RZ 3711 (# 99721) 06/05/08
Computer Science 10 pages

Research Report

Reliability Assurance of RAID Storage Systems for a Wide Range of Latent Sector Errors

Ilias Iliadis and XiaoYu Hu

IBM Research GmbH
Zurich Research Laboratory
8803 Rüschlikon
Switzerland

{ili,xhu}@zurich.ibm.com

LIMITED DISTRIBUTION NOTICE

This report has been submitted for publication outside of IBM and will probably be copyrighted if accepted for publication. It has been issued as a Research Report for early dissemination of its contents. In view of the transfer of copyright to the outside publisher, its distribution outside of IBM prior to publication should be limited to peer communications and specific requests. After outside publication, requests should be filled only by reprints or legally obtained copies of the article (e.g., payment of royalties). Some reports are available at http://domino.watson.ibm.com/library/Cyberdig.nsf/home.

IBM Research
Almaden · Austin · Beijing · Delhi · Haifa · T.J. Watson · Tokyo · Zurich

Reliability Assurance of RAID Storage Systems for a Wide Range of Latent Sector Errors

Ilias Iliadis, Xiao–Yu Hu
IBM Zurich Research Laboratory
Säumerstrasse 4, 8803 Rüschlikon, Switzerland
{ili,xhu}@zurich.ibm.com

Abstract

The low-cost disk drives, which are increasingly being adopted in today's data storage systems, have higher capacity but lower reliability, which leads to more frequent rebuilds and to a higher risk of unrecoverable or latent media errors. An intra-disk redundancy scheme has been proposed to cope with such errors and enhance the reliability of RAID systems. Empirical field results recently reported in the literature, however, suggest that the extent to which unrecoverable media errors occur is higher than the data sheet specifications provided by the disk manufacturers. Our results demonstrate that the reliability improvement due to intradisk redundancy is adversely affected because of the increase in the number of unrecoverable errors. We demonstrate that, by revising the parameter choice of the intra-disk redundancy scheme, we can obtain essentially the same reliability as that of a system operating without unrecoverable sector errors. The I/O and throughput performance are evaluated by means of analysis and event-driven simulations. The effects of the spatial locality of errors and of the error-burst length distribution on the system reliability are also investigated.

1. Introduction

Data storage systems of today use hundreds of hard-disk drives (HDDs) to satisfy the ever increasing demand for more capacity. Protection against disk failures is provided by making use of a RAID (redundant array of independent disks) scheme [3, 11].

A popular RAID scheme is RAID 5, in which disks are arranged in groups (or arrays), each with one redundant disk. RAID 5 arrays can tolerate one disk failure per array. Increased reliability can be provided by a RAID 6 scheme, which requires two redundant disks per array but can tolerate up to two disk failures in an array [2, 4]. However, the increase in reliability significantly reduces the overall throughput performance of RAID 6 arrays as well as the available storage space for a fixed number of total disks in an array.

Another factor that limits the reliability of storage systems is the occurrence of unrecoverable or latent sector errors, i.e., of errors that cannot be corrected by either the standard sector-associated error-control coding (ECC) or the re-read mechanism of the HDD. Unrecoverable media errors typically result in one or more sectors becoming unreadable. This is particularly problematic when combined with disk failures. For example, if a disk fails in a RAID 5 array, the rebuild process must read all the data on the remaining disks to rebuild the lost data on a spare disk. During this phase, a media error on any of the good disks would be unrecoverable and lead to data loss because there is no way to reconstruct the lost data sectors. A similar problem occurs when two disks fail in a RAID 6 scheme. In this case, any unrecoverable sectors encountered on the good disks during the rebuild process also lead to data loss. Typical data storage installations also include a tape-based back-up or a disk-based mirrored copy at a remote location. These mechanisms can be used to reconstruct data lost because of unrecoverable errors. However, there is a significant penalty in terms of latency and throughput.

Techniques such as disk scrubbing [13, 14] and intradisk redundancy [5, 6] have been proposed to enhance the reliability of RAID systems. The established, widely used disk scrubbing scheme periodically accesses disks to detect media-related unrecoverable errors. The scrubbing process identifies unrecoverable sector errors at an early stage and attempts to correct them. Lost data are recovered using the RAID capability, and are subsequently written to a good disk location using the bad block relocation mechanism. Thus, the scrubbing effectively reduces the probability of encountering unrecoverable sector errors. On the other hand, the recently proposed, intra-disk redundancy scheme uses a further level of redundancy inside each disk, in addition to the RAID redundancy across multiple disks. It is based on an interleaved parity-check coding scheme [6], which incurs only negligible I/O performance degradation and has been developed to increase the reliability of disks in general, but especially in the presence of multiple correlated media errors on the same track or cylinder. This method introduces an additional "dimension" of redundancy inside each disk that is orthogonal to the usual RAID dimension based on redundancy across multiple disks. The RAID redundancy provides protection against disk failures, whereas the intra-disk redundancy aims to protect against media-related unrecoverable errors. Note that each of these two schemes can also be applied in conjunction with any other mechanism developed to reduce the number of unrecoverable errors and thereby improve reliability. This implies that the two schemes can also be used simultaneously. A thorough comparison of these two schemes was presented in [9]. It was demonstrated that the reliability improvement due to disk scrubbing depends on the scrubbing frequency and the workload of the system, and may not reach the reliability level achieved by the intra-disk redundancy scheme, which is insensitive to the workload. For this reason, we consider the latter scheme in the remainder of the paper.

The parameters associated with the intra-disk redundancy scheme were chosen in such a way as to ensure sufficient degrees of storage efficiency, performance and reliability [6]. It was demonstrated that, for SATA disk drives, a RAID 5 system enhanced by the intra-disk redundancy scheme achieves a similar reliability as that of a RAID 6 system. The parameter choice was based on the assumption that the unrecoverable sector error probability is the one listed in the data sheet specifications provided by the disk manufacturers. However, by studying empirical field results recently reported in [1], we demonstrate that the actual values can be orders of magnitude higher than the values previously assumed. We subsequently study the reliability of the RAID systems, in terms of the mean time to data loss (MTTDL), and find that the reliability level is adversely affected and is no longer the desired one. This observation suggests that the system parameters need to be revised so as to cover a wide range of values for the unrecoverable sector error probability.

