RELIABILITY ANALYSIS OF A SOFTWARE REDUNDANCY SYSTEM

Anant P. Godbole

Michigan Technological University

Nicholas W. Locantore, Jr.

University of North Carolina, Chapel Hill

Gomathi Sadhasivan

Michigan Technological University

Abstract: The Stein-Chen method of Poisson approximation is used to estimate the reliability of a software system consisting of l independent backups of a portfolio of n files; we would like each file to have at least one readable copy, given information about the numbers of erased files within each backup "column". Exact formulas for the system reliability are also obtained.

Key Words: Backup Systems, Coupling, Poisson Approximation, Redundancy, Stein-Chen Method, System Reliability.

1. INTRODUCTION

The motivation for this problem stems from a discussion of the different backup strategies used for software or information file systems. Available tape backup software, manuals [5] etc. prescribe certain backup plans devised on the basis of empirical evidence, but do not give a formal justification for these plans. With a tape backup strategy, one determines how often to backup, how many tapes to use and what to backup. In this paper, we consider a software redundancy system with a specific backup plan and explore the mathematics, in a reliability context, behind the number of backups that are deemed necessary.

The system under consideration consists of a portfolio of n files that each have l independent backups; throughout this paper, the original file will be considered to be the "first backup". We assume that backups are made in the same order as in the original dossier, and suppose that for each of the l backups, k out of n files get erased, with the positions of these unreadable files being determined at random; k is assumed to be known throughout. Such a model is considered in [5]. We define the reliability of the above redundancy system as the probability that each file has at least one readable copy. One of the questions that one may wish to answer is, "How many backups, l, must we take so that we are $100(1-\alpha)\%$ sure that no file is irretrievable?" Since, under our assumptions, the probability that any copy is unreadable equals k/n, independence implies that the probability that a file has no readable copies equals $(k/n)^l$. If k is small compared to n, it is reasonable to suppose that the probability of having at least one readable backup for each of the n files is approximately

$$\bar{R} = P(X=0) = \left(1 - \left(\frac{k}{n}\right)^l\right)^n,$$

where X equals the number of unsalvageable files in the redundant backup system.

We are interested in finding the minimum l so that the unreliability of the system is less than a small value α, i.e., so that

$$1 - \alpha \leq \left(1 - \left(\frac{k}{n}\right)^l\right)^n$$

or

$$\ln(1-\alpha) \leq n \cdot \ln\left(1-\left(\frac{k}{n}\right)^l\right) \tag{1}$$

holds. Since $\ln(1+x) \approx x$ if x is small, (1) is equivalent to searching for the minimal l such that

$$\alpha \geq n\left(\frac{k}{n}\right)^l \tag{2}$$

holds. In other words

$$l \;\geq\; \frac{\ln\left(\frac{\alpha}{n}\right)}{\ln\left(\frac{k}{n}\right)} \tag{3}$$

provides an *approximate but rather crude* answer to our question.

The main problem with the above solution may be best understood by reformulating our reliability problem as follows: Given an $n \times l$ matrix of zeros and ones for which each column contains k ones in randomly selected positions, how large does l have to be so that the probability that no row contains all ones is larger than $1-\alpha$? In general, what can be said about the probability distribution of X, the number of rows that contain all ones? Clearly, $X = \sum_{j=1}^{n} I_j$ where $I_j = 1$ if the j^{th} row contains no readable version of file j (and $I_j = 0$ otherwise). Under this set up, it is evident that the I_js are highly dependent on each other (I_i and I_j are dependent for all $i \neq j$), and that this global dependence provides the main impediment to the validity of the approximate heuristic argument presented in the previous paragraph.

In Section 2 we use elementary techniques (such as the inclusion-exclusion principle) to provide an exact formula for the reliability of our system. Our exact expression is generalized to the case where there are k_i unreadable files on the i^{th} "day" ($1 \leq i \leq l$). These expressions involve heavy computational effort when the size of the system gets large, however, and in Section 3 we use the celebrated Stein-Chen method to obtain *bounds* on the system reliability. The approximate reliability of the general system (with different numbers of unreadable files for each of the backups) is also studied in detail in Section 3. The error bounds for these approximations are easy to compute and good estimates are provided for the system reliability when n is large and k is small compared to n. For uses of the Stein-Chen method in a reliability context see Chryssaphinou and Papastavridis

[4], Godbole [6], Godbole et al. [7], Koutras et al. [8], Koutras and Papastavridis [9], Papastavridis [10], and Roos [11]; the reader is referred to Barbour et al. [3] for the theoretical aspects of the coupling approach to Stein-Chen approximation, where a problem similar to ours is studied in the context of a matrix occupancy model.

