On the Possibility of Small, Service-Free Disk-Based Storage Systems

Jehan-François Pâris and Thomas J. E. Schwarz, S.J.

Department of Computer Science
University of Houston
480 0 Calhoun
Houston, TX 77204-3010
paris@cs.uh.edu

Department of Computer Engineering
Santa Clara University
500 El Camino Real
Santa Clara, CA 95053
tjschwarz@scu.edu

Abstract

For many storage providers, the cost of providing service calls exceeds the costs of the hardware being serviced. In this paper, we show that zero-maintenance, small disk arrays are too expensive, but that low-maintenance arrays are feasible and describe a possible implementation. Our evaluation technique replaces Mean Time to Data Loss with the lifespan. Our results also show the impact of the assumption of constant failure rates on the results of modeling.

1. Introduction

Today's disks have a less than ten percent probability of failing during any given year of their useful lifetime [PWB07]. While this reliability level is acceptable for applications that only require the storage of a few hundreds of gigabytes of non-critical information over relatively short time intervals, it does not satisfy the needs of applications having to store terabytes of data over many years. With the advent of digital photography, these applications have entered the homes of end consumers.

Given the multiple limitations of backups, the best solution for ensuring the survivability of data over long periods is the use of redundant disk arrays. This goal is typically achieved through either disk mirroring or the use of *m*-out-of-*n* codes, among which RAID level 5 and RAID level 6 [PGK88, BM93, CL+94] are prominent.

Despite all their advantages, redundant disk arrays have their own limitations. First, they require more disk space than non-redundant arrays to store the same amounts of data. Second, they also have higher maintenance costs since the addition of parity drives results in more disk failures. This is less of an issue for large installations having their own maintenance teams. For instance, the Google file system is built from cheap commodity disks that are known to fail often and compensates by maintaining three copies of each file chunk [GGL03]. On the other hand, smaller installations need to replace disks through service calls whose costs can exceed the cost of the hardware being serviced. Storage vendors have also to consider the possibility of seriously embarrassing mistakes occurring during these calls. This made us wonder whether it would be possible to build highly reliable data storage solutions that would not need service during their lifetime. Manufacturers would offer a lifetime warranty to their customers that would offer reimbursement for loss of data due to disk failure. By egregious overprovisioning with disks, we can certainly build these zero-maintenance arrays for the small office and home office (SOHO) market, but that approach is too costly to succeed in the market. In this paper, we use modeling to consider the feasibility of small, zero-maintenance disk arrays.

2. Assumptions and Metrics

Hard disks are complex. They can fail completely, sometimes without warning, and they can loose data in only a few sectors. The recent, excellent overview article by Elerath [El07] gives more details on these various failure mechanisms. To our knowledge, no public data is available on the frequency, at which disks suffer partial data losses. People in the disk industry estimate this rate to be about 10 times higher as that for complete disk failures. In a disk array with redundant data storage, scrubbing (verification that stored data can be read) limits the impact of partial data loss, [SX+04]. For these reasons, we only consider complete device failure in our models.

Many models of disk array reliability assume constant failure rates, even though this assumption is

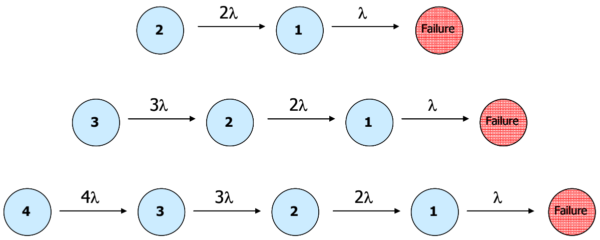

Figure 1: Markov models for mirrored (top), triple, and quadruple data.

known to be false. Empirical studies [El00b, El04, EP07, SG07, YS09] have shown that often, but not always [GGL03], the failure rate of a batch of disks drops over the first year, before it settles on a constant rate. Even assuming constant failure rates, the question arises for the system modeler of its values. Some vintages and batches differ widely from the disk manufacturer specified failure rate of now typically 1/2,000,000 hours [El00b, ES04, EP07].

Estimating the reliability of a storage system means estimating the probability *R*(*t*) that the system will operate correctly over the time interval [0, *t*] given that it operated correctly at time *t* = 0. Computing that function requires solving a system of linear differential equations, a task that becomes quickly unmanageable as the complexity of the system grows, but is feasible for the small arrays that we consider. In fact, we used Mathematica, a symbolic mathematics software tool.