The key contributions of this paper are the following. We explore in further detail the process followed in order to obtain new appropriate parameters. As our results demonstrate, doubling the interleaving depth and segment length of the intra-disk redundancy scheme results in essentially the same reliability as that of a system operating without unrecoverable sector errors. The I/O and throughput performance are evaluated by means of analysis and event-driven simulations. The effect of the spatial locality of errors on the system reliability is also considered. Several measures reported in [1] that are associated with the characteristics of unrecoverable sector errors are analytically obtained here based on the information regarding the distribution of the error-burst length. We also conduct a sensitivity analysis of the system reliability to the error-burst length distribution. The results obtained reveal that the reliability achieved by the intra-disk redundancy scheme is sensitive to the tail of the distribution, but not to the average burst length.

The remainder of the paper is organized as follows. Section 2 considers the unrecoverable or latent errors and investigates the extent to which sectors become erroneous. The basic intra-disk redundancy scheme developed for increasing the reliability of disks in the presence of unrecoverable errors and disk failures is briefly reviewed in Section 3. The relevant performance measures are considered in Section 4. The effect of the spatial locality of correlated media errors on the same track or cylinder of a disk is studied in Section 5. A sensitivity analysis of the system reliability to the error-burst length distribution is also conducted. In Section 6, the I/O performance is evaluated analytically. Section 7 presents numerical results demonstrating the effectiveness of the intra-disk redundancy scheme in improving the reliability of the system. An analytical investigation of the reliability and sensitivity to the various parameters is conducted. The I/O response time and throughput performance are evaluated by means of simulation in Section 8. Finally, we conclude in Section 9.

2. Unrecoverable Errors

According to data sheet specifications, the likelihood of unrecoverable errors occurring in SATA drives is ten times higher than that in SCSI/FC drives [8]. The unrecoverable bit error probability P_{bit} is estimated to be 10^{-15} for SCSI and 10^{-14} for SATA drives. For a sector size of 512 bytes (the default for nearline disks), the equivalent unrecoverable sector error probability is $P_{\text{sec}} \approx P_{\text{bit}} \times 4096$, which is 4.096×10^{-12} in the case of SCSI and 4.096×10^{-11} in the case of SATA drives. In practice, however, and based on the empirical field results recently reported in [1], this probability seems to be much higher. Measurements presented in [1, Fig. 3(b)] show that the average number of latent errors per GB can be as high as 0.0095. As a GB contains 2×10^6 sectors (i.e. 1 GB/512 B), we deduce that the probability of an unrecoverable sector error can be as high as 4.75×10^{-9} (i.e. $0.0095/2\times10^6$), which is more than two orders of magnitude higher than the data sheet specifications for SATA and SCSI/FC disk drives assumed in previous studies [5, 6]. This, in turn, suggests that the reliability of SATA drives should be studied for values of the unrecoverable sector error probability in the range $[4.096\times10^{-11}, 5\times10^{-9}]$ rather than only for the data sheet specification value of 4.096×10^{-11}. As we will see in Section 7, increasing the probability of unrecoverable sector errors in this wide range has a significant impact on the system reliability.

Table 1. Notation of system parameters.

Parameter	Definition
N	Number of disks per array group
n_G	Number of array groups in the system
C_d	Disk drive capacity
S	Sector size
ℓ	Number of sectors in a segment
m	Number of parity sectors in a segment or number of interleaves or interleaving depth
$1/\lambda$	Mean time to failure for a disk
P_{bit}	Probability of an unrecoverable bit error (data sheet specification)
P_{sec}	Probability of an unrecoverable sector error (data sheet specification)
$se^{(\text{RAID})}$	Storage efficiency of the RAID scheme
$se^{(\text{IDR})}$	Storage efficiency of the intra-disk redundancy scheme
$se^{(\text{RAID+IDR})}$	Overall storage efficiency of the system
S_d	Number of sectors in a disk
$1/\mu$	Mean time to rebuild in critical mode for a RAID 5 array
$1/\mu_1$	Mean time to rebuild in degraded mode for a RAID 6 array
$1/\mu_2$	Mean time to rebuild in critical mode for a RAID 6 array
P_s	Probability of an unrecoverable error on a tagged sector at an arbitrary time

3. Intra-Disk Redundancy Scheme

Here we briefly review the intra-disk redundancy (IDR) scheme presented in [5] and developed to increase the reliability of disks in general, but especially to cope with the adverse effect of the spatial locality of errors, such as correlated media errors on the same track or cylinder of a disk [1]. A number of n contiguous data sectors in a strip as well as m redundant sectors derived from these data sectors are grouped together, forming a segment. The redundant parity sectors are obtained using a simple XOR-based interleaved parity-check (IPC) coding scheme [6], which, for small unrecoverable sector error probabilities not exceeding 10^{-8}, is shown to be as effective as the optimum, albeit more complex, Reed–Solomon (RS) coding scheme. The traditional single-parity-check (SPC) coding scheme corresponding to $m = 1$ is also considered. The entire segment, comprising ℓ data and parity sectors, is stored contiguously on the same disk, where $\ell = n + m$. Note that this scheme addresses the issue of spatial locality of errors in that it can correct a single burst of m consecutive sector errors occurring in a segment. However, unlike the RS scheme, it in general does not have the capability of correcting any m sector errors in a segment.

The size of a segment should be chosen such that sufficient degrees of storage efficiency, performance and reliability are ensured. For practical reasons, the strip size should be a multiple of the data-segment size. In addition, the number m of parity sectors in a segment is a design parameter that can be optimized based on the desired set of operating conditions. In general, more redundancy (larger m) provides better protection against unrecoverable media errors. However, it also incurs more overhead in terms of storage space and computations required to obtain and update the parity sectors. Furthermore, for a fixed degree of storage efficiency, increasing the segment size results in an increased reliability, but also in an increased penalty on the I/O performance. Therefore, a judicious trade-off between these competing requirements needs to be made. The storage efficiency $se^{(\text{IDR})}$ of the IDR scheme is given by

$$se^{(\text{IDR})} = \frac{\ell - m}{\ell} = 1 - \frac{m}{\ell} . \qquad (1)$$

The size of a segment and the number of parity sectors in a segment are chosen to be equal to $\ell = 128$ and $m = 8$, respectively, to ensure sufficient degrees of storage efficiency, performance and reliability [6]. The storage efficiency $se^{(\text{IDR})}$ of the IDR scheme is then equal to 94%. The choice of $m = 8$ seems to be reasonable given that recent empirical data indicate that the median number of errors for disks containing one or several errors is 3 [1]. In Section 7, however, we will show that when the unrecoverable sector error probability is higher than the data sheet specification, this choice no longer provides a sufficient degree of reliability, and it therefore needs to be revised.

