# Guaranteeing Service Availability in SLAs; A Study of the Risk Associated with Contract Period and Failure Process

Andrés J. González and Bjarne E. Helvik

Centre for Quantifiable Quality of Service in Communication Systems*
Norwegian University of Science and Technology, Trondheim, Norway
Email: {andresgm,bjarne}@q2s.ntnu.no

***Abstract*—Service Level Agreements (SLAs) are a common means to define the obligations of network/service providers and users in business relationships. The terms that define the guaranteed availability for a given period are an important element of these contracts. The appropriate values selection is difficult due to the large number of variables involved, the complexities of the network and service provision and the computational challenge posed by the transient solution, as opposed to a steady state, that is needed. A common policy taken to solve it, is using the steady state availability as a reference. Nevertheless this simplification may put on risk the contract fulfillment as stochastic variation of the measured availability is significant over a typical contract period. This paper analyzes the relevance that the interval availability analysis has on SLAs, and provides suggestions to the network providers on the selection of adequate availability guarantees. The interval availability of unprotected and shared protected connections is studied under exponential and Weibull failure and repair distributions. It is observed that for a single path scenario, a small reduction of the guaranteed availability below the steady state value improve the probability to meet the requirements considerably. The same is the case for connections with shared backup protection. However performing this analysis in the transient domain is quite demanding. Hence, to simplify it, it is proposed to obtain the steady state results and introduce a *safeguard factor* to control that the availability guarantee is meet. For the Weibull distributed times between failures, where the shape factor is less than one (as observed in operational networks), the probability of meeting a guaranteed availability over a finite contract period, decrease more radically than for the commonly assumed Poisson failure process. This increases the importance of making a transient analysis.**

## I. INTRODUCTION

Network operators and customers use Service Level Agreement (SLA) to define a contracted QoS, where availability is a significant element. Violation of the agreed value may have large economic, performance and reputational consequences for both parts.

This paper studies how to guarantee availability on SLAs, considering the transient working-failed behavior during a contract period, i.e the distribution of the observed interval availability. Under Markov assumptions, the interval availability can be obtained by numerical methods using uniformization techniques as is presented in [8], [15], [17]. Here, it is of special interest analyze the probability of fulfill the contracted interval availability $\alpha$, which will be denoted as the *SLAs success probability*, in a network with link failures and repairs.

This problem was first raised for general systems in [10], and it was observed that if $\alpha$ is larger than the *steady state availability* ($A$) of a service, the success probability decrease continuously with the observation period until reach a value of zero, which means that it will be impossible fulfil the contract. However, it was also shown that there is a considerable risk even when $\alpha < A$. In [16] it was suggested, in a general way, how to use the interval availability for economic planning to maximize the revenue from SLAs. A recent paper [3] studied the transient behavior under adaptive management strategies to control penalties and fairness.

An important issue is the study of connections with shared protection paths. This technique has an increasing application as it offers a considerably increased availability and maintain an efficient resource utilization [5], [6], [2]. On this issue, most of the works have modeled the availability through bounds that approximate the asymptotic solution. This is necessary due to the complexity involved in finding exact solutions due to the process dependencies that exists when links are shared [12], [9], [11], [13], [4]. This paper extends those results by analyzing the transient behavior of these scenarios. A finding is that it is more important to find a conservative solution than to use a huge computational effort to find an almost precise solution with an uncontrolled error, due to the risk that arises form having values of $\alpha$ close to $A$.

During the last years studies have shown that the link inter-failure and reparation/recovery times may be more accurately modelled by a Weibull distribution than by the more commonly used negative exponential distribution, due to the higher occurrence of very short and long times events, e.g. [14]. Hence, the success probability is analyzed under failure and repair processes with Weibull distributed inter-event times. It

*"Centre for Quantifiable Quality of Service in Communication Systems, Centre of Excellence" appointed by The Research Council of Norway, funded by the Research Council, NTNU, UNINETT and Telenor. http://www.q2s.ntnu.no

978-1-4244-4388-8/09/$25.00©2009 IEEE

was observed an interesting relation between the Weibull shape parameter ($\beta$) and the success probability, that emphasizes the importance of using realistic Weibull failure and repair time distributions in setting the interval availability guarantees in the SLAs.

