Journal of Universal Computer Science, vol. 14, no. 5 (2008), 786-799
submitted: 1/11/07, accepted: 28/2/08, appeared: 1/3/08

A Normal Copula Model for the Economic Risk Analysis of Correlated Failures in Communications Networks

Maurizio Naldi
(Università di Roma "Tor Vergata", Italy
naldi@disp.uniroma2.it)

Giuseppe D'Acquisto
(Università di Roma "Tor Vergata", Italy
giuseppe.dacquisto@inwind.it)

Abstract: The reliability of a communications network is often evaluated without taking into account the economic consequence of failures. Here a new approach is proposed to assess the economic consequences of failures as a figure of merit of reliable networks. For this purpose a partition of the network operator's market into service basins is proposed, which includes the presence of correlation between the subsystems needed to serve different service basins as well as within the same service basin. A simulation algorithm, based on the Cross-Entropy method, is fully described to evaluate the probability that the economic loss exceeds a given threshold. An application of the method to a simple scenario is finally reported.

Key Words: Reliability, Communications networks, Risk analysis, Economics of networks

Category: C.2, C.4

1 Introduction

Being reliable is one of the main requirements for a communications networks. The number of activities that are supported by a network (sometimes extremely relevant for human life, such as emergency services or healthcare) is such that we require the network to work properly for an extremely large proportion of time, which in turn means that failures must be rare and quickly fixed. Requirements for a highly reliable network are typically met by a combination of strategies, e.g. by deploying highly reliable network devices, by devising a network architecture which is tolerant to failures of its individual components, and by establishing a capillary and reactive maintenance team. A large body of literature has been devoted to reliability and to the associated issues of availability and dependability (this latter meaning reliability in a broad sense, incorporating both the availability and the reliability concepts, see [ITU 1994] for the pertaining definitions). A sample of the most covered areas of interest is given by the following list (a rather exhaustive survey is provided in [Helvik 2004]):

– Reliability evaluation, i.e. the computation of the failure probability with the aim of providing a performance index for a network/service architecture and driving the design process;

- Fault detection, with the aim of obtaining a fast and precise identification of the occurrence of faults, the devices involved, and their location;
- Fault recovery, with the aim of returning the network to a regular state as fast as possible, and possibly in a totally automated way (self-healing networks).

Despite advances in the design of components and networks, reliability is still a hot issue. In particular, large-scale outages occur more than ever in both data and telephone networks; some examples are reported in [Snow 2001]. A number of factors can be thought of as responsibles for such persistence. In fact, present network design is often driven by the need to deploy services as fast as possible and with costs as low as possible. This is often achieved e.g. through:

- Use of modular software, with an evergoing stream of upgrades and patches;
- Use of COTS (Commercial Off-The-Shelf) products;
- Massive resort to widespread upgrades;
- Employment of cost-abatement procurement procedures;
- Reduction of functional redundancy.

In addition, networks are increasingly built as the connection of a collection of independently designed and often distributed modules, whose number is growing, and with an improper use of diversity mechanisms to achieve reliability [Snow 2001]. As a result failures do occur at the same time on multiple devices and points in the network, which often appear as a giant domino effect. An issue of relevance in reliability analyses is therefore the consideration of correlated failures and of their common causes.

Another issue that is typically missing in reliability studies is the economic impact of a network fault (with rare exceptions, see e.g. [Akeson et al. 1994] and [Kogeda et al. 2005]). When a network device fails, the operator incurs an economic loss, generally even if the service is kept ongoing. The economical losses can in fact be ascribed to these categories:

1. Lost revenues;
2. Penalties for breach of SLA (Service Level Agreement) conditions;
3. Recovery costs;

of which the third one is present even if the customer is not affected by the network failure. However, if we consider the operations of a service provider to be driven by profit (as is typically the case), the economic impact of failures should be the prime factor in determining the way the operator approaches reliability. Under this viewpoint, not all faults are equal: those having a larger economic impact should deserve more attention. Reliability performance should therefore be evaluated not by an index such as the probability that the service is maintained, but rather by the resulting economical losses with their probability.

In this paper we introduce an economical loss function associated to network failures and propose a simulation approach to evaluate the probability that the losses exceed a tolerable threshold. We accomplish that by considering a normal copula model (first established in [Gupton et al. 1997] in the context of credit risk modelling and described in Section 2) to represent the correlation between the failures of different devices. Some considerations on the open issues concerning the concrete application of the normal copula model are provided in Section 4. Resorting to simulation is unavoidable because we are dealing with rare but correlated events concerning a very large number of devices. As a simulation technique we have opted for the Cross-Entropy method, introduced by Rubinstein [Rubinstein and Kroese 2004], whose application to the normal copula model is detailed in Sections 3 and 5 (the latter reporting a toy example).

