Contemporary Engineering Sciences, Vol. 8, 2015, no. 35, 1679 - 1691
HIKARI Ltd, www.m-hikari.com
http://dx.doi.org/10.12988/ces.2015.59270

Computer Simulator of Failures in Super Large Data Storage

Lyudmila Ivanichkina

OOO Proekt IKS, Altufievskoe sh, 44, 127566, Moscow, Russia
MIPT, 9 Institutskiy per., 141700, Dolgoprudny, Moscow Region, Russia

Andrew Neporada

OOO Proekt IKS, Altufievskoe sh, 44, 127566, Moscow, Russia

Abstract

The present paper is concerned with simulation approach to compute data storage reliability. We propose new mathematical model accounting for explicit disk failures, latent sector errors and their detection as well as various policies of data distribution across the storage. The developed algorithm simulates the dynamics of failures and recoveries via Monte-Carlo statistical trials. The reliability is estimated by criterion of mean time to failure (MTTF), when data loss becomes irrecoverable. We present various parametrical dependencies of mean time to data loss (MTTDL) on disk failure rate and mean time to data recovery. The proposed model is compared to the simplified Markov counterpart. Statistical analysis of obtained empirical and analytical results is carried out.

Keywords: mean time to failure (MTTF), mean time to data loss (MTTDL), numerical simulation, Monte-Carlo method, discrete event modelling, distributed data storage, scrubbing procedure, checksum, disk failure, latent sector errors

1 Introduction

One of the distinctive characteristics of a super large data storage is a practical impossibility to measure its real MTTDL experimentally. The main reason comes from a very long expected lifetime of a storage that is estimated from tens and hundreds to billions of years, depending on model parameters. In addition, the experimental studies of storage MTTDL assume nonfailure-operating time, excluding hours of idle power saving mode or standard off-state. Moreover, the consumer demands that storage MTTDL should greatly exceed mean time of data retention of several decades. The experimental studies of such duration are prohibitively expensive and implausible due to onset of obsolescence of deployed technologies.

The data failures and recoveries in a storage have probabilistic nature and are not precisely determined a priory. Therefore, the experimentally obtained storage lifetime to data loss is a stochastic variable. Thus, the sought value is a mean value of this variable. Therefore, to find a mean value one needs a sequence of full-scale trials with in-situ measurements. The numerical simulation becomes the main research technique, when the full-scale experiments on real objects are either implausible or impossible. Thus, the studied object is represented via a model describing its behavior and properties with sufficient accuracy. The corresponding process of model formulation includes determination the most important elements of the simulated system, the analysis of the most significant mutual relationships and the development of the software for numerical experiments.

One of the important properties of numerical simulation model is its level of detail. The considered system can be described on one of the scales depending on aims of research. For data storages, the scale defines the diversity of scenarios, allowed by the model. Thus, the in-depth description of the distributed system can account for spread of individual drive characteristics including its remaining life. However, the higher is the demanded level of detail of the researched object, the bigger are the corresponding computational expenses of the underlining numerical algorithm. For example, the full-scale simulation of occurrence of hidden sector errors on each drive of the storage and their detection by regular scrubbing (checksum validation) procedure are exceptionally expensive in terms of allocated memory and CPU-time.

The methods of the full-scale simulation of data storage reliability gain an increasing priority and are concerned in multiple actual publications. The paper [1] estimates the reliability of disk arrays via the discrete event full-scale simulator RELI, and obtains the empirical cumulative distribution function via the Markov model computational package Sharpe [2, 3]. The authors of work [4] implement the discrete event simulator on the Python library SimPy to estimate the MTTDL of the storage with 1024 drives of 1TB capacity each under various error coding schemes (e.g. [5]). The papers [6, 7] consider the reliability of distributed data storages using the simulation software PARSEC [8]. The generalized model of reliability dependency for simulation analysis of storage MTTDL depending on various laws of statistical distributions of

failures and recoveries is presented in paper [9]. The modern software tools for reliability and performance modeling is the topic of the work [10].
This paper proposes the discrete event model to simulate the storage behavior under specified load up to the sought-for time of data loss (MTTDL). The developed model stands out from existing ones by the more accurate explicit representation of data migrating within a storage.

