The cost for meeting SLA dependability requirements; implications for customers and providers[☆]

Eirik L. Følstad[*], Bjarne E. Helvik[*]

Department of Telematics, Norwegian University of Science and Technology (NTNU), O.S. Bragstads plass 2B, N-7491 Trondheim, Norway.

Abstract

A Service Level Agreement (SLA) describes the service, the service-level objectives (SLOs), the price the customer should pay and the compensation if the SLOs are not met. There is a trade-off for the provider between the costs for improving the deployed service quality vs. probability of paying compensation. We propose how to estimate the provider's optimal service deployment. We show that the optimal deployed service quality is dependent on the SLOs, deployment cost, compensation and observation interval. A service deployment based on cost optimization results in targeted dependability objectives values that are significantly better than stated in the SLOs. The proposed approach provides valuable insight for an aggregator, who buys services from other providers, to negotiate adequate SLOs, price and compensation from the providers to make a valuable offer for its own customers.

Keywords: service level agreement, dependability, service deployment, optimization, evaluation

1. Introduction

Our society is dependent on critical infrastructures where failure free operations are of utmost importance. Examples of critical infrastructures are telecommunications, water supply systems, electrical power systems and banking and finance. Typical for providers of critical infrastructures is that they are part of compound service deliveries, where each of the parties is commercially and technically autonomous. The interdependencies between such parties are discussed in several papers, see e.g. [1–5].

It is important to control the service dependability through the delivery chain of several autonomous parties with legally binding agreements. The delivery of the services between parties and the related economic transactions may be regulated through Service Level Agreements (SLAs), see for instance [6–9]. SLAs for controlling the service dependability have been used in telecommunications and cloud computing, but little in conjunction with compound service deliveries provided by several critical systems and parties.

One aspect in an SLA is the specification of the quality of the service that shall be delivered. This is commonly specified as values of service-level objectives (SLOs). In setting up the SLAs, a fact that is paid surprisingly little attention, is that there are few failures during the typical observation interval for the SLOs, e.g., one year. Hence, due to the stochastic fluctuations of the failure and repair processes, what is observed may deviate significantly from the values representing the asymptotic average behaviour of the service. As a considerable compensation may be given to the customer if the SLOs are not met, neglecting this may have large economic implications. The probability of not meeting the availability SLO requirement during a finite observation interval was first dealt with by Goyal and Tantawi [10].

This paper investigates the relation between the dependability related SLOs, and the asymptotic values of these a service must be designed for. The objective is to provide insight that enables SLAs to be means for a cost quality trade-off beneficial to all parties that may be agreed upon. For the provider, it is a trade-off between the risk of paying compensation to the customer for not meeting the SLOs, and the investment in equipment and operations to improve the service dependability. For the service customer, it is a matter if risk sharing, where the consequences he will experi-

[☆]The work is funded by Telenor and Dept. of Telematics, NTNU.

[*]Corresponding author. Tel.: (+47) 920 44 740.

Email addresses: eirik.folstad@item.ntnu.no (Eirik L. Følstad), bjarne@item.ntnu.no (Bjarne E. Helvik)

ence if the provider does not meet the agreed SLO are compensated by the provider with a penalty stated in the SLA. Dependent on the kind of application or type of infrastructure this compensation may be significant. For the user, this may result in a more dependable and more expensive service than needed. Since a number of interwoven techno-economic relations are sought captured by a few SLOs, there are a number of pitfalls, and the use of SLAs may turn out to be counterproductive, [11, 12].

The objective of this paper is to provide insight that enables SLAs to be means for a cost quality trade-off beneficial to all parties that may be agreed upon. We show how the setting of dependability related SLOs affects the risk of the provider and some possible counteractions the provider may take to reduce the risk. The counteractions are modelled as use of capital expenditures and operational expenses for changing the actual service delivered.

The service delivered is described as a semi-Markov on-off model, being on when the service is delivered according to the temporal performance SLOs (delay, instantaneous loss, etc.) and off otherwise. Regarding the service from the customer's point of view, the naive approach is to derive the parameters of this model directly from the relevant SLOs. The more realistic approach adapted in this paper is to assume that the provider will operate the system in a way that maximizes his profit from the service delivered, which, due to the risk of compensation, is likely to be on the "safe side" of the SLOs. The customer, however will not be aware of the true parameters.

From the following items in the SLA: i) the observation interval, ii) the number of acceptable failures, iii) the maximum number of down times that may exceed a threshold, iv) the accumulated down time, v) the non-compliance compensation paid to the customer, and generic models proving a certain failure intensity and a tightly controlled repair handling time, a more realistic parameterization of a semi-Markov on-off model for the service provision is deduced. The on state defines when the service is delivered according to the temporal performance SLOs (in terms of provided functionality, response times, etc.) and off otherwise. This forms a basis for an understanding of how the SLOs and the cost parameters impact the actual service delivered as well as the economy of the service provision.

Most other quantitative analysis related to the dependability aspects of SLAs concentrate on the service availability. Specifying SLOs for the actual services having control with the failure intensity and down time duration may be just as important and is a salient aspect of this study.

Some work take into account that providers need to over-dimension, i.e. have a safety margin, in the service provision relative to what is specified in the dependability SLOs, to ensure good earnings on a service with penalties. Little work takes into account the provider's need to deploy the service quality to fulfil several dependability SLOs, each with a safety margin, for maximization of the profit. The safety margin is addressed in [13–16] with relation to the known asymptotic unavailability and how the safety margin for the provider is depended on the duration of the observation interval. Both [15] and [16] pinpoint how the safety margin is sensitive to the tail of the repair process. A methodology is presented in [17] to set the unavailability safety margin according to the tolerated customer compensation for a given observation interval. In [18] a framework is proposed for modelling the optimal investment for availability for a two state semi-Markov modelled system where the customer compensation is depending on the down time duration and its variance. A similar trade-off is formalized as an optimization problem in [19] with the objective to minimize the total cost of system improvements and compensations. A two-state Markov model is analysed in [20] where an upper bound for an insurance premium is calculated based upon the dependability related SLOs number of failures, cumulative down time duration and number of down times longer than a defined threshold. In our work we identify the optimal safety margin with respect to the same dependability SLOs as in [20] for an observation interval, investments and operational procedures and customer compensation. We propose to model investments and operational procedures to change the behaviour of the failure and repair processes.

With detailed information of a system numerical methods may be used for dependability analysis for the modelled system with a finite set of states, see e.g. [21, 22]. However, such detailed information and models causes the computing effort to be too demanding for studying the aspects in this paper.

An additional issue not included in this paper is how to manage the system to meet the SLOs for the offered service. Some discussions regarding SLA management may be found in e.g. [19, 23–25].

The rest of the paper is organized as follows. In Section 2 we present the on-off model in more detail. Section 3 deals with the SLA, its items i) - v) listed above and the generic cost models. In Section 4 the probability of violating the SLOs is derived. Based on this, Section 5 describes how a provider may operate the system so it maximizes his profit for the service. Section 6 de-

scribes how aggregated on-off models may be used to deal with compound service delivery systems. In Section 7 we present some case scenarios and discuss the influence of SLO items i)-v) on the inherent/asymptotic behaviour of the service provision and vice versa. Section 8 concludes the paper.

2. The on-off model

The behaviour of the system is modelled by an on-off semi-Markov model. This is of course a gross simplification relative to the behaviour of the real system, but is the most complicated model we may build based on the three dependability related SLOs, i.e., items ii) to iv) as introduced in Section 1 for a given observation interval. For more detailed models, more information has to be included in the SLA combined with in depth knowledge of the provider's system and operational procedures. Note that often just one dependability related SLO is found in SLAs for cloud and communication services, typically the availability, see e.g. [26–30], but here we favour additional dependability related SLOs to better describe the system requirements.

For the on-off semi-Markov model the system alters between the two states described by a failure process and a repair process. The failure process describes the event where the service delivered drops below the satisfactory performance level. Similarly, the repair process describes the event that combinations of fault handling mechanisms and operational procedures have restored the service to a satisfactory performance level. The performances SLOs define the satisfactory service delivery.

Since the failure process is an aggregate of many underlying failures, e.g. originating from hardware and softwareHW and SW in the sub systems, causing a non satisfactory service delivery, we assume that it is a non-homogeneous Poisson process (NHPP), cf. Palm-Khintchine's theorem, with an intensity $\lambda(y)$ where y is the calendar time. The NHPP is a generalization of a homogeneous Poisson process (HPP), which enables us to model seasonal, weekly and daily variations in the failure intensity, as well as trend and change due to change in load or operational conditions. The expected number of failures in an observation interval τ is $\Lambda(\tau) = \int_0^\tau \lambda(y)dy$ and the number of failures observed $N(\tau)$ is Poisson distributed

$$P(N(\tau) = n) = \Lambda(\tau)^n/n! \cdot e^{-\Lambda(\tau)} \quad (1)$$

The repair process is assumed to be any arbitrary process with independent and identically distributed i.i.d. restoration times, D with probability density function (PDF) $h(t)$ and cumulative distribution function (CDF) $H(t) = P(D \leq t) = \int_0^t h(y)dy$. This model of the failure and repair processes defines a semi-Markov process. Fig. 1 illustrates how to deduce a two state model from the SLOs in the SLA between a provider and a customer. Inherent parameters in the system are denoted $\acute{x}$, while the estimated are, to keep the notation simple, untagged.