Analytical/numerical tools existing for analyzing interval availability are found just for trivial systems (single elements) or systems with Markov properties. Hence, here we use simulation to obtain the results that are shown through the paper. Specifically the tool used was Möbius [1] [18], and the simulations in all the scenarios were run with a relative confidence interval of 0.1 and a confidence level of 99% was obtained.

To make the obtain insight easily applicable, a *safe guard factor*, $\sigma$, is introduced to allow the use of steady state availabilities, but taking also into account the stochastic transient effects due to the finite contract period. The results show that though the use of techniques with acceptable computational complexity, the interval availability guaranteed in an SLA may be met with a high probability.

This paper is organized as follows. In Section II issues related to the distribution of the interval availability and their relation to SLAs are introduced and discussed. Section III study the success probability of connections under shared protections schemes, i.e. the probability than the observed availability will meet the guaranteed. Section IV presents the effect of having failure and repair/restore inter event times that follow a Weibull distribution. Section V discusses issues on how "safe" guarantees of availability in SLAs may be obtained. Finally Section VI concludes the paper.

## II. SLA SUCCESS PROBABILITY

Defining a *network connection* as a group of interconnected links that provides an end to end service, maintaining it operational (up/working) is salient to offer a good quality of service. Its performance as a function of time can be modeled according with the random function $\hat{I}(t)$ defined as follows:

$$\hat{I}(t) = \begin{cases} 1 & \text{if the connection is working.} \\ 0 & \text{otherwise.} \end{cases} \quad (1)$$

Specifically on SLAs scenarios, it is important to study the Interval Availability $\hat{A}(\tau)$, which is a stochastic variable that measure the total time that the connection has been working during a defined period $\tau$ and that can be described by the following equation:

$$\hat{A}(\tau) = \frac{1}{\tau}\int_0^{\tau} \hat{I}(t)\, dt. \quad (2)$$

When the transient behavior of a connection is studied, it is common to evaluate the expected interval availability $A(\tau) = E[\hat{A}(\tau)]$. For instance, Figure 1(a) shows that for a five links connection with independent link failures, exponentially distributed, with mean of one year and reparation times with an expectation of 12 hours, $A(\tau)$ converges in few weeks to the steady state availability $A$ defined as:

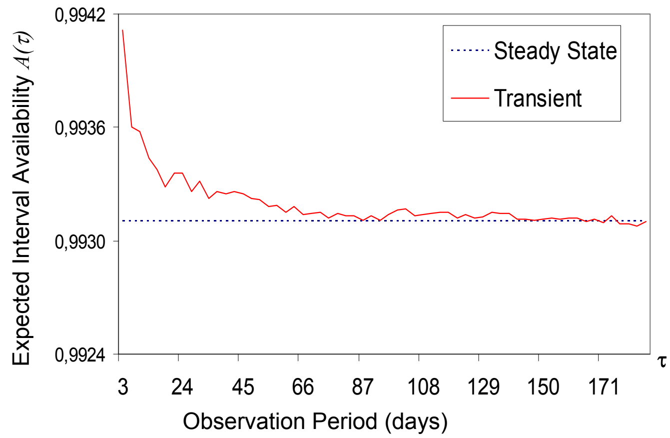


(a) Expected Interval Availability

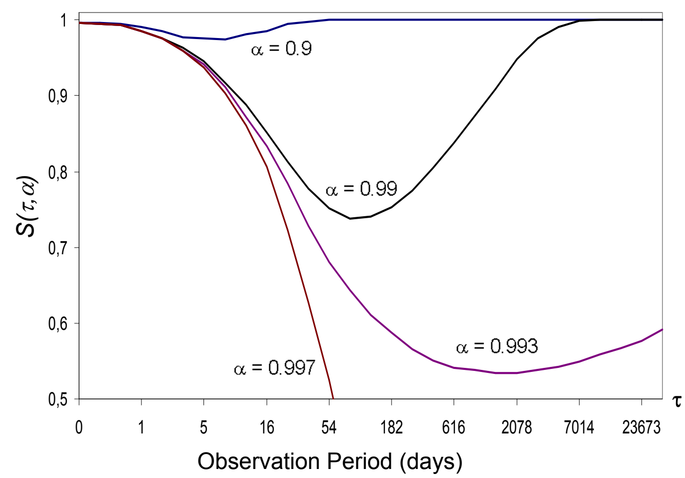


(b) Success Probability under different values of $\alpha$ for a system with $A = 0.9931514$

Fig. 1. Interval Availability Behavior

$$A = \lim_{\tau \to \infty} A(\tau) \quad (3)$$

Nevertheless, the measure of $A(\tau)$ is not descriptive enough since it does not consider the probabilistic variation, that may have a critical influence on the connection quality.