2 A normal copula model for the loss function

In this paper we propose a model for the evaluation of economic losses associated to failures in a communications network. The model revolves around the customers and the services affected by the faults rather than the network elements. The focus is therefore not on the reliability itself but rather on the economical consequences of poor reliability. It can be said that not all faults are equal: those having a deeper impact on the operator's revenues should be addressed first.

For this purpose we introduce the concept of service basin, which can be defined as a rectangular partition of the two-dimensional space of customers and services. A service basin is to be considered disrupted if its customers are not receiving the subscribed services. A representation of a very simple service basin partition is provided by an operator serving two areas and providing triple play services (i.e. voice, TV, and Internet), so that we can identify $3 \cdot 2 = 6$ service basins. Of course the partition can be as fine as desired (or possible), considering even single customers and a more detailed set of services. The state of the i-th service basin is represented by the binary variable B_i, where $B_i = 0$ if the basin is correctly served and 1 otherwise. The disruption of service to the i-th basin causes the economic loss a_i. Though the loss may be actually dependent on the duration of the disruption (which may be random), for the purpose of this model we consider to be able to associate a single loss value to the disruption (e.g. as if the duration of the disruption were known) and to perform the economical loss evaluation over a pre-defined time horizon (e.g. one month). We are not concerned with the details about the frequency of disruptions and their duration but rather describe the overall economic loss due to the whole set of fault events taking place within the time interval of interest. If we have N_b service basins the overall economic loss is therefore

$$L = \sum_{i=1}^{N_b} a_i B_i. \tag{1}$$

We consider that the service to the i-th basin is disrupted if any of a set of subsystems is faulty, though more complex relationship (involving e.g. the simultaneous failure of two or more components) could be envisaged. A given subsystem may however impact several service basins. We indicate by M_i the number of subsystems impacting the i-th service basin. Unlike the approach typically adopted that considers the subsystems as independent, we instead take the much more realistic assumption that the faults of different subsystems may be correlated. We therefore define a model based on the use of M_i latent variables X_{ij} for each service basin (one for each subsystem). Each latent variable is in turn driven by an individual reliability factor η_{ij} and a number of common reliability factors Z_k, $k = 1, \ldots, F$. The common reliability factors allow to model the correlation between the faults of different subsystems due to a common dependence on other exogenous factors, e.g. the influence of the external temperature, of a common control element (whose fault may make all the controlled devices faulty), or the outage of the energy facilities. The relationship between the latent variable of a subsystem and its reliability drivers is the following

$$X_{ij} = \sum_{k=1}^{F} \rho_{ijk} Z_k + \alpha_{ij}\eta_{ij}, \qquad i = 1, \ldots, N_b \qquad j = 1, \ldots, M_i, \tag{2}$$

the weights being subject to the normalizing relation

$$\sum_{k=1}^{F} \rho_{ijk}^2 + \alpha_{ij}^2 = 1. \tag{3}$$

In this model the common reliability factors and the individual reliability factor are i.i.d. standard normal variables, hence the name of normal copula assigned to it, so that the resulting latent variable is itself a standard normal variable. The link between the state of the j-th subsystem in the i-th service basin and its associated latent variable is given by the probability that the subsystem is faulty, equal to

$$P[X_{ij} > b_{ij}] = 1 - G(b_{ij}), \tag{4}$$

where b_{ij} is a suitable threshold and $G(\cdot)$ is the standard Gaussian cumulative distribution function. We can then associate the binary variable Y_{ij} to the state of the single subsystem by the indicator function

$$Y_{ij} = \mathbb{I}\left[X_{ij} > b_{ij}\right], \tag{5}$$

so that $Y_{ij} = 0$ denotes a working subsystem and $Y_{ij} = 1$ a faulty one.
The introduction of the common reliability factors allows us to take into account the correlation between the failures of different subsystems. In fact, the correlation between any two latent variables X_{ij} and X_{lm} is determined by the weights of the common reliability factors in expr. (2):

$$C(X_{ij}, X_{lm}) = \sum_{k=1}^{F} \rho_{ijk}\rho_{lmk}. \tag{6}$$

This model has been introduced first in the context of financial risk analysis where the insolvency of obligors is of interest [Gupton et al. 1997]. It is to be noted that the latent variable approach allows us to consider not just the failure of common causes, but also factors whose influence is gradual (e.g. the rise of temperature or the fluctuations in the energy supply systems) so to make the affected subsystem more prone to failure.
Finally, the service basin state is related to the state of its subsystems by the simple relationship

$$B_i = \max\left(Y_{i1}, \ldots, Y_{iM_i}\right). \tag{7}$$

It is to be noted that, though expr. 7 resembles the reliability relationship valid for a set of independent series-connected devices, it incorporates the correlation between the subsystems through the common reliability factors present in the latent variable associated to each subsystem (see expr. 2).