4. System Analysis

The notation used for the purpose of our analysis is given in Table 1. The parameters are divided into two sets, namely, the set of independent and that of dependent parameters, listed in the upper and lower part of the table, respectively.

The storage efficiency of the RAID scheme chosen is given by

$$se^{(\text{RAID})} = \frac{N-p}{N} = 1 - \frac{p}{N} , \qquad (2)$$

with

$$p = \begin{cases} 1 & \text{for a RAID 5 system} \\ 2 & \text{for a RAID 6 system.} \end{cases} \qquad (3)$$

Note that the above expressions hold for a scheme not using intra-disk redundancy. If an IDR scheme is used, the overall storage efficiency of the entire array (or system) is given by

$$\begin{aligned} se^{(\text{RAID+IDR})} &= se^{(\text{RAID})} \, se^{(\text{IDR})} \\ &= \left(1 - \frac{p}{N}\right)\left(1 - \frac{m}{\ell}\right) . \end{aligned} \qquad (4)$$

The number of sectors in a disk, S_d, is given by

$$S_d = \left\lfloor \frac{C_d}{S} \right\rfloor , \qquad (5)$$

the ratio of disk drive capacity to sector size.

In the remainder we consider SATA disk drives with the parameter values listed in Table 2. The mean time to rebuild (MTTR) a disk, $1/\mu$, is obtained as follows. The time required to rebuild a disk depends on various parameters including the drive capacity and the bandwidth that the drive provides. During a disk rebuild, we assume that the disk array continues to actively service I/O requests with 20% of the time spent on performing rebuilds. We also consider the size of the read/write requests issued to the drive to be 256 KB, the average read/write operations for single randomly chosen sectors to be 150 per second, and the average disk transfer rate to be 60 MB/s. Thus, the average seek and rotational latency for an I/O request is equal to 1/150 = 0.0067 s, its transfer time is equal to (256 KB)/(60 MB) = 0.0043 s, and, therefore, its average time required to complete is 6.7 + 4.3 = 11 ms. Consequently, the effective disk-rebuild bandwidth is (256 KB)/(11 ms) = 23.3 Mb/s, and the mean time to rebuild the 300 GB disk using 20% of the time for rebuilds is equal to (300 GB)/(0.20×23.3 Mb/s) = 17.8 h.

5. Error-Burst Length Distribution

The effect of the spatial locality of correlated media errors on the same track or cylinder of a disk is considered here. Adopting the notation used in [6], let B denote the length (in number of sectors) of a typical burst of consecutive sector errors and $\bar{B}$ the corresponding average length.

Table 2. Parameter values.

Parameter	Value
$1/\lambda$	500,000 h
C_d	300 GB
P_s	4.096×10^{-11}
N	8 (for RAID 5), 16 (for RAID 6)
$1/\mu$	17.8 h
$1/\mu_1$	17.8 h
$1/\mu_2$	17.8 h
S	512 bytes = 4096 bits

Also, let $\{b_j\}$ denote the probability density function of B, i.e., $b_j = P(B = j)$, for $j = 1, 2, \ldots$, and G_n the probability that the length of a burst is greater than or equal to n, i.e., $G_n \triangleq \sum_{j=n}^{\infty} b_j$, for $n = 1, 2, \ldots$. We now consider the following error-burst length distribution, based on actual data collected from the field for a product that is currently being shipped:

$$\begin{aligned} \mathbf{b} =& [0.9812\ 0.016\ 0.0013\ 0.0003\ 0.0003\ 0.0002\ 0.0001 \\ & 0.0001\ 0\ 0.0001\ 0\ 0.0001\ 0.0001\ 0\ 0.0001\ 0\ 0.0001]. \end{aligned} \qquad (6)$$

Then, we have bursts of at most 17 sectors with $\bar{B} = 1.029$, $\overline{B^2} = 1.18$, $G_9 = 0.0005$, and $G_{17} = 0.0001$. As the parameter m of the IPC-based redundancy scheme is chosen to be equal to 8, the probability that this scheme will not be able to correct a single burst of consecutive errors occurring in a segment is equal to G_9, i.e. 0.0005.

The spatial locality issue has been studied in [1]. For the nearline disk family C-1 shown in [1, Fig. 5(a)], the probability P_a that a sector has at least one neighbor sector error within a range of 10 KB (approx. 20 sectors) is equal to 0.09. In the Appendix it is shown that this measure can be obtained as follows using the terminology defined above: $\overline{P_a} = 1 - b_1/\bar{B}$. For the error-burst length distribution given in (6), it follows that $P_a = 0.05$, which is close to the empirical value of 0.09. Furthermore, this figure shows that for all nearline disk families, the probability P_a can be as high as 0.5. This, in turn, implies that b_1 can be significantly smaller than 0.9812, which is the value assumed in (6).

Furthermore, the results shown for the nearline disk family C-1 in [1, Fig. 6(a)] indicate that the average number $\overline{N_a}$ of neighboring errors within a range of 10 KB (approx. 20 sectors) of an existing sector in error is equal to 0.17. In the Appendix it is shown that this measure can be obtained as follows using the terminology defined above: $\overline{N_a} = \overline{B^2}/\bar{B} - 1$. For the error-burst length distribution given in (6), it follows that $\overline{N_a} = 0.15$, which is close to the empirical value of 0.17. Furthermore, this figure shows that for all nearline disk families, the average number $\overline{N_a}$ is less than one, i.e. $\overline{N_a} < 1$. In the Appendix, it is shown that $\bar{B} \le \overline{N_a} + 1$, which, in turn, implies that $1 < \bar{B} < 2$.

