

CALIBRATING A DEPENDENT FAILURE MODEL FOR COMPUTING RELIABILITIES IN TELECOMMUNICATION NETWORKS

Omar Matus
Javiera Barrera
Eduardo Moreno

Faculty of Engineering and Sciences
Universidad Adolfo Ibáñez
Avda. Diagonal Las Torres 2700
Peñalolén, Santiago, CHILE

Gerardo Rubino

INRIA Rennes – Bretagne Atlantique
Campus de Beaulieu
35042 Rennes Cedex, FRANCE

ABSTRACT

In this work, we propose a methodology for calibrating a dependent failure model to compute the reliability in a telecommunication network. We use the Marshall-Olkin (MO) copula model, which captures failures that arise simultaneously in groups of links. In practice, this model is difficult to calibrate because it requires the estimation of a number of parameters that is exponential in the number of links. We formulate an optimization problem for calibrating an MO copula model to attain given marginal failure probabilities for all links and the correlations between them. Using a geographic failure model, we calibrate various MO copula models using our methodology, we simulate them, and we benchmark the reliabilities thus obtained. Our experiments show that considering the simultaneous failures of small and connected subsets of links is the key to obtaining a good approximation of reliability, confirming what is suggested by the telecommunication literature.

1 INTRODUCTION

The topological design and evaluation of telecommunication networks have been addressed in many studies since the 1980s. Among the various definitions of reliability, in this work, we focus on the reliability of a network in terms of the probability that a given set of components will remain in the same connected component of the network at a given time, considering that these components may become inoperative over time.

To evaluate the reliability of a network, a failure model is required. Even for the simplest failure model, in which each component fails independently with some probability, it is difficult to directly evaluate the reliability; more precisely, it is a #P-complete problem (Provan and Ball 1983). Hence, Monte Carlo simulation techniques have become indispensable for computing such reliabilities (see, for instance, Kroese, Taimre, and Botev (2013), or the on-line proceedings of the WSC).

A natural next step is to address topological network design. This problem can also be easily stated; given a set of nodes and potential links, choose a subset of those links to construct a network such that the total cost is minimized and the reliability is maximized. If the evaluation of reliability is already a difficult problem, the design problem of maximizing reliability is even more so; thus, to tackle the latter problem, many simplifications are assumed, the most common being the independent failure of the components.

However, empirical analyses of networks have revealed that a dependence exists among the failures of different components, which could significantly affect reliability computations. Gonzalez et al. (2010) analyzed data from a Norwegian optical network, showing that significant correlations are observed in an important number of pairs of components and that this correlation decays with an increasing distance

between them. Gill, Jain, and Nagappan (2011) found that simultaneous failures of n links become more rare as n increases. Similar conclusions were obtained by Turner et al. (2010) in a study of internet provider networks. These studies show that the dependence among component failures is significant, even for networks based on different technologies.

The question of how to choose a dependency model has been studied by Singpurwalla (2002), who discussed various models, such as the Freund's bivariate exponential distribution, multivariate Weibull/Gamma models and the Marshall-Olkin (MO) copula model. The MO model was proposed by Marshall and Olkin (1967), and it assumes that failures occur in groups of components. This model offers many advantages that make it very flexible and realistic (see the discussion in Botev et al. (2016) and references therein). Botev et al. (2016) proposed rare-event simulation techniques to compute the reliability of a network in which failures arise according to the MO model. For the design problem, Barrera, Cancela, and Moreno (2014) recently proposed an optimization formulation to design networks under MO-type failure models and causal failure models. With the publication of these works, it could be supposed that at least for certain networks, the problem of evaluation and design under mutually dependent failures has been addressed. However, both works assume that the parameters of the MO copula failures are available. Singpurwalla (2002) has already noted that although the MO model is a realistic failure model, it is not scalable because the number of parameters to be calibrated grows exponentially with the number of components. Therefore, the calibration of the MO copula parameters is an important issue to overcome, and this is the problem that we address in this article.

