Modeling Correlated Failures in Survivable Storage Systems

Mehmet Bakkaloglu, Jay J. Wylie, Chenxi Wang, Gregory R. Ganger
Carnegie Mellon University
{mehmetb, jwylie, chenxi, ganger}@ece.cmu.edu

Abstract

The design of survivable storage systems involves inherent trade-offs among properties such as performance, security, and availability. A toolbox of simple and accurate models of these properties allows a designer to make informed decisions. This paper focuses on availability modeling. We describe two ways of extending the classic model with a single "correlation parameter" to accommodate correlated failures. We evaluate the efficacy of the models by comparing their results with real measurements.

1. Introduction

Survivable storage systems [4] encode and distribute data over multiple storage nodes to survive failures and malicious attacks. PASIS is one such system that is configurable to support many data distribution schemes, including threshold schemes [3], for improving storage availability and security. A common property of threshold schemes is that they convert data into n pieces of which only a subset of size m ($m \leq n$) are required to fully reconstruct the data. In general, the n pieces are stored on n different storage nodes; for the data to be available, at least m of these nodes must be reachable. The values of n and m are key design parameters that can affect the availability of the system. Having a simple and accurate availability model allows designers to explore this choice in the context of performance, security, and availability trade-offs [6].

The classic approach to modeling the availability of distributed storage is to assume that storage nodes have identical availabilities and independent failures:

$$f(n, m, Avg.Avail) = \sum_{i=m}^{n} \binom{n}{i} (Avg.Avail)^{i} \quad (1 - Avg.Avail)^{n-i}$$

This model is useful because it enables a system designer to reason about the effect of average node availability on system availability. However, it ignores correlations known to exist among storage node failures.

The most accurate way of modeling correlated failures is to specify exactly the probability of each subset of nodes being unavailable. But, such a model is too complex in practice and provides little intuition to system designers. A better approach would be to extend the classic model with a single correlation parameter that indicates the level of failure correlation of the system. In this paper, we describe two such methods. The first model is based on estimating conditional probabilities. The second model is based on the Beta-Binomial distribution.

There have been many efforts to model correlated failures. Most involve state-based models where all the possible system states and the transition probabilities between the states are identified. Such an approach is not a practical design tool for evaluating high-level design trade-offs. Related work and full details of our work can be found in [1].

2. Modeling Correlated Failures

In our work we use two real datasets. The first one is availability measurements of 99 web servers. The second dataset is from a campus of desktops [5]. For brevity, this paper shows results only from the former.

2.1. Model Based on Conditional Probabilities

For this model, we propose a correlation parameter based on conditional probabilities:

$\forall$ Pair of nodes (X, Y) in the system, where X $\neq$ Y,
Correlation Level= Average (P (X is down | Y is down))

This parameter indicates 2-way correlations. It does not consider failure correlations among more than two nodes (e.g., P (X is down | Y and Z are down)). A close examination of our datasets revealed that (x+1)-way correlations are generally higher than x-way correlations. Based on this observation, we defined a recursive model to estimate x-way correlation probabilities using (x−1)-way and (x−2)-way correlations:

$$P_x = P_{x-1} + \frac{(P_{x-1} - P_{x-2})}{2}$$

(P_1=1−Avg. Avail and P_2=Correlation Level)
where P_x indicates the x-way correlation probability.

2.2. The Beta-Binomial Distribution

A second method for modeling the availability of correlated failures is the Beta-Binomial distribution. This distribution has been used to model other correlated events [2]. The distribution uses the average availability and a correlation level as parameters. The distribution is computed by randomizing the parameter (p=Avg. Avail.) of the classic model using the Beta density function ($f_p(p)$). The probability that i out of n nodes fail is:

$$b_n(i) = \int_0^1 \binom{n}{i} p^i (1-p)^{n-i} f_p(p)\,dp$$

The correlation level for the Beta-Binomial distribution is not intuitive; it determines the degree of variation in p. To compute its value, we use our availability data to determine the value of $b_2(1) + b_2(2)$ empirically (i.e., the case of n=2, m=2). Knowing the measured value of $b_2(1) + b_2(2)$ allows us to solve for the unknown correlation level. Once the correlation level is determined the distribution can be used to estimate the availability of any scheme (n, m). A detailed description of the Beta-Binomial distribution can be found in [2].

