15-440: Distributed Systems
Problem Solving Assignment 3

School of Computer Science
Carnegie Mellon University, Qatar
Fall 2014

Due Date: October 13, 2014

I) Networking (20 Points):

The Internet is far too large for any router to hold routing information for all destinations. How does the Internet routing schemes deal with this issue? Specifically, your answer should describe how Internet routing handles the following issues:

- ***Information storage at routers:*** How do the routing protocols store the information to route packets to all the computers connected to the Internet? What information is stored? Is this information sufficient to route any packet to any computer on the Internet?
- ***Packet forwarding:*** How does the router use the stored routing information for forwarding each packet? What are the challenges while forwarding? What is the fault-tolerance mechanism employed by routing?
- ***Scalability:*** Explain the scalability challenge for routing over the Internet?

II) Naming (35 Points):

a) Discuss the problems raised by the use of *aliases* in a name service serving a distributed system, and indicate how, if at all, these may be overcome? (5 Points)

b) Assume a Server, ***S***, and a Client, ***C***. Suppose that: (1) ***S***'s directory */jewel* is mounted on ***C***'s directory */ruby/red*, (2) */ruby/red/stone* is a symbolic link to */ruby/stone*, and (3) */ruby/red/stone* is stored in ***S***'s name space as */jewel/stone*. If C attempts to resolve */ruby/red/stone*, which naming resolution approach will produce a correct answer, iterative or recursive? Explain. (5 Points)

c) How does caching help a name service's *availability* in a distributed system? (5 Points)