Unfortunately, *R*(*t*) is function and not a single number. Mean Time To Data Loss (MTTDL) is a misleading single figure of merit, since the failure rate of a disk array is far from being constant. For a given reliability level *r*, we therefore calculate the *economic life span* *L*(*r*) of a disk array as the maximum time interval for which data stored on that array will have a probability *r* to survive intact. Our figure of merit becomes dimensionless, if we express it in multiples of the Mean Time Between Failures (MTBF) of the disks.

3. Replicated Single Disk

The first array organization we considered was an array consisting of *n* identical disks each holding an identical copy of the data. We further assumed that we would never repair the array during its useful lifetime. We selected this organization because of its simplicity and its potential for achieving any arbitrary data survival by increasing the number of disks in the array.

Assuming a constant disk failure rate λ, the survival S_1 of a single disk at time *t* is given by the differential equation

$$S_1' = -\lambda S_1, \quad S_1(0) = 1$$

with solution

$$S_1 = \exp(-\lambda t).$$

We model survival of a mirrored disk in the standard Markov model depicted in Figure 1, top. We label the non-failure states by the number of existing disks. The starting state is state 2, from which we transition to state 1 at rate 2λ, whenever one of the two disks fails. We can capture the probability p_i of being in state *i* at time *t* in a system of ordinary differential equations with initial conditions

$$\begin{aligned} p_2' &= -2\lambda p_2 & p_2(0) &= 1 \\ p_1' &= 2\lambda p_2 - \lambda p_1 & p_1(0) &= 0 \end{aligned}$$

The solution is

$$p_2(t) = e^{-2\lambda t}, \quad p_1 = 2e^{-2\lambda t}(e^{\lambda t} - 1)$$

Similarly, the middle Markov model in Figure 1 describes the case of triplicate disks. Using the same convention of writing $p_i(t)$ for the probability of being in state *i*, we now obtain the following system of ordinary differential equations with initial conditions

$$\begin{aligned} p_3' &= -3\lambda p_3 & p_3(0) &= 1 \\ p_2' &= 3\lambda p_3 - 2\lambda p_2 & p_2(0) &= 0 \\ p_1' &= 2\lambda p_2 - \lambda p_1 & p_1(0) &= 0 \end{aligned}$$

This system has solution

$$\begin{aligned} p_3(t) &= e^{-3\lambda t}, \\ p_2 &= 3e^{-3\lambda t}(e^{\lambda t} - 1), \\ p_1 &= 3e^{-3\lambda t}(e^{\lambda t} - 1)^2. \end{aligned}$$

Finally, the final Markov model in Figure 1 gives the Markov model for a quadrupled disk. The system of ODEs is now

$$\begin{aligned} p_4' &= -4\lambda p_4 & p_4(0) &= 1 \\ p_3' &= 4\lambda p_4 - 3\lambda p_3 & p_3(0) &= 0 \\ p_2' &= 3\lambda p_3 - 2\lambda p_2 & p_2(0) &= 0 \\ p_1' &= 2\lambda p_2 - \lambda p_1 & p_1(0) &= 0 \end{aligned}$$

The solution is now

$$\begin{aligned} p_4(t) &= e^{-4\lambda t}, \\ p_3(t) &= 4e^{-4\lambda t}(e^{\lambda t} - 1), \\ p_2(t) &= 6e^{-4\lambda t}(e^{\lambda t} - 1)^2, \\ p_1(t) &= 4e^{-4\lambda t}(e^{\lambda t} - 1)^3 \end{aligned}$$

We capture the chances of survival of the system of *n* replicated disks in the function $S_n(t) = \sum_{i=1}^{n} p_i(t)$.

Table 1: Economic life span for a single data disk and different replication levels

9s	Simple	Mirrored	Triple	Quadr.
1	0.10536	0.38013	0.62392	0.82632
2	0.01005	0.10536	0.24265	0.38013
3	0.00100	0.03213	0.10536	0.19581
4	1.00E-04	0.01005	0.04753	0.10536
5	1.00E-05	0.00317	0.02178	0.05788

We define the economic life span $L(r)$ of an array as the maximum time interval for which data stored on that array will have a probability r to survive intact. More precisely, $L(r)$ is the solution of $S_n(t) = r$, where the reliability level r is 0.9, 0.99, 0.999, 0.9999, and 0.99999, i.e. 1, 2, 3, 4, and 5 nines. After setting $\lambda = 1$, the economic life span is expressed in multiples of the disk MTBF.

