

UC Santa Cruz

UC Santa Cruz Previously Published Works

Title

When MTTDLs Are Not Good Enough: Providing Better Estimates of Disk Array Reliability

Permalink

<https://escholarship.org/uc/item/0q95f2m0>

Authors

Paris, Jehan-Francois

Schwarz, Thomas JE

Long, Darrell

et al.

Publication Date

2008-12-03

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

# When MTDLs Are Not Good Enough: Providing Better Estimates of Disk Array Reliability

Jehan-François Pâris  
Dept. of Computer Science  
University of Houston  
Houston, TX 77204-3010  
paris@cs.uh.edu

Thomas J. E. Schwarz  
Dept. of Computer Engineering  
Santa Clara University  
Santa Clara, CA 95053  
tjschwarz@scu.edu

Darrell D. E. Long<sup>1</sup>  
Dept. of Computer Science  
University of California  
Santa Cruz, CA 95064  
darrell@cs.ucsc.edu

Ahmed Amer<sup>2</sup>  
Dept. of Computer Science  
University of Pittsburgh  
Pittsburgh, PA 15260  
amer@cs.pitt.edu

**Abstract**—While mean time to data loss (MTDL) provides an easy way to estimate the reliability of redundant disk arrays, it fails to take into account the relatively short lifetime of these arrays. We analyzed five different disk array organizations and compared the reliability estimates obtained using their mean times to data loss with the more exact values obtained by directly solving their corresponding Markov model. We observed that the conventional MTDL approach generally provided a good estimate of the long-term reliability of arrays—with the exception of non-repairable arrays—while significantly underestimating the short-term reliability of these arrays.

## I. INTRODUCTION

A critical issue for any large data storage system is how to ensure the survival of the data in the presence of equipment failures. Given the limitations of backup solutions, the best way to achieve this goal is to use redundant storage systems. Two techniques that can be used are *mirroring* and *erasure coding*. Mirroring maintains two or more identical copies of the data on distinct disks. Erasure codes, also known as  $m$ -out-of- $n$  codes, group disks into sets of  $n$  disks that contain enough redundant information to tolerate the loss of  $n - m$  disks. Their best-known implementation is RAID level 3 and 5, which use  $n - 1$ -out-of- $n$  codes [2, 3].

Adding redundancy to a storage system incurs additional costs both in terms of extra storage space and additional update complexity. Since these costs vary widely among all the possible options, it is important to be able to evaluate the survival rate that a given storage system can achieve over time.

The most popular tool for characterizing the reliability of a storage solution is its *mean time to data loss* (MTDL). It

offers the two advantages of being simple to compute and easy to understand, but as we will see, it also has a major limitation. MTDLs characterize fairly well the behavior of a disk array that would remain in service until it fails without being ever replaced for any reason other than a device failure. In reality, disk arrays have much shorter lifetimes and are typically replaced five to seven years after their initial deployment. As a result, most redundant disk arrays will be replaced before they experience any data loss due to a failed disk.

Rather than relying on MTDLs, we should instead focus on characterizing the reliability of disk arrays during their useful lifetime. Two recent studies of disk array reliability [4, 5] have addressed this issue and included data loss predictions obtained through a direct solution of the system of differential equations that characterize the stochastic behavior of each storage system.

We present the first comparison between these new figures and those obtained through the conventional MTDL approach. The disk arrays considered in our study include mirrored and triplicate disks both with and without repairs and a RAID level 5 array consisting of ten disks. The main conclusion of our study is that the conventional MTDL approach generally provides good estimates of the long-term reliability of repairable disk arrays, but significantly underestimates their short-term reliability.

## II. MIRRORED AND TRIPPLICATE DISKS

The first storage organizations we considered were arrays consisting of one, two or three disks, each holding an identical copy of the same data. We distinguished disk arrays that could not be repaired after a disk failure from arrays that would be promptly repaired after each disk failure.

