Topological optimization of reliable networks under dependent failures

Javiera Barrera[a], Héctor Cancela[b], Eduardo Moreno[a]

[a]*Faculty of Engineering and Sciences, Universidad Adolfo Ibáñez, Santiago, Chile.*
[b]*Facultad de Ingeniería, Universidad de la República, Montevideo, Uruguay.*

Abstract

We address the design problem of a reliable network. Previous work assumes that link failures are independent. We discuss the impact of dropping this assumption. We show that under a common-cause failure model, dependencies between failures can affect the optimal design. We also provide an integer-programming formulation to solve this problem. Furthermore, we discuss how the dependence between the links that participate in the solution and those that do not can be handled. Other dependency models are discussed as well.

Keywords: Common-cause failure, Dependent failure, Reliable network design, Sample average approximation.

1. Problem description

The topological design of reliable telecommunications networks has been deeply studied throughout the last 30 years. The problem can be stated as follows: given a set of nodes and a set of potential links between these nodes, we must choose a subset of links to install such that the total cost is minimized and the reliability is maximized. For a good description of this problem and the different types of reliability requirements, we refer to [1].

The reliability of a network can be defined in different ways. The most common measure of reliability is $\mathcal{K}$-connectivity. That is, given a set $\mathcal{K}$ of terminal nodes, the reliability of the network is the probability that there exists a path from every node in $\mathcal{K}$ to every other node in $\mathcal{K}$. When $\mathcal{K}$ includes all nodes in the network then this measure is called the *all-terminal reliability* and when $\mathcal{K}$ is a specific pair of nodes, it is called the *source-terminal reliability* (or *s-t reliability*). To make the computation of reliability affordable, several simplifications over the failures in a network are made. The most common simplifications are as follows: (a) the link failures are independent; (b) the nodes are perfectly reliable; and (c) no repair is allowed. Note that even under

Email addresses: `javiera.barrera@uai.cl` (Javiera Barrera), `cancela@fing.edu.uy` (Héctor Cancela), `eduardo.moreno@uai.cl` (Eduardo Moreno)

these assumptions, to compute the reliability of a given network is a #P-hard problem [2, 3].

Formally, let $\mathcal{G} = (N, E)$ be a graph with node set N and link set E. Let U be a random binary vector taking values in $\{0, 1\}^E$, representing which links are operational. Given U, let E_U be the set of all links that are operational; then, the observed network is given by the graph $\mathcal{G}_U = (N, E_U)$. The $\mathcal{K}$*-reliability* of $\mathcal{G}$ is defined as the probability that $\mathcal{G}_U$ is $\mathcal{K}$-connected.

We study the following problem: given a cost c_e for each $e \in E$ and a budget B, we want to select a subset of links $F \subseteq E$ of total cost less than or equal to the budget, such that the reliability of the selected subnetwork is maximized:

$$\max_{F \subseteq E} \mathbb{P}((N, F_U) \text{ is } \mathcal{K}\text{-connected})$$
$$\sum_{e \in F} c_e \leq B$$

Due to the difficulty associated with the computation of the reliability of a network, the only known methods to exactly solve this problem are based in the enumeration of the possible solutions [4, 5]. This is possible only on small-sized networks and for particular cases, such as when all links have the same probability of failure. Hence, authors have focused on different heuristics and approximation techniques. For example, using Tabu Search [6], Simulated Annealing [7], Genetic algorithms [8, 9] or Neural Networks [1, 10]. These methods give approximate solutions without any guarantee of convergence to optimality. A recent approach, proposed in [11], employs a sample of failure event scenarios to determine the optimal topology. This technique, called the *sample average approximation* (SAA), converges to an optimal solution when the number of samples is sufficiently large.

However, recent studies question the neglect of the dependence between failures. In [12], authors analyze real data from a Norwegian academic network, showing that neighboring links show significant correlation. In particular, for the design of reliable networks, the impact of dependent failures on the resulting reliability is analyzed in [13], showing that the assumption of failure independence can produce an underestimation of the real reliability. There are several approaches to model the failure dependence between components, including causal failure, cascade failure and common-cause failure. For a discussion of the tractability and scalability of these different models, we refer the reader to [14].

In this paper, we study the topological optimization problem under the best-established common-cause failure model [15]. In Section 2, we present the Marshall-Olkin copula model, which supports common-cause failure dependency. In Section 3, we present a SAA model to solve the problem, and in Section 4, we present some extensions of the previous models to consider other dependent failures models. Finally, we present a computational example in Section 5, showing that ignoring correlation can lead to suboptimal solutions.