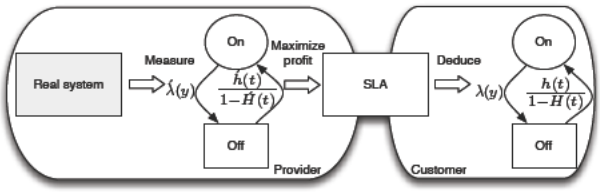

Figure 1: A two state semi-Markov model deduced from the SLA offering from a provider.

As described earlier in this section we have proposed to use three dependability related SLOs where one of them is related to failure intensity. With three dependability related SLOs we can therefore at most describe the repair time distribution with two parameters. In the rest of the paper we have used the gamma distribution with shape parameter α and scale parameter β, i.e., $H(t) = \Gamma(\alpha, t/\beta)/\Gamma(\alpha)$ where $\Gamma(\alpha) = \int_0^\infty e^{-x}x^{\alpha-1}dx$ and $\Gamma(\alpha, t/\beta) = \int_{t/\beta}^\infty e^{-x}x^{\alpha-1}dx$. The gamma distribution may, however, be replaced by any other distribution with two parameters that may be adjusted. Using a gamma distribution, The expected down time is $E[D] = \alpha\beta$ and its coefficient of variation $\sqrt{Var[D]}/E[D] = 1/\sqrt{\alpha}$. If providers or others have a model that better describes the down times of a service, the gamma distribution may in principle be replaced by any distribution with two adjustable parameters, likely at the cost of an increased computational complexity. As described later, the accumulated down time during an observation interval is the convolution of the down time distribution. For i.i.d. gamma distributed down times with parameters α, β, the n-fold convolution is a gamma distribution with parameters $n\alpha, \beta$, whereas in general recursive integrals are needed. Note that if an alternative model for the repair process with two adjustable parameters is available, e.g. based on operational data from a provider, the gamma distribution may be replaced by this, likely at the cost of an increased computational complexity.

Denote the accumulated time spent in the off state during the observation interval τ by $\Omega(\tau)$. In an on-off model, the number of failures $N(\tau)$ and $\Omega(\tau)$ will not be independent. However, in the context of highly depend-

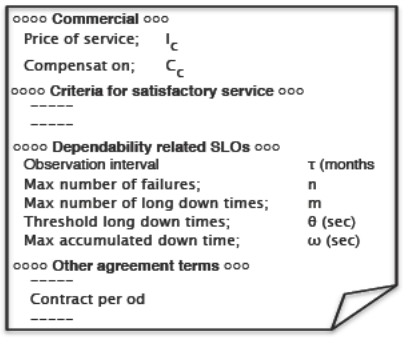

Figure 2: Example of parts of SLA content adapted from [6].

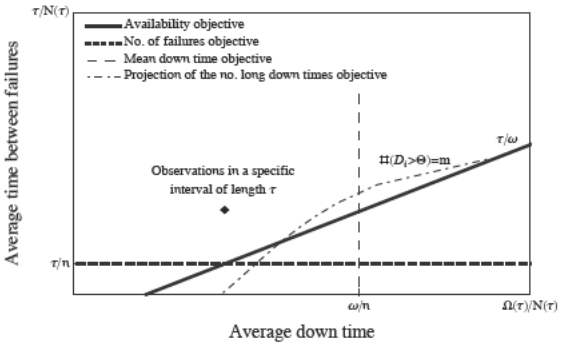

Figure 3: Illustration of dependability related SLOs.

able systems where the accumulated time spent in the off state is very short compared to the observation interval τ, independence is a good approximation. Hence, when

$$\lambda(y) \ll 1/E[D], \forall y \in [0, \tau] \quad (2)$$

holds, we assume that

$$P(\Omega(\tau) \leq \omega, N(\tau) = n) \approx P(N(\tau) = n) \cdot H^{\otimes n}(\omega) \quad (3)$$

where $P(N(\tau) = n)$ is given in (1) and $H^{\otimes x}(t)$ is the CDF of an n-fold convolution of the down time duration. This may in general be derived with the recursion $H^{\otimes x}(t) = \int_0^\infty H^{\otimes(x-1)}(t-y)h(y)dy$. A validation of this approximation relative to the exact joint distribution is given in Appendix A.

3. The SLA and the provider's cost model

In a rational market with competition, a provider has to deliver services that are competitive and attractive with respect to quality and price. The SLA formalizes the service to be delivered, price and compensation, and should reflect the deployed service quality such that the (estimated) business for the provider in the long run is attractive for its stakeholders.

3.1. Dependability related service-level objectives

The SLOs defined in an SLA must be measurable over a finite period. In SLAs and in most research, the contracted unavailability (or availability) has been the most dominant dependability related SLO. In contrast to this, the service requirements and SLOs are most likely different for different usage of the service. For some customers , e.g. end-users performing on-line surgery, the mean time between failures may be the most important dependability attribute, while for e.g. the retail business others restrictions on long outage times might be the most critical requirement , and the finance business is reputed to require high availability.

Fig. 2, adapted from [6], shows parts of an SLA that provides relevant information referred in this paper. Besides the commercial parts, the criteria for satisfactory service are defined. These criteria should describe the performance requirements for the service as well as the usage limitations the customer must comply with. The dependability SLOs define the agreed deviations for the service delivery. The dependability SLOs are independent of the underlying failure and repair processes, and do not require a stationary failure process.

Note that there is a distinct difference between the SLA contract period and the observation interval. The contract period defines the legal commitment of the validity period of the SLA, whereas the observation interval is the agreed time interval(s), during which the dependability SLOs are calculated. Within a contract period there may exist several observation intervals. In this paper the same observation interval is used for all the dependability SLOs.

We denote the number of failures during the observation interval, τ, by $N(\tau)$, the number of down times during the interval that exceeds θ by $M(\tau)$, and accumulated down time during the interval by $\Omega(\tau)$. These are dependent stochastic variables. The dependability related SLOs for the finite observation interval τ are then;

- $N(\tau) \leq n$; maximum number of failures.
- $M(\tau) \leq m_\theta$; maximum number of down times longer than the threshold θ. In the rest of the paper we will use m for m_θ since m is always related to a threshold defined by θ.
- $\Omega(\tau) \leq \omega$; maximum accumulated down time.

Fig. 3 illustrates relations between some dependability related SLOs. During the observation interval τ, we will get an average time between failures $\tau/N(\tau)$ and an av-

erage down time $\Omega(\tau)/N(\tau)$, which represents a point in the plane. The bold dashed line represents the maximum number of failure objective τ/n. The availability objective is represented by the straight bold line representing the inverse of the observed unavailability, i.e., τ/ω. The vertical dashed lines represent requirements for maximum average down times, not discussed further in this paper, and the dash-dot curve is a projection of a limitation of a restriction on the number of long down times for a given m, θ and τ. To comply with the SLOs, the point representing the observations of the system must be above and to the left of the curves representing the objectives. The optimum service deployment point will most likely be in this region, sufficiently away from the constraints represented by the objectives, to accommodate the stochastic fluctuations of the failure and repair processes for the agreed observation interval. This optimum service deployment point will be further investigated and discussed in the following sections.

3.2. Provider's cost model

To be able to offer services the provider has to use capital expenditures (CAPEX) and operational expenses (OPEX) for hardware/softwareHW/SW, implementation, maintenance etc. For simplicity, we denote these costs as deployment cost. The inherent quality in the system is dependent on the deployment cost, and increasing the quality implies increased deployment cost. Given the inherent system quality, the provider has to offer services that comply with SLOs in the SLA. If the SLOs are not fulfilled over the observation interval, the provider has to pay a compensation C_c to the customer. Here, the compensation is modelled as a fixed amount irrespective of which SLO that is violated and the severeness of the violation. We consider this as representative for commercial services, although some service providers use different compensation schemes. For instance the cloud service providers Amazon [26] and Microsoft [27] do only compensate for availability violations. The network communication providers Nextgen [28] and CenturyLink [29] have guarantees for both availability and repair times, but do only compensate for availability violations. Verizon [30] guarantees 100% availability for some Internet access services and compensates violation of repair times in addition to the availability violation. CenturyLink and Verizon also compensate for IP packets performance violations such as jitter and delay. For specific cases, (4) below may be modified to adopt to these.