We tabulate our results in Table 1. A single disk reaches a life span of 1% of the device MBFT with probability 99%. A quoted disk MTBF of 1,000,000 would then be 10,000 hrs or a little bit more than a year. More realistic MTBF estimates would result in even lower economic life spans

4. The 2D+P+Q Disk Array

In this section, we consider disk arrays that store two disks worth of data. Our first and simplest configuration consists of two data disks supplemented by two parity disks, the P and the Q disk, whose contents we calculate with erasure coding. Figure 2 gives the state diagram of this array. Notice that this configuration does not suffer any data loss as long as any two disks can still be read. We represent the initial configuration in the first state. We capture the first disk failure in the state transition out of this state. If the failed disk is a parity disk, then obviously we are left with two data and one parity disks. (To improve the speed of further updates, we could replace the parity with the ordinary parity, but this would not affect the modeling of data survival in our disk array.) If the failed disk was a data disk, then reads to this drive need to be serviced by accessing any two of the three surviving disks and calculate the data, and these reads are therefore slow. However, we can carefully replace the data on one parity drive by the recalculated data that used to be on the failed disk. We take care by storing temporary data so that no additional failure during the reconstruction will loose data. Therefore, the reconstruction has no effect on the survival of our data, but it reestablishes a configuration where all data can be directly read. (If further performance improvements are desired, we can afterwards replace the parity data on the surviving parity drive with ordinary parity.) Another data failure still preserves all information, but each of the remaining disks is now critical and no further disk failure can be tolerated.

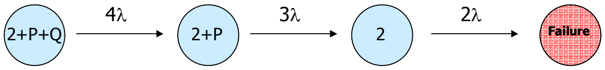

Figure 2. Markov models for 2D+P+Q

Table 2: Economic life span: 2D+P+Q Array

Nines	Mirrored	2+P+Q	Triple
1	0.38013	0.38634	0.623918
2	0.105361	0.151832	0.242637
3	0.0321336	0.0661806	0.105361
4	0.0100503	0.0299014	0.047528
5	0.0031673	0.0137122	0.02178

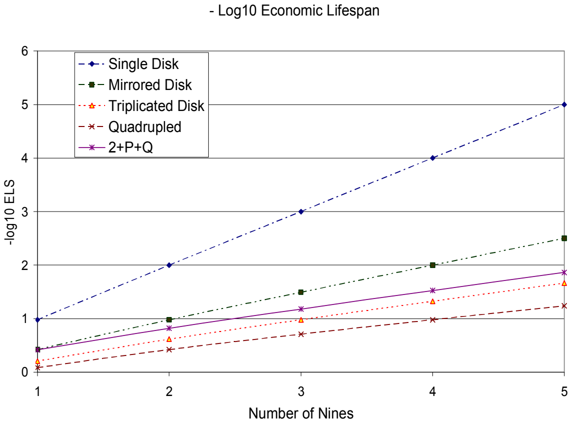

Figure 3: Economic life span comparison

Using our previous conventions, we first solve the system of ODE

$$\begin{aligned} p_4' &= -4\lambda p_4 & p_4(0) &= 1 \\ p_3' &= 4\lambda p_4 - 3\lambda p_3 & p_3(0) &= 0 \\ p_2' &= 3\lambda p_3 - 2\lambda p_2 & p_2(0) &= 0 \end{aligned}$$

to obtain

$p_4(t) = e^{-4\lambda t}, p_3(t) = 4e^{-4\lambda t}(e^{\lambda t} - 1), p_2(t) = 6e^{-4\lambda t}(e^{\lambda t} - 1)^2$

Survival is

$$S(t) = 3e^{-4\lambda t} - 8e^{-3\lambda t} + 6e^{-2\lambda t}.$$

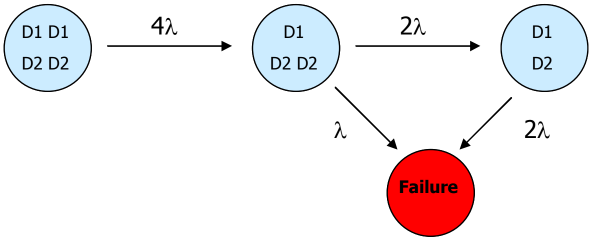

Figure 4: Markov model: Pair of mirrored disks

We present the results in Table 2. Figure 3 compares our schemes so far. The x-axis gives the number of nines. The y-axis contains the negative base 10 logarithm of the economic lifespan of the ensemble described. Recall that chance of data loss during this economic lifespan is given by the level of nines. We give the life spans in multiples of the MTTF of the disks. For instance, our graph for a single disk guaranteed to not loose data with probability 99.99% (four nines) has a y-value of about 4. This means that it should not be used longer than 1/1000 = 10^{-4} of the disk MTTF. In contrast, at 99.99%, the quadrupled disk configuration has an economic lifespan of approximately 1/10 = 10^{-1} of the disk MTTF. The graph illustrates that the 2+P+Q configuration lies between mirroring and triplication.