3 Cross-Entropy simulation algorithm

In this section we describe a method to estimate the probability of large losses by simulation, by applying the Cross-Entropy approach. We will largely draw on the reference formulation of the method by Rubinstein [Rubinstein and Kroese 2004]. Our problem is to estimate the probability that the overall loss exceeds a given threshold

$$\gamma = \mathbb{P}(L > l). \tag{8}$$

Of particular interest is the case where the threshold is quite large. The interest is due to two concurrent reasons: the relevance of large losses for the financial conditions of the operator and the difficulty of estimating the very small probability associated to this rare event. In the financial literature the threshold l is named Value-at-Risk (VaR).

Since the loss L actually depends on a number of random variables, we can group them together in the random vector $\mathbf{S} = \left\{Z_1, \ldots, Z_F, \eta_{11}, \ldots, \eta_{N_b M_{N_b}}\right\}$, made up of $F + \sum_{i=1}^{N_b} M_i$ elements, with the associated vector $\mathbf{u}$ of parameters (in our case made of all the expected values and the variances of the common and individual reliability factors, which are respectively 0 and 1). The probability density function (pdf) for the random vector $\mathbf{S}$ is $f(\mathbf{s}; \mathbf{u})$. An alternative expression for the probability of losses larger than the given threshold is therefore

$$\gamma = \mathbb{P}_{\mathbf{u}}(L > l) = \mathbb{E}_{\mathbf{u}}\left[\mathbb{I}_{\{L(\mathbf{S})>l\}}\right]. \tag{9}$$

After drawing a random sample of size N of the random vector $\mathbf{S}$ we could get a crude MonteCarlo estimation of the probability of interest by the sample average

$$\hat{\gamma}_{MC} = \frac{1}{N}\sum_i \mathbb{I}_{\{L(\mathbf{S}_i)>l\}}. \tag{10}$$

The problem with MonteCarlo estimation is the large variance associated with the low value of the probability to be estimated. Namely the normalized standard error of the MonteCarlo estimator has the well known expression for small values of γ

$$\frac{\sigma_{\hat{\gamma}_{MC}}}{\mathbb{E}\left(\hat{\gamma}_{MC}\right)} = \sqrt{\frac{1-\gamma}{N\gamma}} \simeq \sqrt{\frac{1}{N\gamma}}. \tag{11}$$

This problem can be overcome by resorting to the Importance Sampling (IS) simulation method, where the probability associated to the values of interest (those such that $\mathbb{P}[L > l]$) is artificially increased through the use of a biased pdf $g(\mathbf{s})$. The bias is then recovered by using the IS estimator

$$\hat{\gamma}_{IS} = \frac{1}{N}\sum_i \mathbb{I}_{\{L(\mathbf{S}_i)>l\}} \frac{f(\mathbf{s}; \mathbf{u})}{g(\mathbf{s})}. \tag{12}$$

The extent of the improvements depends on the proper choice of this biased pdf, which has to be such to reduce the variance of the associated estimator. An ideal zero-variance IS estimator would be attained when the biased pdf is

$$g^*(\mathbf{s}) = \frac{\mathbb{I}_{\{L(\mathbf{S}_i)>l\}} f(\mathbf{s}; \mathbf{u})}{\gamma}, \tag{13}$$

which unfortunately depends on the same quantity γ to be estimated.