Before proceeding, we consider two elementary models which are "randomized" or "parametric" simplifications of the general system introduced above. Consider n files that each have l backups with the probability of erasure being p for any of the $n \cdot l$ files. The system reliability, defined as the probability that all the files have a readable backup, can be computed exactly as $(1-p^l)^n$, or, approximately, as $e^{-\lambda}$ where $\lambda = np^l$, with the error bound, using the main result of Barbour and Hall [2], being given by

$$0 \leq e^{-\lambda} - P(X=0) \leq \min\left(\frac{np^{2l}}{2}, p^l\right).$$

In what follows, $\mathrm{Po}(\lambda)$ will denote the Poisson random variable with mean λ. The *total variation* distance $d_{TV}(X, \mathrm{Po}(\lambda))$ between the distribution of X and $\mathrm{Po}(\lambda)$ is defined as

$$d_{TV}(X, \mathrm{Po}(\lambda)) = \sup_{A \subseteq Z^+} |P(X \in A) - P(\mathrm{Po}(\lambda) \in A)|.$$

Barbour and Hall ([2]) proved, more generally, that in the i.i.d. case,

$$d_{TV}(X, \mathrm{Po}(\lambda)) \leq \min(np^{2l}, p^l)$$

Considering the non-identically distributed case where p_i is the probability of erasure of the i^{th} file, the exact formula for the reliability of the system equals $\prod_{i=1}^{n}(1-p_i^l)$; this can be approximated by $e^{-\lambda}$, where $\lambda = \sum_{i=1}^{n} p_i^l$ and the error bound is computed, using Barbour and Hall's [2] result, as

$$0 \leq e^{-\lambda} - P(X=0) \leq \min\left(p_{\max}^l, \frac{p_{\max}^l \cdot \sum_{i=1}^{n} p_i^l}{2}\right)$$

where $p_{\max}$ is the largest erasure probability. Notice how these parametric models involve independent indicators I_j, as opposed to the globally dependent indicators we consider throughout the rest of this paper.

2. EXACT FORMULÆ

In this section we derive exact formulæ for the two cases mentioned in Section 1. Case 1 addresses the case where $k_i = k$ for all $i = 1, 2, \ldots, l$, while Case 2 considers the case where the k_is are not necessarily equal.

Case 1: The exact reliability can be computed on using the principle of inclusion and exclusion: Let A_j be the event that the j^{th} row has all bad files; clearly $P(A_j) = \left(\frac{\binom{n-1}{k-1}}{\binom{n}{k}}\right)^l = \left(\frac{k}{n}\right)^l$ by independence of the entries in the columns. The system reliability, denoted by R, is computed as

$$\begin{aligned} R &= P(\text{each row has at least one good file}) \\ &= 1 - P(\text{at least one row has all bad files}) \\ &= 1 - P(\bigcup_{j=1}^{n} A_j) \\ &= 1 - \sum_{j=1}^{n} P(A_j) + \sum_{i<j} P(A_i \cap A_j) - \ldots + (-1)^n P(A_1 \cap A_2 \ldots A_n). \end{aligned}$$

Consider the probabilities in the above series: We will compute the probabilities up to the typical term $P(A_{i_1} \cap A_{i_2} \ldots \cap A_{i_k})$ of the k^{th} series and notice that since there are just k bad files in each backup, every subsequent term is zero. It is easy, moreover, to see that for $r \leq k$,

$$P(A_{i_1} \cap A_{i_2} \ldots \cap A_{i_r}) = \left(\frac{\binom{n-r}{k-r}}{\binom{n}{k}}\right)^l$$

and thus that an expression for the reliability of the system is given by

$$R = 1 - \sum_{i=1}^{k} \binom{n}{i} (-1)^{i+1} \left(\frac{\binom{n-i}{k-i}}{\binom{n}{k}}\right)^l .$$

Case 2: Once again, we use the principle of inclusion-exclusion to obtain the exact reliability. If A_j is the event that the j^{th} row has all bad files, then the system reliability is given by

$$\begin{aligned} R =& 1 - P(\bigcup_{j=1}^{n} A_j) \\ =& 1 - \sum_{i=1}^{k_{\min}} \binom{n}{i} (-1)^{i+1} \left(\prod_{j=1}^{l} \frac{\binom{n-i}{k_j-i}}{\binom{n}{k_j}} \right) \end{aligned}$$

where $k_{\min}$ denotes the minimum value of $k_i \, (i = 1, 2, \ldots, l)$. The reasoning behind the proof of the above expression is the same as in Case 1, with some minor modifications, and we skip the details. Notice how tedious the computation of the above reliability could be.

