CPSC 441
COMPUTER COMMUNICATIONS

MIDTERM EXAM SOLUTION

Department of Computer Science
University of Calgary
Professor: Carey Williamson

March 2, 2012

This is a CLOSED BOOK exam. Textbooks, notes, laptops, personal digital assistants, and cellular phones are NOT allowed. However, **calculators are permitted**.

It is a 50 minute exam, with a total of 50 marks. There are 12 questions, and 8 pages (including this cover page). Please read each question carefully, and write your answers legibly in the space provided. You may do the questions in any order you wish, but please USE YOUR TIME WISELY.

When you are finished, please hand in your exam paper and sign out. Good luck!

Student Name: ____________________

Student ID: ____________________

Score: _______ / 50 = ______ %

Multiple Choice

Choose the best answer for each of the following 8 questions, for a total of 8 marks.

1 1. In the early days of the Internet, the three primary network applications were:

(a) electronic mail, file transfer, and the Web

(b) network news, file transfer, and the Web

(c) **electronic mail, file transfer, and remote login**

(d) file transfer, ssh, and electronic mail

(e) Facebook, YouTube, and Flickr

1 2. On today's Internet, electronic mail is provided using:

(a) Simple Mail Transfer Protocol (SMTP)

(b) Internet Mail Access Protocol (IMAP)

(c) Post Office Protocol (POP)

(d) Hyper-Text Transfer Protocol (HTTP)

(e) **all of the above**

1 3. The Internet Protocol (IPv4) is an example of a:

(a) **connection-less network layer protocol**

(b) connection-oriented network layer protocol

(c) connection-less transport layer protocol

(d) connection-oriented transport layer protocol

(e) none of the above

1 4. According to the protocol specification, the maximum legal size of an IP packet is:

(a) 512 bytes

(b) 576 bytes

(c) 1024 bytes

(d) 1500 bytes

(e) **64 kilobytes**

1 5. The transmission time required for sending a 512-Byte packet using a 56 kbps modem is:

(a) about 7 milliseconds

(b) about 9 milliseconds

(c) about 40 milliseconds

(d) **about 70 milliseconds**

(e) about 90 milliseconds

1 6. The primary difference between iterated DNS queries and recursive DNS queries is:

(a) iterated queries use port 53, while recursive queries use port 54

(b) recursive queries use port 53, while iterated queries use port 54

(c) iterated queries use UDP, while recursive queries use TCP

(d) recursive queries use UDP, while iterated queries use TCP

(e) **none of the above**

1 7. In the Transmission Control Protocol (TCP), acknowledgements (ACKs) are used for:

(a) error control

(b) flow control

(c) congestion control

(d) error control and flow control

(e) **error control, flow control, and congestion control**

1 8. In typical implementations of TCP (e.g., TCP Reno), the types of ACKs used are always:

(a) selective ACKs only

(b) a mix of selective ACKs and cumulative ACKs

(c) a mix of positive and negative ACKs

(d) **positive, cumulative, lookahead ACKs**

(e) delayed, cumulative, duplicate ACKs

Networking Concepts and Definitions

12 9. For each of the following pairs of terms, **define** each term, and make sure to **clarify** the key difference(s) between the two terms. Give an example, if appropriate.

(a) (3 marks) "circuit-switching" and "packet-switching"

Two different ways of managing the core of a communication network.
Circuit switching: explicit call setup phase; allocates an end-to-end path (and resources); all data flows on same path. Example: classic telephony.
Packet switching: no end-to-end setup; data divided into chunks called packets; each packet independently addressed and routed through network. Example: Internet.

(b) (3 marks) "internetwork" and "the Internet"

An internetwork: general concept; a network of networks (or between networks); more than one network connected together. Example: U of C internetwork.

The Internet: a specific global network of networks, all of which are running the Internet Protocol (IP). Example: the Internet.

(c) (3 marks) "connection-less" and "connection-oriented"

Connection-oriented: a protocol with explicit phases for setup, data transfer, and teardown afterwards; involves state information at nodes. Example: TCP.
Connection-less: a protocol with no advance setup or release of resources, and minimal or no state information involved. Example: UDP.

(d) (3 marks) "pull protocol" and "push protocol"

Refers to data movement paradigm, for example in client-server communication.
Pull: server has data, client must retrieve it. Example: HTTP on the Web.
Push: client has data, must send it to server. Example: SMTP for email.

Application Layer Protocols

10 10. Assume that a typical Web page on the Internet today consists of a base HTML file (about 6 kilobytes) and about 4 embedded objects (with a cumulative total of 12 kilobytes), which may come from the same Web server or from different Web servers.

(a) (3 marks) In class, we argued that retrieving such a Web page using classic HTTP/1.0 would be "slow". Why would it be slow? Identify at least three distinct inefficiencies that contribute to the delay.

```
With HTTP/1.0, there is one object transferred per TCP connection.
1. Repeated TCP connection handshaking incurs extra control packets and RTTs.
2. TCP slow start takes place on each new connection (low throughput).
3. Transfers of different objects are serialized in time (one at a time).
```

(b) (3 marks) Which of these inefficiencies are addressed by the "persistent connection" feature of HTTP/1.1? How effective would it be in this Web page scenario? Why?

```
Persistent HTTP connections can send multibple objects over one TCP connection.
1. Reduces the TCP handshaking overhead.
2. Reduces the repeated TCP slow starts.
Very effective if multiple objects come from the very same server,
but not if they come from different servers.
```

