

READ CAREFULLY BEFORE BEGINNING

1. This exam comprises 25% (400 level) or 20% (500 level) of your total grade. Points are indicated for each question/section.
2. This is an **in-class, open book** exam (which includes open notes and papers).
3. On-campus: You have **50** minutes to finish.
4. Video: You have **55** minutes to finish.
5. Make sure your copy has **6** questions.
6. Read each question completely before answering. Be sure you understand all assumptions and constraints.
7. **Show all work and state all assumptions.** Partial credit can not be given for work not shown. Use the back of the page, if necessary.

REG MAXIMUM	100
DEDUCTIONS	—
TOTAL SCORE	_____

PRINTED NAME _____

1. (20 pts) Draw the Petri-Net describing the following systems. You may omit the inhibitory arcs that halt each timed transition upon failure, if applicable.

(a) A simple RAID (level 1) system consisting of 2 drives total, each with fail-rate λ . The repair/recover rate is indicated by δ .

(b) Again, assume we have a RAID system (level 1), with each of the drives having a failrate λ . However, now each mirror consists of 2 drives. Thus the total number of drives is 4. The repair/recover rate is indicated by δ . Define a single place that constitutes the fail state of the system. To avoid visual clutter you may omit the inhibitory arcs from that state to all timed transitions.

Hint: with two drives on each side of the mirror it *does* matter which of the two drives fail, when each side of the mirror experiences a failure before recovery.

2. (20 pts) Answer the following questions with respect to fault models:

(a) Why does simple majority not solve the oral message based agreement problem?

(b) Assume the Lamport bound of $N = 3m + 1$, with $r = m + 1$ rounds. N nodes participate in the message exchanges of a simplex agreement, i.e. you have the general and $N - 1$ lieutenants. What is the total number of messages sent in round i , $1 \leq i \leq m + 1$ by all processors?

Hint: think of the message tree of a processor. Unlike in our examples in class, you may assume a complete tree, i.e. each level of the tree is full.

3. (10 pts) In Lamport's paper the SM(m) algorithm is discussed.

(a) What is the main advantage of using "signed messages"?

(b) What is the main disadvantage?

4. (20 pts) With respect clock synchronization:

(a) We have discussed what initiated the agreement research, i.e. clock synchronization problems originally observe in SIFT. Given three clocks A, B and C describe how the clocks can be forced by a malicious clock, say A , to drift appart if each clock uses the median value in each round of the synchronization algorithm.

(b) Given the convergence function $CF(p, x_1, x_2, \dots, x_n) = CF(1, 2, 2, 2, 3, 5, 5, 6, 6, 6, 10, 10, 11)$, $\delta = k = 3$, what are the results for the following conversion functions? Note that in general the values x_i are not sorted. However, here they are sorted to make your life easier :-)

i. egocentric average:

ii. fast convergence algorithm:

iii. fault-tolerant midpoint:

iv. fault-tolerant average:

5. (10pts) Still considering clock synchronization: Below is the interval for clock values under the assumption that there is no clock drift at all, based on the interval from the paper by F. Cristian.

$$C_Q(t) \in [T + \min, T + 2D - \min]$$

- (a) Derive an expression for the clock value that should be selected to minimize possible errors?

- (b) Derive an expression for the maximum error?

6. (20 pts) With respect to RAID systems:

(a) Give an example of a realistic MTTF or MTBF for a contemporary single disk drive?

(b) Describe the essential properties of the original five RAID levels,

i. RAID-Level 1:

ii. RAID-Level 2:

iii. RAID-Level 3:

iv. RAID-Level 4:

v. RAID-Level 5:

(c) A firm dealing with important financial transactions is interested in a RAID system. What RAID technology would you suggest to them? Justify your answer.

Extra Workspace