2.3. Evaluation

Figures 1 and 2 compare the results of the three models to the actual availability values for the dataset described. The classic model renders the most inaccurate results. Because the conditional probability model takes the 2-way correlation as a parameter, it is expected to accurately estimate the case where n=2, m=1. The estimation error, shown in Figure 1, is due to the model using only averages; it does not consider variances. For schemes with m=1, if n is low the model under-estimates, but it over-estimates for higher n's. This increase is due to the model under-estimating the higher x-way correlations.

In Figure 2, for m between 1 and (n−1), the conditional probability model first over-estimates and then under-estimates. The Beta-Binomial distribution exhibits the same behavior. This is due to two factors. First, note that a decrease in failure correlation increases the availability of schemes with large (n−m)'s more than schemes with small (n−m)'s. Since the model under-estimates the correlation, availability of schemes with large (n−m)'s tend to be over-estimated. Second, the model assumes that all the nodes are identical and assumes smooth transitions from one x-way correlation to another. In reality, some sets of servers are more correlated to each other than to others; thus, using averages leads to inaccuracies.

For schemes with n=m, the two proposed models always over-estimate. This is because the models consider only the average availability. In reality the schemes with n=m are bound by the server with minimum availability. In summary, a model that uses only average availability is bound to over-estimate the availability of schemes with n=m unless it significantly under-estimates the higher x-way correlations (e.g., independence).

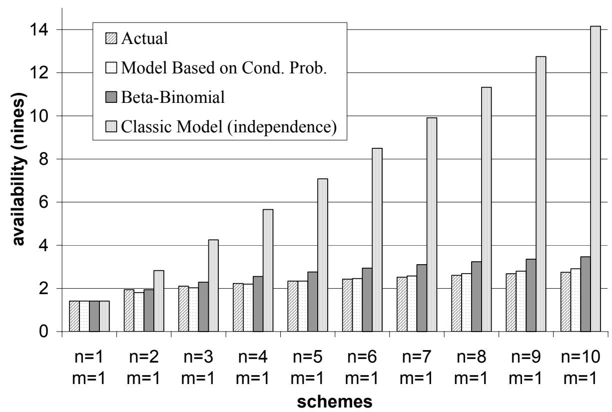

Figure 1. Availabilities as n increases

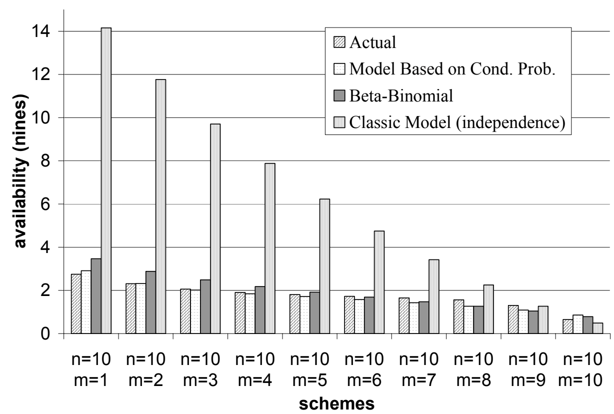

Figure 2. Availabilities as m increases

3. References

[1] M. Bakkaloglu, et. al., "On Correlated Failures in Survivable Storage Systems", Technical Report CMU-CS-02-129, Carnegie Mellon University, May 2002.

[2] V. F. Nicola, et. al., "Modeling of Correlated Failures and Community Error Recovery in Multiversion Software", IEEE Transaction on Software Engineering, Vol.16 No.3, March 1990.

[3] A. Shamir, "How to Share a Secret", Communications of ACM, 22(11):612-613 ACM November 1979.

[4] J. J. Wylie, et. al., "Survivable Information Storage Systems", IEEE Computer, 33(8): 61-68, August 2000.

[5] W. J. Bolosky, et. al., "Feasibility of a Serverless Distributed File System Deployed on an Existing Set of Desktop PCs", Proceedings of the international conference on Measurement and modeling of computer systems: 34-43, 2000.

[6] J. J. Wylie, et. al., "Selecting the Right Data Distribution Scheme for a Survivable Storage System", Technical Report CMU-CS-01-120, Carnegie Mellon University, May 2001.