5. Mirroring Arrays

First, we consider an array consisting of two pairs of mirrored disks. We present the Markov model in Figure 4. A failure from the initial state (left) leads to a state where half of the data only resides on one drive indicated by our label of D1, D2, D2, regardless on what specific disk has failed. From there, failure of one of the mirrored disks left leads to a state where all data is still available, whereas failure of the lone disk leads to data loss.

Our Markov analysis yields the following system of ODEs where we label states and hence state probabilities with the number of surviving disks. Thus, the probability of the system being in the initial state is $p_4(t)$.

$$\begin{aligned} p_4' &= -4\lambda p_4 && p_4(0)=1 \\ p_3' &= 4\lambda p_4 - 3\lambda p_3 && p_3(0)=0 \\ p_2' &= 2\lambda p_3 - 2\lambda p_2 && p_2(0)=0 \end{aligned}$$

The solution is

$$p_4(t)=\exp(-4\lambda t)$$

$$p_3(t)=4\exp(-4\lambda t)(\exp(\lambda t)-1)$$

$$p_2(t)=4\exp(-4\lambda t)(\exp(\lambda t)-1)^2$$

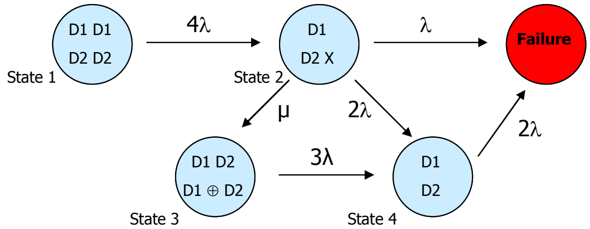

Figure 5: Markov model: Pair of mirrored self-reorganizing disks

Not surprisingly, the resulting economic life spans are smaller than the ones for a single pair of mirrored disks, by about 30%. Compared to the 2D+P+Q organization, the lower life span is caused by the transition to the Failure State from State 3, (D1, D2, D2). Assume that immediately after the first failure, we reorganize the array to a RAID Level 5 so that the disk contents are now of the form D1, D2, D1⊕D2. In this case, the Markov model becomes that of the 2D+P+Q array. The advantage of starting out with a pair of mirrored disks over the 2D+P+Q organization is the better performance. A read to mirrored disks is directed to the disk with shortest service time and a write only involves writing to two disks and completely avoids the cumbersome read-modify-write operation in RAIDs.

As the results of instantaneous reorganization are spectacular, we now model a more realistic reorganization with a non-zero reorganization time 1/μ. Reorganization involves reading two data disks and overwriting the third one with the parity of the two data disks. We can realize it in an idle array by sweeping the disks in parallel. In this optimal case, the rate of reconstruction is given by the time to access a single disk completely. At a sustainable rate of 100MB/sec, this takes 2.8 hours for a 1TB drive. We give the Markov model in Figure 5.

Unfortunately, with the presence of two rates, our simple dimensionless analysis is no longer possible. In addition, the speed of reorganization depends on factors such as array utilization, the size of the disks, and the sustainable bandwidth of the disks. While reorganization takes a constant time, we nevertheless model it exponentially distributed with rate μ. After a failure in the initial state (state 1), reorganization starts immediately, in which we systematically replace the contents of one of the drives in the surviving pair with parity data. We write D1 for the surviving drive in the lost pair, D2 for one of the drives making up the other pair and X for the drive that contains partially data from D2 and partially the parity of D1 and D2. If D1 fails during this time, then we have suffered data loss.

Table 3: Economic life span of a self-reorganizing pair of mirrors

9s	2+P+Q	μ/λ=10^5	μ/λ=10^4	μ/λ=10^3
1	0.38634	0.38633	0.38625	0.38547
2	0.15183	0.15181	0.15161	0.14964
3	0.06618	0.06613	0.06568	0.06120
4	0.02990	0.02979	0.02879	0.01972
5	0.01371	0.01347	0.01132	0.00345

If X fails, then data is preserved and we transition to the unprotected State 4 with a D1 and a D2 disk. If D2 fails, then either we can reconstruct its contents directly by copying from X or we can reconstruct it from the already stored parity data. In either case, after a possible further reorganization, the disk array ends in State 4. This final reorganization has no implications for the reliability of the disk array.

