

Name: _____

USC ID: _____

CSci 555 Midterm Exam

Fall 2006

Instructions:

Show all work. No electronic devices are allowed. This exam is open book, open notes. You have **100 minutes** to complete the exam.

Please prepare your answers on separate sheets of paper. You may write your answers on the sheet of paper with the question (front and back). If you need more space, please attach a separate sheet of paper to the page with the particular question. **Do NOT extend your answer on the back of the sheet for a different question, and do NOT use the same extra sheet of paper to answer more than one question.**

In particular, **each numbered questions must appear on separate pieces of paper so that the exam can be split for grading.**

Be sure to include your **name** and **USC ID** number **on each page.**

There are **100 points** in all and **3 questions.**

	Q1	Q2	Q3	Total Score
Score				

Name: _____

USC ID: _____

1. (25 points) Consistency and Commitment

Consistency mechanisms for several of the systems described in class exhibit the property of commitment, a step in the process of making a change when the change becomes permanent. For the following systems/techniques, describe the point in the process of making a change that the change is committed.

a) In the two phase commit protocol for distributed transactions [5 points]

b) When writing a file in the log structured file system [5 points]

c) When processing events in the Time Warp system (this one requires more thought than the others, you will likely not find the word commit in the Time Warp paper. Think instead about what it means to “commit” a change.). [5 points]

Name: _____

USC ID: _____

d) For part a (distributed transactions) explain the process through which these systems recover from server crashes, discussing both the case where a change was committed, and the case where a change was not committed. [5 points]

e) For part b (the log structured file system) explain the process through which these systems recover from server crashes, discussing both the case where a change was committed, and the case where a change was not committed. [5 points]

Name: _____

USC ID: _____

2. (35 points) Replication and caching

The use of replication and caching in distributed systems improves performance and reliability, but it requires attention to consistency of data.

a. Place the following systems in order (a partial order is OK if multiple systems fall into the same category) in terms of the strength of their consistency semantics (i.e. how strong is the guarantee that data remains consistent. [10 points]

Whole file caching in AFS-1

Whole file caching in AFS-2 and 3

Replica consistency using quorum consensus

Grapevine

Caching in the Domain Name System

File caching in the Sprite File System

b. In our discussion of the domain name system we saw the benefit of caching at multiple levels (i.e. caching on DNS servers and the use of chained queries so that the server cache was shared by multiple clients). Explain how one might apply a similar technique to cache web pages? In your discussion describe the mechanisms that might be used, configuration issues (how one selects the local server), as well as suggested approaches to manage consistency. [15 points]

Name: _____

USC ID: _____

c. In the quorum consensus protocol for replica consistency there is a requirement that the write quorum be greater than half the total votes. Explain the need for this requirement. [10 points]

Name: _____

USC ID: _____

- c. You are designing a content distribution system to support Time Warner's new on demand streaming video feeds. Discuss the techniques that should be used to improve the scalability of the system? In particular, discuss where the videos will be stored, and how users will find and retrieve the videos? What techniques are needed to maintain consistency of the videos in the video library? [15 points].