In the following let $\hat{S}$ denote the contracted service quality, i.e., what is stated in the SLA with the values n, m and ω and S the properties of the deployed service, i.e., estimated with the parameters $\Lambda(\tau)$, α and β. A provider's profit $R(\hat{S}\,|S)$ may be viewed as

$$\begin{aligned} R(\hat{S}\,|S) &= I_{c|\hat{S}}(\tau) - C_{d|S} - C_{c|\hat{S}} P_{\hat{S}\,|S}(\tau) \\ &= I_{c|n,m,\omega}(\tau) - C_{d|\Lambda(\tau),\alpha,\beta} - C_{c|n,m,\omega} P_{n,m,\omega|\Lambda(\tau),\alpha,\beta}(\tau) \end{aligned} \tag{4}$$

where $C_{d|S}$ is the deployment cost, $I_{c|\hat{S}}(\tau)$ is the selling price for the service, i.e., the provider's income, and $P_{\hat{S}\,|S}(\tau) = P_{n,m,\omega|\Lambda(\tau),\alpha,\beta}(\tau)$ is the probability of violating the SLOs. The probability of violating the SLOs will be further explained in Section 4.

The compensation cost for the provider is linked to the probability of violating the SLOs. The compensation cost can be reduced at the expense of the deployment cost that affects the deployed service quality S. For the provider there is an optimum profitable service deployment where the sum of deployment and compensation costs are minimized. The estimation of the optimal deployment, described by the failure and repair processes, will be derived in the following.

4. Probability of violating service-level objectives

The provider's probability of payment of compensation is linked to the probability of violating the SLOs as defined in Section 3.1. The compensation has to be paid if one or more of the requirements are not met during the observation interval. The probability of breaching one or several of the SLOs by the provider can be expressed as

$$\begin{aligned} P_{n,m,\omega|\Lambda(\tau),\alpha,\beta}(\tau) &= P[N(\tau) > n \cup M(\tau) > m \cup \Omega(\tau) > \omega] \\ &= P[N(\tau) > n] + P[M(\tau) > m \cap N(\tau) \leq n] \\ &+ P[\Omega(\tau) > \omega \cap N(\tau) \leq n \cap M(\tau) \leq m] \end{aligned} \tag{5}$$

The probability $P_{n,m,\omega|\Lambda(\tau),\alpha,\beta}(\tau)$ separates into three disjoint sets of events. Each of these disjoint events will be derived in the following.

4.1. Maximum number of failures

Having defined the semi-Markov process, cf. Fig. 1, the number of failures during an observation interval is equal to the on/off transitions. Assuming that the down times may be negligible with respect to the up times, the probability of violating the number of failure occurrences during the observation interval may be expressed as

$$P[N(\tau) > n] = 1 - \sum_{i=0}^{n} \frac{\Lambda(\tau)^i}{i!} e^{-\Lambda(\tau)} \tag{6}$$

where $\Lambda(\tau)$ is the expected number of failures during the observation interval τ.

4.2. Maximum number of long down times

The repair process is modelled as any arbitrary process with independent and identically distributed i.i.d. restoration times as described in Section 2. In an environment with highly dependable systems the provider's probability of violating the maximum number of long down times may be expressed with independent Bernoulli variables yielding the approximation

$$\begin{aligned} &P[M(\tau) > m \cap N(\tau) \leq n] \qquad (7)\\ &= \sum_{i=0}^{n} P[M(\tau) > m | N(\tau) = i] P[N(\tau) = i] \\ &= \sum_{i=m+1}^{n} \frac{(\Lambda(\tau))^i}{i!} e^{-\Lambda(\tau)} \Big(1 - \sum_{j=0}^{m} \binom{i}{j} (1 - H(\theta)^j H(\theta)^{i-j})\Big) \end{aligned}$$

where $H(\theta)$ is the estimated probability of a down time duration less or equal to the threshold θ. In (7) the probability of the number of failures that may violate the maximum number of long times are from $m+1$ to n since (6) already has counted for violating the number of failures. Last summation of (7) expresses the probability of the combinations of down times for the given number of failures that do not violate the maximum number of long down times.

4.3. Maximum accumulated down time

Takacs [31] derived the probability of violating the accumulated down time $P[\Omega(\tau) \leq t]$ for a two state system with $\hat{G}(t)$ and $H(t)$ as the CDFs of times between failures and duration of down times respectively as

$$P[\Omega(\tau) \leq \omega] = \sum_{n=0}^{\infty} H^{\otimes n}(\omega)[\hat{G}^{\otimes n}(\tau - \omega) - \hat{G}^{\otimes n+1}(\tau - \omega)] \qquad (8)$$

where $\otimes$ is the convolution operator and $H^{\otimes x}(t)$ is the x-fold convolution of a given CDF of the distribution of the down times. In the case by Takacs the failure process $\hat{G}(t)$ is homogeneous. In our case, the estimated failure process is non-homogeneous where $G(t) = 1 - e^{-\Lambda(t)}$. Further more, closed form solutions exist only when $H(t)$ is negative exponential or deterministic. Hence, an approximation is needed.

As motivated at the end of Section 2 and validated in Appendix A, we may use (3) to get the following approximation for highly dependable systems

$$P[\Omega(\tau) \leq \omega] \approx \sum_{n=0}^{\infty} H^{\otimes n}(\omega) P[N(\tau) = n] \qquad (9)$$

The conditional probability that the provider violates the maximum accumulated down time becomes

$$\begin{aligned} &P[\Omega(\tau) > \omega \cap N(\tau) \leq n \cap M(\tau) \leq m] \\ &= \sum_{i=0}^{n} \sum_{j=0}^{Min[i,m]} P[\Omega(\tau) > \omega | M(\tau) \leq j \cap N(\tau) \leq i] \\ &\cdot P[M(\tau) \leq j | N(\tau) = i] P[N(\tau) = i] \\ &= \int_{\omega}^{\infty} \Big[\sum_{i=0}^{n} \frac{(\Lambda(\tau))^i}{i!} e^{-\Lambda(\tau)} \sum_{j=0}^{Min[i,m]} \binom{i}{j} (1 - H(\theta))^j H(\theta)^{i-j} \\ &\cdot h^{\otimes j}(t | t > \theta) \otimes h^{\otimes (i-j)}(t | t \leq \theta)\Big] dt \qquad (10) \end{aligned}$$

where $h^{\otimes n}(t)$ is the n-fold convolution of the estimated PDF of duration of down times. For gamma distributed duration of down times with shape α and scale β, $h(t) = (e^{-t/\beta} t^{\alpha-1} \beta^{-\alpha})/\Gamma[\alpha]$. The first summation in (10) expresses the probability of the number of failures that may contribute to violating the accumulated down time, yielding a maximum of n failures, since (6) already has counted for violating the number of failures. Only the probabilities of combinations of down times that do not violate the maximum number of long down times are included, given by the second summation in (10) going from 0 to the minimum of i and m, since (7) counts for violating the number of long down times. The integral of (10) expresses the probability of combinations of down times, longer than the threshold θ and equal to or shorter than θ, that violates the accumulated down time.

The convolutions of mixtures of right-truncated $h(t|t > \theta)$ and left-truncated $h(t|t \leq \theta)$ distributions of down times are numerically obtained by discretization and using the discrete Fourier transform. For efficiency, the summation in (10) is performed in the Fourier (frequency) domain. The accuracy obtained is validated. Note that this approach is flexible, as it allows us to use arbitrary $H(t)$ in the model.

Inserting (6), (7) and (10) into (5) yields the probability $P_{n,m,\omega|\Lambda(\tau),\alpha,\beta}(\tau)$ of not complying with the SLOs.

5. Optimizing providers profit

To maximize (4), i.e., the providers profit, we need a relation between the deployment cost $C_{d|S}$ and the asymptotic service quality S. This will depend on a variety of factors and the specific options that exist for the system, which delivers the service, as well as the options for operating and maintaining the system. To illustrate, to get insights and to obtain indicative results, we introduce a generic relationship, that captures the salient factors, but which is not claimed to be exact. In cases

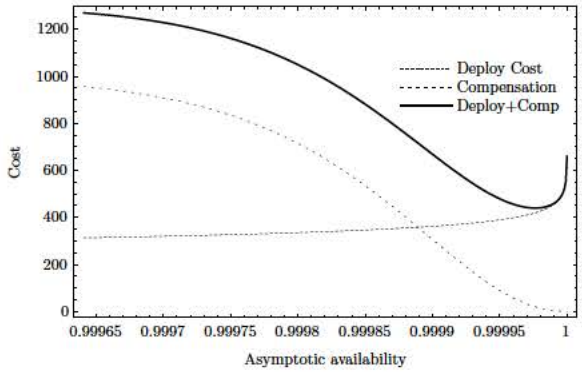

Figure 4: Deployment for asymptotic availability vs. cost for different failures intensities (λ) and gamma distributed duration of down times with $\alpha = 6$ and $\beta = 300$ sec for SLOs and costs found in Table 1. A minimum cost may be observed.

