COSC 3330/6308—COMPUTER ARCHITECTURE
BACK OF THE ENVELOPE COMPUTATIONS OF DISK ARRAY RELIABILITY

We do not always have the time or the resources to do exact computations of disk array reliability. Her are a few simpler techniques that will give you reasonable approximations of the reliability of a disk array.

1. *MTTF and Failure Rate:*

If λ designates the failure rate of any device, we always have λ = 1/MTTF and MTTF = 1/ λ

2. *Probability of a failure during a short interval:*

Consider any device that fails at a constant failure rate λ. Under relatively robust assumptions, the probability that the device will not fail over a time interval of duration *t* is the solution of the differential equation

$$\frac{dR(t)}{dt} = -R(t)\lambda t$$

with the initial condition $R(t) = 1$. The solution of this equation is $R(t) = e^{-\lambda t}$. Hence the probability that the device will fail at least once during the time interval is $1 - e^{-\lambda t}$. Developing that expression into a Taylor series, we obtain

$$1 - e^{-\lambda t} = 1 - 1 + \frac{\lambda t}{1!} - \frac{(\lambda t)^2}{2!} + \frac{(\lambda t)^3}{3!} - \ldots$$

For *very small* values of λt, this expression can be approximated by its first nonzero term λt. Consider, for instance, a disk drive whose MTTF is 200,000 hours, the probability that the disk drive would fail would fail during a given day would be very close to 24/200,000= 1.2×10^{-4} while the actual value would be 1.1999208×10^{-4}.The approximation becomes less accurate as the product λt increases: it would predict that the probability that the same disk drive will fail at least once during a useful lifetime of 5 years is roughly equal to (365×24)/200,000=0.0438 while the exact value is 0.0429.

3. *Probability of a data loss for a mirrored pair of disks:*

Since all data are replicated on both disks, we will lose data if one of the two disks fails and the other disk fails while the first one getting replaced by a new disk containing the same data. This is to say that the repair process will require replacing the failed disk and then copying on it the data stored on the second disk. If T_r designates the average disk repair time, the probability of a data loss over an interval of duration T_d will be given by

$$(2\lambda T_d)(\lambda T_r) = 2\lambda^2 T_d T_r.$$

4. *Probability of a data loss for a RAID array:*

We could extend the same approach to a RAID level 5 array with *n* disks. We obse4rve first that the array will lose data if any of the *n* disks fails and any other disk fails while the first one getting replaced by a new disk containing the same data. We would then have

$$(n\lambda T_d)((n–1)\lambda T_r) = n(n–1)\lambda^2 T_d T_r.$$

but should expect to get increasingly worse approximations of the actual data loss probability as *n* increases.

The main usefulness of these formulae is that we can compute them with a very basic pocket calculator. They are quick but we should expect too much from them.

Updated last on 5/4/2010