## CMPSCI 677: Distributed Operating Systems

**Final Exam**

May 10 2011

- Please typeset your solutions if possible. If you hand-write your solutions, please be legible and scan in your solutions for submission.
- This exam is due in 27 hours (at noon tomorrow). You can upload your exam to SPARK (preferred method) or email a pdf or a plain text document to the instructor (no Word documents please).
- Do not forget to put down your name and student number on the exam books.
- If your answer depends on any specific assumptions, please state them clearly.
- If your answer is based on material other than the textbook or class notes, please cite it in your solution (e.g., From Wikipedia page on Distributed systems at http://URL )
- This is an open book, open notes exam. *You shall not seek help from or discuss the exam with ANY human. If you do, you will be awarded a zero for the exam and will be given a grade of "F" for the course.* Send me email if you need a clarification on any question.
- Explain your answers clearly and be concise. *Do not write long essays.*
- Good luck.

---

1. **Short answer questions**

   (a) Construct an example to show that linearizability is a stronger model than sequential consistency.

   (b) Can you implement the event channel service in CORBA using the message queuing systems that we studied in class? Why or why not?

   (c) Rumor mongering is an epidemic protocol where an update is quickly pushed out to nodes that have not yet received it. It is possible for a small number of nodes to miss an update once infected nodes back off and stop spreading updates. How can you make a simple modification to this approach to improve the chances of spreading an update to all nodes in the system?

   (d) Give a simple example to show why the pay-as-you-use billing model of cloud computing allows you to incur a lower server cost for applications with time-varying workloads.

2. **File Systems**

   (a) Why doesn't NFS version 3 support the `open` operation? Given the table on page 9 of the NFS lecture lising NFS operations, what NFS version 3 calls will be made by the NFS client to the server for the following sequence of operations in an applications:

```
open(foo);  /* assume file foo does not exist */
write(128 bytes);
seek(start of file);
read(64 bytes);
close(foo);
```

(b) NFS employs the notion of a transaction ID to detect lost RPC messages (see page 22 of the NFS lecture). Does this imply NFS provides transactional semantics to clients? Explain your answer.

(c) The CODA file system supports transactional semantics and yet does not employ locks to implement these semantics. Explain why. Assuming these transactional semantics, explain how CODA handles read-write and write-write conflicts.

(d) Explain why the log-structured file system discussed in class is optimized for predominantly write-only disk traffic.

3. **Security**

(a) Suppose that your laptop is stolen and it contains a copy of your RSA private key. How can you prevent a unauthorized person from taking this key and masquerading as you (e.g., sending email signed using your private key)? Assume that the key was issued by a trusted third-party.

(b) Consider a online payment service such as Paypal. When making an online purchase, the merchant redirects you to Paypal, where you authorize payment for the purchase amount, and then return to the merchant web site to complete the payment. What are the benefits of this approach over using credit cards to make payment? Does this approach provide anonymity like the E-cash approach discussed in class?

(c) Is setting up a SSL connection with a web server subject to the man-in-the-middle attacks? Explain your answer.

4. **Middleware**

(a) Design a simple approach to migrate a DCOM object from one server to another? Your solution should take the transient nature of DCOM objects into account.

(b) A Jini application reads data from the shared tuplespace using queries. What kind of scalability problems might an application encounter when reading this data?

(c) The year is 2011 and Gauls have entered the Internet age. Every Gaul has a mobile smart-phone with a high speed Internet connection. The village marketplace is now virtual and all trades are done over the Internet. You are given the task of implementing an online trading post using Jini. Explain how you can use a Jini-based application on the smart phones to automatically discover the nearest trading post. Explain how you can use the concept of JavaSpaces in Jini to implement buy and sell requests. Assume that there are two trading servers, one at the north end of the village and the other at the south end. *You do not have to write code. Explain what concepts of Jini can be used to implement the above features and how.*

5. **World Wide Web**

(a) A replicated web server uses a request redirector that redirects incoming requests to one of the servers in replica pool (similar to the redirector you implemented in your class project). Typically

incoming requests are sent to servers in a round-robin fashion or sent to the least-loaded server. Assume that each replica maintains an in-memory cache of recently requested web pages (the pages are normally stored on disk). What request redirection policy should the request redirector employ to take advantage of caches at each server?

(b) Consider a web proxy cache that uses the proxy-pull approach for maintaining consistency of cached web pages. Assume that the proxy employs the intelligent polling approach discussed in class. Is this approach suitable for frequently updated web pages? Is it suitable for infrequently changing web pages? What about pages from news web-sites that are frequently updated during the day and less frequently during the night? Explain your answer.

(c) In the cooperative caching approach discussed in class, what is the primary advantage of maintaining a directory of all web pages cached at various proxies?