Exercise 2

1 Replication and Erasure Coding

Suppose a storage system *M* stores a payload of size S. The storage system consists of d disks, each of which has size $s \ll S$. There are enough disks so that the data can be stored replicated with redundancy $r = \frac{ds}{S} > 1$. (For simplicity, assume r is an integer and S is a multiple of s.)

An *m-out-of-n erasure code* is a method to split some data into m parts, and to encode the parts into n fragments (each fragment is of the same size as a part), such that the data can be recovered from any m of the n fragments.

For the following calculations, assume that the MTTF of one disk is 10^6 hours (about 0.9% annual failure rate).

a) We store the data using r-fold replication. Given the reliability function $R_{Disk}(t)$ of one disk, derive the reliability function of the storage system *M* and calculate numerically the MTTF of *M*.

b) We consider a storage array with different variants of replication and erasure coding. Let *M* store $S = 60$TB on 600 disks of capacity $s = 300$GB each.

   Compare (numerically) the MTTF of the following coding schemes:

   i. The disks are divided into 3 equal-sized pools, and the data is replicated on to the three pools.

   ii. Three disks are each taken together to form one logical volume with 3-fold replication, and the data is stored on the 200 logical volumes.

   iii. A $(20, 10)$ MDS erasure code is used to create a logical volume across every group of 20 disks (tolerating the failure of 10 disks inside each group of 20), and the data is stored on the 30 logical volumes.

2 Safety and Liveness

On the opposing sides of a canyon, there are two armies, each led by a different general. General A and General B want to attack a third army (led by General Z), which will travel through the canyon. However, the third army is quite strong, and in order to succeed, both generals need to attack at exactly the same time. Thus, both generals devised the following algorithm for a coordinated attack:

- General A chooses an attack time t and sends a messenger m_1 carrying t (denoted $m_1(t)$) to General B.

- General B waits for m_1. Upon the arrival of $m_1(t)$, General B sends a messenger m_2 carrying a confirmation to General A. Then it waits for time t and attacks.
- General A waits for m_2. Upon the arrival of m_2, General A waits for time t and attacks.

Which of the following properties are safety properties, which ones are liveness properties, and which ones are a mixture of the two? (See Sect. 2.1.3 [CGR11].) Explain your answers.

(a) If some general attacks at time t, then the other general attacks at the same time.

(b) If General A attacks at time t, then General B attacks at the same time.

(c) If General B attacks, then there exists a time when General A attacks as well.

(d) Eventually, General B will attack.

(e) If messengers m_1 and m_2 are not intercepted, then eventually both generals attack.

(f) If messengers m_1 and m_2 are not intercepted, then eventually both generals attack at time t.