However, this ideal estimator can be approached by looking for the best biasing pdf within the family $f(\mathbf{s}; \mathbf{v})$, where $\mathbf{v}$ is the so-called tilting parameter vector, such that the distance between this newly defined pdf and the optimal one is minimized. In our case the tilting vector is $\mathbf{v} = \{\mu_{Z_1}, \ldots, \mu_{Z_F}, \mu_{\eta_{11}}, \ldots, \mu_{\eta_{NM_N}}\}$. A suitable measure of distance is the Kullback-Leibler distance, a.k.a. as Cross-Entropy, i.e. the expected value of the logarithm of the ratio of the two pdfs computed under the probability measure provided by the pdf to be approached

$$\mathcal{D}(g,f) = \mathbb{E}_g\left[\ln\frac{g^*(\mathbf{s})}{f(\mathbf{s};\mathbf{v})}\right] = \int g^*(\mathbf{s})\ln g^*(\mathbf{s})ds - \int g^*(\mathbf{s})\ln f(\mathbf{s};\mathbf{v})ds. \quad (14)$$

Since just the latter term depends on the tilting parameters to be optimized, and we know the optimal biasing pdf (13), minimizing the Kullback-Leibler distance is equivalent to choose $\mathbf{v}$ so to solve the following maximization problem

$$\max_{\mathbf{v}} \int \frac{\mathbb{I}_{\{L(\mathbf{S}_i)>l\}} f(\mathbf{s};\mathbf{u})}{\gamma} \ln f(\mathbf{s};\mathbf{v})d\mathbf{s}, \quad (15)$$

which in turn is equivalent to the program

$$\max_{\mathbf{v}} \mathbb{E}_{\mathbf{u}}\left[\mathbb{I}_{\{L(\mathbf{S}_i)>l\}} \ln f(\mathbf{s};\mathbf{v})\right]. \quad (16)$$

By the repeated application of Importance Sampling, using again the pdf family $f(\mathbf{s};\mathbf{w})$ with a reference tilting parameter vector $\mathbf{w}$ the maximization program can finally be written as

$$\max_{\mathbf{v}} \mathbb{E}_{\mathbf{w}}\left[\mathbb{I}_{\{L(\mathbf{S}_i)>l\}} \frac{f(\mathbf{s};\mathbf{u})}{f(\mathbf{s};\mathbf{w})} \ln f(\mathbf{s};\mathbf{v})\right], \quad (17)$$

whose solution is

$$\mathbf{v}^* = \operatorname*{argmax}_{\mathbf{v}} \quad \mathbb{E}_{\mathbf{w}}\left[\mathbb{I}_{\{L(\mathbf{S}_i)>l\}} \frac{f(\mathbf{s};\mathbf{u})}{f(\mathbf{s};\mathbf{w})} \ln f(\mathbf{s};\mathbf{v})\right]. \quad (18)$$

The optimal tilting parameter vector can be estimated by solving the corresponding stochastic program, which uses a simulated sample $\mathbf{S}_1, \ldots, \mathbf{S}_N$ extracted from $f(\cdot;\mathbf{w})$

$$\max_{\mathbf{v}} \frac{1}{N}\sum_{i=1}^{N} \mathbb{I}_{\{L(\mathbf{S}_i)>l\}} \frac{f(\mathbf{S}_i;\mathbf{u})}{f(\mathbf{S}_i;\mathbf{w})} \ln f(\mathbf{S}_i;\mathbf{v}). \quad (19)$$

The j-th component of the tilting vector can therefore be obtained by the following equation

$$\frac{1}{N}\sum_{i=1}^{N} \mathbb{I}_{\{L(\mathbf{S}_i)>l\}} \frac{f(\mathbf{S}_i;\mathbf{u})}{f(\mathbf{S}_i;\mathbf{w})} \frac{\partial}{\partial v_j} \ln f(\mathbf{S}_i;\mathbf{v}) = 0, \quad (20)$$

which can be solved analytically if the distributions of the random variables of interest belong to a natural exponential family.

In our case we decide to bias the distributions of the common reliability factor as well as those of the individual factors. For all these variables we shift the expected value while maintaining the original variance (the reference vector $\mathbf{u}$ is therefore a zero vector of size equal to the number of the reliability factors – individual plus common):

$$Z_i \sim N(0,1) \Longrightarrow Z_i \sim N(\mu_{Z_i}, 1) \qquad i = 1, \ldots, F, \quad (21)$$

$$\eta_{ij} \sim N(0,1) \Longrightarrow \eta_{ij} \sim N(\mu_{ij},1) \qquad i=1,\dots,N \qquad j=1,\dots,M_i. \quad (22)$$

Since all the variables at play are independent, the resulting multidimensional pdf is just the product of the individual pdf's. In the absence of bias we have for the k-th sample element

$$f(\mathbf{S}_k;\mathbf{u}) = \prod_{j=1}^{F} \frac{1}{\sqrt{2\pi}} e^{-z_{jk}^2/2} \prod_{i=1}^{N_b} \prod_{j=1}^{M_i} \frac{1}{\sqrt{2\pi}} e^{-\eta_{ijk}^2/2}. \quad (23)$$