The system reliability, expressed in terms of the MTTDL, depends on the probability P_{seg} that a segment is in error [6]. For the correlated model, which takes the spatial locality of media errors into account, and for P_s values in the interval specified in Section 2, the probability $P_{\text{seg}}^{\text{IPC}}$ corresponding to the IPC coding scheme is given by the following expression [6, Eq. (35)]:

$$P_{\text{seg}}^{\text{IPC}} = \left[1 + \frac{(\ell - m - 1) G_{m+1} - \sum_{j=1}^{m} G_j}{\bar{B}} \right] P_{\text{s}} \, . \quad (7)$$

Next we examine the sensitivity of $P_{\text{seg}}^{\text{IPC}}$, and therefore of the reliability, to the error-burst length distribution. Note that (7) can be rewritten as follows:

$$P_{\text{seg}}^{\text{IPC}} = \frac{(\ell - m)\, G_{m+1} + \sum_{j=m+2}^{\infty} G_j}{\bar{B}} \, P_{\text{s}} \, . \quad (8)$$

We now consider two different distributions $\{b_j\}$ and $\{b_j\}'$ for B, with $1 < \bar{B} < 2$ and $1 < \bar{B}' < 2$. From (8), it follows that the ratio r' of the corresponding probabilities of a segment error is given by

$$r' \triangleq \frac{P_{\text{seg}}^{\text{IPC}\,\prime}}{P_{\text{seg}}^{\text{IPC}}} = \frac{\bar{B}}{\bar{B}'} \cdot \frac{(\ell - m)\, G'_{m+1} + \sum_{j=m+2}^{\infty} G'_j}{(\ell - m)\, G_{m+1} + \sum_{j=m+2}^{\infty} G_j} \, . \quad (9)$$

From the above, and given that $0.5 < \bar{B}/\bar{B}' < 2$, it follows that r' primarily depends on the value of the second fraction at the right-hand side of (9). Note that this value depends only on the values b_j and b'_j for $j > m$. For example, if $b'_j = h\, b_j$ for $j = m+1, m+2, \ldots$, then $r' = (\bar{B}/\bar{B}')\, h$, which is of order $O(h)$, regardless of the values of b_j and b'_j for $1 \le j \le m$.

6. I/O Performance Analysis

The two key components that make up the time required for the processing of an I/O request to a disk are the seek time and the access time [12]. The seek time depends on the current and the desired position of the disk head and is typically specified using an average value corresponding to a seek that requires the head to move half of the maximum possible movement. The access time depends on the size of the data unit requested. The processing time is determined by the type of workload (e.g., random vs. sequential I/O) and the size of the data unit. The processing time of an I/O request normalized to the seek time is expressed by the I/O equivalent metric, denoted by IOE, which was introduced in [7]. In [7] it is shown that the IOE of an I/O request containing k 4-KB chunks is given by

$$IOE = 1 + k/50 \, . \quad (10)$$

For RAID 5 arrays, writing small (e.g., 4 KB) chunks of data located randomly on the disks poses a challenge, the so-called "small-write" problem. This is because each write operation to data also requires the corresponding RAID parity to be updated. A practical way to do this is to read the old data and the old parity from the two corresponding disks, compute the new parity, and then write the new data and the new parity. Hence, each small-write request results in four I/O requests being issued. A RAID 6 array must update two parity units for each data unit being written. This leads to six I/O requests, namely, reading of the old data and two old parity units, and writing of the new data and the two new parity units. Because of the small size of the data units involved, the predominant component of the processing time for each I/O request is the seek time. Based on the above, it follows that the corresponding (normalized) time required for the processing of a small-write request for RAID 5 and RAID 6, expressed through the IOE metric, is given by

$$IOE(n) = \begin{cases} 4\,(1 + n/400) & \text{for RAID 5} \\ 6\,(1 + n/400) & \text{for RAID 6,} \end{cases} \quad (11)$$

where n is the I/O request size expressed in sectors.

Using the intra-disk redundancy scheme requires that the intra-disk parity must also be updated whenever a data unit is written. This imposes some constraints on the design of IDR schemes. For a long write, it is natural to directly compute the new intra-disk parity from the new data, and write it along with the data to the disk, resulting in large I/O request lengths and thus longer access times. For a small write, a practical solution is to read the old data and the corresponding old intra-disk parity as part of a single I/O request. Then the new data and the new intra-disk parity are computed and subsequently written back to the disk by a single I/O request. The size of the requested data increases, thereby increasing the access time. However, for small writes and an appropriately designed IDR scheme, the processing time is still dominated by the seek time. In [6] it is shown that the average length of a single-sector write request when the IPC scheme is used is given by

$$\bar{n} = \begin{cases} 1 + \dfrac{\ell^2}{4(\ell - m)} & \text{for } \ell/m \text{ even} \\ 1 + \dfrac{\ell + m}{4} & \text{for } \ell/m \text{ odd.} \end{cases} \quad (12)$$

From (11), it now follows that the corresponding IOE metrics for RAID 5 and RAID 6 are given by $IOE(\bar{n})$, i.e.,

$$IOE = \begin{cases} 4\,(1 + \bar{n}/400) & \text{for RAID 5} \\ 6\,(1 + \bar{n}/400) & \text{for RAID 6.} \end{cases} \quad (13)$$

Equations (12) and (13) imply that the larger the segment size and the interleaving depth, the higher the IOE

Table 3. I/O Equivalent for Writes.

Request Length	RAID Scheme	IOE	Relative Difference
Small write (1 Sector)	RAID 5	4.01	0.0 %
	IPC (128,8)	4.351	8.5 %
	IPC (256,16)	4.692	17.0 %
	RAID 6	6.015	50.0 %
Long write (480 Sectors)	RAID 5	8.8	0.0 %
	IPC (128,8)	9.12	3.6 %
	IPC (256,16)	9.12	3.6 %
	RAID 6	13.2	50.0 %

metric. Let us now consider an IPC scheme with a segment length of 128 sectors, using 8 redundant sectors for every 120 data sectors. This corresponds to $\ell = 128$, $m = 8$, and is denoted by IPC (128,8). The corresponding IOE metrics for RAID 5 + IPC and RAID 6 are obtained from (13) and listed in Table 3. In the case of no-coding, $\bar{n}$ is equal to 1, whereas in the case of IPC coding, $\bar{n}$ is derived from (12) and is equal to 35.13. It follows that the processing time is dominated by the seek time and that the introduction of the IPC scheme causes the processing time for a single-sector I/O request to increase by approx. 9%. Let us now consider an IPC scheme with a segment length of 256 sectors, using 16 redundant sectors for every 240 data sectors. This corresponds to $\ell = 256$, $m = 16$, and is denoted by IPC (256,16). Note that the storage efficiency of IPC (256,16) is 94%, which is the same as that of IPC (128,8). In this case $\bar{n}$ is equal to 69.26. It follows that the processing time is dominated by the seek time and that the introduction of the IPC scheme causes the processing time for a single-sector I/O request to increase by approx. 17%. In contrast, for a long request of 480 sectors, the increase is much less, namely only about 4% for both IPC schemes considered, as shown in Table 3. This is because the corresponding request length including the intra-disk parity sectors is 512 sectors in both cases. The corresponding IOE for no-coding and IPC coding are derived from (13) by setting $\bar{n} = 480$ and 512, respectively. From the results shown in Table 3, we also deduce that a plain RAID 6 system has an I/O performance penalty of 50% compared with a plain RAID 5 system.

