CS-666 Advanced Distributed System

Homework Assignment #1

Due: 12 February 2018 in class

- Please type the solutions using a word processor such as WORD, or write by hand very neatly and legibly, comparable to typing*.
- Please pay special attention to the due date – no late turn ins or special case consideration.
- Please do the following problems from the text, and submit solutions:

1. Problem 1.1; 1.2; 1.6

2. Problem 2.11; 2.14

3. Problem 4.5

4. Problem 5.22

5. Experimenting with the simulator (group-oriented):

(1) install Matlab on your computer;
(2) download and run the two pieces of Matlab code posted on our class web site;
(3) understand the code and the file assignment paper that everyone has read;
(4) plot and turn in the following two figures by comparing the Greedy algorithm and the Sort Partition algorithm: (a) Average response time for 60/40 skew, one batch (similar to Fig. 9); (b) Average response time for 70/30 skew, one batch (similar to Fig. 8)

Note: The text of HW1 is attached in this document.

1.1 Give five types of hardware resource and five types of data or software resource that can usefully be shared. Give examples of their sharing as it occurs in distributed systems.

1.2 How might the clocks in two computers that are linked by a local network be synchronized without reference to an external time source? What factors limit the accuracy of the procedure you have described? How could the clocks in a large number of computers connected by the Internet be synchronized? Discuss the accuracy of that procedure.

1.6 Use the World Wide Web as an example to illustrate the concept of resource sharing, client and server. What are the advantages and disadvantages of HTML, URLs and HTTP as core technologies for information browsing? Are any of these technologies suitable as a basis for client-server computing in general?

2.11 Consider a simple server that carries out client requests without accessing other servers. Explain why it is generally not possible to set a limit on the time taken by such a server to respond to a client request. What would need to be done to make the server able to execute requests within a bounded time? Is this a practical option?

2.14 Consider two communication services for use in asynchronous distributed systems. In service A, messages may be lost, duplicated or delayed and checksums apply only to headers. In service B, messages may be lost. delayed or delivered too fast for the recipient to handle them, but those that are delivered arrive order and with the correct contents.

Describe the classes of failure exhibited by each service. Classify their failures according to their effect on the properties of validity and integrity. Can service B be described as a reliable communication service?

4.5 The programs in Figure 4.5 and Figure 4.6 are available at cdk5.net/ipc. Modify them so that the client repeatedly takes a line of user's input and writes it to the stream and the server reads repeatedly from the stream, printing out the result of each read. Make a comparison between sending data in UDP datagram messages and over a stream.

5.22 A client makes remote procedure calls to a server. The client takes 5 milliseconds to compute the arguments for each request, and the server takes 10 milliseconds to process each request. The local operating system processing time for each send or receive operation is 0.5 milliseconds, and the network time to transmit each request or reply message is 3 milliseconds. Marshalling or unmarshalling takes 0.5 milliseconds per message.

Calculate the time taken by the client to generate and return from two requests:

(i) if it is single-threaded, and

(ii) if it has two threads that can make requests concurrently on a single processor.

You can ignore context-switching times. Is there a need for asynchronous RPC if client and server processes are threaded?