3. RELIABILITY BOUNDS

Our main results are given in Theorems 1 and 2. Theorem 1 provides lower and upper bounds for the reliability of a software system consisting of n files (each with l backups) that is considered operational if each file has at least one readable backup. For each of the l backups, k out of n files are assumed to be rendered unreadable. We also consider the case where k_i files are unreadable in the i^{th} redundancy backup column ($i = 1, 2, \ldots, l$) and present bounds for the reliability of such a system in Theorem 2.

THEOREM 1.

Let n be the total number of files in the system with each of the l independent backup "columns" having k randomly positioned unreadable files. Let X denote the number of files with no readable backup copies; clearly, $E(X) = \lambda = \sum_{j=1}^{n} (k/n)^l = n \cdot (k/n)^l$. Then the system reliability $R = P(X = 0)$ satisfies the following inequality:

$$|R - e^{-\lambda}| \le d_{TV}(X, \mathrm{Po}(\lambda)) \le \left(1 - e^{-\lambda}\right) \left(\left(\frac{k}{n}\right)^l + n \left(\frac{k}{n}\right)^l \left(\frac{l}{k}\right) \right).$$

PROOF: In order to prove the theorem we must first develop a coupling model for our system $\{I_j\}_{j=1}^n$ of indicator variables, where X, the number of unreadable rows, equals $\sum_{j=1}^{n} I_j$, with $I_j = 1$ if the j^{th} row has all unreadable entries and $I_j = 0$ otherwise. We note that I_j depends on I_i for all $i \ne j$. This global dependence makes it impossible for us to use, e.g., the result of Barbour and Eagleson [1] which hinges on the "dissociatedness" of the I_js that is valid in several reliability situations (but not in ours). The more general Theorem 2.C of Barbour et al. [3] takes this large-scale dependence into account, however, and yields

$$\begin{aligned} |R - e^{-\lambda}| &\le d_{TV}(X, \mathrm{Po}(\lambda)) \\ &\le \left(\frac{1 - e^{-\lambda}}{\lambda}\right) \sum_{j=1}^{n} \left(P^2(I_j = 1) + P(I_j = 1) \sum_{i \ne j} P(I_i \ne J_i) \right) \end{aligned} \qquad (4)$$

where the coupled indicator random variables $\{J_i\}_{i=1}^n = \{J_{i,j}\}_{i=1}^n$ satisfy the condition

$$\mathcal{L}(J_1, J_2 \ldots, J_n) = \mathcal{L}(I_1, I_2 \ldots, I_n | I_j = 1),$$

with $\mathcal{L}(\mathcal{X})$ denoting, as usual, the probability distribution of the random variable $\mathcal{X}$. The idea is for the J_is to mimic the distribution of the I_is when we know that $I_j = 1$; the $\{J_i\}_{i=1}^n$ sequence is defined as follows: Suppose we know that $I_j = 1$, then we let $J_i = I_i$ for all i. If $I_j = 0$, then there exists at least one good file in the j^{th} row. We alter the existing situation very slightly by randomly swapping the status of each good file in the j^{th} row with that of one of the k bad files in its column and let $J_i = 1$ if the i^{th} row has all unreadable files after this interchange is made. In other words, we change each good file in the j^{th} row into a bad one, while keeping the total number of bad files within each column unchanged. It is clear that $\mathcal{L}(J_1, J_2 \ldots, J_n) = \mathcal{L}(I_1, I_2 \ldots, I_n | I_j = 1)$, and that for $i \neq j$, $I_i = 0$ always implies $J_i = 0$, but that $I_i = 1$ might lead to $J_i = 0$. Therefore we have a monotone coupling where $I_i \geq J_i$ for $i \neq j (i = 1, 2, \ldots, n)$. Next consider, for $i \neq j$, the quantity

$$P(I_i \neq J_i) = P(I_i = 1, J_i = 0) = P(I_i = 1)P(J_i = 0 | I_i = 1) = \left(\frac{k}{n}\right)^l P(J_i = 0 | I_i = 1). \quad (5)$$