7. Reliability Results

Here we assess the reliability of the various schemes considered through illustrative examples. The reliability of a RAID system is assessed in terms of the MTTDL, which clearly depends on the size of the system. It turns out that the MTTDL scales with the inverse of the system size. For example, increasing the system size by a given factor will result in an MTTDL decrease by the same factor. Consequently, for the purpose of studying the behavior of the various schemes, the choice of the system size is not essential. Also, the conclusions drawn regarding the performance comparison are independent of the system size chosen. We proceed by considering an installed base of systems using SATA disk drives and storing 10 PB of user data.

From (2), (3), and (4), it follows that the storage efficiency of the entire system is independent of the RAID configuration if the arrays in a RAID 6 system are twice the size of those in a RAID 5 system. For a RAID 5 system with $N = 8$, when no IDR is used, the required number of arrays to store the user data is equal to 4762 (i.e. 10 PB/(7×300 GB)), whereas for a RAID 6 system with $N = 16$, it is equal to 2381 (i.e. 10 PB/(14×300 GB)). The corresponding storage efficiency is equal to 7/8, i.e. 0.875. For the IPC (128,8) and IPC (256,16) IDR schemes, the intra-disk storage efficiency is obtained from (1) and is equal to 0.94. Furthermore, the required number of arrays for a RAID 5 configuration is obtained as the ratio of 4762 to the intra-disk storage efficiency and is equal to 5080. Similarly, for a RAID 6 configuration, the required number of arrays is equal to 2540. The overall storage efficiency is obtained by (4) and is equal to 0.82. Note that the cost of the system is proportional to the number of arrays required and, therefore, inversely proportional to the storage efficiency.

The combined effects of disk and unrecoverable failures can be seen in Figure 1 as a function of the unrecoverable sector error probability. The system reliability is assessed in terms of the MTTDL, which is analytically obtained by the closed-form expressions (Equations (37), (45) and (52)) derived in [6]. The left vertical dashed line indicates the SATA drive specification for unrecoverable sector errors. Note that for small sector error probabilities, the MTTDL remains unaffected because data is lost owing to a disk rather than an unrecoverable failure. In particular, the MTTDL of a RAID 6 system is three orders of magnitude higher than that of a RAID 5 system. However, as the sector error probability increases, the probability of an unrecoverable failure in the critical mode P_{uf} also increases and therefore the MTTDL decreases. This decrease ends when the sector error probability is such that the corresponding P_{uf} is extremely high, i.e., close to one. In this case the rebuild process in critical mode cannot be successfully completed because of an unrecoverable failure. Consequently, the MTTDL is the mean time until the system (i.e. any of the disk arrays) enters the critical mode. In a RAID 5 system, this occurs when the first disk fails after an expected time of $1/(n_G N \lambda)$. In a RAID 6 system, this occurs when a second disk fails while the system is in the degraded mode. Note that this corresponds to the MTTDL of a RAID 5 system without unrecoverable sector errors. This also explains why the RAID 6 curves become flat at about the height of

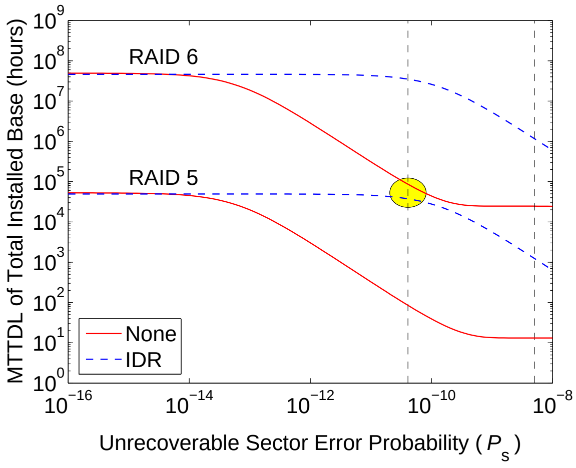

Figure 1. RAID 5 vs. RAID 6 systems under correlated unrecoverable sector errors ($\ell = 128$, $m = 8$).

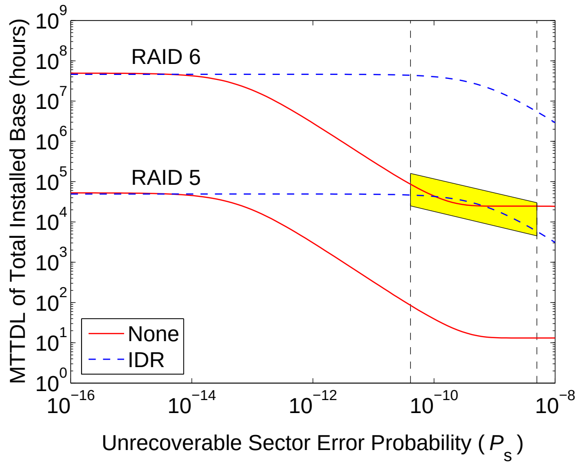

Figure 2. RAID 5 vs. RAID 6 systems under correlated unrecoverable sector errors ($\ell = 256$, $m = 16$).

a RAID 5 system, as can be seen in Figure 1. The MTTDL for RAID 6 is slightly lower than that for RAID 5 because the arrays in a RAID 6 system are larger than those in a RAID 5 system.

The results shown in Figure 1 along the left vertical dashed line reveal that in the practical case of SATA-drive unrecoverable sector errors, the MTTDL is reduced by more than two orders of magnitude. The IPC-based IDR scheme, however, improves the MTTDL by more than two orders of magnitude, therefore eliminating the negative impact of the unrecoverable sector errors.