The joint biased pdf is instead

$$f(\mathbf{S}_k;\mathbf{w}) = \prod_{j=1}^{F} \frac{1}{\sqrt{2\pi}} e^{-\left(z_{jk}-\mu_{Z_j}\right)^2/2} \prod_{i=1}^{N_b} \prod_{j=1}^{M_i} \frac{1}{\sqrt{2\pi}} e^{-\left(\eta_{ijk}-\mu_{\eta_{ij}}\right)^2/2}. \quad (24)$$

The likelihood ratio is therefore

$$\frac{f(\mathbf{S}_k;\mathbf{u})}{f(\mathbf{S}_k;\mathbf{w})} = \exp\left\{ \frac{\sum_{j=1}^{F}\left[\left(z_{jk}-\mu_{Z_j}\right)^2 - z_{jk}^2\right] + \sum_{i=1}^{N_b}\sum_{j=1}^{M_i}\left[\left(\eta_{ijk}-\mu_{\eta_{ij}}\right)^2-\eta_{ijk}^2\right]}{2}\right\}. \quad (25)$$

If we mark the parameters (to be estimated) of the target pdf by a star, the equations to solve are

$$\frac{1}{N}\sum_{k=1}^{N} \mathbb{I}[L_k > l] \frac{f(\mathbf{S}_k;\mathbf{u})}{f(\mathbf{S}_k;\mathbf{w})} \frac{\partial}{\partial \mu_{Z_j}^*} \ln[f(\mathbf{S}_k;\mathbf{v})] = 0 \qquad j=1,\dots,F, \quad (26)$$

where L_k is the overall loss as evaluated at the k-th simulation run, and

$$\frac{1}{N}\sum_{k=1}^{N} \mathbb{I}[L_k > l] \frac{f(\mathbf{S}_k;\mathbf{u})}{f(\mathbf{S}_k;\mathbf{w})} \frac{\partial}{\partial \mu_{\eta_{ij}}^*} \ln[f(\mathbf{S}_k;\mathbf{v})] = 0, \quad (27)$$

the latter with $i=1,\dots,N_b$ and $j=1,\dots,M_i$. The solution provides the following parameter values

$$\mu_{Z_j}^* = \frac{\sum_{k=1}^{N} \mathbb{I}[L_k > l] \frac{f(\mathbf{S}_k;\mathbf{u})}{f(\mathbf{S}_k;\mathbf{w})} z_{jk}}{\sum_{k=1}^{N} \mathbb{I}[L_k > l] \frac{f(\mathbf{S}_k;\mathbf{u})}{f(\mathbf{S}_k;\mathbf{w})}} \qquad j=1,\dots,F, \quad (28)$$

$$\mu_{\eta_{ij}}^* = \frac{\sum_{k=1}^{N} \mathbb{I}[L_k > l] \frac{f(\mathbf{S}_k;\mathbf{u})}{f(\mathbf{S}_k;\mathbf{w})} \eta_{ijk}}{\sum_{k=1}^{N} \mathbb{I}[L_k > l] \frac{f(\mathbf{S}_k;\mathbf{u})}{f(\mathbf{S}_k;\mathbf{w})}} \qquad i=1,\dots,N_b, \qquad j=1,\dots,M_i. \quad (29)$$

At each updating step the probability of interest can be estimated by the IS formula

$$\mathbb{P}[L > l] \simeq \frac{1}{N}\sum_{i=1}^{N} \mathbb{I}_{\{L(\mathbf{S}_i)>l\}} \frac{f(\mathbf{S}_i;\mathbf{u})}{f(\mathbf{S}_i;\mathbf{v})} \quad (30)$$

In order to apply these equations with a sufficient number of useful (i.e. nonzero) terms it is usually expedient to start the algorithm with a threshold $\hat{l} < l$, to be progressively increased as the bias level grows. In the end, the resulting simulation programme can be synthetically described by the following list of steps:

1. Set $\mathbf{v}_0 = \mathbf{u}$, $l_0 = 1$, and the iteration counter $t = 1$;
2. Generate a sample $\mathbf{S}_1, \ldots, \mathbf{S}_N$ from the probability density $f(\cdot; \mathbf{v}_{t-1})$;
3. Obtain the resulting values of the economic loss $L(\mathbf{S}_1), \ldots, L(\mathbf{S}_N)$, sort them in increasing fashion, and extract the 90% percentile to assign it to the threshold l_t;
4. Use the same sample to update the tilting vector $\mathbf{v}_t$ according to equations (28) and (29);
5. If $l_t < l$, set $t = t + 1$ and go back to step 2, else go to the next and final step;
6. Estimate the loss probability by eq. (30) with $\mathbf{v} = \mathbf{v}_t$.