From a SLA point of view, when the observed interval availability distribution is taken into account, and given an availability guarantee $\alpha$, it is of special interest obtain the probability that the availability after some observation period $\tau$ will be larger or equal than the defined guarantee. This will be defined as the Success Probability:

$$S(\tau, \alpha) = Pr[\hat{A}(\tau) \geq \alpha] \quad (4)$$

Additionally the *risk* will be defined as the probability that the specified guarantee availability will not be met, which can be expressed as $1 - S(\tau, \alpha)$.

There are some numerical methods that through the use of uniformization techniques may obtain $S(\tau, \alpha)$ for Markovian systems [8], [15], [17]. Additionally for a single item, the failure/repair process may be regarded as an alternating renewal process when independence are assumed, obtaining interval availability results for general distributions [7]. However, as we address compound systems under non-Markovian assumptions, i.e. Weibull failure and repair processes, $S(\tau, \alpha)$ is obtained by simulation for the various cases studied.

The Success Probability as a function of the observation period was first discussed by Goyal & al. [10]. It has basically two different behaviors dependent on whether the guaranteed value can be met in the asymptotic case or not. See Figure 1(b) for an illustration. In the first case, $\alpha \leq A$, it drops below one for a period, but converges to one. The time until it converges and the depth of the dip, depend on the ratio between the guarantee $\alpha$ and the asymptotic availability $A$. In the second case, $\alpha > A$, it decrease continuously until reach zero.

Taking into account those observations, in this paper is defined a *responsible promise* as the stipulated guarantee on the SLA that lies below $A$. The example in Figure 1(b) shows this for a network connection with five independent links having Poisson failure processes with the expectation of one failure per year and i.i.d. negative exponentially distributed down times with reparation with expectation of 12 hours. The path steady state availability is 0.99315. It is observed that the risk may become very large. This can be avoided if the contract period is very short or very large.

However, these safe observation periods are out of the range of typical SLA contract periods, and therefore, it is necessary to look for another type of solution, like for example use guarantees somewhat poorer than than the asymptotic availability. Note that the worst case occurs when $\alpha$ matches exactly the steady state availability, where we have a null recurrent stochastic process with a risk of 0.5 for an infinitely period.

## III. SHARED PROTECTED CONNECTIONS

To meet high availability requirements and to be able to guarantee it, telecommunications networks implement mechanisms that protect connections trough the reservation of additional resources that can be used when the primary fails. This includes the assignation of a dedicated principal path ($W$) and a shared backup path ($B$). We study the first and more demanding case, where backup links may be shared between several connections. This mechanism is known as *shared backup protection* and allows the combination of dependability improvement via protection and a more efficient resource utilization.

Under this approach, both paths (main and backup) must be failure disjoint. The connection, by default uses the primary path, but if it is affected by a failure, the backup may be used instead, if it is working. However, as it is not dedicated, its availability also depends on its potential use by other connections.

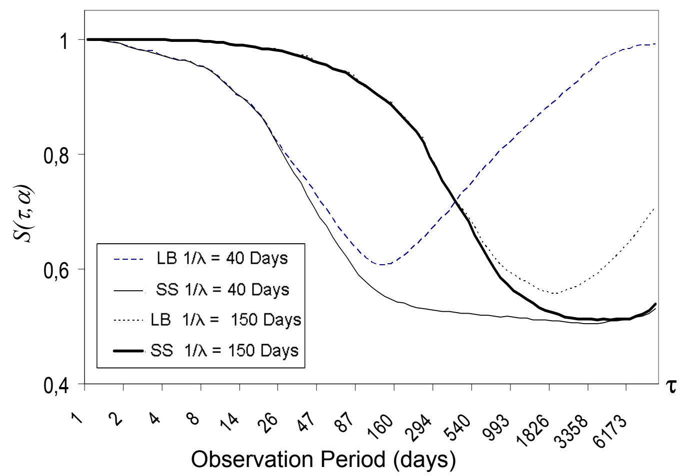


Fig. 2. Success Probability with $\alpha$ =Lower Bound (LB) and $\alpha \approx$ Near to the Steady State (SS) under different links quality.