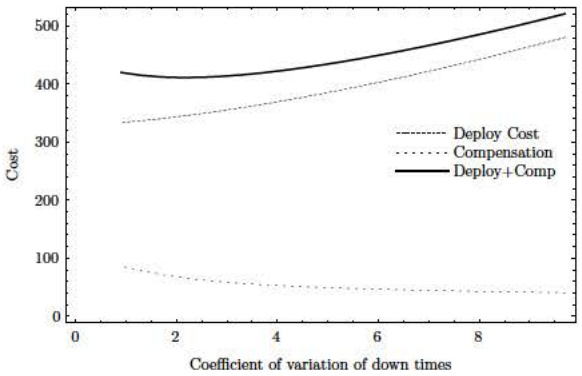

Figure 5: Deployment for coefficient of variation of down times vs. cost for different value of shape parameter of a gamma distributed duration of down times for SLOs and cost found in Table 1. $E[D] = 1800$ sec. A minimum cost may be observed.

where specific design, configuration and/or operation options are available, the generic relations may be replaced by these and a corresponding (integer) optimization performed.

In establishing the generic relations, consider The deployment cost consists of two parts, one related to the asymptotic unavailability, $C_{du|S}$, of the service and one part related to the coefficient of variation of the duration of down times, $C_{dcv|S}$. This yields the deployment cost $C_{d|S} = C_{du|S} + C_{dcv|S}$. The deployment cost is increasing with decreasing unavailability. Likewise, as the coefficient of variation of the duration of down times is decreased, the deployment cost is increasing. To keep the illustration simple, we also use a HPP for the failure process, i.e., $\lambda(y) = \lambda$, $y \in [0, \tau]$.

The profit given by (4) may be refined with the probability of violating the SLOs given by (5) and the cost models to estimate the optimal deployment as

$$R_{n,m,\omega|\lambda,\alpha,\beta}(\tau) = I_{c|n,m,\omega}(\tau) - \underset{\lambda,\alpha,\beta}{\arg\min}\Big(C_{du|\lambda,\alpha,\beta} + C_{dcv|\alpha,\beta} + C_{c|n,m,\omega}P_{n,m,\omega|\lambda,\alpha,\beta}(\tau)\Big) \quad (11)$$

where we, to keep the illustration simple, use a HPP for the failure process, i.e., $\lambda(y) = \lambda$, $y \in [0, \tau]$.

To establish an aggregated cost vs. availability model, we introduce a reference unavailability U_0 and a degree of replication δ. In simple single element systems, duplication corresponds to $\delta = 2$, triplication to $\delta = 3$, etc. The obtained unavailability U is then modelled as $U = U_0^{\delta}$, which yields

$$\delta = \frac{\ln U}{\ln U_0} \approx \frac{\ln(\lambda E[D])}{\ln U_0} \quad (12)$$

The deployment cost typically increases slightly faster than δ, say δ^{ν}, since spare routes etc. tend to be more costly than the primary, i.e., ν is close to, but a little larger than 1. For instance, Telenor [32] offers a spare route at a higher price than the primary while Amazon [26] offers virtual computing instances at the same price, but if two instances are in different availability zones an additional cost is introduced for data transfer between the independent zones. Hence, a simple aggregated relation between asymptotic unavailability cost and model parameters becomes

$$C_{du|\lambda,\alpha,\beta} = \Big(\frac{-\ln U}{-\ln U_0}\Big)^{\nu} C_0 = \Big(-\ln(U)\Big)^{\nu} C_{du0} \approx \Big(-\ln(\lambda\alpha\beta)\Big)^{\nu} C_{du0} \quad (13)$$

where C_{du0} is the reference deployment cost for the reference unavailability U_0. If the provider wants to reduce the variance of the duration of down times, more money has to be put in deployment and operations (CAPEX) and (OPEX). The variance of the duration of the down times expresses the providers' knowledge, preparations and processes for keeping the duration of the down times within certain limits, as well as the deployed systems' mechanisms for failure corrections. To model the aggregated cost vs. variance of the duration of down times we introduce the reference cost C_{dcv0} for a reference coefficient of down time durations and a degree of reduction η. The deployment cost increases with reduced variation, and deployment cost for reducing the coefficient of variation of the service interruption may be modelled as

$$C_{dcv|\alpha,\beta} = C_{dcv0}\Big(\frac{E[D]}{\sqrt{Var[D]}}\Big)^{\eta} = C_{dcv0}\alpha^{\frac{\eta}{2}} \quad (14)$$

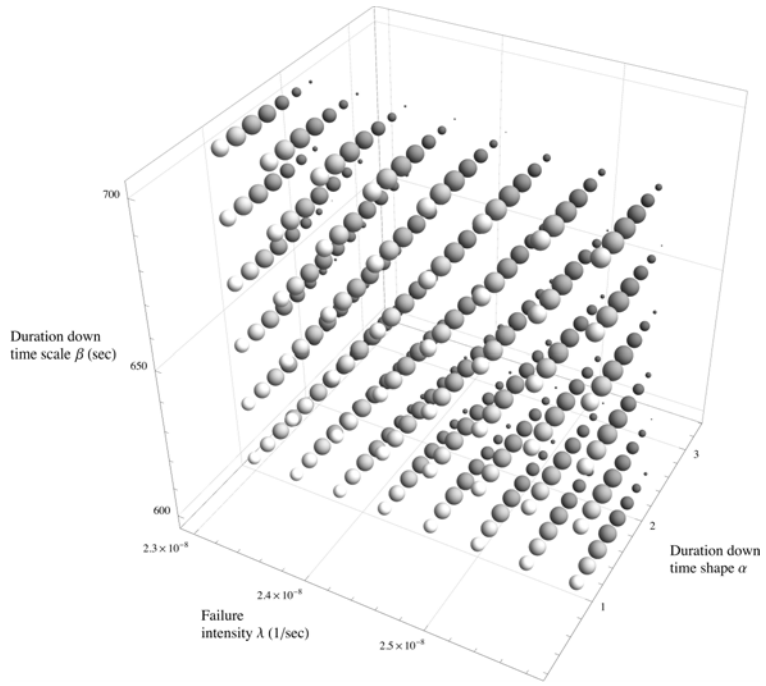

Figure 6: Providers' profit is dependent on failures intensity (λ) and shape (α) and scale (β) parameters of the distribution down time. The volumes of the bubbles are proportional to the profit. The whitest bubble indicates the highest availability and darkest indicates the lowest availability.

Telstra [33] offers customers to buy different SLA premium restoration services to reduce the time to repair dependent on time of the day. This example illustrates how a change of the repair process with the aim to reduce the not only the mean, but also the variance has a cost.

To illustrate how the compensation and deployment cost for asymptotic availability, $A = (\lambda\alpha\beta + 1)^{-1}$ is depending on different failures intensities, Fig. 4 depicts an example where the down times are gamma distributed with $\alpha = 6$ and $\beta = 300$ sec for SLOs and costs found in Table 1. Note that there is a distinct minimum of the total cost. Likewise, Fig. 5 illustrates an example of compensation and deployment cost for different coefficient of variation for gamma distributed down time durations where $E[D] = 1800$ sec.

6. Aggregation of several on-off models

In this section we describe how to find the properties of an aggregated system from a number of SLAs. An aggregated system provides services composed of sub services from the underlying systems. Each underlying system is operated by an autonomous service provider that delivers its sub service in accordance with an SLA with the set of SLOs defined in Section 3.1.

A structure function may be used to describe the aggregated system as composed of underlying systems A structure function in minimal product-of-sum form focuses on the combinations of failed underlying systems that make the aggregated system to fail. Each of the maxterm in a minimal product-of-sum corresponds to a minimal cut set, see e.g. [34] for an introduction. A minimal cut set contains a number of underlying systems. In the following denote the set of underlying systems as $\mathcal{J}$, the number of minimal cuts sets as k and the set of underlying systems in minimal cut set x as J_x where $x \in 1, ..., k$.

In [35] the Palm distribution of the duration of down times for an aggregated system is derived. With the assumptions that the underlying systems are independent, and that only one cut-set of the aggregated system yields system failure at any time, which is permissible for a highly available system, the Palm distribution of the aggregated system's down time duration is

$$1 - H_A(t) = \sum_{x=1}^{k} \frac{\Lambda_x}{\sum_{r=1}^{k} \Lambda_r} [1 - H_x(t)] \qquad (15)$$

where Λ_x and $H_x(t)$ represent failure intensity and cu-

mulative duration of down time distribution for minimal cut-set x respectively. In [35] the derivations of Λ_x and $H_x(t)$ may be found as

$$\Lambda_x^{-1} = \frac{\left(\prod_{i\in J_x} E[D_i]/(E[D_i]+E[S_i])\right)^{-1} - 1}{\sum_{i\in J_x} E^{-1}[D_i]} \quad (16)$$

and

$$1 - H_x(t) = \sum_{i\in J_X} \frac{E^{-1}[D_i]}{\sum_{j\in J_x} E^{-1}[D_j]} [1 - H_i(t)] \cdot \prod_{j\in J_x\setminus\{i\}} \frac{\int_t^\infty 1 - H_j(s)\ ds}{E[D_j]} \quad (17)$$

where $E[S_i]$ is the expected time between failures for the underlying system i and $E[D_i]$ is its expected duration of down times with CDF $H_i(t)$. See Section 2 for relations to the model parameters.