As noted by Hagstrom and Ross (2001), correlation data for existing network are scarce; hence, suitable models should consider this lack of data. In addition, the techniques discussed above are intended to be applied for the design of a network before its construction. Therefore, it is natural to assume that only partial information can be gathered, mostly from similar existing networks. Our main result is a methodology for calibrating an MO failure model to approximate the reliability of a network. In our methodology, we assume that only the marginal failure probabilities and correlations are available for each component. We evaluate the proposed methodology by means of an experiment in which the reliabilities obtained using two failure models are compared: we directly simulate a phenomenon that affects the network, causing dependent failures, and we compare the results with simulations using the calibrated MO model. The article is organized as follows. In Section 2, we introduce the notation and formally introduce the MO model; we discuss the cases of serial and parallel components under the MO failure model to quantify the errors that can be produced by the assumption of independence and the differences that can arise between two different MO models with the same marginal failure probabilities and correlations. In Section 3, we present an optimization problem for obtaining an MO model with the desired calibration parameters. We also discuss the behaviors observed in different networks and how to obtain MO models that emulate these behaviors. In Section 4, we introduce a physical failure model and compare it with various MO models. Finally, in Section 5, we present the conclusions of our study.

2 NOTATION AND MARSHALL-OLKIN MODEL

2.1 Marshall-Olkin (MO) Copula Model for Dependent Failures

Let \mathcal{G} be a graph with set of nodes \mathcal{N} , and let $\mathcal{C} = \{1, 2, \dots, n\}$ denote the set of links or components connecting distinct pairs of nodes. We assume that nodes cannot fail. Links, however, are subject to failure, and at any time $t \geq 0$, they can be in one of two states, either operational (up) or not operational (down). At time $t = 0$, all components are operational. After random lifetimes in the up state, links fail, moving into the down state, in which they remain forever (that is, we consider non-repairable systems). For each link i , we define its lifetime T_i as the random time at which it fails. In the standard model of network reliability, the random variables T_1, \dots, T_n are assumed to be independent. Here, the model is much more general. Let \mathcal{P} be a collection of subsets of \mathcal{C} , which are not necessarily disjoint. For each subset $V \in \mathcal{P}$, we define a positive random variable W_V that represents the instant at which all links in V fail simultaneously. We

refer to such an event as a *shock*. The result of a shock associated with subset V is that at time W_V , all links in V are in the down state (if they were already down, they remain down, that is, nothing happens). Hence, the lifetime of a link i corresponds to the earliest time at which a shock that affects i occurs, that is, $T_i = \min_{i \in V} \{W_V\}$. In the MO copula model, the random variables $\{W_V\}_{V \in \mathcal{P}}$ are mutually independent with exponential distributions. Hence, T_i also follows an exponential distribution for all $i \in \mathcal{C}$.

At any time $t \geq 0$, the entire network is either up or down, depending on whether a certain connectivity property is satisfied. In this paper, we illustrate our approach using the basic s, t -connectivity: two nodes in \mathcal{N} , denoted by s and t , are marked, and the network is up if and only if there exists a path connecting them that is composed only of operational links. We assume that the graph \mathcal{G} is connected, so, at time $t = 0$ the network is operational. When all links are down, that is, at any time $t \geq \max_{i \in V} \{W_V\}$, the network is down.

The metric of interest here is the probability that the network will be up at time t , $R(t)$. In some cases, the main instant t of interest is the so-called “mission time”, the maximal time value until which we wish the network to continue operating without problems. Because t is mathematically arbitrary here, we choose to develop our technique using a value of 1 for simplicity. Equivalently, we can also think of t as the mission time and consider that the time unit we are using is precisely equal to that mission time.