It can be said that the total availability of a given connection depends on the availability of the main paths of the other connections which share links with on the backup. Hence the availability prediction becomes more demanding than for the unprotected or dedicated path protection case. This is further complicated by the fact that connections in a network are continuously established and closed, so the set of "partners" sharing a backup link is dynamic [4]. For this study we use a lower bound that approximates the asymptotic solution.

To predict the total asymptotic availability of a connection ($C_n$) that uses shared protection scheme, the unavailability of its main path is considered, which contains a number of links ($L_i$) with independent availabilities $A_i$, and the unavailability of the backup path with links ($L_j$) with availabilities $A_j$ and that they only can be used when all the working paths of the group of connections ($C_s$) that share that link are available. This yields.

$$A^{C_n} \geq 1-(1-\prod_{\forall i|L_i \epsilon W} A_i).(1-\prod_{\forall j|L_j \epsilon B} A_j.[\prod_{\forall s|C_s} A_W^{C_s}]) \quad (5)$$

For the transient analysis the problem becomes more complex and therefore it is more appropriate the use of simulation. In the case shown in Figure 2 two connections are considered, each with a dedicated principal path and an independent and disjoint backup path. Both backup paths have one link in common that have to be shared in case of failure. This link can be used by the first connection that need it, if it is not being used, otherwise, the current connection that request the service has to wait until the working path of the other become again available, or if another link (different from the shared) on the backup of the connection that is currently using it, goes down, making useless the reservation of the shared link.

Figure 2 shows the results obtained in the simulation, for links with two different expected link interfailure times. First using a guarantee $\alpha$ as close as possible to the steady state availability of the system, secondly an $\alpha$ equal to the lower bound obtained from equation (5).

In general it was observed that the properties of $S(\tau, \alpha)$ remains for the shared protected scenarios as well.

As was expected, according to the results in the previous section, the lower bound is always a safer decision for a network provider. Specially the safety margin is more pronounced when the links do not have a good quality. Even though the lower bound is safer, from a SLA success probability point of view, the risk of not meeting a guarantee is significant and should be dealt with.

## IV. WEIBULL ANALYSIS

Measurements of operational systems show higher occurrence of very short and long link interfailure and reparation/recovery times than what is properly described by a negative exponentially distribution, e.g. [14], [19].

Hence these are more accurately modelled by a Weibull distribution with a scale parameter, $\beta$, less than one, i.e. we assume i.i.d times $T_x$ between events of type $x$, where $x$ may be failure of a link or a repair of link and

$$P(T_x > t) = 1 - F_{T_x}(t) = e^{-(\frac{t}{\theta})^{\beta}}, \forall t \geq 0, \qquad (6)$$

where $\theta$ is the scale parameter. It is seen that when $\beta = 1$ the distribution becomes exponential and the results obtained previously is a special case.

As above the the success probability is obtained for unprotected and shared protected connections. The value of $\theta$ is set in order always to have an expected time between failures of one year, for all the values of $\beta$ used. The repair times used were also Weibull distributed with the shape parameter less than one and an expectation equal to 12 hours for all cases. Hence, the asymptotic values are the same in all cases and identical to those presented in section II.

Figure 3 shows some results, given a value of $\alpha = 0.98$ for a responsible promise and $\alpha = 0.997$ for the other case. As can be seen, exponentially distributed link failure times ($\beta = 1$) may incur a smaller risk than the cases with a Weibull process with shape factor below one. Figure 3(a) shows that if the value of $\beta$ is very small, the failure process may generate a very steep and large dip in $S(\tau, \alpha)$. Additionally the time needed to return to an acceptable low risk may be longer than the exponential case as well as the typical contract periods, making less likely that the SLA agreement is met. When $A < \alpha$, it can be observed from Figure 3(b) that the effect of $\beta$ is also significant for the case where the contract can not be met in the asymptotic case.