Assume that a duration of down time caused by cut-set J_x is followed by the state where all systems in the cut-set are again functional. This implies an approximation of the intensity of failures for the cut-set in concern. With this approximation the distribution of the time between failures for the aggregated system can be given as

$$g_A(t) \approx \Lambda_A e^{-\Lambda_A t}, \text{where } \Lambda_A = \sum_{x=1}^{k} \Lambda_x \quad (18)$$

With the time between failures and duration of down time distributions derived for the aggregated system, an aggregator may form relations between its own SLA and the SLAs for the set of underlying sub systems $\mathcal{J}$ as

$$R_{n,m,\omega|\Lambda_A,H_A(t)}(\tau) = I_{c|n,m,\omega}(\tau) - \sum_{\forall j\in\mathcal{J}} \left(I_{c_j|n_j,m_j,\omega_j}(\tau) - C_{c_j|n_j,m_j,\omega_j} P_{n_j,m_j,\omega_j|\lambda_j,\alpha_j,\beta_j}(\tau)\right) - C_{c|n,m,\omega} P_{n,m,\omega|\lambda,\alpha,\beta}(\tau) \quad (19)$$

Note that the deployment cost in (19) is related to the price for the services delivered by underlying systems reduced with the expected compensations.

7. Case scenarios

In this section we exemplify the on-off model deduced from the SLA and other parameters as described in Sections 4, 5 and 6. A discussion of the sensitivity of the parameters is provided.

First we want to recall the main approximations and assumptions used. Appendix A validates the approximation of the independence of number of failures and the accumulated down times and discusses the insensitivity to fluctuations in the failure intensity. The main assumptions are related to the dependability SLOs given in Section 3.1. Note that SLAs typically define the maximum number of failures. In Sections 7.1 and 7.2 we deal with the problem under the assumption of a constant failure intensity, i.e., a homogeneous failure process. In Section 7.4 the result from a simulation study is shown, demonstrating that the results obtained are insensitive to fluctuations in the failure intensity. The model of the deployment cost is as given in Section 5. A compensation is assumed to be paid if one or several of the SLOs are violated during an observation interval as described Section 4.

Table 1: SLOs and commercial terms as regulated in an example SLA with the assumed deploy cost parameters.

Group	Parameter	Symbol	Value
Commercial	Cost (income for provider)	$I_\tau(n,m,\omega)$	400
	Compensation Cost	C_c	1000
SLO	Observation interval (months)	τ	12
	Max number of failures	n	3
	Max number long down times	m	1
	Threshold long down time (sec)	θ	1800
	Max acc. down time (sec)	ω	4500
Deploy cost	Deploy Cost, unavail reference	C_{du0}	18
	Deploy Cost, unavail cost factor	ν	1,25
	Deploy Cost, CV reference	C_{cv0}	5
	Deploy Cost, CV cost factor	η	3
Optimized	Failure intensity, Possion (1/sec)	λ	$2.46 \cdot 10^{-8}$
	Duration outage, Gamma	α (shape)	1.40
	Duration outage, Gamma (sec)	β (scale)	642

7.1. Reference scenario

In Table 1 the assumed values of the SLOs, commercial terms and deployment cost parameters are given as a reference scenario. The values of the SLOs are realistic for a highly dependable system, whereas the deployment cost parameters are examples to show how these impact the estimated deployed dependability quality of the system.

Mathematica [36] is used to solve the numerical optimization. The optimal case, i.e., assuming operator behaviour from (11) is obtained, corresponding to the values for λ, α and β are included in Table 1.

To study the parameters' sensitivity of estimated failure and repair processes on the operators profit a 3D plot for the profit for a range of values of λ, α and β is given in Fig. 6. The plotted ranges enclose the optimal values yielding the maximum profit 26.2. In the figure the profit is proportional to the volume of the bubble, i.e., the bubble with the largest volume is the most profitable combination of λ, α and β. As indicated by Fig. 6 the

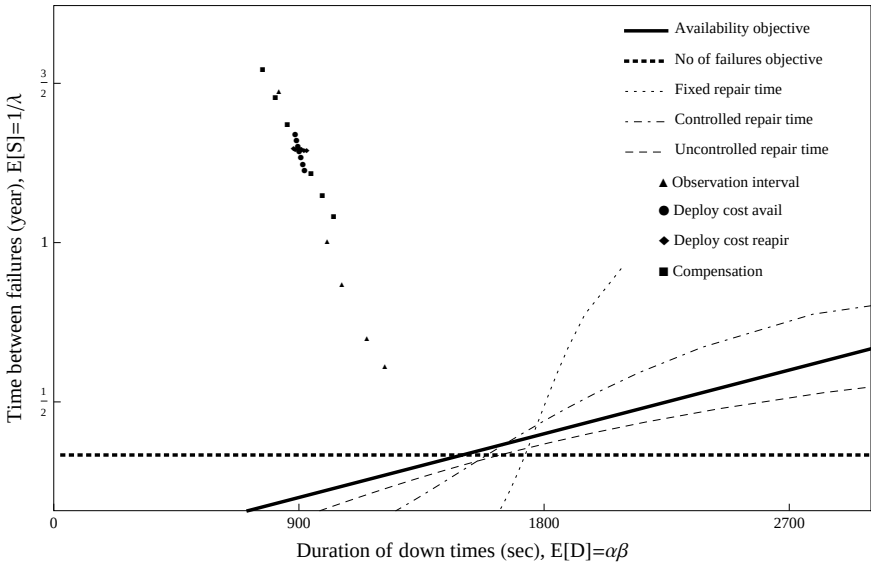

Figure 7: Optimal system deployment for the given SLOs with variations in cost parameters and observation interval. Cost parameters are changed with ±15% compared with values given in Table 1. Observation interval = {8, 12, 24, 36, 72, 120} months with scaled requirements from Table 1.

profit has a global optimum and is not very sensitive to the plotted range of parameter values. The asymptotic availability based on the parameter values of the failure and repair processes is given by the grey tones of the bubbles in Fig. 6. As may be found in the figure, a high profitable deployed service does not correspond to the highest availability.

7.2. Deployed vs. required quality

Examples of relations between contracted SLOs are depicted in Fig. 3. Inherent system parameters that satisfy these objectives are in the upper left area enclosed by the y-axis and the SLOs.

A refinement of Fig. 3 is provided in Fig. 7 where the requirements are the contracted SLOs given by Table 1. The maximum number of failures objective and maximum unavailability objective are represented in the figure, derived from n/τ and $1/U = \tau/\omega$ respectively. There is no mean duration of down time SLO, but the maximum down time threshold θ puts constraint on the mean duration of down times. In the figure three lines are illustrating how θ influences the mean duration of down times depending on the providers control of the repair process. These lines are named; fixed, controlled and uncontrolled repair time to associate the provider's capability to manage the repair process in terms of coefficient of variation, i.e., $1/\sqrt{\alpha}$, of the repair time. As an example the values of $\alpha = \{100, 5, 1\}$ have been used for illustrating the fixed, controlled and uncontrolled repair time and for each the following is solved with respect to β

$$1 - \frac{\Gamma(\alpha, \theta/\beta)}{\Gamma(\alpha)} = \frac{\gamma}{\tilde{\lambda}} \tag{20}$$

where $\gamma = m/\tau$ is the intensity of the long down times and $\tilde{\lambda}$ is the intensity of service failures. The equations are solved for different values of $\tilde{\lambda} > \gamma$. The higher the shape parameter α gets, representing a fixed repair time, the closer will $E[D]$ get the to threshold θ.

To investigate how the deployed dependability qualities are dependent on the SLOs and cost parameters a number of different values are studied for the parameters. As a reference scenario the parameters and cost as given in Table 1 is used. From this reference point one of the parameters is changed and the corresponding optimal service quality deployment is derived and depicted in Fig. 7. For the parameters compensation (C_c), deployment cost unavailability factor (ν) and deployment cost repair (η) the changes are ±15% in steps of ±5%, while for the observation interval (τ) we have used $\{4, 8, 12, 24, 36, 72, 120\}$ months. The SLOs n, m and ω and costs are scaled in proportion with the observation interval. For the number of long outages, $m = 1$ at 12 months, this is not feasible. Hence, for $\tau = \{4, 8\}$ we use $m = 1$.