For a row i with all unsalvageable files to change into one for which a backup is readable, we must have swapped at least one good file in row j with a bad file from row i. The probability of this event is clearly $1 - \left(\frac{k-1}{k}\right)^Y$, where Y equals the number of good files in row j. It follows from (5) that

$$P(I_i \neq J_i) = \left(\frac{k}{n}\right)^l \cdot \left(1 - \left(\frac{k-1}{k}\right)^Y\right) \leq \left(\frac{k}{n}\right)^l \left(1 - \left(\frac{k-1}{k}\right)^l\right),$$

and thus from (4) that

$$\begin{aligned}
|R - e^{-\lambda}| &\leq d_{TV}(X, \text{Po}(\lambda)) \\
&\leq \left(\frac{1 - e^{-\lambda}}{\lambda}\right) n \left(\left(\frac{k}{n}\right)^{2l} + \left(\frac{k}{n}\right)^l \cdot n \left(\frac{k}{n}\right)^l \left(1 - \left(\frac{k-1}{k}\right)^l\right)\right) \\
&= (1 - e^{-\lambda}) \left(\left(\frac{k}{n}\right)^l + n \left(\frac{k}{n}\right)^l \left(1 - \left(\frac{k-1}{k}\right)^l\right)\right) \\
&\leq (1 - e^{-\lambda}) \left(\left(\frac{k}{n}\right)^l + n \left(\frac{k}{n}\right)^l \left(\frac{l}{k}\right)\right),
\end{aligned}$$

proving Theorem 1.

THEOREM 2.

Let n be the total number of files in the system with the i^{th} independent backup column $(i = 1, 2, \ldots, l)$ having k_i randomly positioned unreadable files. Let X denote the number of files with no readable backup copies; we have $E(X) = \lambda = \sum_{j=1}^{n} \prod_{i=1}^{l} (\frac{k_i}{n}) = n \prod_{i=1}^{l} (\frac{k_i}{n})$. Then the system reliability $R = P(X = 0)$ satisfies the following inequality:

$$|R - e^{-\lambda}| \leq d_{TV}(X, \text{Po}(\lambda)) \leq \left(1 - e^{-\lambda}\right) \left(n \left(\frac{k_{\max}}{n}\right)^l - (n-1) \left(\frac{k_{\min} - 1}{n - 1}\right)^l \right).$$

PROOF: The proof of this theorem could have been developed using reasoning similar to that used in the proof of Theorem 1. For variety, however, we exploit the monotonicity of the coupling described below and use Corollary 2.C.2 in Barbour et al. [3] instead; the use of this alternate method gives us almost the same bounds as would have been obtained had we adopted the approach taken in Theorem 1. The coupled random variables $\{J_i\}_{i=1}^n = \{J_{i,j}\}_{i=1}^n$ are defined as in the proof of Theorem 1: If $I_j = 1$, we let $J_i = I_i$ for each i. If $I_j = 0$, on the other hand, we randomly swap the status of each good file in the j^{th} row with that of one of the bad files in its column. In other words, if the r^{th} backup of the j^{th} file is readable, we "change" it into an unreadable file, while maintaining the total number of bad files in the r^{th} column by randomly choosing one of the k_r bad files and changing it into a readable document. Finally, we let $J_i = 1$ if the i^{th} file is unsalvageable *after* the abovedescribed changes are implemented. It is clear that the condition

$$\mathcal{L}(J_1, J_2 \ldots, J_n) = \mathcal{L}(I_1, I_2 \ldots, I_n | I_j = 1)$$

is satisfied and that $J_i \leq I_i$ for each $i \neq j$. Corollary 2.C.2 in Barbour et al. [3] yields

$$|R - e^{-\lambda}| \leq d_{TV}(X, \text{Po}(\lambda)) \leq (1 - e^{-\lambda}) \left(1 - \frac{V(X)}{\lambda}\right). \tag{6}$$

The variance, $V(X) = V(\sum_{j=1}^{n} I_j)$, can be written as

$$V(\sum_{j=1}^{n} I_j) = \sum_{j=1}^{n} V(I_j) + 2 \sum_{j=1}^{n} \sum_{i<j} \text{Cov}(I_i, I_j). \tag{7}$$

Furthermore,

$$\begin{aligned}\operatorname{Cov}(I_i, I_j) = E(I_i \cdot I_j) - E^2(I_j) =& P(I_i \cdot I_j = 1) - P^2(I_j = 1) \\ =& \prod_{i=1}^{l} \left(\frac{k_i(k_i - 1)}{n(n-1)} \right) - \prod_{i=1}^{l} \left(\frac{k_i}{n} \right)^2\end{aligned}$$

and

$$V(I_j) = E(I_j^2) - E^2(I_j) = E(I_j) - E^2(I_j) = \prod_{i=1}^{l} \left(\frac{k_i}{n} \right) - \prod_{i=1}^{l} \left(\frac{k_i}{n} \right)^2.$$

