

- 1a What is the difference between data-centric consistency models and client-centric consistency models? 5pt
- 1b Researchers often distinguish caching from replication. Explain the difference between the two. 5pt
- 1c Content Delivery Networks in the Web such as Akamai often use existing caching mechanisms to establish replication of Web documents. Explain how this works. 5pt
- 2a A  $k$ -fault tolerant process group can withstand the failure of  $k$  processes before failing to meet its specifications. Discuss the relation between the size of such a group and the assumptions that are made with respect to failure semantics and process behavior. 5pt
- 2b Why is it so hard to implement *exactly once semantics* in the presence of server crashes and reliable RPCs? Assume a client receives an acknowledgement when its request has been delivered at the server. 10pt
- 2c Explain how an orphan process comes into existence when dealing with message-logging systems. 5pt
- 3a How can we protect a certificate from being successfully used by an unauthorized user? 5pt
- 3b Public-key cryptosystems are claimed to scale better than shared-key systems. Does this claim really hold? *Hint*: think of certificate revocation. 5pt

**Grading:** The final grade is calculated by accumulating the scores per question (maximum: 45 points), and adding 5 bonus points. The maximum total is therefore 50 points.