There is another reason that a value of 1 is interesting. Consider the classical setting with independent components and a *static* model in which time is not an explicit variable. Components and systems are either up or down; we are given the graph and the individual probabilities for the components to be, say, down, denoted by p_i for component i , and we measure the network dependability using some measure of the capability of the network to provide communication services, for instance, the probability that s and t will be connected in the implicit random partial graph of \mathcal{G} . This measure is typically called the *network reliability*. Then, consider our initial dynamic model, with $\mathcal{P} = \{\{1\}, \dots, \{n\}\}$, where the random variable $W_{\{i\}}$ has a rate $\lambda_i = -\ln(1 - p_i)$. It can be seen that the probability that link i will have failed at time t is $\Pr(T_i \leq t) = 1 - \exp(-\lambda_i t)$, and $\mathbb{P}(T_i \leq 1) = 1 - \exp(-\lambda_i) = 1 - (1 - p_i) = p_i$. Thus, in this case, the static s, t -network reliability is equal to $R(1)$. This is essentially the Creation Process of Elperin, Gertsbakh, and Lomonosov (1991), or rather, its complement, considering that in the formulation of Elperin, Gertsbakh, and Lomonosov (1991), the system’s components all start in the failed state and are repaired until the system becomes operational.

As explained in the introduction, we are interested in calibrating an MO copula model that represents the behavior of a communication network based on the marginal failure probabilities and the correlations between failures. In principle, these values are easy to measure in an operating communication infrastructure.

Lemma 1 Let X_i denote the state of link i at time 1. Let p_i be the marginal failure probability of link i , that is, $p_i = \mathbb{P}(X_i = 0)$, and let $\rho_{i,j}$ be the correlation between the failures of links i and j , that is, $\rho_{i,j} = \text{Cov}(X_i, X_j) / (\sqrt{\text{Var}(X_i)} \sqrt{\text{Var}(X_j)})$. Then, an MO copula model with these marginal distributions and correlations must satisfy the following set of equations, where the set of unknown variables is $\{\lambda_V : V \subseteq \mathcal{C}\}$:

$$\sum_{V \subseteq \mathcal{C}: \{i,j\} \in V} \lambda_V = \ln \left(\frac{\rho_{i,j} \sqrt{p_i} \sqrt{p_j}}{\sqrt{(1-p_i)(1-p_j)}} + 1 \right), \quad \forall i, j \in \mathcal{C}, i \neq j, \quad (1a)$$

$$\sum_{V \subseteq \mathcal{C}: i \in V} \lambda_V = -\ln(1 - p_i), \quad \forall i \in \mathcal{C}, \quad (1b)$$

$$\lambda_V \geq 0, \quad \forall V \subseteq \mathcal{C}. \quad (1c)$$

Proof. This result is obtained directly by expressing the covariances and marginal failure probabilities in terms of the links' lifetimes,

$$\begin{aligned} \mathbb{E}(X_i) &= 1 - p_i = \exp\left(-\sum_{V \subseteq \mathcal{C}: i \in V} \lambda_V\right) \quad \text{and} \\ \mathbb{E}(X_i \cdot X_j) &= \Pr(T_i > 1 \wedge T_j > 1) = \exp\left(-\sum_{V \subseteq \mathcal{C}: i \in V \vee j \in V} \lambda_V\right) \\ &= (1 - p_i)(1 - p_j) \cdot \exp\left(-\sum_{V \subseteq \mathcal{C}: \{i, j\} \in V} \lambda_V\right), \end{aligned}$$

and isolating the corresponding terms. □

Note that this linear system contains $n(n + 1)/2$ equations and $2^n - 1$ variables, so in general, it will have an infinite number of solutions. For this reason, it is important to know the accuracy of estimation of the system's reliability for each of these solutions.

Henceforth, we will denote the network reliability at $t = 1$ simply by R , such that $R = R(1)$. Moreover, we will abuse the notation and refer to the subsets V as "copulas", as the shock times $\{W_V\}$ couple the links' lifetimes.