Both the plain RAID 6 and the RAID 5 + IDR systems improve the reliability over the plain RAID 5 system, with the respective gains shown in Figure 1. Note that in the case of SATA drives the resulting MTTDLs for these two systems are of the same order (indicated by the ellipse). Therefore, the RAID 5 + IDR system is an attractive alternative to a RAID 6 system, in particular because its I/O performance is better than that of a RAID 6 system, as we shall see in Section 8. However, for larger unrecoverable sector error probabilities, this no longer holds.

The interval $[4.096\times10^{-11}, 5\times10^{-9}]$ of practical importance for P_s is indicated in Figure 1 between the two vertical dashed lines. As P_s increases in this interval, the MTTDL corresponding to the plain RAID 6 system remains practically constant. In contrast, the MTTDL corresponding to the RAID 5 system enhanced by the intra-disk redundancy scheme is reduced, despite the fact this scheme can correct burst of errors in 99.95% of the cases. This percentage corresponds to the probability that the length of a burst does not exceed 8. Therefore, for larger unrecoverable sector error probabilities, the reliability improvement offered by the IPC-based IDR scheme with an interleaving depth of 8 does not reach the reliability level achieved by RAID 6. Consequently, to further improve the MTTDL achieved by the IPC-based IDR scheme, the interleaving depth should be increased. Furthermore, to maintain the storage efficiency, the segment length should also be proportionally increased. We therefore consider a system in which both the segment size and the interleaving depth are twice as long, i.e. $\ell = 256$ and $m = 16$, such that the storage efficiency and the required number of arrays remain the same. The IPC-based IDR scheme has now the capability of correcting any burst of up to 16 errors and therefore the MTTDL improves, as can be seen in Figure 2. Furthermore, the resulting MTTDL is of the same order as the MTTDL of the RAID 6 system (indicated by the shaded area in the figure). In fact, for unrecoverable sector error probabilities in the interval $[10^{-10}, 7\times10^{-10}]$, the reliability level achieved by the IDR scheme is higher than that of the RAID 6 system.

As can be seen in Figures 1 and 2, in the range of sector error probabilities of interest, the MTTDL increases as the interleaving depth m increases. This is to be expected because the larger the m the higher the likelihood that a burst of errors can be corrected. In contrast, the MTTDL is practically insensitive to the segment length ℓ because, regardless of the segment length, an unrecoverable failure within a segment is essentially caused by a single burst of errors. A judicious selection of ℓ can be made by considering that increasing ℓ results in an increased storage efficiency (i.e. reduced cost), but also in an increased penalty on the I/O performance, according to (4), (12) and (13). In the case

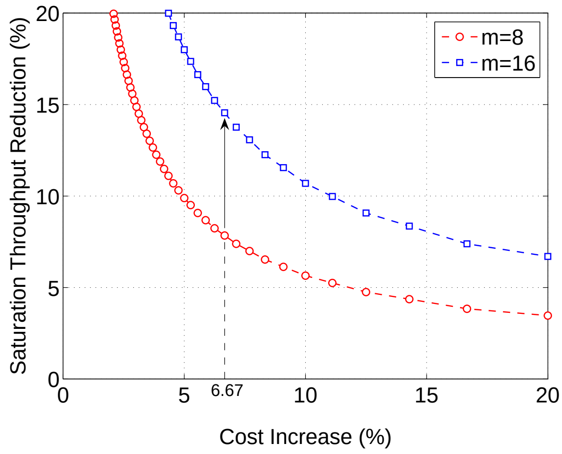

Figure 3. Saturation throughput reduction as a function of cost for $m = 8, 16$ (small writes).

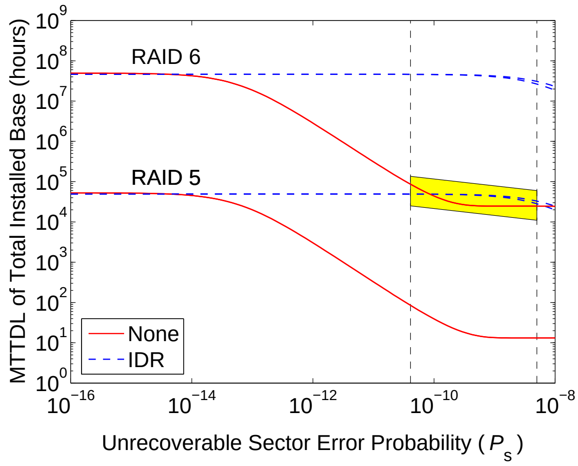

Figure 4. MTTDL for various burst-length error distributions ($\ell = 256$, $m = 16$).

of small writes, and because the saturation throughput of a RAID system is inversely proportional to the IOE metric, the saturation throughput reduction, r_d, due to the IDR scheme is given by

$$r_d = \frac{IOE(1)^{-1} - IOE(\bar{n})^{-1}}{IOE(1)^{-1}} = \frac{IOE(\bar{n}) - IOE(1)}{IOE(\bar{n})} \ , \tag{14}$$

where $\bar{n}$ is given by (12).

From the above, we obtain the tradeoff between the reduction of saturation throughput and the cost in terms of the number of disks deployed, as shown in Figure 3, with the values for these two measures obtained by considering various ℓ that are multiples of m. The vertical dashed line corresponds to $\ell/m = 16$, which implies a cost increase of $m/(\ell - m) = 6.67\%$. The vertical arrow shows the additional reduction of the saturation throughput from 7.8% to 14.5% when we change from the IPC (128,8) to the IPC (256,16) scheme.

We now consider the IPC redundancy scheme employed in conjunction with RAID 5 and RAID 6 systems, with $\ell = 256$ and $m = 16$, and examine the sensitivity of the reliability to the error-burst length distribution. First, we consider that the burst length distribution specified in (6) slightly changes, with the probability of encountering a burst of 17 consecutive sector errors becoming ten times smaller, from $b_{17} = 10^{-4}$ to $b'_{17} = 10^{-5}$, and the probability of encountering a burst of 15 consecutive sector errors almost doubling, from $b_{15} = 10^{-4}$ to $b'_{15} = 1.9{\times}10^{-4}$. Note that for the new distribution, the values corresponding to the error-related measures considered in Section 5 remain practically unchanged, i.e., $\bar{B}' = 1.029$, $P'_a = 0.05$, and $\overline{N_a}' = 0.147$.