2. The Marshall-Olkin copula model for common-cause failures

Common-cause failures are a subset of dependent events in which two or more component fault states exist at the same time and are either direct results or a shared cause [16]. These failures arise naturally in several contexts, for example, in an overlay (virtual) network that is connected through an underlying physical network, so a failure in the physical layer could affect several components of the overlay network. Another example is a network in which the components share equipment that is essential for their function. The Marshall-Olkin (MO) copula introduced in [17] is one of the best-established models for common-cause failures. In this model, events cause one or more components to fail simultaneously, but the lifetime of each link remains exponentially distributed. This model was used in [18] for evaluating the reliability of a network using an importance sampling technique to generate samples of correlated failures.

Formally, let $\mathcal{G} = (N, E)$ be a network. At time zero, we assume that all links are operational. As time passes, links start to fail (alone or simultaneously). Because no repair is allowed, fewer links operate over time. For each link, we can define the lifetime V_e that is the instant at which link e fails. Let $\mathcal{P}_0^E$ be a collection of non-empty subsets of E and let $(W_D)_{D\in\mathcal{P}_0^E}$ be a family of independent positive random variables. The time W_D represents the instant in which a failure that affects all links in D occurs. Therefore, the lifetime of link e is the first time that a set D containing e fails. That is,

$$V_e = \min_{D:e\in D}\{W_D\}.$$

When W_D's are exponential random variables, then this coupling is known as the MO copula. Note that the marginal distribution of the lifetime V_e is also exponential for all e. Let $U_e(t)$ be the state of link e at time t; hence, we have $U_e(t) = \mathbb{I}_{(V_e\leq t)}$. Let $U(t) = (U_e(t))_{e\in E}$, then the graph $\mathcal{G}_U(t) = (V, E_{U(t)})$ is the observed network at time t. Because we are interested in a static model, we take a snapshot of the network at time 1 and evaluate the reliability at that instant (this can also correspond to a dynamic model with a fixed mission time). We fix the status of each link $U_e := U_e(1)$.

Note that this is a natural extension of the independent failure model, setting $\mathcal{P}_0^E$ as the collection of singleton sets, and V_e as exponential random variables of parameter $\log(1-p_e)$. In this particular case, we have independent failure probabilities between links, and at time 1, the failure probability is p_e for link e. We can also recover the model in which the nodes and links fail independently, adding to the previous model the collection of sets $\{D_n : n \in N\}$ such that D_n is the set of all links with n as an end node. In this case, if all nodes fail independently with probability q and the links fail independently with probability p, then the marginal probability failure for each link is $(1-p)(1-q)^2$, and the node failures induce a correlation between adjacent links of $\frac{q(1-p)}{1-(1-p)(1-q)^2}$.

3. An integer programming model using sample average approximations

In this section, we present an integer programming model that considers the dependencies between links, given by a MO copula. This model is related to the ideas of [11], which is based on a sample average approximation (SAA) of the probability that the network is $\mathcal{K}$-connected.

SAA is a popular technique for approximating the stochastic objective function. Some of the first applications of this technique appeared in [19], and the approach was empirically studied in [20]. Recalling that the probability of an event is equal to the expected value of the indicator function of the event, the basic idea is to approximate the expected value by its sampled average. That is, we sample a set of scenarios S, and we approximate the objective function by

$$\begin{aligned}\mathbb{P}\left((N,F_x) \text{ is } \mathcal{K}\text{-connected}\right) &= \mathbb{E}\left(\mathbb{I}_{(N,F) \text{ is } \mathcal{K}\text{-connected}}\right) \\ &\approx \frac{1}{|S|}\sum_{s\in S}\mathbb{I}_{(N,F_x^s) \text{ is } \mathcal{K}\text{-connected}},\end{aligned}$$

where F_x^s is the subgraph obtained after removing from F_x the failed links in scenario s.