Substituting the above facts in (7), we get the following expression for the variance of X:

$$V(X) = n \left(\prod_{i=1}^{l} \left(\frac{k_i}{n} \right) - \prod_{i=1}^{l} \left(\frac{k_i}{n} \right)^2 \right) + n(n-1) \left(\prod_{i=1}^{l} \left(\frac{k_i(k_i - 1)}{n(n-1)} \right) - \prod_{i=1}^{l} \left(\frac{k_i}{n} \right)^2 \right). \quad (8)$$

On substitution of (8) in (6) and simplifying, we get

$$|R - e^{-\lambda}| \leq d_{TV}(X, \operatorname{Po}(\lambda)) \leq \left(1 - e^{-\lambda}\right) \left(n \prod_{i=1}^{l} \left(\frac{k_i}{n} \right) - (n-1) \prod_{i=1}^{l} \left(\frac{(k_i - 1)}{(n-1)} \right) \right). \quad (9)$$

Finally, we can upper bound $\prod_{i=1}^{l} \left(\frac{k_i}{n} \right)$ by $\left(\frac{k_{\max}}{n} \right)^l$ and the quantity $-\prod_{i=1}^{l} \left(\frac{(k_i - 1)}{(n-1)} \right)$ by $-\left(\frac{k_{\min} - 1}{n-1} \right)^l$, to obtain the total variation estimate

$$|R - e^{-\lambda}| \leq d_{TV}(X, \operatorname{Po}(\lambda)) \leq \left(1 - e^{-\lambda}\right) \left(n \left(\frac{k_{\max}}{n} \right)^l - (n-1) \left(\frac{k_{\min} - 1}{n-1} \right)^l \right) \quad (10)$$

as asserted.

Remarks:

Since Theorem 1 and Theorem 2 both provide total variation bounds, they may each be used, e.g., in situations where the redundancy system is considered operative provided that there are m *or fewer* totally erased files. Computation of the exact system reliability $P(X \leq m)$ is completely intractable in this case, but a Poisson approximation is provided by $\sum_{j=0}^{m} \frac{e^{-\lambda} \lambda^j}{j!}$ and, more importantly, the same error bound as that provided by Theorem 1 or Theorem 2 governs the closeness of the approximation. Notice moreover, that (9) yields a sharper bound for the error, but is somewhat less streamlined than the estimate (10).

We next tabulate the error bound for the "equal k" case. Here we use the bound (10) which specializes to

$$|R - e^{-\lambda}| \le d_{TV}(X, \mathrm{Po}(\lambda)) \le \left(1 - e^{-\lambda}\right) \left(n \left(\frac{k}{n}\right)^l - (n-1) \left(\frac{k-1}{n-1}\right)^l \right)$$

on setting $k_{\max} = k_{\min} = k$. The reason why we prefer Theorem 2 to Theorem 1 is that for equal k_is, the ratio

$$\frac{\text{Error bound of Theorem 1}}{\text{Error bound of Theorem 2}} \to \frac{l}{k(1-(1-\frac{1}{k})^l)} \ge 1$$

as $n \to \infty$, so that Theorem 2 actually outperforms Theorem 1. The tables that follow are for $n = 100$ and $n = 1000$, and for values of k and l that show off our bounds to advantage. It should be pointed out, however, that our error bounds are not always small. For example, with $k = 8, l = 3$, the approximate reliabilites for $n = 10, 20$ are 0.005976 and 0.278037, with error bounds of 0.88 and 0.24 respectively! The point is that we need n to be large, and k to be small compared to n for our errors to be small. Moreover, given such values of k and n, larger values of l will lead to progressively lower error bounds. It should also be mentioned that we are really not considering huge (and very reliable) databases – such as the file systems for telephone calls or air passengers in a given period. For example, even with millions of passengers travelling by air each day, a person's records over a five-year period, say, can be traced with little difficulty. k is clearly so low in this case that building a totally automatic backup system is not a problem.