The Markov model in Figure 5 gives rise to the following ODE.

$$p_1' = -4\lambda p_1 \qquad p_4(0) = 1$$
$$p_2' = 4\lambda p_1 - (3\lambda + \mu) p_2 \qquad p_3(0) = 0$$
$$p_3' = \mu p_2 - 3\lambda p_3 \qquad p_2(0) = 0$$
$$p_4' = 2\lambda p_2 + 3\lambda p_3 - 2\lambda p_4 \qquad p_1(0) = 0$$

The solutions are

$$p_1(t) = e^{-4\lambda t},$$
$$p_2(t) = 4\frac{\lambda}{\lambda - \mu} e^{-4\lambda t} (e^{(\lambda-\mu)t} - 1)$$
$$p_3(t) = \frac{4e^{-4\lambda t}\left(\mu + (\lambda - \mu)e^{\lambda t} - e^{(\lambda-\mu)t}\right)}{\lambda - \mu}$$
$$p_4(t) = 2e^{-4\lambda t} (2(1 - 3e^{\lambda t} + e^{2\lambda t} + e^{(\lambda-\mu)t})\lambda^2 +$$
$$(e^{2\lambda t} - 1)\lambda\mu - 3(e^{\lambda t} - 1)^2 \mu^2)/(\lambda^2 - \mu^2)$$

We tabulate the life spans of the reorganizing disk array in Table 3. The right columns give the life spans of the reorganizing pair of mirrors for various assumptions of the ratio μ/λ that equals the ratio of the disk MTTF over the reorganization time. Since the reorganization time is somewhere between a few hours and maybe 100 hours, and disk MTTF is somewhere between 100,000 hours and 2,000,000 hours, the true ratio lies in the spectrum covered by the table. From the numbers, it is clear that the speed of reorganization only matters for very stringent demands on disk array longevity.

6. Impact of Variable Disk Failure Rates

The assumption of constant disk failure rates enables Markov modeling, but is often not realistic.

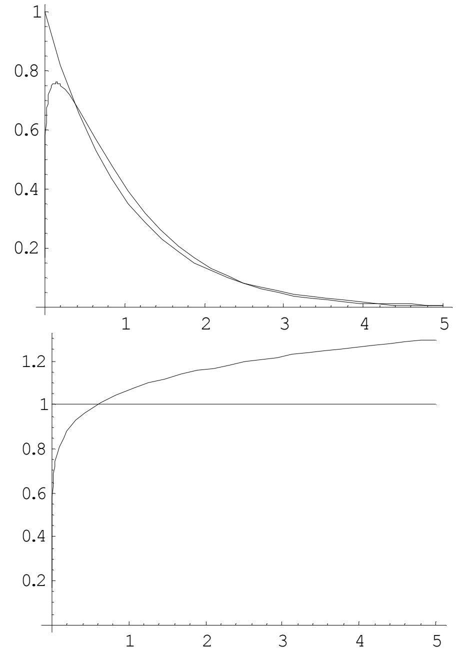

Figure 6: Probability density and hazard rate

A two-parameter Weibull distribution with probability density function usually fits field disk operational failure data much better. Exceptions typically arise when a disk drive family suffers from two or more competing early failure causes. The probability density function of the two-parameter Weibull distribution is

$$p_{\text{Weibull}}(t) = \frac{\beta}{\eta} \cdot \left(\frac{t}{\eta}\right)^{\beta - 1} \exp(-\left(\frac{t}{\eta}\right)^{\beta})$$

with *shape parameter* β and *characteristic lifetime* η. Since the mean time to failure of a device following the Weibull distribution is $\text{MTTF} = \eta \cdot \Gamma(\frac{1}{\beta} + 1)$, we set $\eta = \Gamma(\frac{1}{\beta} + 1)^{-1}$. The shape parameter is typically a value between 0.9 and 1.5. A typical value is $\beta = 1.12$ [EP07]. Let $F(t)$ = Prob($X \leq t$) be the cumulative probability distribution of the failure event X. The failure rate or *hazard rate* $h(t)$ is the probability that a device fails in a given unit of time. We always have the relationship $h(t) = p(t)/(1\text{-F}(t))$. For the Weibull distribution, this amounts to

$$h_{\text{Weibull}}(t) = \frac{\beta}{\eta} \cdot \left(\frac{t}{\eta}\right)^{\beta - 1}.$$

The failure rate for the exponential distribution with parameter λ is simply constant λ. The Weibull failure rate decreases if $0 < \beta < 1$, is constant if $\beta = 1$, and