To formulate an integer programming model that uses the previous approximation, we define a binary variable w_s that indicates the event that the graph (N, F_x^s) is $\mathcal{K}$-connected (or not) under scenario s. A $\mathcal{K}$-cut is a subset $M \subseteq N$ such that $\mathcal{K} \cap M \neq \emptyset$ and $\mathcal{K} \setminus M \neq \emptyset$. Recall that a graph $G = (N, E)$ is $\mathcal{K}$-connected if and only if for every $\mathcal{K}$-cut M the cut-set induced by M, $\delta(M) = \{uv \in E : u \in M, v \notin M\}$, is not empty. Finally, let $\{W_D^s \ : \ D \in \mathcal{P}_0^E,\ s \in S\}$ be a sampling of the MO copula. Then, we solve the following integer programming problem:

$$\begin{aligned}
&\max \sum_{s\in S} z_s && && (1)\\
&\sum_{e\in E} c_e x_e \le B && && (2)\\
&\sum_{e\in\delta(M)} u_e^s \ge z_s && \forall M\ \mathcal{K}\text{-cut}, \forall s\in S && (3)\\
&u_e^s \le x_e && \forall e\in E, \forall s\in S && (4)\\
&u_e^s \le 0 && \forall e\in D, \forall s\in S \text{ such that } W_D^s < 1 && (5)\\
&x_e \in \{0,1\} && \forall e\in E &&\\
&z_s \in \{0,1\} && \forall s\in S &&\\
&u_e^s \in \{0,1\} && \forall s\in S, \forall e\in E &&
\end{aligned}$$

Binary variables x_e for each link $e \in E$ determine the resulting network, where constraint (2) bounds the total cost of these links. Variables u_e^s represent whether or not link e is operative under scenario s. To do so, constraints (4) and (5) force $u_e^s = 0$ in the case that link e is not chosen ($x_e = 0$), or whether one of the copula associated to link e indicates that this link fails in scenario s ($W_D^s < 1$). Finally, constraint (3) verifies whether the resulting network is $\mathcal{K}$-connected under scenario s or not. To verify this, note that binary variables z_s can take a value of 1 (which is desired by the objective function), only if for every $\mathcal{K}$-cut there is at least one link operative in scenario s. Note that it is possible to eliminate variables u_e^s by replacing constraints (3)-(5) by $\sum_{e\in\delta(M)} \mathbb{I}_{(W_D^s\geq 1)} x_e \geq z_s$, obtaining a similar model to the one presented in [11]. However, we adhere to the model including explicitly the u_e^s variables because we will discuss some extensions in which these variables are required in the model.

It should be noted that, in this IP formulation, dependence between the chosen and unchosen links in the solution still exists, which is reasonable for many cases. Nevertheless, there are other cases in which dependence with unchosen links should be ignored. An example case is when common-cause failures arise from the interactions between two components. Therefore, if one of these components is not constructed, this dependence should be ignored. In a MO copula model in which dependence is observed only between a pair of links $\{e, e'\}$, the failure event $\{W_{\{e,e'\}} < 1\}$ for this pair will be ignored if one link has not been constructed. This behavior can be added to the IP formulation by replacing constraint (5) for $D = \{e, e'\}$ by

$$u_e^s + x_{e'} \leq 1 \quad \forall e, e' \in E, e \neq e', \forall s \in S \text{ such that } W^s_{\{e,e'\}} < 1. \tag{6}$$

These cover-inequalities can be generalized for sets D of arbitrary sizes, by

$$u_e^s + \sum_{e'\in D\setminus\{e\}} x_{e'} \leq |D| - 1 \quad \forall e \in D, \forall s \in S \text{ such that } W_D^s < 1. \tag{7}$$

Even if these formulations result in an exponentially sized model, they can be implemented efficiently and be used to solve mid-sized instances of the problem. For example, we can relax the integrality condition of variables u_e^s and z_s, obtaining a problem with only $|E|$ binary variables. Additionally, constraint (3) can be added in a cutting-plane scheme, where the separation problem is to find, for a given scenario s, a $\mathcal{K}$-cut of G with edge-weights equal to the incumbent solution u_e^s, such that the cut-weight is less than the value of z_s. This can be done solving a min $\mathcal{K}$-cut, which can be done in $\mathcal{O}(|N||E| + |N|^2 \log |N|)$ time [21]. Moreover, this separation problem can be applied only when an integer solution has been found in the branch-and-bound tree, by solving a breadth-first search over the incumbent solution, which can be done in linear time. Additionally, complex constraints such as (6) and (7) appear in small numbers because these events are rare. Code that incorporates these features can be downloaded from `http://emoreno.uai.cl/code/saa-reliability/` for testing purposes. Note that in this model it is also possible to consider lifetimes with distributions different from the exponential distribution.