The corresponding MTTDLs for RAID 5 and RAID 6 systems are shown in Figure 4. Compared with Figure 2, in the interval of interest, the MTTDL corresponding to the RAID 5 system enhanced by the IDR scheme is significantly higher, despite the fact that the above-mentioned error-related measures practically have not changed. In particular, for $P_s = 5{\times}10^{-9}$ the MTTDL increases along the right vertical dashed line by an order of magnitude, from 5818 to 28,232 hours. This behavior is explained as follows. From (9), and given that $\bar{B}' \simeq \bar{B}$, $G_{17} = 10^{-4}$, and $G'_{17} = 10^{-5}$, it follows that the probability of a segment error is reduced by one order of magnitude, which in turn results in an increase of the MTTDL. In particular, for unrecoverable sector error probabilities in the interval of interest that are larger than 10^{-10}, we observe that the reliability level achieved by the RAID 5 system enhanced by the IDR scheme is higher than that of the RAID 6 system (indicated by the shaded area in the figure) and essentially the same as that of a system operating without unrecoverable sector errors. Note also that the MTTDLs corresponding to the plain RAID 5 and RAID 6 systems remain practically unchanged.

In Section 5 it was mentioned that for some disk families, the probability of having a single sector in error may not be very high. We therefore consider a modified error-burst length distribution with a significant likelihood of having two and three consecutive sectors in error. We obtain this distribution by changing the values of the first three elements of the above distribution as follows: from $\{b'_1, b'_2, b'_3\} = \{0.9812, 0.016, 0.0013\}$ to $\{b''_1, b''_2, b''_3\} = \{0.5812, 0.316, 0.1013\}$. Note that G''_{17} remains the same, i.e. $G''_{17} = G'_{17} = 10^{-5}$, but the above-mentioned error-related measures change significantly, with the new values being $\bar{B}'' = 1.529$, $P_a'' = 0.62$, and $\overline{N_a}'' = 0.88$. Nevertheless, as shown in Figure 4, the MTTDL changes only for the IDR enhanced systems and only slightly. In particular, for $P_s = 5\times10^{-9}$ the MTTDL for the RAID 5 + IDR system increases along the right vertical dashed line slightly from 28,232 to 32,832 hours. This behavior is due to the fact that the resulting probability of a segment error is of the same order because it is further reduced, according to (9), by a factor of only $\bar{B}''/\bar{B}' = 1.529/1.029 = 1.486$.

8. Performance Results

Here we study the performance impact of the RAID and IDR schemes on the response time and the saturation throughput of the corresponding systems by using event-driven simulations. Most modern RAID controllers have a large battery-backed cache that boosts the overall system performance by reducing the I/O requests to the disks and performing aggressive read-ahead and write-behind. The response time of an array as experienced by the end user can be dramatically shortened by increasing the size of the array cache and selecting the replacement strategy based on the characteristics of workloads. As our main interest in the simulation is the difference in performance of the RAID schemes considered, rather than caching mechanisms or characteristics of workloads, we start measuring the response time of requests after caching, i.e., from the instant when they are sent to the disks. Therefore, the saturation throughput measures the maximum throughput between the front-end (cache) and the back-end (disk array), assuming sufficient bandwidth in between. The higher the saturation throughput, the better the performance of the underlying RAID mechanism. We use the lightweight event-driven simulator developed in [6] to obtain the results. We assume a first-come first-served (FCFS) scheduling policy for serving the I/O requests at each disk. Actually, we have tested several other disk-scheduling policies such as SSTF, LOOK, and C-LOOK, and found that the scheduling policy does not change the relative performance of the schemes considered.

We obtain the mean response times of a RAID 5 array consisting of 8 SATA disks and a RAID 6 array consisting of 16 SATA disks as a function of the I/O requests per disk per second. We also evaluate the performance of the RAID schemes enhanced by the IDR schemes. For the IDR schemes, we employ the IPC (128,8) and the IPC (256,16) schemes. We consider the small-read/write scenario and use synthetic workloads generating aligned 4-KB-small I/O requests with uniformly distributed logical block addresses. The ratio of read to write is set to be 1:2, i.e., there are 33.33% reads and 66.67% writes. The request inter-arrival times are assumed to be exponentially distributed.

Figure 5 shows that the increase observed in the response times when RAID 5 is enhanced by the IPC-based IDR scheme is minor for mean arrival rates of less than 25 I/O requests per second. However, the increase in the response times for RAID 6 is significant. Furthermore, the saturation throughput for RAID 6 is about 66% of the saturation throughput

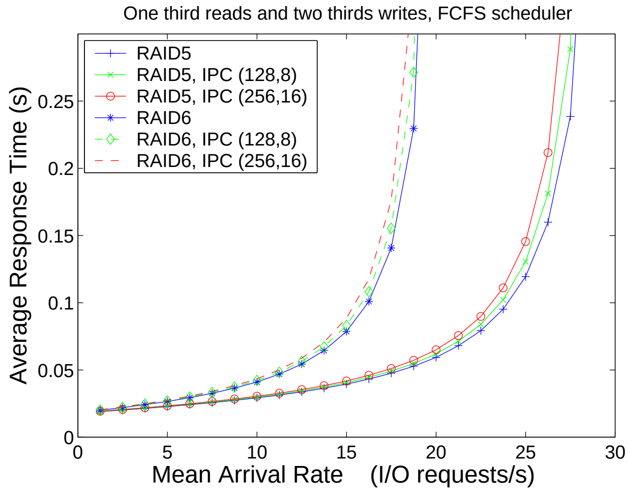

Figure 5. Response time of various RAID systems (synthetic workload, small read/write requests, one third reads, two thirds writes).

Also the saturation throughput for RAID 5 is 30.25 I/O requests per disk per second, whereas for RAID 5 enhanced by the IPC-based IDR scheme it is 29.75 I/O requests per disk per second. This represents a minor, 2%, degradation in saturation throughput due to the IPC-based IDR scheme. However, for RAID 6 the saturation throughput is 20.16 I/O requests per disk per second, which represents a substantial degradation of 33%.