2.2 Relevance of Choosing Different MO Copula Parameters

Before imposing specific criteria on which to choose an MO model from among all possible solutions for system (1), we will demonstrate the impact of choosing different MO copula parameters on the reliability. We consider the extreme cases of a system with n components: one is a system in which all components are in series, and the other is a system with all components in parallel. For each system, we also consider the two extreme cases of sets of copulas that satisfy the equation given in (1): one in which shocks affect only one or two components at once and one in which only shocks that affect either one component or all components simultaneously are considered.

We are interested in showing that for different sets of copulas that satisfy the system of equations given in (1), the resulting reliabilities can be very different even if correlations are not that high, indicating that it is important not only to satisfy these equations but also to choose sets of copulas that are representative of the possible failures in the system.

Let G_1 and G_2 denote networks composed of n links in series (G_1) and in parallel (G_2). Let p be the marginal unreliability of any link, and let ρ be the correlation between the states of any pair of links. For both networks, we obtain two different solutions to Equation (1): the solution obtained using only copulas of a size of at most 2, and the solution obtained using copulas of size 1 and one copula that includes all n links.

Then, by solving the system of equations given in (1) using the restricted set of valid copulas for each case, we can compute the exact reliability R between the two terminal nodes of each network, yielding the following formulas:

1. G_1 : n links in series.
 - (a) Copulas of one or two components:

$$R = (\rho p + 1 - p)^{\frac{n(n-1)}{2}} (1 - p)^{\frac{n(3-n)}{2}}. \tag{2}$$

- (b) Copulas of one or n components:

$$R = (\rho p + 1 - p)^{n-1} (1 - p). \tag{3}$$




Figure 1: Reliabilities of series and parallel networks with $n = 5$ components, as functions of the correlation between any pair of link states, for various values of the elementary reliability p and for two choices of copulas: either only copulas of size ≤ 2 or only copulas of size 1 or n (all components).

2. G_2 : n links in parallel.
 - (a) Copulas of one or two components:

$$R = \sum_{k=1}^n \left[\binom{n}{k} (-1)^{k+1} (\rho p + 1 - p)^{\frac{k(k-1)}{2}} (1 - p)^{\frac{k(k-3)}{2}} \right]. \quad (4)$$

- (b) Copulas of one or n components:

$$R = \frac{1 - p^n (1 - \rho)^n}{\rho p + 1 - p} (1 - p). \quad (5)$$

Figure 1 shows the difference between the two solutions for various values of p and ρ and for the two types of networks with $n = 5$ components. Continuous lines represent the reliabilities for one- or all-component copulas, and dashed lines represent the corresponding values obtained using one- or two-component copulas. Only small correlations can be considered using copulas of size 1 or 2. However, note that the difference between the results for the different copula subsets can be considerable even for small correlations (greater than 10% for correlations smaller than 0.1). Moreover, if we ignore the correlations, we obtain the reliability calculated for a correlation of $\rho = 0$ (the intersection with the Y axis in the figures). Hence, ignoring the correlation produces an underestimation of the reliability for the series case and an overestimation of the reliability for the parallel case, and the difference grows as the correlation increases.

3 CHOOSING AN APPROPRIATE SET OF COPULA PARAMETERS

3.1 An Optimization Problem for Choosing MO Copula Parameters

In this section, we discuss how to obtain a set of copula parameters that satisfies the system of equations given in (1). Note that this is not a simple problem, because the system has an exponential number of variables. To find a solution to this system, we rewrite (1) as the following optimization problem:

$$\min \sum_{i,j \in \mathcal{C}} (t_{ij}^+ + t_{ij}^-) \tag{6a}$$

$$\sum_{V \subseteq \mathcal{C}: \{i,j\} \in V} \lambda_V + t_{ij}^+ - t_{ij}^- = \ln \left(\frac{\rho_{i,j} \sqrt{p_i} \sqrt{p_j}}{\sqrt{(1-p_i)(1-p_j)}} + 1 \right), \quad \forall i, j \in \mathcal{C}, i \neq j \tag{6b}$$