2 Basic mathematical model

We propose reliability model to simulate storage operations during prolonged timeframe and obtain stochastic estimation of storage lifetime up to data loss. The main elements of the model are disks and blocks. The blocks define the distributed representation of files. Each file recorded to the storage undergoes sequential partitioning into nonintersecting fix-sized blocks, processed into the error coding algorithms. The numbers of disks N_{disks} and blocks N_{blocks} are computed from the macroscopic parameters of the model, namely: the total maximum size of useful data in a storage, the quotient of storage space utilization, and error coding parameters, as well as block size and disk capacity. We consider the number of disks to be proportional and lesser to the number of blocks $N_{disks} < N_{blocks}$. One can estimate the total number of blocks via relation $N_{blocks} = \omega C/S_{Block}$, where C is the maximum allocation size of useful data, ω is the quotient of storage space utilization, and S_{Block} is a block size, defined by the storage architecture. The main parameters for the simulated disk are disk failure rate λ and recovery rate μ for the data, stored on the failed disk.
The storage reliability is achieved by redundant representation of stored data. The common variants are replication and error coding of blocks. The latter requires less disk space but more computational time. The reliability level depends on the number of replicas or parameters of error coding. One of the most common and widely uses schemes is The ReedSolomon (n, k) coding (e.g. [11]), including the variant $(6, 4)$ with four blocks of initial data and two parity blocks. This set of parameters is popular, since two parity blocks are sufficient for reliability and the redundancy rate $6/4 = 1.5$ provides saving of disk space. The further increase of parameter n reduces both the reliability and performance.
The classical analytical model based on Markov chains (e.g. [12]) assumes the independence of data loss events for different blocks of the system. In this case the preliminary estimation of storage MTTDL is described by analytical expression, depending on data block MTTDL for (n, k) error coding scheme: $MTTDL_{Storage} \approx MTTDL_{Block}/N_{blocks}$. Additionally, the value of the data block lifetime $MTTDL_{Block}$ for arbitrary n and excluded negligibly small terms has the following analytical approximation: $MTTDL_{Block} =$

$\mu^{n-k} / \left(\lambda^{n-k+1} \prod_{l=0}^{n-k} (n-l)\right)$. As one can see, the classical model estimates the reliability of a single block in the first place. At the same time, the proposed simulation model yields macroscopic estimations for a large storage.

The simulation discrete event model, comparing to analytical one, allows one to model more rich set of cases occurring in real storages. The simulation model do not use two sets of undistinguishable operational and undistinguishable failed disks. Instead, the model keeps the list of all data storing devices. The transition from one system state to another is defined not by ratio of elements in corresponding sets, but by the ensemble of events, taking place at every disk.

There are two events for a disk, namely: the failure and the disk data recovery. The events possible for a block include latent sector error within the block fragment and recovery of the given block fragment after checksum validation. The statistical distributions and failure rates, the processes of data validation and recovery are the input parameters for the model. The initial conditions for a problem of statistical modeling of reliability correspond to some distribution of block fragments across the storage disks. The program implementation of computational algorithm permits to simulate various policies of fragment placement, such as fully random distribution or distribution within a group of n disks. This explicit fragment distribution can increase the accuracy of simulation modeling by accounting data migration dynamics within a storage. The main program cycle contains pseudorandom generators, simulating the events of disk failures, bit sector errors for blocks and data recovery after failures, according to respective statistical distributions. Each iteration proceeds as follows:

1. Determine event that is nearest to the current moment of simulation time;
2. Update dependent data structures (e.g. add disk to the list of failed drives);
3. Check for irrecoverable data loss if the latest event was a failure;
4. Generate new failure or recovery event.

The sought-for variable, computed during the runtime, is a simulation time of onset of the irrecoverable data loss event.

3 Computational complexity of the model

The proposed simulation model at the considered level of detail is computationally demanding for RAM and CPU time. The numerical algorithm is implemented on data structures, mutually matching blocks and disks. Thus, each disk corresponds to the index table of blocks with fragments stored on the disk. At the same time, each block corresponds to the table of disks, holding fragments of the block. In addition, the algorithm uses a dynamical list of failed disks and damaged block fragments as well as an index data structure with future system events sorted in ascending order. The mean size of the list of blocks having fragments with irrecoverable bit sector errors depends on the ratio between the irrecoverable bit error rate and the rate of continuous checksum validation. The average size of this list can be estimated as $O(1)$ to the number of disks in the system. Therefore, the RAM volume, required for the simulation experiment, equals $O(N_{blocks})$, i.e. is linearly proportional to N_{blocks}.