4 Open issues: Model Identification and Calibration

So far we have assumed that the model is completely defined, i.e. that all its parameters are known. However, the actual application of the model to a concrete situation requires that we are able to completely define the model on the basis of a set of measurements on the system of interest (typically measurement of failure rates). Though this is still an open issue, we provide here some indication. In our case the definition of the model requires the following operations:

1. Identifying the common reliability factors;
2. Setting the values of the $M = \sum_{i=1}^{N_b} M_i$ thresholds b_{ij} ($i = 1, \ldots, N_b$, $j = 1, \ldots, M_i$);
3. Setting the values of the $F \cdot M$ correlation parameters ρ_{ijk} ($i = 1, \ldots, N_b$, $j = 1, \ldots, M_i$, $k = 1, \ldots, F$).

As to the first item, this requires a deep knowledge of the system at hand, so to identify those factors that may influence more than a single subsystem. However, the number F of the common reliability factors has an impact on the capability of the model to correctly represent the system: too low a number of common reliability factors may result in the impossibility to fit the model to the observed failure rates.

The thresholds b_{ij} are easily derived from the observed indidivual failure rates. In fact if we take the observed failure rate r_{ij} of the generic subsystem as an approximation of its failure probability $\mathbb{P}[X_{ij} > b_{ij}]$ the thresholds are obtained as

$$b_{ij} = G^{-1}\left(1 - r_{ij}\right). \tag{31}$$

Instead, the determination of the correlation parameters may be quite difficult. We have to determine $F \cdot M$ parameters. A mathematical view of the problem would lead us soon to consider the correlation between the latent variables as constraints (though these correlations are not observable). The minimum number of constraints is therefore the number of couples of latent variables $\binom{M}{2} = \frac{M(M-1)}{2}$. Both the number of unknowns and the number of constraints grow with M, but the former grows linearly and the latter grows quadratically. By equating the two quantities we can derive some bounds on F of practical interest for the model determination. When $M < 2F + 1$ we may have (if we don't introduce additional constraints) an underconstrained system, so that there is room to fit the model to the observed failure rates. Instead, if we exceed this bound, the resulting system of equations is overconstrained, with the practical consequence that we won't be able to model all of the possible correlation conditions. As an example we can consider the case where we have four subsystems (represented by the latent variables X_1, X_2, X_3, X_4) and a single common reliability factor. In this case we won't be able e.g. to obtain the correlation parameters such that the following conditions on the correlation between the latent variables are met $C\left(X_1, X_3\right) = C\left(X_2, X_4\right) = 0.8$, $C\left(X_1, X_2\right) = C\left(X_3, X_4\right) = 0.1$. In those cases a way out could be represented by the addition of common reliability factors (i.e. increasing F) so to make the resulting system again underconstrained. This is not tantamount to introducing dummy variables, but is rather related to the identification of real common factors that were hidden in the first exploration of the system at hand.

However, in practical cases the determination of the correlation parameters ρ_{ijk} can't be arrived at through the correlation between the latent variables, since these are essentially unobservable. Instead, we can measure fault correlations, i.e. the correlation between the random variables Y_{ij} indicating the state of the subsystems. Indicating by $p_{ij} = \mathbb{P}[X_{ij} > b_{ij}]$ the marginal probability that the j-th subsystem in the i-th service basin is out of service, and by $p_{ij,lm} = \mathbb{P}[X_{ij} > b_{ij}, X_{lm} > b_{lm}]$ the joint probability of failure of two subsystems, the correlation between the faults of these two subsystems is

$$c_{ij,lm} = \frac{p_{ij,lm} - p_{ij}p_{lm}}{\sqrt{p_{ij}(1 - p_{ij})p_{lm}(1 - p_{lm})}}. \tag{32}$$

