

# Distributed Systems

## Exam

The exam is on 3 pages. Solve the exam in anyway that is easy for you. Upload a PDF, write on your laptop, on a piece of paper then scan it, or whatever is appropriate for you. I expect your submission to make it easy for me to know what you are answering, so please be organized and neat.

The exam is graded from 70 points. Choose any set of questions that you would like such that you acquire 70 points. You can mix-and-match whichever questions you like!  
Contact me with questions by email or on Canvas.

### **Question1: (25 points)**

Design a distributed system for electronic voting for the next parliamentary election in Sweden. Start with thinking about how many machines will you need if one voter takes on average 2 minutes in-front of the machine to vote, and that there are 7.5 Million eligible voters in Sweden. Voting will take place during 2 days. The system has the following functionality. It allows a voter to pick among different candidates. A voter can only vote once. A voter can modify her vote as long as she has not clicked the submit button. A voter can later using her personal number check online if she has already voted, and to whom did she vote. Discuss how such a system should be built, including how the CAP theorem should be taken into account.

### **Question 2: (10 points)**

Write a Map-Reduce program (in pseudo-code) that counts the total number of votes each candidate has gotten in the above election.

### **Question 3: (10 points)**

Write a Map-Reduce program (in pseudo-code) that searches if a voter exists in the voters list.

### **Question 4: (25 points)**

Design a distributed social network that is not governed by a single entity, i.e., where no company controls the social network. The only requirement is that the system can not be governed by any single entity. In your design, discuss the different distributed system aspects that you as a system designer need to consider. Make your own decisions on the types of data that users are allowed to post. Expect at least 100 Million users to be online.

### **Question 5: (10 points)**

The above distributed social network platform became a cesspool of illegal activity. Write a Map-Reduce program (in pseudo-code) that searches for a list of keywords in all the posts in the network, returning posts that include these keywords, along with their poster identity.

### **Question 6: (5 points)**


The above network became even worse with law-enforcement unable to . Think of two attacks that you can launch on this network to bring it down or at least make it unreliable.

**Question 7 (25 points)**

a- At 10:05.700, a client sends the server a request for the time. The server response arrives at 10:05.900 containing a time stamp of 10:05.850. Using Cristian's algorithm, to what value does the client set its clock?

b- The coordinator's clock in a Berkeley algorithm reads 1:00. The clocks on the other systems read 1:02, 1:04, 1:10. After running the algorithm, to what value will the other systems be set?

c- In the following diagram, give an example of one consistent cut, and one inconsistent cut.



d- Discuss the ordering and possible consistency model for the following event diagrams. Note that the box shows the length an operation happens.