The CPU time required for a computational experiment to obtain storage lifetime to data loss is proportional to simulated lifetime and, therefore, is not deterministic. To generate block distribution across disk the implemented algorithm takes $O(N_{blocks})$ arithmetic operations in average. The CPU time expenses for one main cycle iteration consist of next event processing and validation of data loss. The processing of next event includes the search for a nearest event, taking O(1) operations, the generation of a new event with an insertion to the index and an update of list of failed disks or list of damaged blocks, taking $O(N_{disks})$ operations.

The check for data loss demands a traversal of all blocks having fragments on failed disks and fragments with irrecoverable bit sector errors. The mean number of blocks requiring validation depends on a specified model parameters and can be estimated as $O(N_{blocks\ per\ disk})$, where $N_{blocks\ per\ disk}$ is a mean number of blocks with fragments contained on a specified single disk.

4 Parameters estimation for a real storage

The frequency of irrecoverable read errors is a random variable, which is calculated by means of statistical analysis of large samples from produced disk models. This stochastic property is declared in HDD data sheet provided by the manufacturer and is about one error per 10^{14} bits or one error per 11 TB of data read from disk. The proposed model considers how often such errors occur in a data storage lifetime. The frequency estimation relies on a mean volume of data read from a single HDD at the datacenter during typical timeframe.

Mean HDD access rate can be reasoned from mean HDD load rate reported in the respective device specification. Let expected drive load be about 20 per cent. In other words, disk is accessed 20 per cent of uptime and is idle for the remaining 80 per cent. If the specified [13] maximum sustained transfer rate equals 140 Mbps then the size of annually read data is

$$(140 \cdot 0.2 \cdot 3600 \cdot 24 \cdot 365)\,/1024 = 862312 \text{ GB}$$

or approximately 842 TB. Assuming that, one can approximate the value of irrecoverable bit errors frequency for one drive at 77 errors per year. Let the disk capacity be 2 TB, data fill ratio be 50 per cent, and the typical storage segment size be equal 50 MB. According to this, the mean error frequency while reading specified segment is about 0.0367 errors per year or one error within 272 years. This frequency, equal to 272 years^{-1}, is comparable to disk failure rate estimated as 23 years^{-1} for MTTF of 200000 hours. Therefore, the storage MTTDL should consider this parameter to achieve higher fidelity.
It should be noticed that this parameter was calculated for the scenario of relatively high disk load. In general, the frequency of irrecoverable read errors depends significantly on disk usage rate. The mean expected disk load, provided in product data sheet, is essentially only tentative value for estimations. The more preferable way is to measure mean disk load for the given storage, since it is a function of access rate and access patterns. For example, the archive data storage can exhibit very low disk load. Nevertheless, the obtained value of the parameter yields both the lower bound for storage MTTDL and the estimated impact of bit errors on reliability.

5 Computation Results

We studied the dependence of super large distributed data storage MTTDL on model parameters. The reference input has the following values: 0.5 PB of useful data size in a cluster, 2 TB capacity of a single disk, 300 MB size of a data block, (6,4)-scheme of error coding, 200000 hours of a disk MTTF, 4 hours of a mean recovery time for a disk data. The sequence of computational experiments yielded the storage lifetimes to data loss with one model parameter chosen to vary and other parameters fixed at reference values.
Next, we plotted the MTTDL values, obtained from both computer simulation and analytical expressions, as parametrical dependencies on disk failure rates, mean time of data recovery after failure, and size of a data block (see fig. 1-3). The shape of the curves are qualitatively conform to the results, found independently from Markov simulation models for petabyte storages (e.g. [14]). The presented plots show that analytical and simulation models mutually agree in the parametrical dependencies of MTTDL values. Some difference between corresponding mean values of both models under the same parameters can be explained by a large variance in the results of simulation modeling.
It should be noted that the reasonableness of accuracy level for the considered model depends on an applied usage of simulation results. For real data storages, the significant value of MTTDL has the order of hundreds and thousands