Table 4: Economic life span for various Weibull shape parameters

9s	$\beta = 0.8$	$\beta = 0.9$	$\beta = 1.0$	$\beta = 1.1$	$\beta = 1.2$
Single Disk					
1	0.0530	0.0780	0.1054	0.1340	0.1630
2	0.0028	0.0058	0.0101	0.0158	0.0230
3	0.0002	0.0005	1.0E-3	0.0019	0.0034
4	8.8E-6	3.4E-5	1.0E-4	0.0002	4.9E-4
5	5.0E-7	2.6E-6	1.0E-5	3.0E-5	7.2E-5
Mirrored Disk					
1	0.2634	0.3245	0.3801	0.4302	0.4748
2	0.0530	0.0780	0.1054	0.1340	0.1630
3	0.0120	0.0208	0.0321	0.0455	0.0606
4	0.0028	0.0057	0.0101	0.0158	0.0230
5	0.0007	0.0016	0.0032	0.0055	0.0088
Triplicate Disk					
1	0.4894	0.5627	0.6239	0.6749	0.7175
2	0.1503	0.1970	0.2426	0.2860	0.3266
3	0.0530	0.0780	0.1054	0.1340	0.1630
4	0.0196	0.0322	0.0475	0.0650	0.0840
5	0.0074	0.0135	0.0218	0.0320	0.0438
2D+P+Q					
1	0.2688	0.3304	0.3863	0.4365	0.4813
2	0.0837	0.1170	0.1518	0.1868	0.2210
3	0.0296	0.0465	0.0662	0.0878	0.1106
4	0.0110	0.0192	0.0299	0.0426	0.0571
5	0.0041	0.0081	0.0137	0.0210	0.0298

increases if $1<\beta$. Figure 6 compares the probability density and failure rate for a Weibull with λ=1, β=1.12, η=1.04238 and an exponential distribution with λ=1. In this example, the Weibull hazard rate first increases and then very shortly afterwards steadily decreases. This behavior is true for many, but not for all disk populations studied by Elerath and Shah [El00b, ES04, EP07]. To estimate the impact of the shape parameter β on the economic life span of the array, we replace the constant failure rate λ in our systems of ODE's with the non-constant hazard rate of the Weibull distribution.For example, when calculating the life span for a system of mirrored disks, we obtain the ODE

$$p_2'(t) = -2h_{\text{Weibull}}(t)p_2(t)$$

$$p_1'(t) = 2h_{\text{Weibull}}(t)p_2 - 2h_{\text{Weibull}}(t)p_1$$

We can easily solve these equations for given Weibull parameters. We similarly proceed to reproduce the results for triplicate disks. We present our results in Table 4. As we compare the results, we find the strong influence of the shape parameter. The discrepancy is particularly strong at larger numbers of nines, where it attains more than one decadic order of magnitude.

These results force us to qualify our previous conclusions. They also show clearly that the Markovian assumption of constant disk failure rates is a dangerous one. The effect is more pronounced where we are interested in the reliability of very young disks, namely when we define the lifespan by a high level of nines survivability of the ensemble.

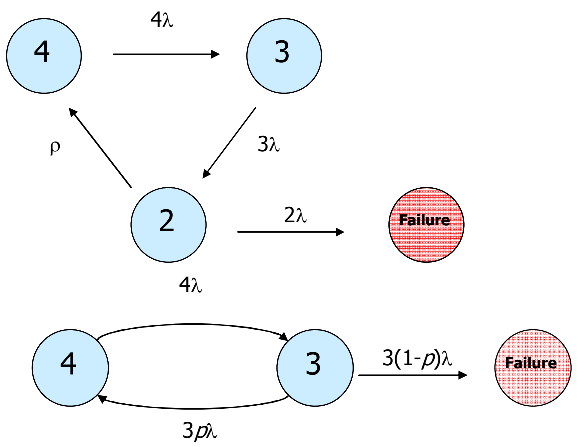

Figure 7: Markov model: Repairable disk array

7. Small Disk Arrays with Repair

Since small repair-free disk systems cannot provide sufficient survivability guarantees, we now look at small disk arrays that are repaired only in an emergency. Basically, we assume a 2D+P+Q layout that is repaired whenever two disks have failed. When modeling such a system, we run into the problem of how to model repair times. In our settings, we cannot assume a service technician call to the home office or small business. Therefore, repair involves shipping to the manufacturer, replacement by a new system, and copying data from the old array.

Under these circumstances, we also have to accommodate other reasons for data loss such as shipping damage and service provider error. Assume that for the moment, we assume an exponential repair time with rate ρ. We can then model our system as in Figure 7, where the initial state is state 4, representing the system without disk failure. The other states are states 3 and 2, labeled by the number of available disks. In state 2, we have an additional failure transition to the failure state, taken with rate 2λ and a repair transition back to the initial state 4, taken with rate $\rho\lambda$.