$$\sum_{V \subseteq \mathcal{C}: i \in V} \lambda_V + t_{ii}^+ - t_{ii}^- = -\ln(1-p_i), \quad \forall i \in \mathcal{C} \tag{6c}$$

$$\lambda_V \geq 0, \quad \forall V \subseteq \mathcal{C} \tag{6d}$$

$$t_{ij}^+, t_{ij}^- \geq 0, \quad \forall i, j \in \mathcal{C}. \tag{6e}$$

That is, we add slack variables t_{ij}^+ and t_{ij}^- for each equation, and we minimize the sum of these slack values. If we obtain a solution with an objective value equal to 0, then it is an exact solution to equation (1). A solution to this linear optimization problem with an exponential number of variables can be obtained using a technique known as column generation (Dantzig and Wolfe 1960). That is, we solve the problem with a reduced set of variables λ_V , called the “master” problem. To verify whether an additional variable should be included in the problem, we need to find a variable with a negative reduced cost. The reduced cost associated with a variable λ_V is given by

$$\bar{C}_V = -\sum_{i \in V} \mu_i - \sum_{i,j \in V} v_{ij},$$

where μ_i and v_{ij} are the dual variables of equations (6c) and (6b), respectively.

We can formulate an integer programming model to find a variable λ_V with the minimum reduced cost by solving the following “pricing” problem:

$$\min -\sum_{i \in \mathcal{C}} \mu_i x_i - \sum_{i,j \in \mathcal{C}} v_{ij} y_{ij} \tag{7a}$$

$$y_{ij} \leq x_i \quad \forall i, j \in \mathcal{C} \tag{7b}$$

$$y_{ij} \leq x_j \quad \forall i, j \in \mathcal{C} \tag{7c}$$

$$x_i + x_j \leq y_{ij} + 1 \quad \forall i, j \in \mathcal{C} \tag{7d}$$

$$x_i, y_{ij} \in \{0, 1\} \quad \forall i, j \in \mathcal{C}. \tag{7e}$$

The binary variable x_i indicates whether link i should be included in the copula V . The variable y_{ij} takes a value of 1 if and only if x_i and x_j are included in the copula. The objective function is equal to the reduced cost associated with V . An optimal value of the objective function smaller than 0 indicates that the new copula defined by $\{x_i\}_{i \in \mathcal{C}}$ should be included in the master problem. Then, we solve the master problem again and iterate until all variables with a negative reduced cost have been included.

Note that this pricing problem is NP-hard (by reduction of the Independent Set problem), but our implementation shows that optimization software can solve large-scale instances of this pricing problem in a few seconds by virtue of the structure of the pricing problem that is exploited by such solvers.

Note also that additional constraints can be included in the pricing problem. For example, we can limit the size of the copula V by adding a cardinality constraint ($\sum_{i \in \mathcal{C}} x_i \leq k$). This allows us to impose additional criteria on the set of copulas used to calibrate the model. Note that these additional criteria could lead us to an optimal solution that does not satisfy the system of equations given in (1), but such a solution does provide a good approximation of the desired marginal distributions and correlations.

3.2 Empirical Failure Behaviors in Telecommunication Networks

As mentioned in the introduction, empirical analyses of networks have shown that the probability of simultaneous failure in a network decreases as the number of components involved in the simultaneous failure increases. For the data center studied in Gill et al. (2011), §4.6, the authors showed that failures involving up to four links accounted for 90% of all failures. Gonzalez et al. (2010), §IV.B, showed that the number of simultaneous failure events decreases with increasing distance between the components involved. Turner et al. (2010), §4.2.2, reported that simultaneous failures occur in links that share a node, or more precisely, a router. A similar behavior was reported by Markopoulou et al. (2008) for IP backbone networks. These studies provide evidence that among all models that satisfy equation (1), we should preferentially choose a model with copulas that include only links that share a node, thereby limiting the maximum copula size. To do so, we restrict the set of available copulas in model (6) to consider only subsets V that satisfy the desired conditions. Note that in this case, the solution to the problem (6) could be different from zero (i.e., the resulting copulas would not satisfy the system of equations given in (1)). Nevertheless, as we show in the following section, this solution still provides a good approximation of the reliability of the network.