In turn the probability of a joint failure can be computed through the joint pdf of the two latent variables (which is a two-dimensional normal one $g_{\beta}(\cdot, \cdot)$ with a correlation coefficient $\beta = C(X_{ij}, X_{lm})$)

$$p_{ij,lm} = \int_{b_{ij}}^{+\infty} \int_{b_{lm}}^{+\infty} g_{\beta}(s,t) ds dt, \tag{33}$$

providing a link between the correlation of faults (observable but not explicitly present in the copula model) and the correlation of the latent variables (unobservable but directly related to the weights of the common reliability factors in the normal copula model). As noted in [Frey et al. 2001] fault correlations are much lower than the corresponding latent variable correlations.
Since faults, rather than latent variables, are observable, a better calibration procedure involves the maximization of the likelihood function, following the approach of [Demey et al. 2004], which we briefly introduce here. We can first write the fault probability of the individual subsystem conditioning on the individual reliability factor of that subsystem

$$p_{ij} = \mathbb{P}[X_{ij} > b_{ij}] = G\left(\frac{b_{ij} - \sum_{k=1}^{F} \rho_{ijk} Z_k}{\sqrt{1 - \sum_{k=1}^{F} \rho_{ijk}^2}}\right). \tag{34}$$

If we observe the network for a number T of periods and introduce the observed variable $D_{ij}(t)$, equal to 1 if we observe the fault of that subsystem at time t and 0 otherwise, the conditional likelihood function is

$$\mathcal{L}_c = \prod_{t=1}^{T} \prod_{i=1}^{N_b} \prod_{j=1}^{M_i} \left[p_{ij} D_{ij}(t) + (1 - p_{ij})(1 - D_{ij}(t))\right]. \tag{35}$$

Since we have conditioned on the individual reliability factors, which are independent and identically distributed according to a standard normal law, the log likelihood function is finally

$$\mathcal{L} = \sum_{t=1}^{T} \sum_{i=1}^{N_b} \sum_{j=1}^{M_i} \int \left[p_{ij} D_{ij}(t) + (1 - p_{ij})(1 - D_{ij}(t))\right] g(z) dz, \tag{36}$$

where $g(z)$ is the standard normal probability density function. The log-likelihood function depends, through the marginal probability values p_{ij}, on the correlation parameters ρ_{ijk} of interest as expressed by expr. (34). We can therefore calibrate the ρ_{ijk} values by maximizing, by numerical procedure, this likelihood function. As the maximization problem can be quite complex, since there are many parameters and in many instances the observed fault indicator D_{ij} will be zero, a constrained maximization approach has been proposed e.g. in [Demey et al. 2004], where some conditions are imposed on the correlation between the latent variables (e.g. that they are all equal or that take some value within a small set).

Once the correlation parameters have been determined the weights of the individual reliability factors α_{ij} can be readily derived by the normalizing relationship (3).

5 A toy example

For the purpose of illustrating the application of the method described so far we consider in this section a simple model, made of 100 service basins, with each service basin made of a single subsystem and a single reliability factor, common to all the subsystems. The latent variable of the i-th subsystem is therefore

$$X_i = \rho Z + \sqrt{(1-\rho)^2}\eta_i, \qquad i = 1, \ldots, 100. \tag{37}$$

The loss associated to the disruption of a single service basin is unitary, so that the overall loss ranges from 0 to 100 in steps of 1. We apply the method previously described to evaluate the probability of a large loss, corresponding to the simultaneous failure of many service basins. As a simple case we set the failure threshold for each latent variable $b = 1$, so that the individual failure probability is $1 - G(1) \simeq 0.16$. We examine how the probability of a large loss varies with the Value-at-Risk (i.e. the threshold l as defined in Section 2) and with the correlation weight ρ. We examine the former issue for a case of mild dependance, setting $\rho = 0.1$ and letting the loss treshold vary in the range $l = 30 \div 70$. As can be seen in Figure 1 the fall of the loss probability is roughly exponential. In order to assess the effect of the correlation we can compute the loss probability when all subsystems are independent by the binomial distribution, obtaining for the extreme cases ($l = 30$ and $l = 70$) a loss probability respectively equal to $1.697 \cdot 10^{-4}$ and $1.774 \cdot 10^{-33}$. The distance from the values reported in the picture shows that the effect of correlation is particularly relevant for large values of VaR.

The effect can be highlighted in greater detail by setting a value for the VaR (say $l = 50$) and varying the correlation weight ρ in the range $0.05 \div 0.5$. Now the loss probability, reported in Figure 2, grows much more rapidly when the correlation weight first grows, approaching then its limiting value (i.e. the individual failure rate, in this case $1 - G(1) \simeq 0.16$). Again we can derive the limiting value on the other side, i.e. when all the subsystems are independent ($\rho = 0$), obtaining $\mathbb{P}(L > 50) \simeq 1.88 \cdot 10^{-15}$.