The resulting system of differential equations is:

Table 5: Economic life span for various repair parameters

Exponential Repair Rate				
9s	ρ = 10	ρ = 100	ρ = 1000	ρ = 10000
1	0.6061	3.3224	30.9729	307.5440
2	0.1756	0.4472	3.0845	29.4659
3	0.0702	0.1129	0.4294	3.0620
4	0.0307	0.0385	0.1033	0.4276
5	0.0139	0.0154	0.0309	0.1022

Constant Repair Rate				
9s	ρ = 10	ρ = 100	ρ = 1000	ρ = 10000
1	0.7778	6.3051	61.6189	614.762
2	0.1565	0.7303	6.0071	58.7713
3	0.0440	0.1523	0.7258	5.9793
4	0.0134	0.0430	0.1519	0.7253
5	0.0042	0.0131	0.0429	0.1518

$$\begin{aligned} p_4' &= -4\lambda p_4 + \rho p_2 & p_4(0) &= 1 \\ p_3' &= 4\lambda p_4 - 3\lambda p_3 & p_3(0) &= 0 \\ p_2' &= 3\lambda p_3 - (2\lambda + \rho) p_2 & p_2(0) &= 0 \end{aligned}$$

The general solution (obtained with Mathematica) involves solving polynomial equations and is too involved to represent here. However, we obtain the economic life span for various values of the repair time parameter. The results, in Table 5, are quite encouraging, even for very low repair rates – only 10 times faster than the disk failure rate (first column). The columns in Table 5 are indexed by the repair rate ratio ρ. A value of ρ = 10000 means that the average repair time is 10000 times less than the disk MTTF.

Our Markov model depicted in Figure 7 top uses an exponentially distributed repair time. This assumption is obviously false, since it assigns a positive probability to a repair time of a minute and similarly a positive probability to repair times longer than a decade. To gauge the influence of this assumption, we now assume that the repair time is a predetermined value of 1/ρ. If we consider a system in state 3 in our Markov model in Figure 7 that is suffering a transition to state 2, which triggers a repair. Assuming a constant repair time, our system will transit to state 4 if the repair is successful and to the failure state if another disk failure intervenes before the repair becomes effective. The chance for the former is $p = \exp(-\lambda / \rho)$. If we decide to neglect the time that the system is in state 2, we e obtain a transition from state 3 to state 4 taken at rate $3p\lambda$ and from state 3 to the failure state at rate $3(1-p)\lambda$. We depict the resulting approximate model in Figure 7 bottom. Assuming that the time between disk failures is exponentially distributed and setting $\lambda = 1$ to obtain normalized values, we can capture the state of the system in the following system of ordinary differential equations:

$$\begin{aligned} p_4'(t) &= -4p_4(t) + 3p \cdot p_3(t) & p_4(0) &= 1 \\ p_3'(t) &= 4p_4(t) - 3p_3(t) & p_3(0) &= 0 \end{aligned}$$

with p = exp(-1/ρ). The explicit solution for the survival rate $R(t) = p_4(t) + p_3(t)$ is somewhat involved and we do not write it down here. Because we do not take the repair time into account, it is a pessimistic (lower) bound. When we tabulate the results, we find a significant difference. For very small life spans (the lower right corners of Table 5), our neglect of the time spent during repair shows up, whereas for longer ones, our life spans are close to twice as good. Our results show that modeling repair time with an exponential distribution can also lead to significant errors.

8. Conclusions and Related Work

We have investigated the feasibility of building data storage solutions that would not need any servicing during their useful life span, as these solutions would eliminate the need of costly on-site repairs. We found that cost-effective solutions simply do not guarantee sufficient data survival. Based on our investigation, we propose instead low-maintenance disk arrays that would not be repaired until they reach a state where they cannot tolerate any additional failure. We have shown that one such solution is actually feasible and have described a possible implementation.

Our proposal is only directed to a situation where a disk failure would result in a service call, whose costs might approach the costs of the hardware being serviced. Other strategies should apply to larger installations where the possibility of a human error during the disk replacement process is a major concern. Anecdotal evidence amply suggests that mistakes in repairs constitute one of the major causes of data loss in disk arrays. From the manufacturer's perspective, the loss to the customer is compounded by the appearance of incompetence of the staff. The main advantage of low maintenance solutions will then be a reduction in the likelihood of these mishaps.