(c) (4 marks) In class, we identified at least three other mechanisms that could be used to reduce the latency for retrieving typical Web pages. List **any two** of these mechanisms, and comment briefly on their potential effectiveness (or not) for this particular Web page download scenario.

```
Parallel connections: open multiple TCP connections concurrently, either to the
same server or to different servers; allows objects to be downloaded in parallel.
Caching: store copies of recently retrieved objects, either in browser cache or
in proxy cache on LAN; use conditional GET; avoid repeated retrievals.

Increase TCP initial window size (IWS): transfer small objects quickly.
gzip: compress objects before transfer, to reduce volume of bytes exchanged.
```

Domain Name Service (DNS)

10 11. The output on the following page was generated by running the `dig` command recently on `ict602c.cpsc.ucalgary.ca`. Use this output, and your knowledge of DNS, to answer as many of the following questions as you can.

(a) (1 mark) What well-known port number is used for DNS queries?

`53`

(b) (1 mark) What is the IP address of `cs.usask.ca`?

`128.233.101.16`

(c) (1 mark) What is the IP address of the DNS server that answered the query in (b)?

`136.159.2.9`

(d) (1 mark) Was this response authoritative?

`No (came from DNS cache at U of C)`

(e) (1 mark) How many mail servers are evident in the `cs.usask.ca` domain?

`Zero (no MX records are evident)`

(f) (1 mark) What is the IP address of `cs.ubc.ca`?

`142.103.6.5`

(g) (1 mark) What is the IP address of the DNS server that answered the query in (f)?

`136.159.2.9`

(h) (1 mark) For how long can the query result from (f) be cached?

`3600 seconds (1 hour)`

(i) (1 mark) Which of the two queries shown generated the faster response?

`cs.usask.ca (1 msec versus 195 msec)`

(j) (1 mark) Which of the two queries shown generated the larger response?

`cs.usask.ca (141 bytes versus 117 bytes)`

```
[carey@ict602c]$ dig cs.usask.ca

; <<>> DiG 9.7.3-RedHat-9.7.3-2.el6 <<>> cs.usask.ca
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 32104
;; flags: qr rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 4, ADDITIONAL: 0
;; QUESTION SECTION:
;cs.usask.ca.                   IN      A
;; ANSWER SECTION:
cs.usask.ca.            12201   IN      A       128.233.101.16
;; AUTHORITY SECTION:
usask.ca.               12201   IN      NS      hymie.usask.ca.
usask.ca.               12201   IN      NS      access.usask.ca.
usask.ca.               12201   IN      NS      delphi.usask.ca.
usask.ca.               12201   IN      NS      zeppelin.sasknet.sk.ca.
;; Query time: 1 msec
;; SERVER: 136.159.2.9#53(136.159.2.9)
;; WHEN: Sat Feb 18 08:15:36 2012
;; MSG SIZE  rcvd: 141

[carey@ict602c]$ dig cs.ubc.ca

; <<>> DiG 9.7.3-RedHat-9.7.3-2.el6 <<>> cs.ubc.ca
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 60274
;; flags: qr rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 2, ADDITIONAL: 2
;; QUESTION SECTION:
;cs.ubc.ca.                     IN      A
;; ANSWER SECTION:
cs.ubc.ca.              3600    IN      A       142.103.6.5
;; AUTHORITY SECTION:
cs.ubc.ca.              3600    IN      NS      fs1.ugrad.cs.ubc.ca.
cs.ubc.ca.              3600    IN      NS      ns1.cs.ubc.ca.
;; ADDITIONAL SECTION:
fs1.ugrad.cs.ubc.ca.    3600    IN      A       198.162.35.1
ns1.cs.ubc.ca.          3600    IN      A       142.103.6.6
;; Query time: 195 msec
;; SERVER: 136.159.2.9#53(136.159.2.9)
;; WHEN: Sat Feb 18 08:18:03 2012
;; MSG SIZE  rcvd: 117
```

Transmission Control Protocol

10 12. One of the unique features of TCP as a transport-layer protocol is its implementation of end-to-end congestion control. Although it was not part of the original TCP protocol design, congestion control was added in 1988 because of the inadequacies of network-layer congestion control in the Internet.

(a) (2 marks) What two new TCP sender state variables are used in TCP congestion control? What is the purpose of each of these state variables?

```
slow start threshold (ssthresh): remembers a recent estimate of a safe operating
point for the congestion window size
congestion window (cwnd): limits the maximum number of segments transmitted
into the network, based on implicit feedback received from the network
```

(b) (6 marks) What two new algorithms were added to provide TCP congestion control? What is the purpose of each? Give a brief overview of the operation of each.

```
Slow Start (SS): Start small, with a minimal window size of 1 MSS.
Increase cwnd by 1 MSS for each successful ACK received. This results in
doubling cwnd every RTT (exponential growth). Continue until ssthresh is
reached, or a loss occurs, or done. Provides estimate of good window size.

Congestion Avoidance (CA): Cautious continuous bandwidth probing for additional
network capacity. Expand cwnd by 1 MSS for each full window of data exchanged
successfully. Linear growth. Upon loss detection, update ssthresh and cwnd.
```

(c) (2 marks) What is the fundamental underlying control principle during the steady-state operation of TCP congestion control? Why is this principle a good one?

```
Additive Increase Multiplicative Decrease (AIMD):
Speed up slowly and cautiously when extra capacity seems to be available.
Slow down quickly when problems occur, to give the network time to clear out.

A safe design principle (from control theory) for stable network operation.
```

*** THE END ***