The effect of the shape of the repair distribution, i.e. its $\beta$ parameter, has also been thoroughly investigated. It is found that if the value of $\beta$ is very small the dip in $S(\tau, \alpha)$ is smoothed and therefore the risk may be reduced. This effect has inverse tendency as for the time between failures, but it is, for typical parameter values, approximately on order of magnitude, i.e. ten times, less, making the failure process dominant.

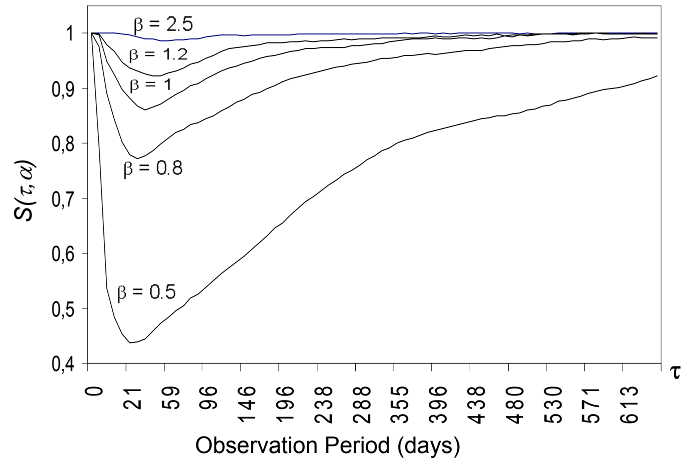


(a) Success Probability under different Failure Link $\beta$ values and $A \geq \alpha$

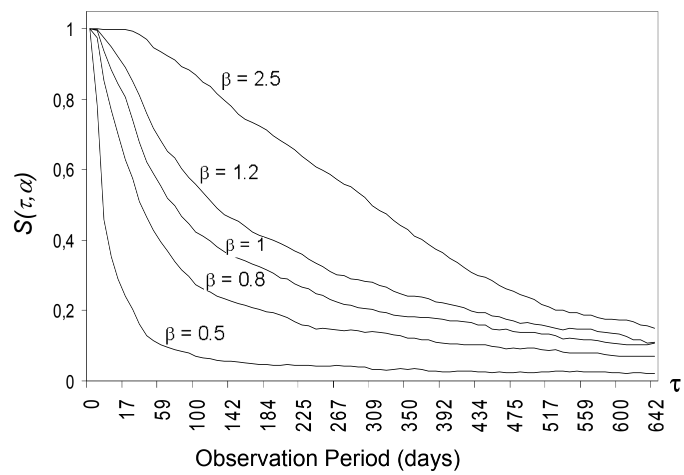


(b) Success Probability under different Failure Link $\beta$ values and $A < \alpha$

Fig. 3. $S(\tau, \alpha)$ Behavior under Weibull Process

The above findings show the importance of a proper characterization of the failure processes, because of their influence on the distribution of the observed interval availability and hence, on the risk of not meeting an availability contracted in an SLA.

## V. MEETING AVAILABILITY GUARANTEES

The analysis made on the last sections, shows that the stochastic behavior of the interval availability may put the contract fulfillment at serious risk, and actions should be taken to counteract these effects.

Knowing that is very demanding for a network provider to have a success equal to one, it must be planned a risk value that takes into account the commercial and technical costs implied. For this propose a guarantee value $\alpha^*$ is specified in the SLA, to have a controlled risk.

To obtain this value a *Safe-Guard Factor* $\sigma$ is suggested as the correction value that multiply the steady state availability in order to obtain a desired success that take into account different consequences that imply the violation of the contracted parameters, i.e.

$$\alpha^* = \sigma \cdot A \tag{7}$$

In this way is possible to have a solution with the asymptotic simplicity and the transient safety.

To have a better idea about the magnitudes implied to have an adequate SLA definition, the connection interval availability for a fixed observation period, on paths with multiple links, that follow independent failure and reparation Weibull distributed processes was simulated.

In Figure 4 can be seen the results obtained for the Safe Guard Factor tendency in an SLA with one year duration, using different network conditions. The values of $\beta$ shown on the figure correspond to the failure process.