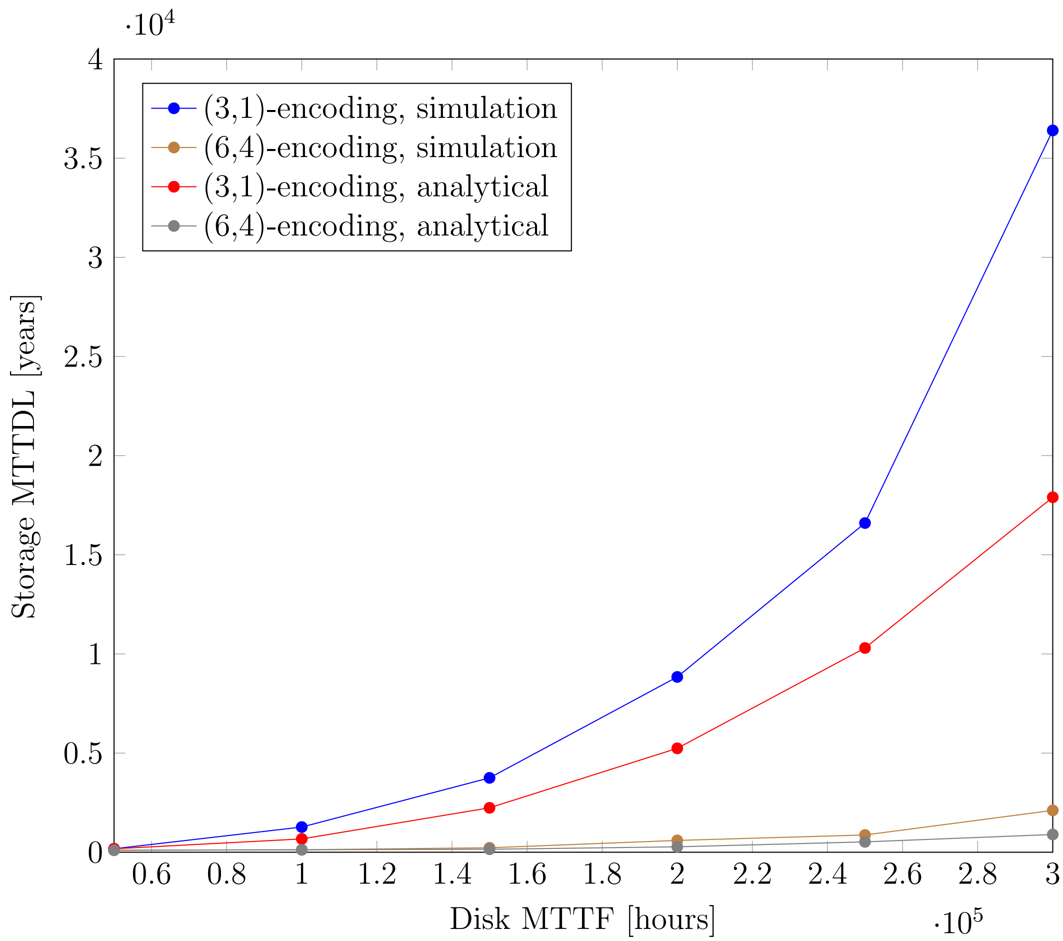

Figure 1: The dependence of data storage MTTDL on disk MTTF

years. In practice, exactly that high values are strongly sought-for, since the storage time to data loss is a random variable with both large variance and mean. The large mean indicates the decrease of probability of data loss during first 5–20 years of operation up to complete software and hardware renewal. One of the widely used methods to estimate model accuracy is a statistical test on series of experiments. The test compares an empirical distribution of real data with an assumed theoretical distribution law. A reasonable null hypothesis suggests that in the limiting case the mean of an experimental sample matches the mean value for analytical model of MTTDL. To test the hypothesis one can use Student's one-sample parametric t-test. Generally, this test expects the normal distribution of analyzed random variable. The assumption of normality provides precise probability values to accept or reject the considered hypothesis. However, t-test is also suitable for various non-normal distributions, since the corresponding t-statistic asymptotically converges to