4. Extension to the causal-failure model

In this section, we discuss how to extend the previous model to incorporate other kinds of dependency models between links, in particular, the causal failure model. Causal failures between two components arise when the failure of a first component affects, starting from that moment on, the chances that the surviving component fails. Many approaches are available to formalize this concept, such as the one discussed in [15], although there is no consensus on which model is the best. As suggested in the previous work, Freund's model for dependence can capture this effect. In this model, we have that, while link e is operational, link e' fails with rate $\lambda_{e'}$. After link e fails, the failure rate of the surviving link e' increases to $\lambda_{e'} + \Delta\lambda_{e \Rightarrow e'}$. For details on the joint distribution and properties of the model, see [22]. This model can be easily simulated with three independent exponential random variables W_e, $W_{e'}$ and $W_{e \Rightarrow e'}$ with rates λ_e, $\lambda_{e'}$ and $\Delta\lambda_{e \Rightarrow e'}$, respectively. Hence, the lifetime of link e is

$$V_{e'} = \min\{W_{e'}, W_e + W_{e \Rightarrow e'}\}.$$

We can extend the IP formulation presented in the previous section to represent this causal-failure model. As in the previous section, if the causal effect of link e on link e' is to be considered even when e is not included in the network, then we add the following constraint to the IP formulation:

$$u^s_{e'} \leq 0 \quad \text{if } W^s_e + W^s_{e \Rightarrow e'} < 1.$$

If causal failure should be ignored when link e is not in the solution, then we add this other constraint to the IP formulation:

$$u^s_{e'} \leq 1 - x_e \quad \text{if } W_e + W^s_{e \Rightarrow e'} < 1.$$

Note that these pairwise constraints can be added for several pairs, allowing the incorporation of complicated one-to-many cause failures or failures that are similar in the same way.

5. Computational Example

In this section, we present a computational implementation of the IP model over a medium-sized network. We also show that ignoring the dependence between link failures can lead to suboptimal solutions of our network design problem.

We tested our model over the Italian WDM backbone network, extracted from [23], where the marginal failure probabilities p_e are computed as a function of the length of each link. Inspired by [12], we introduced a correlation using the distance d_{ef} between the midpoint of a pair of links e and f, given by $\rho_{ef} = \exp(-\theta d_{ef}))$ rounded to the first decimal and using $\theta = 2$. This formula obtains correlations up to 0.5. In order to compute parameters for the MO copulas λ_s that support these correlations and marginal probabilities of failures, we solved the following system of linear equations

$$\sum_{D\supseteq\{e\}} \lambda_D = -\log(1-p_e) \qquad \forall e\in E$$

$$\sum_{D\supseteq\{e,f\}} \lambda_D = \log\left(1+\rho_{ef}\sqrt{\frac{p_e p_f}{(1-p_e)(1-p_f)}}\right) \qquad \forall e,f\in E,\ e\neq f$$

$$\lambda_D \geq 0 \qquad \forall D\subseteq N,$$

obtaining a set of 114 copulas with positive parameters containing between 1 and 3 links. We set the cost of each link equal to its length in kilometers and the budget to $B = 3500$ kms. We ran our IP formulation assuming independence between failures and using the MO copulas computed previously, with 3000 samples for each case. Finally, we repeated the execution of these models 10 times. The average running time of each instance was 29.3 minutes.

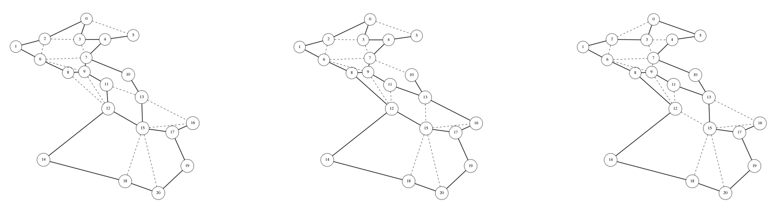

(a) Instance I0 (b) Instance I1 (c) Instances I2,I3 and I4

Figure 1: Optimal solutions obtained by the IP model for the different instances.