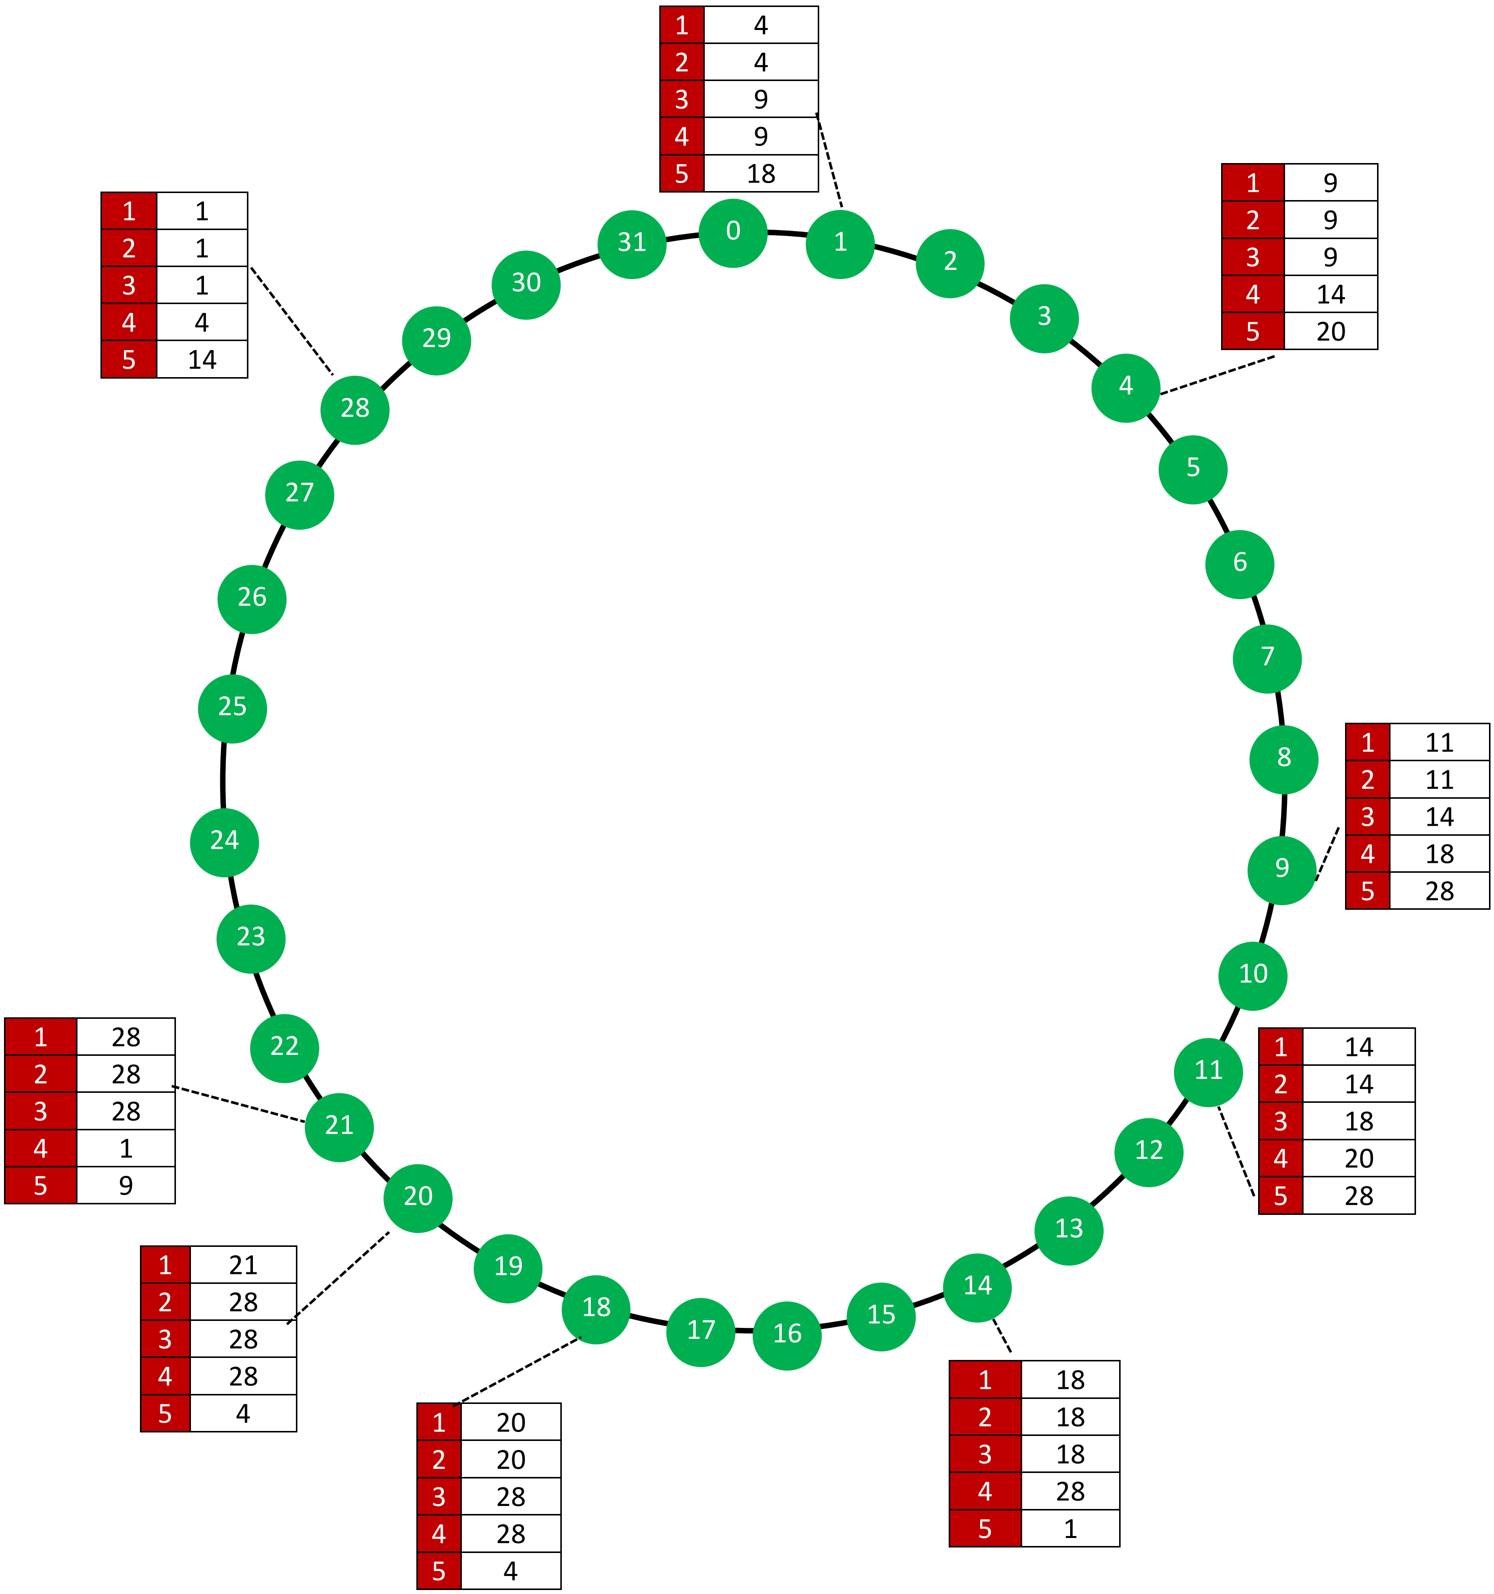

0
1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
1 4
2 4
3 9
4 9
5 18
1 9
2 9
3 9
4 14
5 20
1 11
2 11
3 14
4 18
5 28
1 14
2 14
3 18
4 20
5 28
1 18
2 18
3 18
4 28
5 1
1 20
2 20
3 28
4 28
5 4
1 21
2 28
3 28
4 28
5 4
1 28
2 28
3 28
4 1
5 9
1 1
2 1
3 1
4 4
5 14

d) Consider an entity, ***E***, moving from a source location, ***A***, to a destination location, ***B***. While passing across intermediate locations, ***E*** will reside for only a relatively short time. When arriving at ***B***, ***E*** does not move for a while. Assume the existence of a *hierarchical location service,* HLS. Suppose that changing an address in HLS may take a relatively long time to complete, and should therefore be avoided when visiting an intermediate location. How can you locate ***E*** at an intermediate location assuming the existence of HLS? (5 Points)

e) Consider the Chord system shown in the previous page. (15 Points)

1) Assume that node 10 joins the system. Fill in the values of its finger table accordingly. Make sure to reflect any resultant changes in any of the other finger tables (if any).
2) If a change occurs in a Chord system due to a node joining or leaving, is it required to instantly update the affected finger tables? Discuss.
3) Consider a Chord DHT-based system for which ***k*** bits of an ***m***-bit identifier space are reserved for assigning to superpeers. If identifiers are randomly assigned, how many superpeers can one expect to have in an ***N***-node system?

