Improving Disk Array Reliability Through Faster Repairs

(Extended Abstract)

Jehan-François Pâris
Department of Computer Science
University of Houston
Houston, TX
jfparis@uh.edu

Thomas Schwarz, S. J.
Department of MSCS
Marquette University
Milwaukee, WI
thomas.schwarz@marquette.edu

Darrell D. E. Long[1]
Department of Computer Science
University of California, Santa Cruz
Santa Cruz, CA
darrell@cs.ucsc.edu

I. INTRODUCTION

Magnetic disk capacities have grown over the last decades by a factor of at least ten thousand. While the minicomputer disk drives of the late eighties could only store 600 MB of data [5], 8TB or 10TB disk drives are common today. The same is not true for disk transfer rates: they are just one hundred times higher than those of the late eighties. As a result, copying the entire contents of a disk will take now considerably more time than thirty or forty years ago.

This development has now a critical impact on the reliability of disk arrays. Magnetic disks are the least reliable component of most computer systems, with failure rates that can sometimes reach 25 percent per year [2]. As a result, all medium-size to large-size disk arrays include some redundancy and provisions for the quick replacement of failed units. Mirroring, RAID level 5 [5] and RAID level 6 [3] [6] are the best known exemplars of these redundant disk array organizations. In all cases, the level of protection that these organizations offer depends on both the prompt detection of disk failures and the time it will take to reconstitute the contents of failed disks on spare devices. In particular, any increase in this reconstruction time will have a negative impact on the reliability of the array.

A well-known solution to this problem is *declustering* [1] [4]. Declustering partitions each physical disk into k physically contiguous parts, which we will call *disklets* [8], and combines these disklets into distinct reliability stripes. The scheme requires that each of these parity stripe have no more than one disklet located on any physical disk.

The main advantage of declustering disk array organizations is their faster recovery. Once a disk has failed, we can recover the contents of its k disklets in parallel through their own parity stripes. This will allow us to reconstruct the lost data k times faster as long as no two of the k parity stripes has to access disklets located on the same physical disk.

To the best of our knowledge, none of the recent papers analyzing the fault-tolerance of disklet-based storage systems have ever tried to estimate their five-year reliability, preferring instead to evaluate the probability of a data loss in the presence of a fixed number of failures [7].

We present here a preliminary study of the effect of declustering on the five-year reliability of disk arrays. We chose for our investigation an array consisting of ten RAID level 6 reliability stripes with ten disks each for a total of one hundred disks. Such organization is large enough to demonstrate the benefits of declustering while remaining small enough to let us neglect the impact of non-fatal quadruple disk failures on the array reliability.

II. OUR MODEL

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 intractable as the complexity of the system grows. A simpler option is to use instead the five-year reliability of the array. As this value is typically very close to one, we will express it in "nines" using the formula $n_n = -\log_{10}(1 - R_d)$, where R_d is the five-year reliability of the array. Thus four nines corresponds to a five-year reliability of 99.99 percent.

We develop first a generic Markov model that will apply to both declustered and non-declustered disk arrays. The specific behavior of each disk array will be represented by the three parameters N, β and μ, where N is the number of disks in the array, β is the probability that the array will lose data in the presence of a simultaneous failure of three disks, and μ is the disk repair rate. For the sake of simplicity, we will neglect the probability that the array will tolerate quadruple failures without data loss, assuming that this probability is small enough to be ignored as a first approximation.

Our model consists of an array of drives with independent failure modes. Whenever a drive fails, a repair process is immediately initiated for that drive. Should several drives fail, this repair process will be performed in parallel on those drives. We assume that drive failures are independent events and are exponentially distributed with mean λ. In addition, we require repairs to be exponentially distributed with mean μ. Both hypotheses are necessary to represent our system by a Markov process with a finite number of states.

Fig, 1 displays our state transition probability diagram. State <0> is the initial state where all N drives are operational and no drive has failed. Should any of the drives fail, the system would move to state <1> with an aggregate failure rate $N\lambda$. A second failure would bring the system to state <2> with

---

[1] Supported in part by National Science Foundation Grants CCF-1219163, CNS-1528179, and IIP-1266400, by the Department of Energy under Award Number DE-FC02-10ER26017/DE-SC0005417, and by the industrial members of the Center for Research in Storage Systems.

978-1-5090-5252-3/16/$31.00 ©2016 IEEE

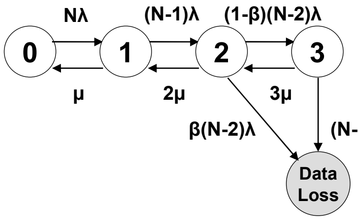

Fig. 1. State–transition probability diagram of a disk array with N disks that tolerates all double and most, but not all, triple failures.

rate $(N-1)\lambda$. Since some but not all triple failures will result in a data loss, we have two failure transitions from state <2> namely:

1. A transition to the data loss state with rate $\beta(N-2)\lambda$ where the actual value of the β parameter will depend on the specific storage organization.
2. A transition to state <3> with rate $(1-\beta)(N-2)\lambda$.