Both compensation and deployment cost affect the deployed dependability qualities as indicated in Fig. 7. When the compensation (C_c) is decreased, the deployed dependability qualities approach the SLOs. Similarly,

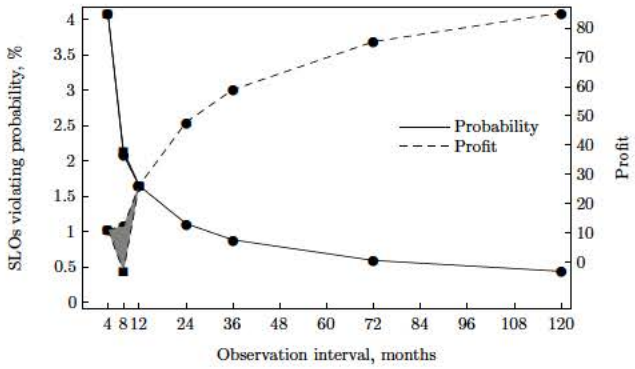

Figure 8: Sensitivity of the observation interval for violating SLOs and the profit per year for the values given in Table 1.

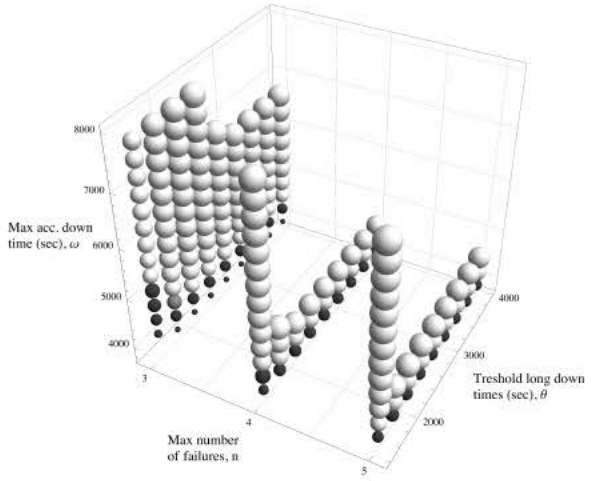

Figure 9: The maximum provider's profit, represented as the volume of the bubbles, for a range of values of the SLOs n, ω and θ. All other parameters kept at the values in Table 1. The medium grey bubble indicates the profit of the reference scenario in Table 1. Dark grey bubbles indicate less profit than this scenario, light grey larger.

the higher the deployment cost gets, the further away is the deployed dependability quality the SLOs. The interpretation is that the lower deployment cost gets relative to the expected compensation, the better it is for the provider to deploy better dependability qualities for the service. For the observation interval (τ) the deployed dependability quality tends to get closer to the SLOs when the interval gets longer as indicated in Fig. 7. When the observation interval is four months, the service failure intensity gets very low ($4.02 \cdot 10^{-9}$), while the mean duration of down time gets very high (more than 10.000 sec), not shown in the figure. The model provides a good insight on how the observation interval, in particular how a relative short interval, affects the optimal deployment. In the example with $\tau = 4$ months the probability of a failure during the interval gets very low, but the consequence in terms of expected outage duration gets very long.

The observation interval is known to have a significant effect on the probability of violating the availability SLO [13, 37]. In the more complex setting discussed here, Fig. 8 shows probability of violating the SLOs and the provider's profit for optimized values of λ, α and β for the actual observation interval. It is clearly indicated that the probability of violating the SLOs is higher for a short observation interval than for the longer intervals. In addition, the figure illustrates that the profit is higher with longer observation interval, which implies that the provider should use an observation interval that is equal the contract period of the SLA.

The shaded area in Fig. 8 provides a probability zone caused by the scaling of the maximum number of long down times. For the intervals of $\tau = 4$ and $\tau = 8$ months curves for both $m = 0$ and $m = 1$ are plotted. As can be observed, the $m = 0$ requirement is a more demanding than $m = 1$ for an observation interval of 8 months, while it has no effect when the observation interval is 4 months, since the maximum accumulated down time is the most demanding requirement in both cases.

7.3. *Negotiation of service-level objectives*

It may be unfeasible in practice to control the failure rates and repair times distributions in a continuous way for a system as assumed in the previous sections. However, this kind of analysis may also be a useful tool in the negotiations between the service provider and the customers. Fig. 9 depicts the provider's profit for a range of values of the SLOs parameters n, ω and θ, while the income from the service is kept constant. Starting with the reference scenario of Table 1 (the medium grey bubble), it may be seen how the potential profit from the service may change as the service requirements are strengthened or weakened by altering the SLO parameters n, ω and θ. If the customer has a similar relation between the value of the service to him, i.e., what he is willing to pay as $I_{c|n,\omega,\theta}$, and the parameter values, this may help to find a set of SLO parameters that yields a better trade-off for both, tentatively Pareto optimality.

7.4. *Non-homogeneous Poisson failure process*

Earlier in Section 7 the probability of violating the dependability SLOs, i.e., the quantity $P_{n,m,\omega|\lambda,\alpha,\beta}(\tau) \rightarrow P_{\text{opt}}(\tau)$ in (5), was obtained as a key quantity in finding the optimal operational parameters for a system. These

parameters were obtained for an HPP failure process. SLAs do not typically have any requirement for the failure process, just the values observed over the given interval. In this section, we will investigate whether fluctuations in the failure process has any significant influence on the result obtained, as long as the requirement of short down times, i.e., (2) is met. This is done by performing a simulation study where the probability of violating the SLOs is compared between HPP and NHPP failure processes for the reference scenario as given in Table 1 with the optimal operational parameters. If this probability, $P_{\text{obs}}(\tau)$, is nearly the same as $P_{\text{opt}}(\tau)$, cf. (11), the optimal operation point is insensitive to fluctuations in the failure process. The simulations are performed without neglecting down times in the failure generation process, so the robustness of the assumption in (3) is demonstrated as well.

An NHPP failure process may have a variety of different time varying failure intensities. For instance, in [38] a data set for covering more than 1000 consecutive days from a cellular network operator was analyzed and it was found that the failure intensity had strong cyclic effects of 12 hours, 24 hours and 7 days. The following model may be used for the variations of failure intensity for an NHPP failure process with multiple cyclic and trend effects

$$\lambda(y) = \sum_{i=0}^{s} \psi_i y^i + \sum_{j=1}^{k} \rho_j \sin(\varpi_j y + \varrho_j) \qquad (21)$$

where y is the calendar time and the periods of the cyclic effects are given by $2\pi/\varpi_j$.

For the reference scenario as given in Table 1 the optimized HPP failure intensity, λ, is $2.46 \cdot 10^{-8}$ for gamma distributed down times with the parameters $\alpha = 1.40$ and $\beta = 642$. As may be found in Figure 8 the probability of violating the SLOs is $P_{\text{opt}}(\tau) = 1.652\%$. To compare this result with an NHPP with the same expected number of failures in the observation interval τ the failure intensity fulfills $\Lambda(\tau) = \int_0^\tau \lambda(y)dy = \lambda\tau$.

The simulator was implemented in Mathematica 8 [36] running on a Mac with 2.5GHz Intel Core 2 Duo CPU with 4GB memory and OS X version 10.6.8. A total of 400000 years were simulated, i.e., equal 400000 observation intervals.

In Table 2 the estimated probabilities, P_{obs}, with the 95% confidence intervals are found by simulation for the HPP case without neglecting down times (first line) and three different NHPP cases where the cyclic effects are given in hours or days. Note for the three NHPP cases the phase shifts, ϱ_j, are 0. The probability of violating the SLOs by simulation for the three NHPP cases and HPP are found to be very similar and close to the calculated optimal value. This demonstrates the insensitivity to fluctuations in the failure process, as well as the approximation in (3).

Table 2: Simulation results for the probability of violating the SLOs for the reference scenario in Table 1 for the HPP (first line) and three NHPP failure processes with the same expected number of failures during the observation interval τ. The $\varrho_j = 0$ for all NHPP.

ψ_0	ψ_1	ρ_1	ϖ_1	ρ_2	ϖ_2	P_{obs} [%]
$2.46 \cdot 10^{-8}$	0	0	0	0	0	1.655 ± 0.038
$3.0 \cdot 10^{-8}$	$-3.43 \cdot 10^{-16}$	$0.5 \cdot 10^{-8}$	24h	$0.2 \cdot 10^{-8}$	7d	1.667 ± 0.051
$3.0 \cdot 10^{-8}$	$-3.49 \cdot 10^{-16}$	$1.5 \cdot 10^{-8}$	30d	$0.2 \cdot 10^{-8}$	7d	1.608 ± 0.040
$2.8 \cdot 10^{-8}$	$-2.18 \cdot 10^{-16}$	$1.5 \cdot 10^{-8}$	90d	$0.5 \cdot 10^{-8}$	7d	1.665 ± 0.035
$2.46 \cdot 10^{-8}$	From optimization procedure: $P_{\text{opt}}(\tau)$					1.652

7.5. Aggregated systems

Assume a system consisting of several underlying systems as depicted in Fig. 10. This scenario describes how an aggregator may combine offerings from several providers to be able to offer a new service. The aggregated system consists of a certain structure of services provided by independent network operators, named N_1 and N_2, and independent data centre providers, named C_1, C_2 and C_3 respectively. As indicated in Fig. 10, the on-state of the aggregated system depends on one data centre and its connected network to be in the on-state. The aggregated system is described by the following structure function

$$\begin{aligned} \Phi = & \ N_1 \cap C_1\big) \cup \ N_2 \cap (C_2 \cup C_3)\big) \\ & = (N_1 \cup N_2) \cap (N_2 \cup C_1) \\ & \cap (N_1 \cup C_2 \cup C_3) \cap (C_1 \cup C_2 \cup C_3) \end{aligned} \qquad (22)$$

The aggregator's mechanisms for providing a fault tolerant system of the underlying systems are not counted for. For instance, a replication procedure may be needed to ensure that data is replicated to all the data centres. The SLAs for the underlying systems are defined in Table 3. In the table the estimated optimized values for the underlying systems as λ_i, α_i and β_i are also included.

