CMPSCI 677: Distributed Operating Systems

Final Exam

December 16, 2009

General instructions:

- Please typeset your solutions if possible. If you hand-write your solutions, please be legible.
- This exam is due in 27 hours (at noon tomorrow). You can email your solutions in pdf or plain text format (no Word documents please) or drop off a hard-copy in my office (CS building room 336).
- 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.
- This is an open book, open notes exam. *You may 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.*
- All solutions, even if they are based on a reference, should be your own, i.e., paraphrased and not taken verbatim from the source. Cite all references clearly whether they are research papers or Internet sources.
- Good luck.

---

1. **Short answer questions** (25 points)
   Answer the following questions in brief.

   (a) **Two Phase Locking** Construct an example to show how two phase locking can lead to *cascaded aborts.* A cascaded abort occurs when an abort of a transaction triggers a sequence of aborts of other transactions.

   (b) 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?

   (c) Is linearizability is a stronger or a weaker model than sequential consistency? Give an example to justify your answer.

   (d) **EJBs and Middleware.** Modern web sites such as Amazon.com can maintain the state of the shopping cart of a user across browser sessions. For example, if you put an item in your shopping cart on amazon and then shut down your browser/user session, the item remains in your cart when you return to Amazon at a later time. Assuming that the amazon web site uses Enterprise Java Beans (EJBs), what kind of EJBs would you use to maintain this type of state at the server? Explain your choice.

(e) **Jini and Middleware.** Would you use partitioned or replicated tuple-spaces for read-intensive Jini applications? How would your answer change for write-intensive Jini applications?

2. **Security** (20 points)

(a) Can the RSA public key cryptography algorithm discussed in class be compromised using a brute-force attack? If so, explain under what scenarios this is possible. If not, explain why it is not practical to launch such an attack. Assume that the public key used for encryption is known to the intruder.

(b) Amazon wishes to implement a digital rights management (DRM) scheme for e-books on their popular kindle device. Kindle is a low-end tablet-like computer for readings electronic books. Each book purchased by Alice should have the following property:

- The e-book should be encrypted so that only Alice can decrypt and read it.
- Alice should be able to prove that she purchased the book from Amazon and that she did not download a bootlegged copy from a site such as The Pirate Bay.
- The encryption should be efficient so that low-end processors on Kindle can decrypt and render the book in real-time.
- There should not be any way for Amazon to revoke rights to read the book once Alice has purchased it (e.g., in a recent incident, Amazon revoked rights to certain books from kindle devices of users after they had been legitimately purchased).

Design a simple technique to implement DRM-based encryption with the above properties.
Note: Focus on the concepts taught in class to solve this problem rather than looking up more complex DRM schemes on the Internet or studying Kindle's DRM mechanisms.

3. **Fault tolerance** (10 points)

- In the three phase commit protocol, if a participant is in the READY state, is it possible for a crashed participant to recover to a COMMIT state? If so, construct an example scenario where a participant crashes and then recover in this state. If not, argue why this is impossible.

4. **World Wide Web** (10 pts)

(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. Design a request redirection policy that the request redirector should employ to take advantage of caches at each server?

5. **Distributed File Systems** (20 pts)

(a) Given the figure on lecture 20, page 9, what NFS version 3 calls will be made by the NF S 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) 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.

(c) What are the advantages of distributing parity blocks across all disks in RAID-5 over keeping all parity blocks on the same disk in RAID-4? Does distributed parity in RAID-5 provide better fault tolerance properties than the centralized parity in RAID-4?

6. **Distributed snapshots** (10 pts)

The distributed snapshot algorithm studied in class can capture the global state of a distributed computation. The algorithm presented in class assumed reliable FIFO channels. Modify this algorithm to record global state for non-FIFO channels. A non-FIFO channel can deliver messages out of order. You may assume that all channels are reliable (i.e., no packets are lost). Explain all steps of your algorithm clearly. Be sure to state any assumptions you make.

7. **Cloud Computing** (15 pts)

You are an IT manager responsible for setting up the server infrastructure for your company's web application. Assume that the web application sees a peak workload of 2000 requests/second and that each request takes 1ms of CPU processing time (which is equivalent to 10^6 cpu cycles on a 1GHz processor).

You are considering two options: (i) purchasing servers to setup your own infrastructure to run this application, or (ii) renting servers from a cloud platform to run the application. You must perform a cost analysis to determine which option is cheaper.

Assume that a 1GHz Dell servers costs $1000 and has a life-time of 2 years. Assume that a 1GHz cloud server costs 10 cents/hour to rent.

(a) How many servers are needed for your application? How much annual costs are incurred for the renting versus the buying option?

(b) Next assume that the application sees a peak workload of 2000 requests/second for 11 months of the year and a higher peak workload of 5000 requests/second in December (e.g.,m due to higher online traffic in the holiday season). What are the annual costs for this scenario?

(c) Based on this analysis, what do you think are the key advantages of the cloud platform and it's pay-as-you go model?