4 COMPUTATIONAL EXPERIMENTS

To compare the performances of different methods of selecting the MO copula parameters, we must benchmark our results against another simulation model that naturally induces correlations between failures. To do so, we simulate a physical model of failures similar to that proposed by Agarwal et al. (2013), and we attempt to replicate the obtained reliabilities using the MO copula model based on only the empirical marginal probabilities of failure and their correlations.

In the physical model, the network can be affected by physical disasters such as earthquakes, storms or electromagnetic pulses. These disasters occur at random locations and with random intensities. When an event occurs, we assign to each component a probability of failure that decays with increasing distance from the location of the disaster.

In this model, the disaster position q is given by a uniform random variable over the square area of the network and the intensity of the event is given by an independent exponential random variable D , as suggested by Gutenberg and Richter (1954).

The probability of failure for each component, as a consequence of the disaster, is inversely proportional to its distance from the event location and directly proportional to the intensity of the event. Hence, given a disaster located in a position q with intensity D , the failure probability of a component in a position u is given by

$$f_u = \max \left\{ 0, 1 - \frac{d(q, u)}{D} \right\}, \quad (8)$$

where $d(q, u)$ is the Euclidean distance between u and q .

We apply the model described above to the network shown in Figure 2, which consists of 12 nodes and 16 links. We sample the position q of the event from a uniform distribution over $[0, 16] \times [0, 16]$, and we independently sample the intensity D from an exponential distribution with the parameter $\lambda = 5$. Using these parameters, we compute the failure probability of each link f_u , where u is the midpoint of the link, and we sample the state of each link (operative or failed) using a Bernoulli random variable with the corresponding probability f_u . This process represents one simulation of the physical model. We denote this model by PHYSICALM.

From the resulting simulations, we compute the empirical marginal failure probabilities p_i and the empirical correlations between failures ρ_{ij} for all $i, j \in \mathcal{C}$, which will be used to calibrate the MO copula parameters.

We calibrate six different MO copula models based on the marginal probabilities and correlations, as follows:




Figure 2: Example network for computational experiments.

Model INDEPM This model considers only copulas of size 1, with the corresponding marginal failure probabilities obtained from the physical model. This model mimics the assumption of independence between failures.

Model COLGENM This model considers both the empirical marginal failure probabilities and the correlations between failures from the physical model. We obtain an exact solution for system (1) using the proposed column generation algorithm, as explained in Section 3.1.

Models SC-2, SC-3 and SC-4 Following the discussion in Section 3.2, we solve the problem defined in (6) by considering only copulas that share a node, with sizes of at most 2, 3 and 4 components, respectively.

Model COLGENM-4 To demonstrate the relevance of the proximity between the links in each copula, we solve problem (6) via column generation while including a constraint in the pricing problem that limits the maximum copula size to 4 links. Note that this approach produces small copulas (with few components), but the distances between the components of each copula are not necessarily small.

For all models, we simulate multiple scenarios and evaluate the reliability by computing the existence of an operative path between each pair of nodes in the network for each scenario. A total of 10^6 simulations yield errors bounded by 10^{-3} .

The results are presented in Figure 3, where the color of each cell in the figure represents the resulting reliability between each pair of nodes in the network. Additionally, in Table 1, we compute the mean difference and the maximum difference (in both absolute and relative terms) between the physical model and each other model.