It is also important to analyze the effect that the individual link quality has. On Figure 4 can be also appreciated that in order to have correction values very close to one, it is necessary to have links with extremely low failure rate (high expected time for link failure). In the practice, obtain a link with this characteristics may imply a huge effort and therefore for most of the existing networks the deviation that put $\sigma$ may be considerable.

A particular example that may illustrate better the effect that the use of the Safe-Guard Factor has in the SLA specification may be obtained, assuming a contract for a given connection defined for a fixed observation time of one year, where it is previously known that it's steady state availability is 363 days (maximum tow days of cumulative time-out per year). If the links failures on the network follow a weibull distribution with $\beta$ equal to 0.5 and expectation of one year, using Figure 4, can be obtained that the optimal $\alpha$, to have a success probability of 0.99, imply around 16 days of connection unavailability. Hence, if this is analyzed under a business model, it is possible that the commercial image may be considerably affected. On the other hand if the provider offers guarantees that are not fulfill in most of the case the reputational consequences may be even worst in the long time. This dilemma raise the challenge of having a model able to balance this two important tendencies.

As was shown before, there are many variables that have an implication on and SLA success in terms of availability. From the business point of view, there are three facts that are relevant for the Network Providers regarding the network dependability and SLA specification. The first is the cost implied in the installation, configuration and assignment of network resources in order to provide high quality connections, which include the selection of a high reliable infrastructure and the assignation of resources with the appropriate protection. The second is the commercial impact that may have to offer a guaranteed availability that may be improved for the competitors (as was shown before, the difference may be of several days). Finally the third one is the cost generated by the refund payment caused when the requirement are not met, additionally to the

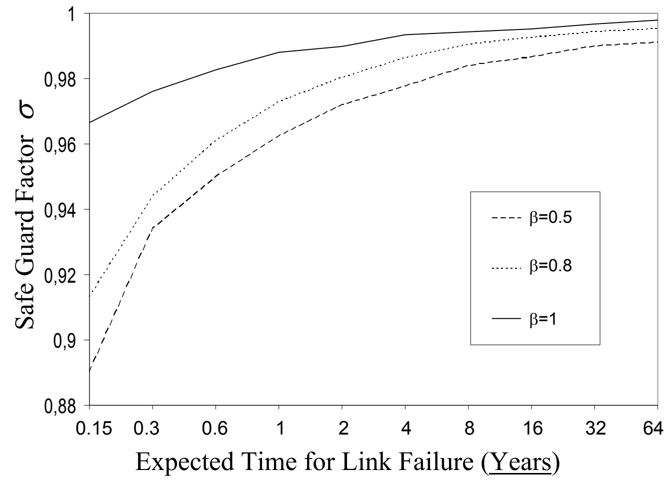


Fig. 4. Safe-Guard Factor needed to obtain a 99% Success Probability

image impact that is generated when the customers assimilate those events as a weakness quality.

One of the ideas behind the Safe-Guard factor and all the others analysis made previously; is to provide tools in order to have a realistic joint model, able to include and correlate the three facts already enounced, in order to offer a good quality service and at the same time obtain the best profit. The precise setting of the model variables may differ for each provider. However they can be easily obtained, based on the particulars technical and marketing conditions that the network companies have.

## VI. Conclusions

This paper demonstrates the importance of taking the distribution of the interval availability into account for defining the SLA specification.

It was shown that when a provider offers an availability equal to the steady state availability, the dip in $S(\tau, \alpha)$ is extended for an infinite period with a value of 0.5. To have a guarantee in the SLA that is likely to be met, it must be below the steady state value. For the case of shared protection, it was confirmed that the lower bound obtained in previous works is useful and sufficiently accurate to be used to decide an appropriate guarantee $\alpha$. It was shown that $S(\tau, \alpha)$ has similar behavior under Weibull failures and reparation process, but additionally was found that the shape parameter of the failure process has a major impact on the SLA success. For values of $\beta$ less than 1, as found in operational data, the risk increase considerably fast and the duration at the high risk period is larger than the typical SLA contract period.

Finally, it was demonstrated that it is feasible to control the probability of success using easy implementable procedures.