9. Conclusion

Owing to increasing disk capacities and the adoption of low-cost disks in modern data storage systems, unrecoverable or latent media errors are becoming a significant cause of user data loss. An intradisk redundancy scheme developed to enhance the reliability of RAID systems in the presence of unrecoverable errors was considered. Based on empirical field results, we have shown that the extent to which unrecoverable sector errors occur is substantially higher than the data sheet specifications provided by the disk manufacturers. Our results demonstrate that the reliability level achieved by this intradisk redundancy scheme is adversely affected owing to the increased number of unrecoverable errors in this wide range. The effect of the spatial locality of errors on the system reliability is also considered. A sensitivity analysis to the error-burst length distribution reveals that the reliability achieved by the intra-disk redundancy scheme is sensitive to the tail of the distribution. The I/O and throughput performance of the RAID 5 and RAID 6 systems enhanced by the intra-disk redundancy scheme was evaluated by means of analysis and simulations. The trade-off between I/O performance and cost, in terms of capacity, was also assessed. For the SATA drives considered, the results obtained demonstrate that by appropriately revising the parameters of the intra-disk redundancy scheme, the reliability can be improved substantially over a wide range of sector error probabilities. This scheme offers the maximum possible improvement on reliability, with only a minimal penalty on the I/O performance. Therefore a RAID 5 system enhanced by an intra-disk redundancy scheme is an attractive alternative to a RAID 6 system, as its reliability is similar to and its I/O performance better than that of a RAID 6 system.

APPENDIX
NEIGHBORING SECTOR ERRORS

Let us consider a randomly chosen sector in error, SEC, and let us denote its position by α. The length B_s of the corresponding burst is distributed as follows: $P(B_s = j) = j\, b_j / \bar{B}$ for $j = 1, 2, \dots$ [10]. Let also N_a denote the number of additional unrecoverable sector errors in the interval $[\alpha - 20, \alpha + 20]$ around SEC. Assuming that the burst lengths do not exceed 20, i.e. $B_s \le 20$, the number of other sector errors within a radius of 20 sectors (i.e. 10 KB) of SEC is equal to the burst length reduced by one, i.e. $N_a = B_s - 1$.

As the likelihood that a second burst occurs within the specified interval is negligible, the probability P_a that SEC has at least one neighbor sector in error within the specified interval is therefore equal to the probability of the burst length exceeding one. From the above it follows that

$$P_a = P(B_s > 1) = 1 - P(B_s = 1) = 1 - \frac{b_1}{\bar{B}} \; . \qquad (15)$$

Also, the average number of neighboring errors $\overline{N_a}$ is given by

$$\begin{aligned} \overline{N_a} &= \sum_{i=0} i\, P(N_a = i) = \sum_{i=0} i\, P(B_s = i+1) \\ &= \sum_{i=0} i\, \frac{(i+1)\, b_{i+1}}{\bar{B}} = \sum_{i=0} \frac{[(i+1)^2 - (i+1)]\, b_{i+1}}{\bar{B}} \\ &= \sum_{j=1} \frac{j^2\, b_j}{\bar{B}} - \sum_{j=1} \frac{j\, b_j}{\bar{B}} = \frac{\overline{B^2}}{\bar{B}} - \frac{\bar{B}}{\bar{B}} = \frac{\overline{B^2}}{\bar{B}} - 1 \; . \end{aligned} \qquad (16)$$

Note also that the following inequalities hold: $E[(B - \bar{B})^2] \ge 0 \Rightarrow E(B^2) \ge \bar{B}^2 \Rightarrow \frac{\overline{B^2}}{\bar{B}} \ge \bar{B}$. This, together with (16), implies that $\bar{B} \le \overline{N_a} + 1$.

References

[1] L. N. Bairavasundaram, G. R. Goodson, S. Pasupathy, and J. Schindler. An analysis of latent sector errors in disk drives. *ACM SIGMETRICS Performance Evaluation Review*, 35(1):289–300, June 2007. (*Proc. ACM SIGMETRICS 2007*, San Diego, CA).

[2] M. Blaum, J. Brady, J. Bruck, and J. Mennon. EVENODD: An efficient scheme for tolerating double disk failures in RAID architectures. *IEEE Trans. Comput.*, 44(2):192–202, Feb. 1995.

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

[4] P. Corbett, R. English, A. Goel, T. Grcanac, S. Kleiman, J. Leong, and S. Sankar. Row-diagonal parity for double disk failure correction. In *Proceedings of the 3rd USENIX Conference on File and Storage Technologies (FAST)* (San Francisco, CA), pages 1–14, Mar.-Apr. 2004.

[5] A. Dholakia, E. Eleftheriou, X.-Y. Hu, I. Iliadis, J. Menon, and K. Rao. Analysis of a new intra-disk redundancy scheme for high-reliability RAID storage systems in the presence of unrecoverable errors. *ACM SIGMETRICS Performance Evaluation Review*, 34(1):373–374, June 2006. (*Proc. ACM SIGMETRICS 2006/Performance 2006*, Saint Malo, France).

[6] A. Dholakia, E. Eleftheriou, X.-Y. Hu, I. Iliadis, J. Menon, and K. Rao. A new intra-disk redundancy scheme for high-reliability RAID storage systems in the presence of unrecoverable errors. *ACM Trans. Storage*, 4(1), 2008.

[7] J. L. Hafner, V. Deenadhayalan, T. Kanungo, and K. Rao. Performance metrics for erasure codes in storage systems. IBM Research Report, RJ 10321, Aug. 2004.

[8] Hitachi Global Storage Technologies, Hitachi Disk Drive Product Datasheets. http://www.hitachigst.com/. 2007.

[9] I. Iliadis, R. Haas, X.-Y. Hu, and E. Eleftheriou. Disk scrubbing versus intra-disk redundancy for high-reliability RAID storage systems. In *Proceedings of the ACM SIGMETRICS Conference on Measurement and Modeling of Computer Systems* (Annapolis, MD), June 2008.

[10] L. Kleinrock. *Queueing Systems, Volume 1: Theory*. Wiley, New York, 1975.

[11] D. A. Patterson, G. Gibson, and R. H. Katz. A case for redundant arrays of inexpensive disks (RAID). In *Proceedings of the ACM SIGMOD International Conference on Management of Data* (Chicago, IL), pages 109–116, June 1988.

[12] C. Ruemmler and J. Wilkes. An introduction to disk drive modeling. *IEEE Computer*, 27(3):17–28, Mar. 1994.

[13] D. C. Sawyer. Dependability analysis of parallel systems using a simulation-based approach. *NASA-CR-195762*, Feb. 1994.

[14] T. J. E. Schwarz, Q. Xin, E. L. Miller, D. D. E. Long, A. Hospodor, and S. Ng. Disk scrubbing in large archival storage systems. In *Proceedings of the 12th Annual International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunications Systems (MASCOTS)* (Volendam, The Netherlands), pages 409–418, Oct. 2004.