Table 1: Differences in reliability between the physical model and the other models.

| | INDEPM | COLGENM | SC-2 | SC-3 | SC-4 | COLGENM-4 |
|---------------------|--------|---------|--------|--------|--------|-----------|
| Mean absolute diff. | 0.0742 | 0.0116 | 0.0267 | 0.0163 | 0.0115 | 0.0539 |
| Mean relative diff. | 0.0848 | 0.0135 | 0.0302 | 0.0184 | 0.0132 | 0.0621 |
| Max. absolute diff. | 0.1006 | 0.0326 | 0.0631 | 0.0369 | 0.0257 | 0.1087 |
| Max. relative diff. | 0.1198 | 0.0388 | 0.0689 | 0.0403 | 0.0285 | 0.1293 |

As seen, both COLGENM and SC-4 provide very good approximations of the reliabilities obtained using the physical model, with mean absolute differences in reliability of close to 1%. This indicates




Figure 3: Reliabilities between all node pairs for the different models.

that a model with copulas composed of a small number of connected components can approximate the physical model, although the parameters only approximate the right-hand side of the equations in (1). Moreover, the reliabilities produced by COLGENM-4 indicate that small-sized copulas are not a guarantee of a good approximation. This suggests that the connectivity of the components in the copulas is probably the key property that determines whether good approximations can be obtained, which is consistent with the observations reported in the literature discussed in Section 3.2. We can also see that the INDEPM model, which ignores the correlations, overestimates the reliability of the network, with differences of up to 10% compared with the physical model. Finally, we note that increasing the size of the connected copulas from 2 to 4 (models SC-2, SC-3 and SC-4) also enhances the quality of their approximations.

5 CONCLUSIONS

It is important to consider the dependencies between component failures in networks. We show that ignoring these correlations can considerably affect computations of network reliability. We also show that by using techniques such as column generation, it is possible to calibrate an MO copula model to obtain the required marginal failure probabilities and failure correlations. However, different calibration solutions can lead to considerable differences in the resulting reliabilities. By calibrating a physical model, we find that, as suggested by the literature, considering the simultaneous failures of small and connected sets of components provides a good approximation of the reliability, even if the obtained correlations are not exactly identical to those in the physical model. Further experiments should be conducted using different networks and failure models to confirm this finding. Finally, we remark that knowledge of the correlations is equivalent to knowledge of the probability that two links will fail at the same time. The proposed methodology could be easily extended if more information were available, such as the probability that three or more links will fail at the same time, but the empirical values required to calibrate an MO model of this type can be more difficult to obtain in real networks.

ACKNOWLEDGMENTS

The authors acknowledge the financial support from Fondecyt grants 1130681 and 1161064, INRIA-CIRIC and Programa Iniciativa Científica Milenio NC130062.