As we assumed that all quadruple failures were fatal, there is a single failure transition leaving state <3>.

Recovery transitions are more straightforward: they bring the array from state <3> to state <2>, then from state <2> to state <1> and finally from state <1> to its initial state <0>.

We derive first the Kolmogorov system of differential equations that describes the behavior of the array then use it to compute the mean time to data loss (MTTDL) of the system using:

$$MTTDL = \sum\nolimits_{i=0}^{3} p_i^*(0) \; ,$$

where $p_i^*(s)$ is the Laplace transform of the probability $p_i(t)$that the system is in state $<i>$ at time t. We then convert this MTTDL into five-year reliability, using:

$$R_d = \exp(-d/MTTDL)$$

where d is a five-year interval expressed in the same units as the MTTDL. Observe that the above formula implicitly assumes that long-term failure rate 1/MTTDL does not significantly differ from the average failure rate over the first five years of the array.

We applied our model to a set of m RAID level 6 parity stripes that contain n disks each; thus $N = nm$. As long as the array is not declustered, the only triple failures that will cause a data loss are the failures of three disks that belong to the same parity stripe. Given that there are m parity stripes and the total number of triple failures is $\binom{mn}{3}$, the fraction β of triple failures that will result in a data loss is

$$\beta = m\binom{n}{3}\Big/\binom{mn}{3}.$$

Consider now what would happen if each disk is partitioned into k disklets. A data loss will now occur whenever three disks fail and three of their disklets are on the same parity stripe. As a result, the fraction β of triple failures that result in a data loss can be approximated by $\beta' = km\binom{n}{3}\Big/\binom{mn}{3}$. At the same time, the repair rate of the new array will be $\mu' = k\mu$.

TABLE I. FIVE-YEAR SURVIVAL RATES AND DATA LOSS PROBABILITIES FOR A DISK ARRAY CONSISTING OF TEN RAID LEVEL 6 RELIABILITY STRIPES WITH TEN DISKS EACH FOR DIFFERENT NUMBERS OF DISKLETS PER PHYSICAL DISK, A DISK MTTF OF 100,000 HOURS AND A FULL DISK MTTR OF 24 HOURS.

Disklets per disk	**Five-year reliability (nines)**	**Five-year data loss probability**	**Data loss probability ratios**
1	3.74	0.01799%	1.00
2	4.25	0.00568%	0.31
4	4.62	0.00260%	0.14
8	4.94	0.00115%	0.06

Table I summarizes our results. We assumed a disk MTTDL of 100,000 hours and a full disk mean time to repair (MTTR) of 24 hours. As we can see, partitioning each disk into 2, 4 or 8 disklets results in a significant improvement of the five-year reliability of the array. The best results were obtained for $k = 8$ and resulted in a reduction of 94 percent of the cumulative probability of a data loss over a five-year interval. We did not consider higher values of k given the limited size of the array because it would have increased the likelihood that the reconstruction process would have to access two or more disklets located on the same physical disk.

The improvement became much less significant when we increased the disk MTTDL from 100,000 to 200,000 hours and much more significant when we reduced it to 35,000 hours. Space limitations prevent us from discussing these results.

III. CONCLUSION

We have presented a preliminary study of the impact of declustering on the five-year reliability of a disk array consisting of ten RAID level 6 reliability stripes with ten disks each for a total of one hundred physical disks. Our results show that partitioning each disk into eight disklets that belong to separate parity stripes could reduce by 94 percent the cumulative probability of a data loss over a five-year interval.

REFERENCES

[1] G. A. Alvarez, W. A. Burkhard, L. J. Stockmeyer, and F. Cristian, Declustered disk array architectures with optimal and near-optimal parallelism, *ACM SIGARCH Computer Architecture News*, 26(3):109–120, 1998.

[2] B. Beach, What hard drive should I buy? http://blog.backblaze.com/2014/01/21/, January 21, 2014, retrieved September 24, 2016.

[3] W. A. Burkhard and J. Menon, Disk array storage system reliability. *Proc. 23rd FTC Symposium*, pp. 432–441, June 1993.

[4] M. Holland and G. A. Gibson, Parity declustering for continuous operation in redundant disk arrays, *Proc.5th ASPLOS Conference,* pp. 23–35, Oct. 1992.

[5] D.A. Patterson, G. Gibson, and R. H. Katz, A case for redundant arrays of inexpensive disks (RAID). *Proc. SIGMOD 1988 Int. Conference,* pp. 109–116, June 1988.

[6] T. J. E. Schwarz, *Reliability and Performance of Disk Arrays*, Ph.D. Thesis, Department of CSE, U. C. San Diego, 1994.

[7] T. Schwarz, S. J., A. Amer, T. Kroeger, E. L. Miller, D. D. E. Long and J.-F. Pâris, Reliable Storage at Exabyte Scale, *Proc. 24th MASCOTS Symposium*, Sep. 2016.

[8] A. Wildani, T. Schwarz, S, J., E. L. Miller and D. D. E. Long, Protecting against rare event failures in archival systems, *Proc. 17th MASCOTS Symposium*, pp. 246–256, Sep. 2009.