6 Conclusions

A new approach, based on the economic loss, has been proposed to evaluate the reliability behaviour of a communications network. The model is based on the partition of the network operator's market in a number of service basins,

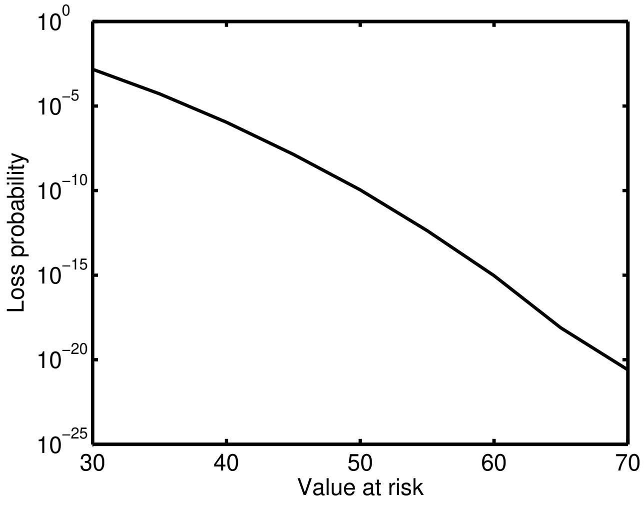

Figure 1: Probability of losses larger than a given Value-at-risk ($\rho = 0.1$)

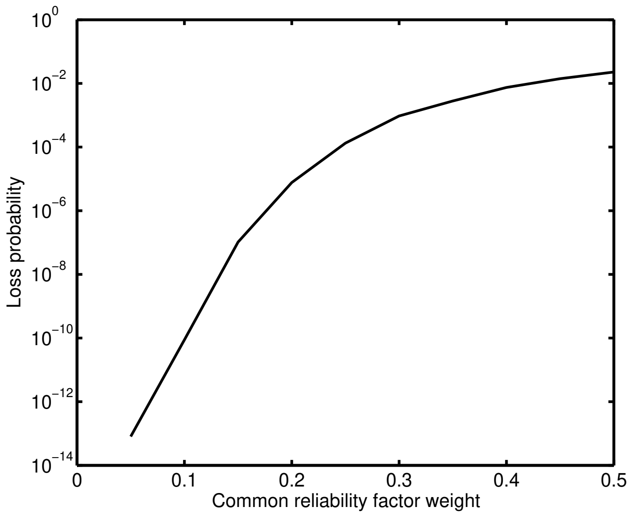

Figure 2: Effect of the correlation weight ($l = 50$)

each relying on a number of subsystems. The dependence among different subsystems, largely ignored in most reliability models, has been taken into account by introducing a normal copula model. In order to solve the model and obtain the probability of a given level of losses a simulation algorithm, based on the Cross-Entropy approach, has then been proposed. The full description of the algorithm is finally accompanied by a toy example to show its usability in a simple scenario.

References

[Akeson et al. 1994] Akeson, J.K., Casey, B.K., Dagli, A.A., Glickman, S.E., Kalisz, T.: "Cost Impact of Telecommunications Switch Software Problems on Divested Bell Operating Companies"; IEEE Journal Sel. Areas in Communications, 12, 2 (February 1994)

[Demey et al. 2004] Demey, P., Jouanin, J.F., Roget, C., Roncalli, T.: "Maximum likelihood estimate of default correlations"; Risk, 17, 11 (November 2004) 104-108

[Frey et al. 2001] Frey, R., McNeil, A.J., Nyfeler, M.A.: "Modelling Dependent Defaults: Asset Correlations Are Not Enough!"; Working paper (9 March 2001)

[Gupton et al. 1997] Gupton, G., Finger, C., Bhatia, M.: "CreditMetrics Technical Document"; J.P. Morgan and Co. (2 April 1997)

[Helvik 2004] Helvik, B.:"Perspectives on the dependability of networks and services"; Telektronikk, 3 (2004) 27-44.

[ITU 1994] ITU-T: "Recommendation E.800: Terms and definitions related to quality of service and network performance including dependability"; Geneva (1994).

[Kogeda et al. 2005] Kogeda, O.P., Agbinya, J.I., Omlin, C.W.: "Impacts and Cost of Faults on Services in Cellular Networks"; Proceedings of the International Conference on Mobile Business (ICMB'05) (2005).

[Rubinstein and Kroese 2004] Rubinstein, R.Y., Kroese, D.P.: "The Cross-Entropy Method"; Springer, New York

[Snow 2001] Snow, A.P.: "Network Reliability: The Concurrent Challenges of Innovation, Competition, and Complexity"; IEEE Trans. Reliability, 50, 1 (March 2001) 38-40.