REFERENCES

- Agarwal, P. K., A. Efrat, S. K. Ganjugunte, D. Hay, S. Sankararaman, and G. Zussman. 2013. “The resilience of WDM networks to probabilistic geographical failures”. *IEEE/ACM Transactions on Networking (TON)* 21 (5): 1525–1538.
- Barrera, J., H. Cancela, and E. Moreno. 2014. “Topological optimization of reliable networks under dependent failures”. *Operations Research Letters* 32 (2): 132–136.
- Botev, Z. I., L. Pierre, R. Simard, and B. Tuffin. 2016. “Static network reliability estimation under the Marshall-Olkin copula”. *ACM Transactions on Modeling and Computer Simulation* 26 (2): 14.
- Dantzig, G. B., and P. Wolfe. 1960. “Decomposition principle for linear programs”. *Operations Research* 8 (1): 101–111.
- Elperin, T., I. Gertsbakh, and M. Lomonosov. 1991. “Estimation of network reliability using graph evolution models”. *IEEE Transactions on Reliability* 40 (5): 572–581.
- Gill, P., N. Jain, and N. Nagappan. 2011. “Understanding network failures in data centers”. In *Proceedings of the ACM SIGCOMM 2011 Conference*, Volume 41, 350–361. Association for Computing Machinery.
- Gonzalez, A. J., B. E. Helvik, J. K. Hellan, and P. Kuusela. 2010. “Analysis of dependencies between failures in the UNINETT IP backbone network”. In *2010 IEEE 16th Pacific Rim International Symposium on Dependable Computing*, 149–156. Institute of Electrical and Electronics Engineers, Inc.
- Gutenberg, B., and C. Richter. 1954. *Seismicity of the Earth and Associated Phenomena*. 2nd ed. Princeton University Press.
- Hagstrom, J. N., and S. Ross. 2001. “Component State Dependence and Error in Reliability Computation”. Available in <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.20.8460>.
- Kroese, D. P., T. Taimre, and Z. I. Botev. 2013. *Handbook of Monte Carlo Methods*, Volume 706. John Wiley & Sons.
- Markopoulou, A., G. Iannaccone, S. Bhattacharyya, C.-N. Chuah, Y. Ganjali, and C. Diot. 2008. “Characterization of failures in an operational IP backbone network”. *IEEE/ACM Transactions on Networking (TON)* 16 (4): 749–762.
- Marshall, A. W., and I. Olkin. 1967. “A multivariate exponential distribution”. *Journal of the American Statistical Association* 62 (317): 30–44.
- Provan, J. S., and M. O. Ball. 1983. “The complexity of counting cuts and of computing the probability that a graph is connected”. *SIAM Journal on Computing* 12 (4): 777–788.
- Singpurwalla, N. D. 2002. “Dependence in network reliability”. In *Proceedings of the 5th International Conference on Information Fusion, FUSION 2002*, edited by M. Morelli and A. J. DeSimone, Volume 2, 981–985. International Society of Information Fusion.
- Turner, D., K. Levchenko, A. C. Snoeren, and S. Savage. 2010. “California fault lines: understanding the causes and impact of network failures”. In *Proceedings of the ACM SIGCOMM 2010 Conference*, Volume 40, 315–326. Association for Computing Machinery.

AUTHOR BIOGRAPHIES

OMAR MATUS is a Project Engineer at Universidad Adolfo Ibáñez and a part-time lecturer. He holds a Master of Science degree in Industrial Engineering from Universidad Adolfo Ibáñez. His master’s thesis studied dependent failures in telecommunication networks, which is the basis of this paper. His email address is omar.matus@uai.cl.

JAVIERA BARRERA is an Assistant Professor with the Faculty of Engineering and Sciences, Universidad Adolfo Ibáñez, Chile. She received her Ph.D. from Université René Descartes-Paris V and Universidad de Chile in 2006. Her research interests lie in the asymptotic behavior of stochastic models and, more recently, in stochastic optimization. Applications in telecommunications, computer science, physics and

finance motivate the models addressed in her research. Her email address is javiera.barrera@uai.cl.

EDUARDO MORENO is an Associate Professor with the Faculty of Engineering and Sciences, Universidad Adolfo Ibáñez, Chile. He received his Ph.D. from Université Paris-Est Marne-la-Vallée and from Universidad de Chile in 2005. His research focuses on combinatorial optimization, algorithms and computational methods, with applications in network design, telecommunications, transportation and open-pit mining. His email address is eduardo.moreno@uai.cl.

GERARDO RUBINO is a senior researcher at INRIA (the French National Institute for Research in Computer Science and Control), where he is the leader of the DIONYSOS (Dependability, Interoperability and performance analysis of networks) team. His research interests lie in the quantitative analysis of computer and communication systems, mainly using probabilistic models. He also works on the quantitative evaluation of perceptual quality of multimedia communications over the Internet. With Bruno Tuffin, he co-edited the book *Rare event simulation using Monte Carlo methods*, published by John Wiley & Sons in 2009, and co-authored several of its chapters. He has published more than two hundred papers in journals and conference proceedings, in several fields of applied mathematics and computer science, and has performed various editorial tasks and managing activities in research. He is a member of the IFIP WG 7.3. His email address is gerardo.rubino@inria.fr.