

Name:

CS461, Final
Spring 2003
May 20, 2003

Pledge:

Directions:

- Please answer each question in the space provided. The amount of space should be sufficient for a correct answer. If you need more space, please use the backs of pages, and make a note to that effect. If you run out of space, exam books are provided at the front of the room.
- This exam is closed-book, closed-notes, and is covered by the Honor Code. Please write and sign the pledge after you finish your exam.
- There are a total of five sections, with the number of points for each shown by the question. While it is not the intent for the exam to be a race, spending too much time on a single question may preclude finishing the exam. Budget your time wisely.
- To be fair, I will try to avoid answering content-related questions during the exam, unless it's to correct a mistake on my part.
- If you feel that a question **requires** additional assumptions or information to answer, please state them. Your guiding principle should be *Occam's razor*, which loosely translated states that you should allow as few assumptions as necessary to explain the situation.
- Unless otherwise stated/implied, assume a C-like language running on a Unix-like operating system.
- Please first read over the entire exam and then begin to answer questions. I will wait outside the exam room for the first 5-10 minutes, and then will be available in my office (room 322).
- Please write legibly and bring your exams to my office as you finish.

1. True or False (10 pts) For each statement, write “true” if the statement is true, or “false” if the statement is false. If you believe the statement does not have a clear answer, give whichever choice is more appropriate and explain why.

- RSVP requires more state than DiffServ.
- BLAST-like protocols are appropriate for shared wide-area networks.
- Consistent Hashing is a common technique for building peer-to-peer systems.
- Quorums can be used to implement Byzantine fault tolerance.
- Head-of-line blocking is prevalent in output-buffered systems

2. Short Answer (16 pts) Answer each item *well* in no more than 4 sentences.

- What is “marshalling” and why is it needed in Remote Procedure Call (RPC) systems?

- What sort of performance issues/implications are associated with various marshalling approaches?

- What is the relation between JPEG and MPEG?

- What is a “bidirectional prediction” frame in MPEG and why have it in addition to a “forward prediction” frame?

3. Security (10 pts)

- Describe how “private key” encryption systems work, and explain how they differ from “public key” systems.

- The use of a public key system eliminates a problem in private key systems, but introduces a new problem. Explain.

4. Routing and routers (16 pts)

- What are the basic parts of a router and what do they do?

- What sorts of issues limit the scalability of high-end routers?

Some form of “tag switch” routing has been attempted/implemented by various companies, and is becoming a popular option for high-end routers.

- What is the basic concept behind these kinds of routers, and what are the technical/cost considerations that drive their design?

- Describe what additional mechanism is necessary in these kinds of routers that are not needed in pure IP routers.

5. Content Distribution Networks (CDNs) and Peer-To-Peer (P2P) Systems (18pts)

CDNs tend to be built by individual companies, either for their own use or to host/serve content for other companies. In contrast, P2P systems tend to consist of large numbers of end users in a more loosely-coupled environment.

- Describe what kinds of concerns are shared in the two environments, and any mechanisms that can be used in both contexts.

- What areas of difference exist between the two systems, and what are the origins of such differences?

One “common” way of implementing CDNs is to use/abuse the DNS system to provide a layer of indirection between the clients and the servers from which they fetch content.

- Describe a (good) plausible system for using DNS in this manner, and explain its technical operation.

- What sort of considerations are important in deciding which server should be used by a client when fetching a particular object? Explain their operation/interrelation.