III) <u>Synchronization (45 Points):</u>

a) Consider the behavior of two machines in a distributed system. Both have clocks that are supposed to tick 1000 times per millisecond. One of them actually does, but the other ticks only 990 times per millisecond. If UTC updates come in once a minute, what is the maximum clock skew that will occur? (5 Points)

b) A client attempts to synchronize with a time server. It records the round-trip times and timestamps returned by the server in the table shown below (i.e., Table I). (10 Points)

1. Which of these times should the client use to adjust its clock? To what time should the client set its clock? Estimate the accuracy of such a setting with respect to the server's clock.
2. If it is known that the time between sending and receiving a message in the system concerned is at least 8 milliseconds, do your answers (i.e., the time setting and the accuracy) change?

Round-Trip (ms)	Time(hr:min:sec)
22	10:54:23.674
25	10:54:25.450
20	10:54:28.342

Table I

c) Show how Lamport's algorithm can be used to accomplish *totally-ordered multicasting* in distributed systems. (15 Points)

d) Consider the centralized mutual exclusion algorithm that we discussed in the class. Suppose that the coordinator crashes. Does this always bring the system down? If not, under what circumstances does this happen? Is there any way to avoid the problem and make the system able to tolerate coordinator crashes? (8 Points)

e) Suggest how to adapt the bully algorithm to deal with temporary network partitions and slow processes. (7 Points)