## References


[1] *Mobius Manual. Version 2.2.1. November 20, 2008.*

[2] B.E. Helvik A. Mykkeltveit. Comparison of schemes for provision of differentiated availability-guaranteed services using dedicated protection. In *Proceedings of the Seventh International Conference on Networking (ICN)*, April 2008.
[3] B.E. Helvik A. Mykkeltveit. Adaptive management of connections to meet availability guarantees in slas. In *Proceedings of the IM 2009 Mini-Conference*, June 2009.
[4] B.E. Helvik A. Mykkeltveit. Comparison of schemes for provision of differentiated availability-guaranteed services using dedicated protection. In *Proceedings of the 12th Conference on Optical Network Design and Modelling (ONDM)*, March 2008.
[5] B.E. Helvik A. Mykkeltveit. Provision of connection-specific availability guarantees in communication networks. In *Proceedings of the 6th International Workshop on Design of Reliable Communication Networks (DRCN)*, Oct, 2007.
[6] P. Cholda, A. Mykkeltveit, B. E. Helvik, O. J. Wittner, and A. Jajszczyk. A survey of resilience differentiation frameworks in communication networks. *IEEE Communications Surveys & Tutorials*, 9(4):32–55, Fourth Quarter 2007.
[7] D.R. Cox. *Renewal Therory*. Methuen, 1967.
[8] H.R. De Souza E Silva, E. Gail. Calculating cumulative operational time distributions of repairable computer systems. *IEEE Transactions on Computers*, Volume: C-35, Issue: 4:322–332, April 1986.
[9] A. Fumagalli, M. Tacca, F. Unghvary, and A. Farago. Shared path protection with differentiated reliability. In *Proc. IEEE International Conference on Communications ICC 2002*, volume 4, pages 2157–2161, 28 April–2 May 2002.
[10] A.N. Goyal, A.; Tantawi. A measure of guaranteed availability and its numerical evaluation. *IEEE Transactions on Computers*, Volume 37, Issue 1:25 – 32, 1988.
[11] Pin-Han Ho, J. Tapolcai, and A. Haque. Spare capacity reprovisioning for shared backup path protection in dynamic generalized multi-protocol label switched networks. *IEEE TRANSACTIONS ON RELIABILITY*, 57(4):551–563, Dec. 2008.
[12] Hongfang Yu Hongbin Luo, Lemin Li. Routing connections with differentiated reliability requirements in wdm mesh networks. *IEEE/ACM TRANSACTIONS ON NETWORKING,*, VOL. 17, NO. 1:253–266, FEBRUARY 2009.
[13] Changcheng Huang, Minzhe Li, and A. Srinivasan. A scalable path protection mechanism for guaranteed network reliability under multiple failures. *IEEE TRANSACTIONS ON RELIABILITY*, 56(2):254–267, June 2007.
[14] A. Markopoulou, G. Iannaccone, S. Bhattacharyya, C. N. Chuah, Y. Ganjali, and C. Diot. Characterization of failures in an operational ip backbone network. *IEEE/ACM Transactions on Networking*, 16(4):749–762, Aug. 2008.
[15] D.A. Waldman H. Mello, D.A.A. Schupke. A matrix-based analytical approach to connection unavailability estimation in shared backup path protection. *IEEE COMMUNICATIONS LETTERS*, VOL. 9, NO. 9:844–846, 2005.
[16] G.S. Waldman H. Schupke D.A. Mello, D.A.A. Quiterio. Specification of sla survivability requirements for optical path protected connections. In *Optical Fiber Communication Conference, 2006 and the 2006 National Fiber Optic Engineers Conference. OFC 2006*, 2006.
[17] B Rubino, G.; Sericola. Interval availability analysis using denumerable markov processes: application to multiprocessor subject to breakdowns and repair. *IEEE Transactions on Computers*, Volume 44, Issue 2:286 – 291, Feb. 1995.
[18] W. H. Sanders and II Obal, W. D. Dependability evaluation using ultrasan. In *Proc. Twenty-Third International Symposium on Fault-Tolerant Computing FTCS-23. Digest of Papers*, pages 674–679, 22–24 June 1993.
[19] The Norwegian Research Network UNINETT. Downtime statistics. Available at: http://drift.uninett.no/downs/.