TABLE 1

n	l	k	APPROX. REL.	ERROR BOUND
100	3	5	0.987578	7.41612e-05
		6	0.978632	0.000189
		7	0.966282	0.000413
		8	0.950089	0.000809
	4	5	0.999375	2.256569e-07
		6	0.998705	8.442746e-07
		7	0.997602	2.554787e-06
		8	0.995912	6.628101e-06
	5	5	0.999969	6.434259e-10
		6	0.999922	3.516801e-09
		7	0.999832	1.464107e-08
		8	0.999672	5.003369e-08
	6	5	0.999998	1.768422e-12
		6	0.999995	1.41021e-11
		7	0.999988	8.069327e-11
		8	0.999974	3.628861e-10
	12	5	1	5.505428e-28
		6	1	4.145512e-26
		7	1	1.579282e-24
		8	1	3.660013e-23

TABLE 1 (Contd.)

n	l	k	APPROX. REL.	ERROR BOUND
1000	3	35	0.958031	0.000147
		36	0.954416	0.000168
		37	0.950608	0.000193
		40	0.938005	0.000283
	4	35	0.998501	2.403277e-07
		36	0.998322	2.928136e-07
		37	0.998128	3.548135e-07
		40	0.997443	6.126038e-07
	5	35	0.999947	3.626148e-10
		36	0.99994	4.676124e-10
		37	0.999931	5.987858e-10
		40	0.999898	1.209692e-09
	6	35	0.999998	5.252177e-13
		36	0.999998	7.168087e-13
		37	0.999997	9.699137e-13
		40	0.999996	2.292156e-12
	12	35	1	3.218729e-30
		36	1	6.310054e-30
		37	1	1.173905e-29
		40	1	7.139341e-29

4. QUESTIONS FOR FURTHER RESEARCH

There are several questions that could be looked into and studied in greater detail; we mention two: Consider the proof of Theorem 1. Notice that there are at most k indices i for which I_i and J_i could differ. Can a better coupling be defined for which one can exploit this fact? More specifically, our analysis in Theorem 1 used the estimate $\sum P(I_i \neq J_i) \leq nP(I_s \neq J_s)$ where $s \neq j$. Can this bound be bettered to $kP(I_s \neq J_s)$ by using a more judicious coupling? Next, consider the following question: Suppose it is observed that unreadable files are far more prevalent in some "rows" than in others. One would then certainly not wish to create backups in the same order each time. What is an optimal backup strategy under this scenario? How does it compare to the obvious (but sub-optimal) strategy of randomizing the order of backup files within each column?

Acknowledgements

The authors would like to thank Daniel S. Moak, whose questions motivated them to study this problem. The research of the first- and second-named authors was partially supported by U.S. National Science Foundation Grants DMS-9200409 and DMS-9322460.

REFERENCES

[1] Barbour, A. D., and Eagleson, G. K., "Poisson convergence for dissociated statistics," *Journal of the Royal Statistical Society, Series B,* **46**, 397–402 (1984).

[2] Barbour, A. D., and Hall, P., "On the rate of Poisson convergence," *Mathematical Proceedings of the Cambridge Philosophical Society*, **95**, 473–480 (1984).

[3] Barbour, A. D., Holst, L., and Janson, S., *Poisson Approximation*, Oxford University Press, Oxford, 1992.

[4] Chryssaphinou, O., and Papastavridis, S., "Limit distribution for a consecutive-k-out-of-$n : F$ system," *Advances in Applied Probability*, **22**, 491–493 (1990).

[5] Colorado Memory Systems Inc., *COLORADO BACKUP Tape Backup Software*, 1993.

[6] Godbole, A., "Approximate reliabilities of m-consecutive-k-out-of-n : failure systems," *Statistica Sinica*, **3**, 321–327 (1993).

[7] Godbole, A., Khunger, S., and Ososanya, E., "Exact and approximate hypercube reliabilities," in *Runs and Patterns in Probability: Selected Papers*, A. Godbole and S. Papastavridis, eds., Kluwer Academic Publishers, Boston, 1994, pp. 163–172.

[8] Koutras, M., Papadopoulos, G., and Papastavridis, S., "Reliability of 2-dimensional consecutive-k-out-of-$n : F$ system," *IEEE Transactions on Reliability*, **R42**, 658–661 (1993).

[9] Koutras, M., and Papastavridis, S., "Application of the Stein-Chen method for bounds and limit theorems in the reliability of coherent structures," *Naval Research Logistics*, **40**, 617–631 (1993).

[10] Papastavridis, S., "A Weibull limit for the reliability of a consecutive k-within-m-out-of-n system," *Advances in Applied Probability*, **20**, 690–692 (1988).

[11] Roos, M., "Stein-Chen method for compound Poisson approximation," Doctoral Dissertation, Universität Zürich (1993).