---

<sup>1</sup> Supported in part by the Petascale Data Storage Institute under Department of Energy Award DE-FC02-06ER25768.

<sup>2</sup> Supported in part by the National Science Foundation under Award Number 0720578.




Fig. 1. State transition probability diagram for mirrored and triplicate disks that are never repaired.

TABLE I

ECONOMIC LIFESPANS OF SINGLE DISKS, MIRRORRED DISKS AND TRIPPLICATE DISKS WITHOUT REPAIRS.

| <i>Nines</i> | <i>Simple</i> | <i>Mirrorred</i> | <i>Triple</i> |
|--------------|---------------|------------------|---------------|
| 2            | 0.01005       | 0.10536          | 0.24265       |
| 3            | 0.00100       | 0.03213          | 0.10536       |
| 4            | 1.00E-04      | 0.01005          | 0.04753       |
| 5            | 1.00E-05      | 0.00317          | 0.02178       |

TABLE II.

THE SAME ECONOMIC LIFESPANS COMPUTED USING THE MTTDLs OF THE ARRAYS.

| <i>Nines</i> | <i>Simple</i> | <i>Mirrorred</i> | <i>Triple</i> |
|--------------|---------------|------------------|---------------|
| 2            | 0.01005       | 0.01508          | 0.01843       |
| 3            | 0.00100       | 0.00150          | 0.001831      |
| 4            | 1.00E-04      | 0.00015          | 0.000183      |
| 5            | 1.00E-05      | 1.501E-05        | 1.83E-05      |

#### A. First Case: The Disks Cannot Be Repaired

As in almost all studies of disk array reliability, we will assume that disk failures are independent events, exponentially distributed with failure rate  $\lambda$ .

The survival function  $S_1(t)$  of a single disk at time  $t$  is given by the differential equation

$$S_1'(t) = -\lambda S_1(t)$$

with the initial condition  $S_1(0) = 1$ .

Its solution is

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

We model the survival function  $S_2(t)$  of a pair of mirrored disks using the standard Markov model depicted on the top of Fig. 1. We label the non-failure states by the number of operational disks. The starting state is state 2, from which we transition to state 1 at rate  $2\lambda$ , whenever one of the two disks fails. We can capture the probability  $p_i(t)$  of being in state  $i$  at time  $t$  in a system of ordinary differential equations

$$\begin{aligned} p_2'(t) &= -2\lambda p_2(t) \\ p_1'(t) &= 2\lambda p_2(t) - \lambda p_1(t) \end{aligned}$$

with initial conditions

$$\begin{aligned} p_2(0) &= 1 \\ p_1(0) &= 0 \end{aligned}$$

whose solution is

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

The survival  $S_2(t)$  of our pair a mirrored disks is

$$S_2(t) = p_2(t) + p_1(t) = e^{-2\lambda t}(2e^{\lambda t} - 1).$$

To estimate the MTTDL of our system, we first compute the Laplace transforms of its system of differential equations

$$\begin{aligned} sp_2^*(s) - 1 &= -2\lambda p_2^*(s) \\ sp_1^*(s) &= 2\lambda p_2^*(s) - \lambda p_1^*(s) \end{aligned}$$

Observing that the mean time to data loss (MTTDL) of the array is given by

$$\text{MTTDL} = \sum_i p_i^*(0),$$

we obtain

$$\text{MTTDL} = p_2^*(0) + p_1^*(0) = \frac{1}{\lambda} + \frac{2}{\lambda} = \frac{3}{2\lambda}.$$

The bottom Markov model in Fig. 1 describes the case of triplicate disks. Its system of differential equations is

$$\begin{aligned} p_3'(t) &= -3\lambda p_3(t) \\ p_2'(t) &= 3\lambda p_3(t) - 2\lambda p_2(t) \\ p_1'(t) &= 2\lambda p_2(t) - \lambda p_1(t) \end{aligned}$$

with initial conditions

$$\begin{aligned} p_3(0) &= 1 \\ p_2(0) &= 0 \\ p_1(0) &= 0 \end{aligned}$$

Solving the system gives us its survival function

$$S_3(t) = p_3(t) + p_2(t) + p_1(t) = e^{-3\lambda t}(3e^{2\lambda t} - 3e^{\lambda t} + 1).$$