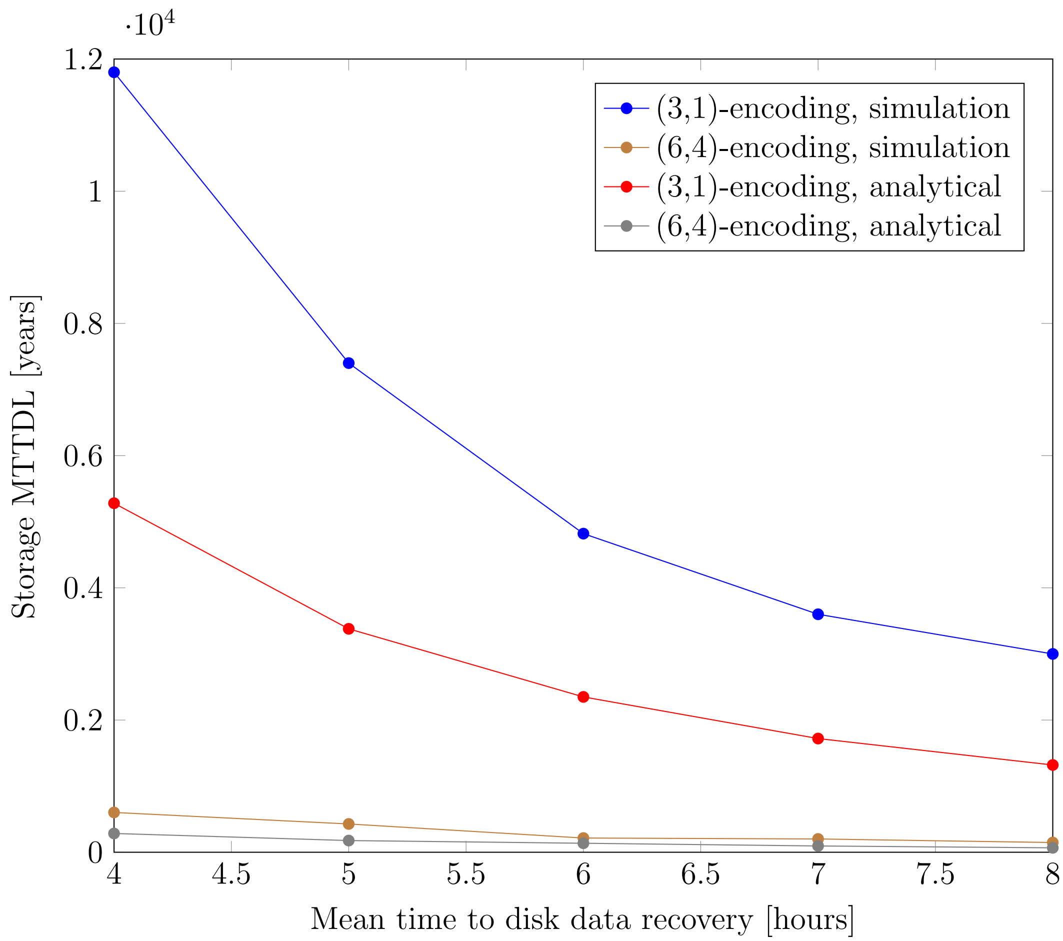

Figure 2: The dependence of data storage MTTDL on mean time to data recovery after failure

standard normal distribution $N(0, 1)$.

In the considered case, the storage lifetime to data loss has non-normal distribution; so, the Students t-test gives only indicative results. Nevertheless, these results were deemed acceptable for model accuracy estimation, since the qualitatively valid order of magnitude for MTTDL is sufficient for a storage design.

We employed one-sample t-test to validate the hypothesis of equality between the sample mean of simulated data storage lifetimes and analytically obtained MTTDL for the same model parameters. The analytical model was considered in its simplified variant without irrecoverable bit sector errors.