In the informed case the aggregated system's failure intensity and its mean down time may be found using the estimated optimized values for each of the underlying systems by using (15) and (18) as described in Section 7.5. As described in Section 7.2, each of the providers of the underlying systems has added a safety margin for not violating their agreed SLOs. The aggregated system's failure intensity is found to be in the range of $1.9 \cdot 10^{-11}$ with mean duration of down time of

Table 3: SLOs and costs with the assumed deploy cost parameters as regulated by multiple SLAs for a scenario as depicted in Fig. 10.

	Parameter	N_1	N_2	C_1	C_2	C_3
Commercial	$I_\tau(n,m,\omega)$	400	400	300	200	200
	C_c	1000	1000	650	400	400
SLO	τ (months)	12	12	12	12	12
	n	3	3	21	30	30
	m	1	1	7	10	10
	θ (sec)	1800	1800	1800	1800	1800
	ω (sec)	4500	4500	32500	40000	40000
Deploy cost	C_{du0}	18	18	15	15	15
	ν	1,25	1,25	1,25	1,25	1,25
	C_{cv0}	5	5	4	4	4
	η	3	3	3	3	3
Optimized	$\lambda \cdot 10^{-7}$ (1/sec)	0.246	0.246	3.435	5.504	5.504
	α	1.404	1.404	0.935	0.802	0.802
	β (sec)	642.2	642.2	1281.7	1414.9	1414.9

approximately 508 seconds (using Palm's identity for $\int_0^\infty (1-H(t))dt$. In the naive case, where the SLOs form the basis for the computation, failure intensity is found to be $1.6 \cdot 10^{-10}$ and mean down time to be 1070 seconds using the n_i and ω_i for each of the underlying systems only. The optimized system is an undoubted over dimensioned system. An aggregator may use this as an opportunity to obtain lower prices for the services provided by the underlying providers combined with lower SLOs or reduced compensations. This in turn will be the aggregator's added value to his customers, since the aggregated service price gets lower than the sum of the prices for all underlying services.

8. Conclusion

It is shown how an on-off model of a service from a provider may be obtained from the parameters of a service level agreement (SLA) between provider and customer. SLAs are expected to be an increasingly important and used means, when a chain of autonomous partners provides critical services. The price of the service, observation interval, service level objectives (SLOs), deployment cost and compensation in case the SLOs are not met, are taken into account. The SLOs considered are number of failures, accumulated down time and number of severe outages over an observation period. It is also shown how aggregated systems may be handled. This approach provides valuable insights into the cost vs. quality (in terms of dependability SLOs) trade-off for the provider and the customer. One such insight is obtaining the risk sharing in providing important infrastructure services.

It is also demonstrated that having short observation periods, as a part the agreement, will incur a significantly increase in cost relative to longer observation intervals. The model may be used by the service provider to find an approximately cost optimal inherent deployed dependability. This will be significantly different from the naive approach using the dependability related SLOs directly. However, the approach for obtaining the optimal inherent deployed dependability is based on a continuous set of solutions using idealised relations between cost, failure intensities and repair processes. In a real system, there will be a discrete set of options for providing the service, both in terms of system designs and/or configurations, as well as maintenance strategies. How to deal with this discrete set of solutions in obtaining the optimal designs and strategies are for future research.

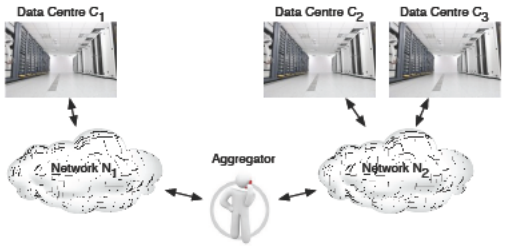

Figure 10: An aggregated system consisting of underlying systems described by the structure function $\Phi = \big(N_1 \cap C_1\big) \cup \big(N_2 \cap (C_2 \cup C_3)\big)$.

Acknowledgment

We acknowledge the input and discussions with the researchers Astrid Undheim and Andres González at Telenor and the professors Yuming Jiang and Poul Heegaard at Institute of Telematics at NTNU. We would like to thank the anonymous reviewers for their comments and suggestions which have lead to a significant improvement of the paper.

Appendix A. Assessing the error of approximation

In this appendix, we validate the approximation introduced in (3). First we validate the approximation of neglecting the down times when counting the failures from an HPP failure process and then when the failure process is NHPP.

A.1 The effect of neglecting down times

To obtain the simultaneous probability for the accumulated down time and the number of failures $P[\hat{\Omega}(\tau) \leq \omega, N(\tau) = x]$ we follow the line of reasoning in Takacs'

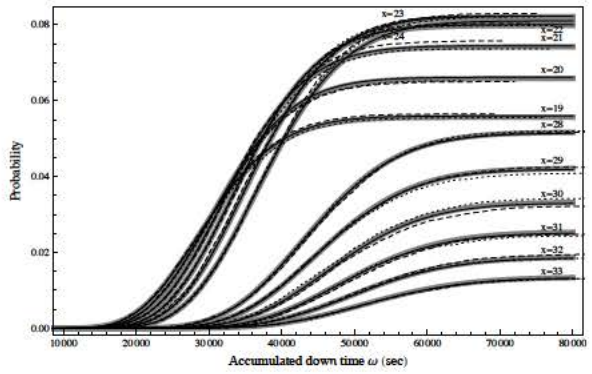

Figure A-1: $P[\hat{\Omega}(\tau) \leq \omega, N(\tau) = x]$ (thin lines) and $P[\Omega(\tau) \leq \omega, N(\tau) = x]$ (thick lines) are depicted for the parameters $\lambda = 7.504 \cdot 10^{-7}\text{sec}^{-1}, \alpha = 0.8, \beta = 2000$ sec and $\tau = 12$ months for an HPP failure process. Note that differences are small for all cases though only some are shown in the figure. $P[\Omega(\tau) \leq \omega, x]$ is less than $P[\hat{\Omega}(\tau) \leq \omega, N(\tau) = x]$ for $x \leq 23$ while higher otherwise. The simulated HPP (dashed) and NHPP (dotted) show also very small effect of neglecting the down time and that the variations in the failure intensity of an NHPP is very small as well.

[31]. Let $\{\xi_1, \xi_2, \xi_3, \ldots\}$ and $\{\vartheta_1, \vartheta_2, \vartheta_3, \ldots\}$ denote the times spent on-state and off-state respectively where $\zeta_x = \sum_1^x \xi_x$ and $\chi_x = \sum_1^x \vartheta_x$. The system starts in on-state at time $t = 0$ and alternates between the on-state and off-state until ending in either on-state or off-state at the end of the observation interval and we may write

$$\begin{aligned} &P[\hat{\Omega}(\tau) \leq \omega, N(\tau) = x] \\ &= P[\chi_x \leq \omega, \zeta_x + \chi_x \leq \tau < \zeta_{x+1} + \chi_x] \\ &+ P[\zeta_x \geq \tau - \omega, \zeta_x + \chi_{x-1} \leq \tau < \zeta_x + \chi_x] \\ &= P[\zeta_x \leq \tau, \zeta_x < \tau - \omega] - P[\chi_x \leq \omega, \zeta_{x+1} < \tau - \omega] \\ &+ \Upsilon(\tau, x, \omega) - \Upsilon(\tau, x+1, \omega) \end{aligned} \quad \text{(A-1)}$$

where

$$\begin{aligned} \Upsilon(\tau, x, \omega) &= \int_{y=0}^{\omega} P\Big[\zeta_x \geq \tau - \omega, \zeta_x \leq \tau - y | \chi_{x-1} = y\Big] \\ &\cdot P[y \leq \chi_{x-1} < y + dy] dy \\ &= \int y = 0^{\omega} h^{\otimes(x-1)}(y) \cdot \Big(\hat{G}^{\otimes x}(\tau - y) - \hat{G}^{\otimes x}(\tau - \omega)\Big) dy \\ &= \Theta(\tau, x, \omega) - \hat{G}^{\otimes x}(\tau - \omega) \cdot H^{\otimes(x-1)}(\omega) \end{aligned} \quad \text{(A-2)}$$