Similarly, we find out that the MTTDL of the triplicate disks is

$$\text{MTTDL} = p_3^*(0) + p_2^*(0) + p_1^*(0) = \frac{11}{6\lambda}.$$

We define the *economic lifespan*  $L(r)$  of an array as the maximum time interval for which data stored on that array will have a probability  $r$  to survive intact. In other words,  $L(r)$  is the solution of  $S_n(t) = r$ . In addition we express  $L(r)$  in multiples of the disk MTTF ( $1/\lambda$ ) to obtain dimensionless values. Table I displays the economic lifespans of simple, mirrored and triplicate disks for reliability levels  $r$  equal to 0.99, 0.999, 0.9999, and 0.99999, that is, 2, 3, 4, and 5 *nines*.

To obtain the same values from our estimates of the array MTTDLs, we assume that the array failure rate  $\lambda_n$  will remain equal to  $1/\text{MTTDL}$  over the lifetime of the array and obtain

$$S_n(t) = \exp\left(-\frac{t}{\text{MTTDL}}\right).$$

As we can see from Table II, the economic lifespans of mirrored and triplicate disks computed from their MTTDL are much lower than those obtained by directly solving the model. From Equations (1), we see that the actual hazard rate of the pair of mirrored disks

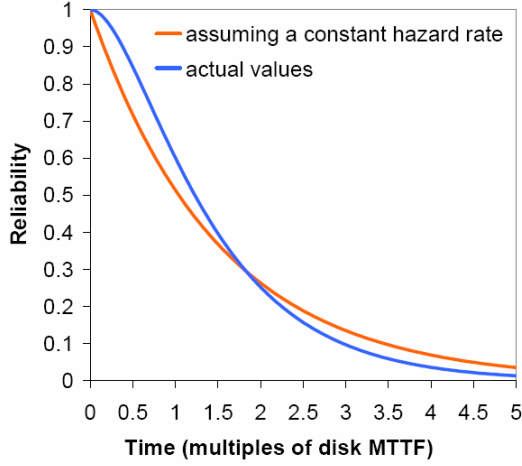


Fig. 2. Reliability of a pair of mirrored disks that cannot be repaired.

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

is initially equal to zero and reaches its maximum for  $t = (\log 2)/\lambda$  while the hazard rate we computed from its MTDDL

$$\frac{1}{MTDDL} S_n(t) = \frac{1}{MTDDL} \exp\left(-\frac{t}{MTDDL}\right),$$

is a monotonically decreasing function.

Fig. 2 compares the reliabilities obtained through the two techniques. As we can see, reliability figures that we obtained assuming a constant hazard rate are much lower than the actual figures during an initial period roughly equal to the disk MTTF. This is a fairly large time interval if we consider that most estimates of disk MTTFs fall between ten and thirty years depending on their working environments. Very few disk arrays last that long before getting replaced.

One way to obtain better estimates of the economic lifespan of disk arrays would be to change the way we compute their MTDDLs. Instead of considering the whole potential lifetime of the array, we should compute the array MTDDL over a shorter observation period  $T_o$ . To achieve this goal, we will assume that disk arrays will be replaced at a rate  $\nu = 1/T_o$ .

Fig. 3 displays the state transition probability diagrams of mirrored and triplicate disks that are never repaired but periodically replaced at a rate  $\nu$ . These diagrams are identical to the diagrams of Fig. 1 but for the new transitions of rate  $\nu$  returning to state  $\langle 0 \rangle$  from all other states.