The numerical simulation was carried out with respect to the above-mentioned model parameters. The simulation consisted of 25 experiments and resulted in the following values of data storage lifetime to data loss (converted to years

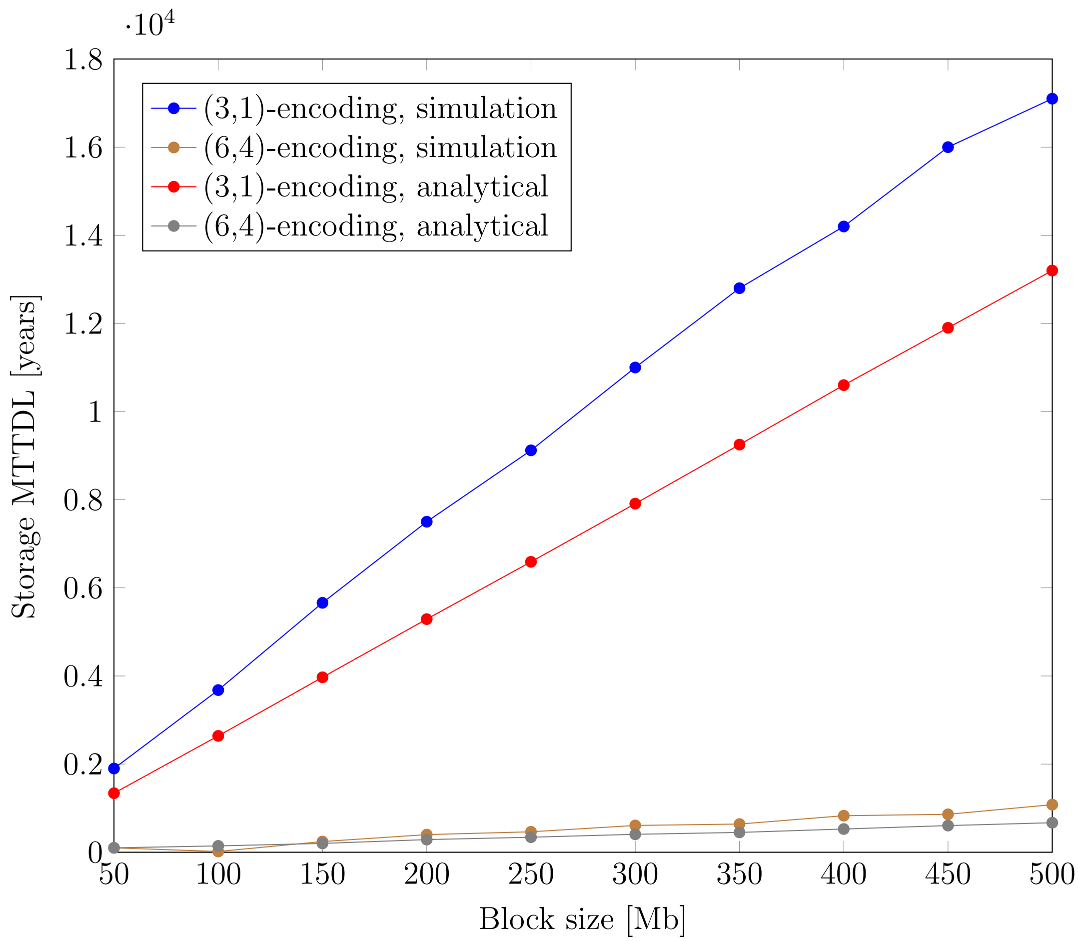

Figure 3: The dependence of data storage MTTDL on data block size

and rounded up): [467, 212, 404, 356, 46, 241, 359, 16, 235, 181, 626, 392, 381, 68, 1026, 317, 4641, 310, 879, 10, 808, 68, 394, 31, 459]. The analytical calculation of MTTDL yielded the value of 266 years.
According to null hypothesis, the sample mean is equal to MTTDL. The alternative hypothesis states the inequality of corresponding values. The test statistic is computed as $t = (\bar{x} - MTTDL)\sqrt{n}/s$, where $\bar{x} = (1/n)\sum_{i=1}^{n} x_i$ is a sample mean, $s^2 = (1/(n-1))\sum_{i=1}^{n}(x_i - \bar{x})^2$ is a sample variance, and n is a sample size. The statistic equals 1.396. The tabulated critical value of t-statistic for Students distribution with $n - 1 = 24$ degrees of freedom and confidential interval of 99 per cent equals to 2.797. Since the calculated value of criterion is lesser than the critical value, then the alternative hypothesis of mismatch between the sample mean and MTTDL is rejected with probability 99 per cent. Therefore, the Students t-test confirms that the analytically derived MTTDL describes the reliability of a super large data storage with

good accuracy. To verify an applicability of the parametric t-test the simulated sample underwent the nonparametric one-sample Wilcoxon test [15] with null hypothesis of medians equivalence. To use this test successfully it is sufficient to assume the symmetry of frequency distributions with respect to the a sample median $\tilde{x}$. The statistic W_{stat} is calculated as a minimum of sums of positive and negative ranks R_i for nonzero absolute errors $|x_i - \tilde{x}|$.

$$W_{stat} = \min\left(\sum_{i=1}^{n} Z_i{\cdot}R_i, \sum_{i=1}^{n}(1 - Z_i)\cdot R_i\right),$$

$$Z_i = \begin{cases} 0, & x_i - \tilde{x} < 0, \\ 2, & x_i - \tilde{x} > 0. \end{cases}$$

The calculated differences $|x_i - \tilde{x}|$ equal: [201, -54, 138, 90, -220, -25, 93, -250, -31, -85, 360, 126, 115, -198, 760, 51, 4375, 44, 613, -256, 542, -198, 128, -235, 193], and the ranks R_i corresponding to their absolute values are: [16, 5, 12, 7, 17, 1, 8, 19, 2, 6, 21, 10, 9, 14.5, 24, 4, 25, 3, 23, 20, 22, 14.5, 11, 18, 13]. Following this, the statistic value equals $W_{stat} = \min(117, 208)$. The tabulated critical value $T_{\alpha,n}$ for Wilcoxon distribution for the two-sided significance level $\alpha = 0.01$ and n degrees of freedom equals $T_{0.01,25} = 68$. Since $W_{stat} > T_{0.01,25}$, then the null hypothesis is accepted. It means that the difference in values of empirical and analytical medians are statistically insignificant.

In case, when the symmetry assumption is not fulfilled, to test null hypothesis of medians equality one can employ the sign test [16], trading statistical power for higher robustness and applicability.

The above-mentioned tests were also used to estimate accuracy of analytical model accounting for both irrecoverable bit sector errors and continuous checksum validation. The simulation model ran the 25 experiments with the same parameters as those that in the analytical model, additionally including irrecoverable bit error rate of 272 years^{-1} in each data block for the 20 per cent disk load, and the expected time of complete data check of one week. The numerical simulation resulted in the following values of data storage time to data loss: [67.0, 38.5, 7.5, 5.6, 2.0, 5.9, 28.0, 9.0, 39.5, 4.8, 13.2, 19.9, 9.1, 21.1, 42.1, 12.4, 15.8, 21.7, 3.7, 54.0, 21.1, 42.1, 56.1, 51.9, 26.3]. Analytically calculated MTTDL for this set of parameters equals to 15.4 years.

The value of t-test statistic for this sample equals 2.441 and is below the critical value of 2.797. According to this, an alternative hypothesis of inequality between sample mean and MTTDL is rejected. The Wilcoxon criterion for this sample is 99.5, and the threshold $T_{0.01,25} = 68$. This also leads to rejection of inequality between empirical and analytical medians. Therefore, the analytical MTTDL matches with a high probability the MTTDL for a simulation model, closely representing a real data storage.

6 Conclusion

We proposed new discrete event simulation model of disk failures in super large distributed data storage, accounting for latent bit sector errors, data validation and recovery. The model provides the statistical representation of operational properties, including time of operations up to lifetime. The developed mathematical description can be extended to simulate disks of variable failure rate and correlated failures of several disks during breakdown of a server or disconnection of a rack. The considered model also permits to investigate, in particular, various policies of data placement (random distribution of data chunks across disk or data distribution based on joining disks into RAID arrays) and other techniques to increase reliability and minimize data loss.
The one-sampled statistical tests confirmed the null-hypothesis of matching between the mean of experimental values of a storage lifetime and the analytically obtained MTTDL for the same set of model parameters. Thereby, it is demonstrated that analytical expressions for the MTTDL estimate the reliability of super large data storage with a good accuracy. The obtained results are the first stage of research on reliability of super large data storage. The proposed simulation model is intended to use in problems that are hard to solve by analytical means. In particular, we plan to extend the developed simulator to research non-exponential distributions of failures and various policies of data placement.

Acknowledgements. The article is published within applied scientific research performed with financial support of the Ministry of Education and Science of the Russian Federation. Subsidy provision agreement 14.579.21.0010. Universal identifier of the agreement is RFMEFI57914X0010.

References

[1] G.A. Gibson and D.A. Patterson, Designing Disk Arrays for High Data Reliability, *Journal of Parallel and Distributed Computing,* **17** (1993), 4 – 27. http://dx.doi.org/10.1006/jpdc.1993.1002

[2] R. A. Sahner and K. S. Trivedi, *Sharpe: Symbolic Hierarchical Automated Reliability and Performance Evaluator, Introduction and Guide for Users,* Department of Computer Science, Duke University, 1986.

[3] R. A. Sahner and K. S. Trivedi, Reliability Modeling using SHARPE, *IEEE Transactions on Reliability,* **R-36** (1987), no. 2, 186 – 193. http://dx.doi.org/10.1109/tr.1987.5222336

[4] A. Wildani, T. Schwarz, E. L. Miller and D. D. E. Long, Protecting against rare event failures in archival systems, *In Proceedings of the 17th IEEE International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems (MASCOTS),* (2009), 1 – 11. http://dx.doi.org/10.1109/mascot.2009.5366825

[5] L. Ivanichkina and A. Neporada, Mathematical methods and models of improving data storage reliability including those based on finite field theory, *Contemporary Engineering Sciences,* bf 7 (2014), no. 28, 1589 – 1602. http://dx.doi.org/10.12988/ces.2014.411236

[6] Q. Xin, T. J. E. Schwarz and 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 05),* (2005), 125 – 134. http://dx.doi.org/10.1109/mascots.2005.27

[7] Q. Xin , E. L. Miller and T. J. E. Schwarz, Evaluation of Distributed Recovery in Large-Scale Storage Systems, *Proceedings of the 13th IEEE International Symposium on High Performance Distributed Computing,* (2004), 172 – 181. http://dx.doi.org/10.1109/hpdc.2004.1323523

[8] R. A. Meyer and R. Bagrodia, PARSEC user manual, release 1.1.

[9] V. Venkatesan and I. Iliadis, A general reliability model for data storage systems, *in Proceedings of the 9th International Conference on Quantitative Evaluation of Systems (QEST),* (2012), 209 – 219. http://dx.doi.org/10.1109/qest.2012.32

[10] R. N. Calheiros, R. Ranjan, A. Beloglazov, C. A. F. De Rose and R. Buyya, Cloudsim: a toolkit for modeling and simulation of cloud computing environments and evaluation of resource provisioning algorithms, *Software: Practice and Experience,* **41** (2011), 23 – 50. http://dx.doi.org/10.1002/spe.995

[11] J. S. Plank, J. Luo, C. D. Schuman, L. Xu and Z. Wilcox-O'Hearn, A performance evaluation and examination of open-source erasure coding libraries for storage, *Proccedings of the 7th conference on File and storage technologies,* (2009), 253 – 265.

[12] D. A. Patterson, G. Gibson, R. H. Katz, A case for redundant arrays of inexpensive disks (RAID), *ACM SIGMOD Record,* **17** (1988), no. 3, 109 – 116. http://dx.doi.org/10.1145/971701.50214

[13] Constellation ES, High-capacity storage designed for seamless enterprise integration, DataSheet, Seagate.
http://www.seagate.com/docs/pdf/datasheet/disc/ds_constellation_es.pdf

[14] K. Greenan, E. Miller, T. Schwarz and D. Long, Disaster recovery codes: increasing reliability with large-stripe erasure correcting codes, *Proceedings of the 2007 ACM workshop on Storage security and survivability,* (2007), 31 – 36. http://dx.doi.org/10.1145/1314313.1314322

[15] F. Wilcoxon and R.A. Wilcox, *Some Rapid Approximate Statistical Procedures, 2 ed.,* Pearl River, Lederle Laboratories, New York, 1964.

[16] G. W. Corder, D. I. Foreman, *Nonparametric Statistics: A Step-by-Step Approach (2nd ed.),* John Wiley & Sons, 2014.

Received: September 29, 2015; Published: December 12, 2015