	Independent		Copula	
Instance	Object.	Estimated	Object.	Estimated
I0	0.90542	0.87381 ± 0.00207	0.87591	0.87381 ± 0.00207
I1	0.90542	0.86473 ± 0.00207	0.87569	0.87776 ± 0.00206
I2	0.90542	0.85401 ± 0.00208	0.87343	0.87407 ± 0.00206
I3	0.90542	0.84588 ± 0.00209	0.87320	0.87282 ± 0.00207
I4	0.90542	0.79534 ± 0.00211	0.83106	0.83083 ± 0.00210

Table 1: Average objective values and Wilson's 95% confidence interval for the reliability of the solutions obtained by the IP model using original and modified copula parameters.

Different samples could lead to different solutions, but the most common solution obtained (and the one with best objective) coincides for both models, and it is presented in Figure 1a. We also present in Table 1 (instance I0) the average of the obtained objective function of the model, and a 95% Wilson's confidence interval of reliability of the solution, computed using 1 million samples of the MO copula. Note that even if both models obtain the same solution, the obtained objectives differ considerably. Moreover, the model that only considers the marginal failure probabilities overestimates the reliability of the network, but the model using the MO copulas obtains objectives very close to the estimated reliability of the resulting network.

In order to show that the dependencies can affect the final topology, we construct new instances by sequentially modifying correlations. First, we construct instance I1 from instance I0 by increasing the correlation between links $(7, 10)$ and $(9, 11)$ to the maximum allowable, keeping the same marginal failure probabilities for all links. We repeat this procedure increasing correlation between links $(0, 2)$ and $(0, 3)$, links $(6, 8)$ and $(3, 7)$, and links $(12, 14)$ and $(14, 18)$. We denote these instances as I2, I3 and I4, respectively. Note that, for all instances, the model assuming independence between failures is exactly the same than before, obtaining the same solution and objective. Solutions obtained by our model are presented in Figure 1b-1c. As we can see, the new solution of I1 avoids using link $(7, 10)$ in order to avoid this increased correlation, and it replaces two other links to keep the three-cycles structure of the solution within budget. The same behaviour is observed for I2. In Table 1, we also present the obtained objective and estimated reliability of the solution obtained by each model, under the modified MO copula parameters. As we can see, the model assuming independence between failures obtains the same solution as before, but its actual reliability again differs considerably from the average objective value. On the other hand, the model using MO copulas obtains better solutions (up to 5% more reliable), with objective values very close to the estimated reliability for all instances.

This example shows that if dependency among failures is neglected, then it could lead to suboptimal solutions. Thus, in order to attain a larger reliability, we may choose less reliable routes if doing so avoids significant correlation.

6. Conclusions

Reliable network design models in past literature have ignored the dependency between link failures, mainly due to the difficulty that dependency represents when solving problems. However, as shown in this work, dependency between failures can affect the optimal design, and ignoring this dependency can lead to suboptimal solutions. This is a key aspect of the design because empirical evidence shows that correlation between link failures exists and it is significant. Considering this point, we introduce an integer programming model in which this type of dependency between failures can be addressed using simulation, allowing us to solve this problem efficiently for a variety of failure dependency models. To implement this IP model there is still a necessary step: to choose and to calibrate a correct failure model. This crucial decision, that should be addressed before using our IP model, depends in the specific problem that need to be analyzed, and can be a challenging task by itself.

Acknowledgments

Authors acknowledge the financial support of CIRIC (INRIA-Chile), Anillo ACT-88, Programa Iniciativa Cientifica Milenio NC130062 (J.B) and FONDE-

CYT's grants 1100618 (J.B) and 1130681 (E.M.).

References

[1] A. Konak, A. E. Smith, Network reliability optimization, in: M. G. C. Resende, P. M. Pardalos (Eds.), Handbook of Optimization in Telecommunications, Springer US, 2006, pp. 735–760. doi:10.1007/978-0-387-30165-5_26.

[2] M. Ball, Computational complexity of network reliability analysis: An overview, IEEE Transactions on Reliability 35 (3) (1986) 230–239. doi:10.1109/TR.1986.4335422.

[3] J. Provan, M. Ball, The complexity of counting cuts and of computing the probability that a graph is connected, SIAM Journal of Computing 12 (4) (1983) 777–788. doi:10.1137/0212053.

[4] K. Aggarwal, S. Rai, Reliability evaluation in computer-communication networks, IEEE Transactions on Reliability 30 (1) (1981) 32–35. doi:10.1109/TR.1981.5220952.

[5] R. Jan, F. Hwang, S. Cheng, Topological optimization of a communication-network subject to a reliability constraint, IEEE Transactions on Reliability 42 (1) (1993) 63–70. doi:10.1109/24.210272.