Relatively little work has been dedicated to disk arrays that self-reorganize to adjust to device failures. *Sparing* is one such form of adaptation to disk failures. Adding a spare disk to a disk array provides the replacement disk for the first failure. Distributed sparing [TM97] gains performance benefits in the initial state and degrades to normal performance after the first disk failure. The authors *et al.* [PSL06] have recently presented a mirrored disk array organization that adapts itself to successive disk failures. When all disks are operational, all data are mirrored on two disks. Whenever a disk fails, the array starts using (n – 1)-out-of-n codes in such a way that no data are left unprotected.

Our modeling proposes a new measure of disk array reliability, the economic life span. In our context, we can express the economic life span of a disk array as a factor of the expected life span of an individual disk and the required data survival probability. In addition, our study has shown that the Markovian assumption that failure rates are exponentially distributed can lead to misleading results, even for small arrays. This complements earlier results on the impact of disk infant mortality for large disk arrays by Xin et al. [XS+05] and a more recent investigation by Elerath and Pecht [EP07], who use an alternative technique based on field data for disk reliability.

References

[BM93] W. Burkhard and J. Menon. "Disk array storage system reliability," *Proc. 23rd International Symposium on Fault-Tolerant Computin*g (FTCS-23), pp. 432-441, Toulouse, France, Aug. 1993.

[CL+94] P. M. Chen, E. K. Lee, G. A. Gibson, R. Katz, and D. Patterson. "RAID, High-performance, reliable secondary storage," *ACM Computing Surveys*, 26(2):145–185, 1994.

[El00a] J. G. Elerath. "AFR: problems of definition, calculation and measurement in a commercial environment," *Proc. Annual Reliability and Maintainability Symposium*, January 2000.

[El00b] J. G. Elerath. "Specifying reliability in the disk drive industry: No more MTBFs," *Proc. Annual Reliability and Maintainability Symposium*, January 2000.

[ES04] J. G. Elerath and S. Shah. "Server class drives: How reliable are they?" *Proc. Annual Reliability and Maintainability Symposium*, 2004.

[EP07] J. G. Elerath and M. Pecht. "Enhanced Reliability Modeling of RAID Storage Systems", *Proc. 37th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DNS'07).*

[El07] J.G. Elerath: "Hard disk drives: The good, the bad, and the ugly!," *Queue*, vol. 5(6), 2007.

[GGL03] S. Ghemawat, H. Gobioff, and S. Leung. "The Google file system," *Proc. 19th ACM Symposium on Operating Systems Principles (SOSP)*, pp. 29–43, Oct. 2003.

[PSL06] J.-F. Pâris, T. J. Schwarz and D. D. E. Long. "Self-Adaptive Disk Arrays," *Proc. 8th International Symposium on Stabilization, Safety, and Security of Distributed Systems* (SSS 2006), Dallas, TX, pp. 469–483, Nov. 2006.

[PGK88] D. A. Patterson, G. A. Gibson, and R. H. Katz. "A case for redundant arrays of inexpensive disks (RAID)," *Proc. SIGMOD 1988 International Conference on Data Management*, Chicago, IL, pp. 109–116, June 1988.

[PWB07] E. Pinheiro, W.-D. Weber, L. A. Barroso. "Failure trends in a large disk drive population," *Proc. 5th USENIX Conference on File and Storage Technologies* (FAST 2007), San Jose, CA, pp. 17–28, Feb. 2007.

[SG07] B. Schroeder and G. A. Gibson. "Disk Failures in the Real World: What Does an MTTF of 1,000,000 Hours Mean to You?" *Proc. 5th USENIX Conference on File and Storage Technologies (FAST 2007)*, San Jose, CA, February 2007.

[SX+04] T. J. E. Schwarz, Q. Xin, E. L. Miller, D. D. E. Long, A. Hospodor, S. Ng. „Disk Scrubbing in Large Archival Storage Systems," *Proc. 12th IEEE International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunications Systems* (MASCOTS), October 2004.

[TM97] A. Thomasian and J. Menon. "RAID 5 performance with distributed sparing." *IEEE Transactions on Parallel and Distributed Systems* 8(6):640–657, June 1997

[XM+04] Q. Xin, E. Miller, T. Schwarz. "Evaluation of Distributed Recovery in Large-Scale Storage Systems," *Proc. Thirteenth IEEE International Symposium on High Performance Distributed Computing* (HPDC-13), Honolulu, HI, pp. 172–181, June 2004.

[XS+05] Q. Xin, T. J. E. Schwarz, S.J., E. L. Miller. "Disk Infant Mortality in Large Storage Systems," *Proc. 13th IEEE International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunications Systems* (MASCOTS), Atlanta, GA, pp.125–134, August 2005.

[YS09] J. Yang and F.-B. Sun. "A comprehensive review of hard-disk drive reliability." *Proc. Annual Reliability and Maintainability Symposium,* 1999.