Substitution of (A-2) into (A-1) yields

$$\begin{aligned} P[\hat{\Omega}(\tau) \leq \omega, N(\tau) = x] &= \Big(H^{\otimes x}(\omega) - H^{\otimes(x-1)}(\omega)\Big) \\ &\cdot \hat{G}^{\otimes x}(\tau - \omega) + \Theta(\tau, x, \omega) - \Theta(\tau, x+1, \omega) \end{aligned} \quad \text{(A-3)}$$

where $\hat{G}(t)$ is a homogeneous failure process. Further, $\Theta(\tau, x, \omega) = \int_0^{\omega} h^{\otimes(x-1)}(y) \hat{G}^{\otimes x}(\tau - y) dy$, $\Theta(\tau, 0, \omega) = 1$ and $\Theta(\tau, 1, \omega) = \hat{G}(\tau)$. For gamma distributed down times and an HPP failure process this may be written as

$$\begin{aligned} \Theta(\tau, x, \omega) &= H^{\otimes(x-1)}(\omega) - \frac{e^{-\lambda\tau}}{\beta^{(x-1)\alpha}\Gamma(x-1)\alpha} \\ &\cdot \sum_{i=0}^{x-1} \frac{\lambda^i}{i!} \int_0^{\omega} y^{(x-1)\alpha-1} \cdot (\tau - y)^i e^{(-\beta+\lambda)y} dy \\ &= H^{\otimes(x-1)}(\omega) - \frac{e^{-\lambda\tau} \sum_{i=0}^{x-1} \lambda^i / i!}{\Gamma((x-1)\alpha)\beta^{(x-1)\alpha}} \\ &\cdot \sum_{j=0}^{i} \tau^j (-\omega)^{i-j} \omega^{(x-1)\alpha} \cdot (\omega(1/\beta - \lambda))^{-i+j+(1-x)\alpha} \binom{i}{j} \\ &\cdot \Big(\Gamma(i - j + (x-1)\alpha) \\ &- \Gamma(i - j + (x-1)\alpha, \omega(1/\beta - \lambda))\Big) \end{aligned} \quad \text{(A-4)}$$

where $\hat{G}(\tau) = 1 - e^{-\lambda\tau}$. We now have an expression (A-3) of the exact simultaneous probability and we may assess the error of the approximation of neglecting the down times as the difference between $P[\hat{\Omega}(\tau) \leq \omega, N(\tau) = x]$ and $P[\Omega(\tau) \leq \omega, N(\tau) = x]$ where $P[\Omega(\tau) \leq \omega, N(\tau) = x] = H^{\otimes x}(\omega) P[N(\tau) = x]$. The differences will be higher the higher the density of failures become and with increased expected down time. In Fig. A-1 $P[\hat{\Omega}(\tau) \leq \omega, N(\tau 2 = x]$ (thin lines) and $P[\Omega(\tau) \leq \omega, N(\tau) = x]$ (thick lines) are depicted for the parameters $\lambda = 7.504 \cdot 10^{-7}\text{sec}^{-1}, \alpha = 0.8, \beta = 2000$ sec and $\tau = 12$ months representing a system behaviour that provoke higher differences than used in this paper. The difference is hardly noticeable for the depicted number of failure in the figure and is less than 2% for all cases. As may be observed, $P[\Omega(\tau) \leq \omega, N(\tau) = x]$ is less than $P[\hat{\Omega}(\tau) \leq \omega, N(\tau) = x]$ for $x \leq 23$ while higher otherwise. The approximation may also be shown to be valid for a NHPP as long as time varying failure intensity $\lambda(y) \ll 1/E[D], \forall y \in [0, \tau]$, where for gamma distributed down times $E[D] = \alpha\beta$.

A.2 The effect of neglecting variations in the failure intensity

For an NHPP failure process with multiple cyclic and trend effects, the model in (21) may be used for the variations of failure intensity. In (A-3) an HPP failure process is assumed. To investigate the effects on the approximation when there is an NHPP failure process we have performed a simulation study.

The on-off system is simulated without neglecting the down period. For the intensity we have used (21) for $s = 1$ and $k = 2$ where $\psi_0 = 7.504 \cdot 10^{-7}, \psi_1 = 2.55 \cdot 10^{-16}, \rho_1 = 1.5 \cdot 10^{-7}, \rho_2 = 0.2 \cdot 10^{-7}$ and ϖ_1

and ϖ_0 represent periods of 24 hours and one week respectively. The phase shifts, ϱ_1 and ϱ_2 are both set to 0. The excepted number of failures during the observation interval τ is equal for the HPP and NHHP failure processes where $\Lambda(\tau) = \int_0^\tau \lambda(y)dy = \lambda\tau$.

A total of 60000 years were simulated, i.e., equal 60000 observation intervals. Results from this simulation study are shown in Figure A-1 for HPP and NHPP failure processes. These results combined with the result in Table 2 confirm that the effect of neglecting down times, cf. (3) is very small, and that variations in the failure intensity, i.e., the effect of an NHPP failure process, is very small as well.

[1] Bobbio, A., Bonanni, G., Ciancamerla, E., Clemente, R., Iacomini, A., Minichino, M., et al. Unavailability of critical {SCADA} communication links interconnecting a power grid and a telco network. Reliability Engineering & System Safety 2010;95(12):1345 – 1357. 19th European Safety and Reliability Conference.

[2] Markovic, D.S., Zivkovic, D., Branovic, I., Popovic, R., Cvetkovic, D.. Smart power grid and cloud computing. Renewable and Sustainable Energy Reviews 2013;24:566 – 577.

[3] Yigit, M., Gungor, V.C., Baktir, S.. Cloud computing for smart grid applications. Computer Networks 2014;70:312 – 329.

[4] Ouyang, M.. Review on modeling and simulation of interdependent critical infrastructure systems. Reliability Engineering & System Safety 2014;121:43 – 60.

[5] Ntalampiras, S., Soupionis, Y., Giannopoulos, G.. A fault diagnosis system for interdependent critical infrastructures based on {HMMs}. Reliability Engineering & System Safety 2015;138:73 – 81.

[6] E.860; framework of a service level agreement. Telecomunication Standardization Sector (ITU-T): 2006;.

[7] Hartley, K.L.. Defining effective service level agreements for network operation and maintenance. Bell labs technical journal 2005;9(4):139–143.

[8] Wu, L., Buyya, R.. Service level agreement (SLA) in utility computing systems. Arxiv preprint arXiv:10102881 2010;.

[9] Trienekens, J., Bouman, J., van der Zwan, M.. Specification of service level agreements: Problems, principles and practices. Software Quality Journal 2004;12(1):43–57.

[10] Goyal, A., Tantawi, A.. A measure of guaranteed availability and its numerical evaluation. Computers, IEEE Transactions on 1988;37(1):25–32.

[11] Sauvé, J., Bartolini, C., Moura, A.. Looking at business through a keyhole. In: Integrated Network Management-Workshops, 2009. IM '09. IFIP/IEEE International Symposium on. 2009, p. 48–51.

[12] Taylor, R., Tofts, C.. Death by a thousand SLAs: a short study of commercial suicide pacts. Forschungsbericht, Hewlett-Packard Labs 2005;.

[13] Zhou, L., Grover, W.. A theory for setting the "safety margin" on availability guarantees in an SLA. In: Design of Reliable Communication Networks, 2005. (DRCN 2005). Proceedings.5th International Workshop on. 2005, p. 7 pp.

[14] Maciejewski, H., Caban, D.. Estimation of repairable system availability within fixed time horizon. Reliability Engineering & System Safety 2008;93(1):100 – 106.

[15] González, A.J., Helvik, B.E.. Guaranteeing service availability in slas; a study of the risk associated with contract period and failure process. In: Communications, 2009. LATINCOM'09. IEEE Latin-American Conference on. IEEE; 2009, p. 1–6.

[16] Snow, A., Weckman, G., Gupta, V.. Meeting SLA availability guarantees through engineering margin. In: Networks (ICN), 2010 Ninth International Conference on. 2010, p. 331–336.

[17] Schulz, F.. Decision support for business-related design of service level agreements. In: Software Engineering and Service Science (ICSESS), 2011 IEEE 2nd International Conference on. 2011, p. 35–38.

[18] Franke, U.. Optimal IT service availability: Shorter outages, or fewer? Network and Service Management, IEEE Transactions on 2012;9(1):22–33.

[19] Leitner, P., Hummer, W., Dustdar, S.. Cost-based optimization of service compositions. Services Computing, IEEE Transactions on 2013;6(2):239–251.

[20] Mastroeni, L., Naldi, M.. Network protection through insurance: Premium computation for the on-off service model. In: Design of Reliable Communication Networks (DRCN), 2011 8th International Workshop on the. IEEE; 2011, p. 46–53.

[21] Reibman, A., Trivedi, K.. Numerical transient analysis of markov models. Computers & Operations Research 1988;15(1):19 – 36.

[22] Rubino, G., Sericola, B.. Interval availability analysis using denumerable markov processes: application to multiprocessor subject to breakdowns and repair. Computers, IEEE Transactions on 1995;44(2):286–291.

[23] Bhoj, P., Singhal, S., Chutani, S.. SLA management in federated environments. