: $P_1, P_2, P_3, \ldots, P_{100}$. Each piece has 10 blocks (e.g. piece P_1 has blocks $B_{1,1}, B_{1,2}, \ldots, B_{1,10}$; piece P_2 has blocks $B_{2,1}, B_{2,2}, \ldots, B_{2,10}$). Peers N_2 and N_3 already have the following pieces:

N_2: $P_5, P_{10}, P_{11}, P_{13}, P_{20}, P_{23}, P_{30}, P_{39}, P_{64}$

N_3: All pieces *except* pieces P_5, P_{20}, P_{30}

In the Peer Exchange Protocol, after an initial *Handshake*, each peer exchanges a *Bitfield* message which indicates the pieces they have available. Then a peer may send a *Request* message to request a specific block, and receive a *Piece* message containing a specific block.

(e) Assuming only peers N_1, N_2 and N_3 are in the swarm, what is the availability of the torrent? Explain you answer. [2 marks]

Answer. *The availability is 1.06; of the 100 pieces in the torrent, all 100 pieces are available, and there are 6 pieces that have two copies*

(f) If there were another two seed peers in the swarm, N_4 and N_5, what would the availability be? [1 mark]

Answer. *As they are seed peers, they both have all 100 pieces. The availability will be 3.06.*

(g) In the *Bitfield* message sent from N_2 to N_1, what values will be included? [1.5 marks]

Answer. *The list of pieces that N_2 has: $P_5, P_{10}, P_{11}, P_{13}, P_{20}, P_{23}, P_{30}, P_{39}, P_{64}$*

(h) Assume N_1 uses the "rarest-piece first" algorithm to select the ordering of pieces to download. If connected to N_2 and N_3 which pieces will N_1 NOT download first? [1.5 marks]

Answer. *There is one copy of many pieces, and only two copies of the following pieces: $P_{10}, P_{11}, P_{13}, P_{23}, P_{39}, P_{64}$. These 6 pieces will not be chosen first.*

Each peer maintains four variables for each other peer it is connected to: *am_choking*, *am_interested*, *peer_choking*, *peer_interested*. Consider the values that peer N_1 maintains for the other two peers:

N_2**:** *am_choking=False*, *am_interested=True*, *peer_choking=False*, *peer_interested=True*

N_3**:** *am_choking=False*, *am_interested=True*, *peer_choking=True*, *peer_interested=False*

(i) Will N_1 send a *Request* message to N_2? Explain your answer. [1 mark]

Answer. *Yes, because N_1 is interested in pieces that N_2 has.*

(j) Will N_3 send a *Request* message to N_1? Explain your answer. [1 mark]

Answer. *No, because N_3 is not interested in pieces that N_1 has.*

(k) Can N_1 download pieces from N_3? Explain your answer. [1 mark]

Answer. *No, because N_3 is choking N_1.*

Question 4 [22 marks]

(a) Explain the difference between flow control and congestion control. [2 marks]

Answer. *Flow control is used to ensure the receiver buffer is not overflowed, whereas congestion control is usde to ensure the network (routers) is not overflowed.*

(b) Using any of the following variables, write an equation that gives the approximate sending rate, S, of a TCP source (when not packets are lost): round trip time, RTT; congestion window, $cwnd$; advertised window, $awnd$; slow start threshold, $ssthresh$; maximum segment size, MSS. [2 marks]

Answer. $S = \min(cwnd, awnd)/RTT$

(c) If a link between two computers has a capacity of 10Mb/s and RTT of 5ms, what is the value its bandwidth-delay product? [2 marks]

Answer. *50,000 bits or 6,250 Bytes*

(d) Ignoring congestion control and packet headers, if the receiver computer in part (c) had a maximum buffer size of 5KB, what do you think the approximate throughput will be for the TCP connection? [2 marks]

Answer. *The advertised window of 5,000 Bytes will limit throughput, allowing 5,000 Bytes per RTT giving a throughput of 8Mb/s*

(e) If the maximum buffer size in part (d) was doubled to 10KB, then explain how it will impact on the TCP throughput. [2 marks]

Answer. *The BDP will limit throughput, i.e. we will achieve the capacity of 10Mb/s.*

(f) What method does a TCP source use to reduce its sending rate for congestion control? [1 mark]

Answer. *The source reduces the congestion window,* cwnd.

(g) What event(s) does a TCP source assume indicates decreased congestion? [1 mark]

Answer. *Receipt of an ACK*

(h) Explain the difference between how a TCP source increases its sending rate in the additive increase phase compared to the slow start phase. Also indicate when does the TCP source change from one of these phases to the other. [3 marks]

Answer. *TCP starts in the slow start phase where it increases its sending rate exponentially, doubling the rate every RTT. When the congestion window reaches the slow start threshold TCP switches the additive increase where it increases the sending linearly, increasing the rate by one MSS every RTT.*

(i) In TCP, packet losses are detected by the source when either a timeout occurs or 3 duplicate ACKs are received. Explain how the TCP congestion control algorithm responds to each of these packet loss events. (You don't need to give exact algorithms, but should indicate how and why the congestion control algorithm responds differently to the two different events). [3 marks]

Answer. *In both cases TCP assumes the events indicate increased congestion and reduces the sending rate. Normally 3 duplicate ACKs will be receive before a timeout the congestion is small (only 1 packet lost), and hence this results in a small decrease in sending rate. If a timeout occurs then it is assumed congestion is large and results in a large decrease in sending rate.*

(j) Explain what is meant if TCP is described as "fair". [2 marks]

Answer. *If multiple TCP connections share the same bottleneck link, the capacity of that link will be equally divided amongst the connections.*

(k) Computers A, B and C share a bottleneck link with capacity 8Mb/s to the Internet. On computer A is a Bittorrent application that is downloading from 6 peers in a swarm. On computer B is a web browsing application that is downloading a file from a web server. Computer C also has a web browsing application downloading from a different web server. What is the approximate download rate achieved by the applications on each computer? (that is, give the download rate for application on A, application on B and application on C) [2 marks]

Answer. *There are 8 TCP connections sharing the bottleneck link. As TCP is fair between connections computer A obtains 6/8 of the capacity, i.e. 6Mb/s, while computers B and C obtain 1/8 of capacity, i.e. 1Mb/s each.*

Question 5 [10 marks]

Two challenges of using P2P systems for sharing resources are: searching and data transfer. This question is only about searching.

(a) In Napster-like P2P systems an index is stored on a central server. What is the index (that is, what important information does it contain)? [2 marks]

Answer. *The index maps resources or keys to peers. That is, it contains a list of keys of resources and the corresponding peers that maintain the resource.*

(b) Where is the index information stored in a FastTrack P2P system? [2 marks]

Answer. *A super-peer stores the index information for its local peers.*

(c) Explain how searching works in a FastTrack P2P system. As a guide, you should clearly explain the steps for all typical cases of where the queried resource may be located (e.g. the conditions when a response to a query is returned). You should state which nodes send the queries, and to what destinations and using what method. [3 marks]

Answer. *A peer sends a query to its local super-peer. If the super-peer maintains an index for the resource, it immediately returns a response. If the super-peer does not have the resource in the index, then the super-peer floods the query to all other super-peers. The super-peers that maintain the index will send a response.*

(d) Explain an advantage that FastTrack has compared to Gnutella. [1.5 marks]

Answer. *Searching can be faster because the super-peers maintain some index information and can immediately respond, whereas in Gnutella no index information is stored about other peers. As a result of the above, another advantage is that the number of messages (overhead) is reduced.*

(e) Explain a disadvantage that FastTrack has compared to Gnutella. [1.5 marks]

Answer. *In FastTrack if a super-peer fails, then its local peers cannot search. Whereas in Gnutella if one peer fails then the others can still search.*

Question 6 [11 marks]

You are designing a P2P system for file sharing within your company using a Chord Distributed Hash Table. The maximum number of users in the company is 1,000, while it is expected that there will be no more than 1,000,000 different resources (files) to be shared.

(a) You use a hash function that produces a k-bit hash value. What should the minimum value of k be in your network, and why? [2 marks]

Answer. *The hash function should produce unique values for each user (10 bits gives 1024) and for each resource (20 bits gives 1048576). Hence k should be at least 20.*

Despite the minimum value given in part (a), the software you are using requires k to be 32. Assume this value for the following parts. Also assume of the 1,000 users, there are 5 users with the IDs: 205, 443, 444, 500 and 680. There are no other users with IDs between 205 and 680 (that is, the other 995 users have IDs outside this range).

(b) What is the maximum number of neighbours a peer may have in its routing table? [1 mark]

Answer. *32*

(c) Which keys may peer 500 index? [1 mark]

Answer. *Key 500 down to 445*

(d) Which, if any, of the other four peers (443, 444, 500 680) does peer 205 have a route to? Explain your answer. [3 marks]

Answer. *Peer 205 should have routes to neighbours that are 1, 2, 4, 8, ... positions away. The closest neighbour is 443 (238 positions away). Peer 205 has a route to 443 (which covers positions 1, 2, ... 128). The next neighbour 256 positions away is peer 500. The next neighbour 512 positions away is greater than 680 (we don't know its ID). Hence 205 maintains routes to 443 and 500.*

(e) If a new peer with ID 300 joins the system, explain what needs to happen. [2 marks]

Answer. *Peer 205 will need to update its neighbours (peer 300 will be a neighbour). Also the index of keys needs to be updated: keys 300 to 206 previously indexed at 443 will be indexed at peer 300.*

(f) Chord uses a specific algorithm for determining which neighbour to maintain routes to. An alternative would be for each peer to maintain routes to the next closest neighbour. Explain both an advantage and disadvantage of the Chord approach, compared to this alternative. [2 marks]

Answer. *Chord requires less neighbours to maintain routes to (updating routes is complex and introduces communications overhead). The disadvantage of Chord against the alternative is Chord requires more queries to search.*

Reference Material

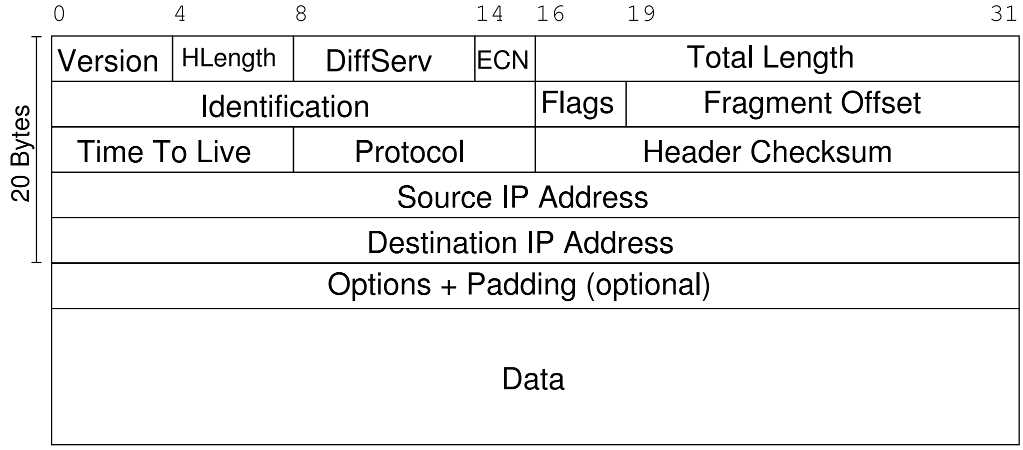

Figure 1: IP Datagram Format. Flags: Reserved, Don't Fragment, More Fragments

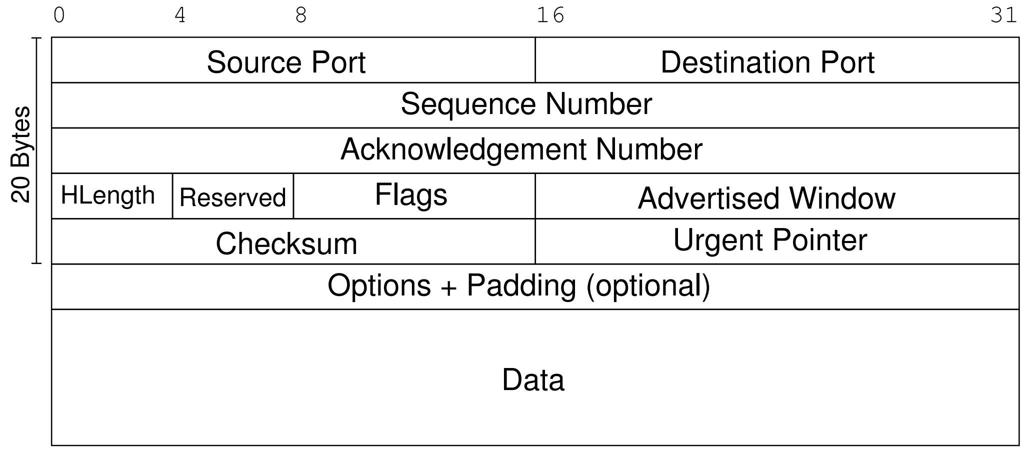

Figure 2: TCP Segment Format. Flags: CWR, ECE, URG, ACK, PSH, RST, SYN, FIN

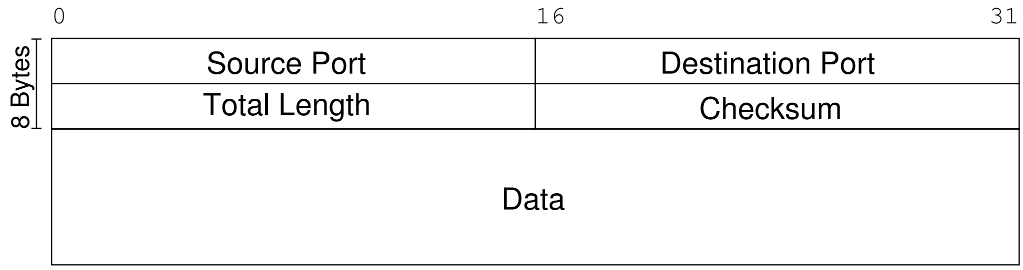

Figure 3: UDP Datagram Format

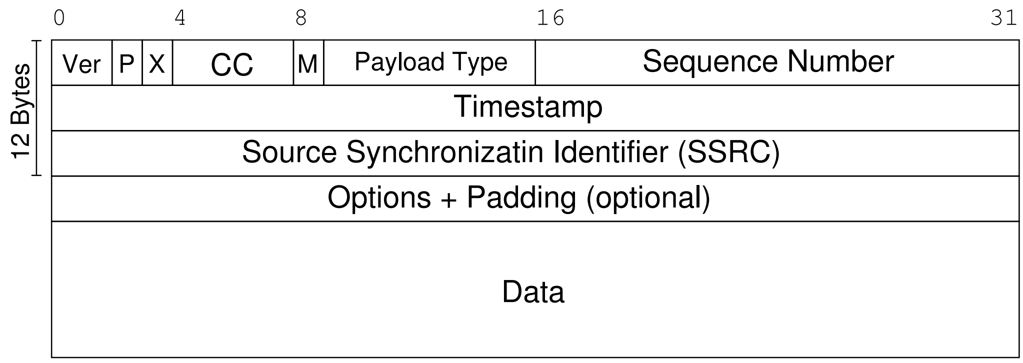

Figure 4: RTP Packet Format. P: Padding; X: Extension; CC: CSRC count; M: Marker