Solving the systems of differential equations of the two models, we obtain:

$$S_2(t) = \frac{3\lambda + \nu}{2\lambda^2}$$

for the pair of mirrored disks and

$$S_3(t) = \frac{11\lambda^2 + 6\lambda\nu + \nu^2}{6\lambda^3}$$




Fig. 3. State transition probability diagram for mirrored and triplicate disks that are never repaired and periodically replaced.

TABLE III.

MTTDL-BASED ESTIMATES OF THE ECONOMIC LIFESPANS OF SINGLE DISKS, MIRRORING DISKS AND TRIPPLICATE DISKS ASSUMING THAT FAILED DISKS DO NOT GET REPLACED BUT ARRAYS GET PERIODICALLY REPLACED AT A RATE EQUAL TO THE INVERSE OF THEIR ECONOMIC LIFESPANS.

| <i>Nines</i> | <i>Simple</i> | <i>Mirrored</i> | <i>Triple</i> |
|--------------|---------------|-----------------|---------------|
| 2            | 0.01005       | 0.07883         | 0.15299       |
| 3            | 0.00100       | 0.02313         | 0.06175       |
| 4            | 1.00E-04      | 0.00715         | 0.02691       |
| 5            | 1.00E-05      | 0.00224         | 0.01223       |

for the triplicate disk organization.

We then used these expressions to compute the economic lifespans  $L(r)$  of the two arrays. Since we wanted to focus on the behaviors of the arrays within their economic lifespans, we set their replacement rates equal to  $1/L(r)$  using a simple iterative technique. Table III summarizes our results. As we can see, our estimates of the economic lifespans of mirrored and triplicate disks still underestimate the values we obtained by directly solving their differential equations but are now much closer to these values. The new estimates of the economic lifespans of the pair of mirrored disks are 25 to 30 percent lower than their true values while the new estimates of the economic lifespans of the triplicate disks are 35 to 45 percent lower than their true values.

### B. Second case: the disks can be repaired

Let us consider the case where failed disks are promptly replaced after each disk failure and assume that disk repair times are exponentially distributed with rate  $\mu$ .

We model the survival function  $S_2(t)$  of a pair of repairable mirrored disks using the Markov model depicted in Fig. 4. Solving its system of differential equations, we obtain:

$$S_2(t) = p_2(t) + p_1(t) = \frac{\exp(-(1/2)t(3 + \mu + R)) \cdot (-3 - \mu + R + \exp(tR)(3 + \mu + R))}{2R}$$




Fig. 4. State transition probability diagram for a pair of mirrored disks when failed disks are promptly replaced.

TABLE IV

ECONOMIC LIFESPANS OF REPAIRABLE MIRRORRED DISKS FOR VARIOUS REPAIR RATE RATIOS  $\mu/\lambda$ .

| <i>Nines</i> | $\mu/\lambda = 10^3$ | $\mu/\lambda = 10^4$ | $\mu/\lambda = 10^5$ |
|--------------|----------------------|----------------------|----------------------|
| 2            | 5.04123              | 50.2669              | 502.532              |
| 3            | 0.50275              | 5.00410              | 50.0265              |
| 4            | 0.05115              | 0.50028              | 5.00041              |
| 5            | 0.00601              | 0.05012              | 0.50003              |
| 6            | 0.00120              | 0.00510              | 0.05001              |

TABLE V

THE SAME ECONOMIC LIFESPANS COMPUTED USING THE MTDDLs OF THE ARRAYS.

| <i>Nines</i> | $\mu/\lambda = 10^3$ | $\mu/\lambda = 10^4$ | $\mu/\lambda = 10^5$ |
|--------------|----------------------|----------------------|----------------------|
| 2            | 5.040243             | 50.26675             | 502.5319             |
| 3            | 0.501751             | 5.004002             | 50.02652             |
| 4            | 0.050150             | 0.500175             | 5.000400             |
| 5            | 0.005015             | 0.050015             | 0.500018             |
| 6            | 0.000502             | 0.005001             | 0.050002             |

where  $R = \sqrt{1 + \mu(6 + \mu)}$ .

Table IV displays the economic lifespans for various numbers of nines and selected values of  $\mu/\lambda$  between  $10^3$  and  $10^5$ . A closer look at Table IV shows that increasing this  $\mu/\lambda$  ratio by 100 yields about the same increase in the economic lifespan.

Table V displays the same economic life spans computed using the MTDDL of the array. As we can see, they are in fairly good agreement with the values obtained by directly solving the model. The sole significant discrepancies occur when computing economic lifespans less than 0.01 times the disk MTTF. Given that disk MTTFs are likely to vary between 100,000 and 300,000 hours [6-8], this threshold corresponds to economic lifespans varying between one thousand and three thousand hours.

Similarly, the MTDDL of the mirrored disks is

$$\text{MTDDL} = \frac{3\lambda + \mu}{2\lambda^2}.$$

The main reason for this better agreement is the existence of a repair transition from state <1> to state <2>, which rapidly brings the array back to its original state. Since this process is likely to be repeated many times over the array




Fig. 5. State transition probability diagram for a RAID level 5 array with ten disks.

TABLE VI

ECONOMIC LIFESPANS FOR A REPAIRABLE RAID LEVEL 5 DISK ARRAY WITH TEN DISKS ASSUMING  $\lambda = 1$  AND VARIOUS REPAIR RATE RATIOS  $\mu/\lambda$ .

| <i>Nines</i> | $\mu/\lambda = 10^3$ | $\mu/\lambda = 10^4$ | $\mu/\lambda = 10^5$ |
|--------------|----------------------|----------------------|----------------------|
| 2            | 0.114800             | 1.119000             | 11.16920             |
| 3            | 0.012300             | 0.111500             | 1.111890             |
| 4            | 0.001984             | 0.011230             | 0.111100             |
| 5            | 0.000512             | 0.001213             | 0.011120             |
| 6            | 0.000153             | 0.000197             | 0.001121             |

TABLE VII

THE SAME ECONOMIC LIFESPANS COMPUTED USING THE MTDDLs OF THE ARRAYS.

| <i>Nines</i> | $\mu/\lambda = 10^3$ | $\mu/\lambda = 10^4$ | $\mu/\lambda = 10^5$ |
|--------------|----------------------|----------------------|----------------------|
| 2            | 0.113792             | 1.118826             | 11.16916             |
| 3            | 0.011328             | 0.111378             | 1.111878             |
| 4            | 0.001137             | 0.011133             | 0.111137             |
| 5            | 0.000113             | 0.001113             | 0.011113             |
| 6            | 0.000011             | 0.000111             | 0.001111             |

lifetime, it brings the average hazard rate of the mirrored disks closer to the rate directly computed from the differential equations of the system.

In addition, we can see that the most serious discrepancies occur when  $\mu/\lambda = 10^3$ , which is a fairly low repair rate to failure rate ratio. Observe first that the MTTFs of individual disks are at least 100,000 hours, which correspond to a disk failure rate of  $10^{-5}$  failures per hour. A repair rate to failure rate ratio of  $10^3$  implies a disk repair rate of  $10^{-2}$  repairs per hour, which corresponds to a mean disk repair time of slightly more than four days.

### III. RAID ARRAYS

We can apply the same approach to estimate the reliability of a RAID level 5 array consisting of ten devices. Solving the Markov model of Fig. 5, we obtain its survival function

$$S_{10}(t) = p_{10}(t) + p_9(t) = \frac{\exp(-\frac{1}{2}t(19 + \mu + R)(-19 - \mu + R + \exp(tR)(19 + \mu + R)))}{2R}$$

with

$$R = \sqrt{1 + \mu(38 + \mu)}$$

and its mean time to data loss

$$MTTDL = \frac{19\lambda + \mu}{90\lambda^2}.$$

Table VI displays the economic life spans of our RAID array obtained by directly solving the model and Table VII displays the same results computed using the MTTDL of the array. As we observed for the repairable mirrored disks, the sole significant discrepancies between the two sets of values occur for very small economic lifespans and are the most pronounced when  $\mu/\lambda = 10^3$ . This should be expected as the error introduced by assuming a constant hazard rate is the greatest for very small lifespans and low repair rate to failure rate ratios tend to produce these lifespans.

#### IV. DISCUSSION

We have considered very simple disk arrays so far and would like to know how the difference between the true and the constant hazard rate reliability is affected by the size of the disk arrays. As an example, we consider a system of 20 identical components that can withstand up to 3 failures and thus needs 17 functioning components. Proceeding just as before, we obtain the survival rate

$$S_{17/20} = e^{-20\lambda t} \left( 1 + 20(e^{\lambda t} - 1) + 190(e^{2\lambda t} - 1) + 1140(e^{3\lambda t} - 1) \right)$$

and for the MTTDL

$$MTTDL_{17/20} = \frac{12617}{58140\lambda}.$$

We present the now widely different survival rates in Fig. 6.

We now investigate our intuition further by investigating systems of  $n$  components such that failure of 1, up to 2, up to 3, or up to 4 components constitute system failure. Accordingly, we call these systems 1, 2, 3, or 4 components failure-resilient. Examples for such systems are declustered disk arrays [1, 9]. Fig. 7 gives the Markov model for the case of a 4 components failure-resilient system. We calculated the survival probability  $S_{i,n}$  for these systems where  $i$  stands for the degree of resilience and  $n$  for the number of components. We then developed the Taylor series of the resulting expressions and obtained:

$$S_{1,n}(t) = e^{-n\lambda t} = 1 - n\lambda t$$

$$S_{2,n}(t) = 1 + \frac{1}{2}(n\lambda^2 - n^2\lambda^2)t^2$$

$$S_{3,n}(t) = 1 + \frac{1}{6}(-2n\lambda^3 + 3n^2\lambda^3 - n^3\lambda^3)t^3 + O(t)^4$$

$$S_{4,n}(t) = 1 + \frac{1}{24}(6n\lambda^4 - 11n^2\lambda^4 + 6n^3\lambda^4 - n^4\lambda^4)t^4 + O(t)^5$$

We calculate economic lifetimes by solving for  $t$  in  $S_{i,n}(t) = r$  with  $r \approx 1$ . As a result, the absence of terms of rank smaller than  $i$  in the Taylor series explains why we cannot approximate  $S_{i,n}(t)$  with a survival function with constant hazard rate for  $i \geq 2$ , since that survival function would have a non-zero linear component in the Taylor series.




Fig. 6. True and constant hazard rate survival rate of the 17/20 disk array.




Fig. 7: Markov model of a four-component failure-resilient system without repairs.

TABLE VIII

ECONOMIC LIFESPAN AT FOUR NINES OF A THREE-COMPONENT FAILURE-RESILIENT REPAIRABLE DISK ARRAY WITH TEN DISKS.

| $\mu/\lambda$ | <i>Economic Lifespan</i> |
|---------------|--------------------------|
| 10            | 0.009853                 |
| 100           | 0.012771                 |
| 1000          | 0.283207                 |
| 10000         | 27.81820                 |

Graphically, the absence of terms of rank smaller than  $i$  in  $S_{i,n}(t)$  shows itself in the flatness of the curve near  $t = 0$ .

These results also hold for repairable systems. In fact, the *same* Taylor series developments hold for repairable systems that are 1, 2, or 3 components failure-resilient. However, in this case, the magnitude of the coefficient of  $t^{i+1}$  is so large that using only the constant and the lowest-order member of the Taylor series is only justified for a reliability level  $r$  so close to 1 that it is of little practical use. To illustrate the fact, we give the economic lifespan of a 3 component failure-resilient system of 10 components in Table VIII, assuming a survival probability of 99.99% (four nines). As we can see, the lifespan depends heavily on the repair rate. To keep our table dimensionless, we set the component failure rate  $\lambda$  to 1. In fact, the Taylor series of  $S_{3,n}^{\text{repair}}(t)$  up to a  $t^4$  is:

$$1 + \frac{1}{6}(-2n\lambda^3 + 3n^2\lambda^3 - n^3\lambda^3)t^3 + \frac{1}{8}(-2n\lambda^4 + 5n^2\lambda^4 - 4n^3\lambda^4 + n^4\lambda^4 + 2n\lambda^3\rho - 3n^2\lambda^3\rho + n^3\lambda^3\rho)t^4$$

The following coefficients are even larger, showing that approximation by Taylor series is not a feasible approach to calculate economic lifespans.

## V. CONCLUSION

While MTTDLs provide an easy way to estimate the reliability of redundant disk arrays, they do not take into account the relatively short lifetime of these arrays and tend to overestimate the probability of a data loss over this lifetime. We analyzed five different disk array organizations and compared the reliability estimates obtained using their MTTDLs with the more exact values obtained by directly solving their Markov model. We observed that the MTTDL approach grossly underestimated the reliability of non-repairable redundant disk arrays and proposed a technique reducing the margin of error by assuming that the disk array was replaced at frequent intervals. In contrast, we found out that the same MTTDL approach provided fairly good estimates of the reliability of repairable redundant disk arrays as long as the individual disk repair rate remained well above one thousand times the individual disk failure rate.

These findings raise the important question of how to evaluate the reliability of complex redundant disk arrays that are not promptly repaired within hours of a disk failure. Since the systems of differential equations describing the behavior of these systems are likely to be intractable, extant analytic tools will not be able to evaluate the correctness of any MTTDL-based estimate of their reliability. The best alternative will be to turn to discrete simulation techniques, a step we plan to take in the near future.

## REFERENCES

- [1] G. Alvarez, W. Burkhard, L. Stockmeyer and F. Cristian, Declustered Disk Array Architectures with Optimal and Near-Optimal Parallelism, *Proc. 25th Annual International Symposium on Computer Architecture (ISCA '98)*, Barcelona, Spain, 1998
- [2] 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.
- [3] 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.
- [4] J.-F. Pâris, D. D. E. Long and T. J. E. Schwarz, Ensuring Data Survival on Solid-State Storage Devices, *Proc. 1<sup>st</sup> International Workshop on Storage and I/O Virtualization, Performance, Energy, Evaluation and Dependability (SPEED 2008)*, Salt Lake City, UT, Feb. 2008.
- [5] J.-F. Pâris and T. J. Schwarz. On the Possibility of Small, Service-Free Disk Based Storage Systems, *Proc. 3<sup>rd</sup> International Conference on Availability, Reliability and Security (ARES 2008)*, Barcelona, Spain, Mar. 2008.
- [6] E. Pinheiro, W.-D. Weber and L. A. Barroso. Failure Trends in a Large Disk Drive Population, *Proc. 5<sup>th</sup> USENIX Conference on File and Storage Technologies (FAST '07)*, San Jose, CA, Feb. 2007.
- [7] 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. 5<sup>th</sup> USENIX Conference on File and Storage Technologies (FAST '07)*, San Jose, CA, Feb. 2007.
- [8] B. Schroeder and G. A. Gibson. A Large-Scale Study of Failures in High-Performance-Computing Systems, *Proc. International Conference on Dependable Systems and Networks (DSN2006)*, Philadelphia, PA, June 2006.
- [9] T. J. E.; Schwarz, and W. A. Burkhard. Almost Complete Address Translation (ACATS) Disk Array Declustering, *Proc. 8<sup>th</sup> IEEE Symposium on Parallel and Distributed Processing (SPDP '96)*, New Orleans, LA, pp. 324–331, Oct. 1996.