[6] S. J. Koh, C. Y. Lee, A tabu search for the survivable fiber optic communication network design, Computers & Industrial Engineering 28 (4) (1995) 689–700. doi:10.1016/0360-8352(95)00036-Z.

[7] S. Pierre, J.-M. Bourjolly, M.-A. Hyppolite, Application of the method of simulated annealing to the topological design of remote computing networks, in: Electrical and Computer Engineering, 1993. Canadian Conference on, 1993, pp. 729–732 vol.2. doi:10.1109/CCECE.1993.332400.

[8] B. Dengiz, F. Altiparmak, A. Smith, Efficient optimization of all-terminal reliable networks, using an evolutionary approach, IEEE Transactions on Reliability 46 (1) (1997) 18–26. doi:10.1109/24.589921.

[9] D. L. Deeter, A. E. Smith, Economic design of reliable networks, IIE Transactions 30 (1998) 1161–1174. doi:10.1023/A:1007584419484.

[10] H. AboElFotoh, L. Al-Sumait, A neural approach to topological optimization of communication networks, with reliability constraints, IEEE Transactions on Reliability 50 (4) (2001) 397–408. doi:10.1109/24.983401.

[11] Y. Song, J. Luedtke, Branch-and-cut approaches for chance-constrained formulations of reliable network design problems, Mathematical Programming Computation 5 (4) (2013) 397–432. doi:10.1007/s12532-013-0058-3.

[12] A. Gonzalez, B. Helvik, J. Hellan, P. Kuusela, Analysis of dependencies between failures in the UNINETT IP backbone network, in: Dependable Computing (PRDC), 2010 IEEE 16th Pacific Rim International Symposium on, 2010, pp. 149–156. doi:10.1109/PRDC.2010.12.

[13] O. Lapcevic, M. Lackovic, B. Mikac, Impact of dependent failures on the availability of the optical network, in: Communication Systems, Networks and Digital Signal Processing, 2008. CNSDSP 2008. 6th International Symposium on, 2008, pp. 423–427. doi:10.1109/CSNDSP.2008.4610799.

[14] N. D. Singpurwalla, Dependence in network reliability, in: Proceedings of the Fifth International Conference on Information Fusion, Vol. 2, IEEE, 2002, pp. 981–985. doi:10.1109/ICIF.2002.1020918.

[15] D. V. Lindley, N. D. Singpurwalla, et al., On exchangeable, causal and cascading failures, Statistical Science 17 (2) (2002) 209–219. doi:10.1214/ss/1030550862.

[16] A. Mosleh, K. Fleming, G. Parry, H. Paula, D. Worledge, D. M. Rasmuson, Procedures for treating common cause failures in safety and reliability studies, Tech. rep., Electric Power Research Insttute, Palo Alto, CA (USA) (1988).

[17] A. W. Marshall, I. Olkin, A multivariate exponential distribution, Journal of the American Statistical Association 62 (317) (1967) 30–44.

[18] Z. I. Botev, P. L'Ecuyer, B. Tuffin, Dependent failures in highly reliable static networks, in: Simulation Conference (WSC), Proceedings of the 2012 Winter, IEEE, 2012, pp. 1–12. doi:10.1109/WSC.2012.6465033.

[19] A. J. Kleywegt, A. Shapiro, T. Homem-de-Mello, The sample average approximation method for stochastic discrete optimization, SIAM Journal on Optimization 12 (2) (2002) 479–502. doi:10.1137/S1052623499363220.

[20] J. Linderoth, A. Shapiro, S. Wright, The empirical behavior of sampling methods for stochastic programming, Annals of Operations Research 142 (1) (2006) 215–241. doi:10.1007/s10479-006-6169-8.

[21] M. Stoer, F. Wagner, A simple min-cut algorithm, Journal of the ACM 44 (4) (1997) 585–591. doi:10.1145/263867.263872.

[22] J. E. Freund, A bivariate extension of the exponential distribution, Journal of the American Statistical Association 56 (296) (1961) pp. 971–977.

[23] Z. Pandi, M. Tacca, A. Fumagalli, L. Wosinska, Dynamic provisioning of availability-constrained optical circuits in the presence of optical node failures, Journal of Lightwave Technology 24 (9) (2006) 3268–3279. doi:10.1109/JLT.2006.879505.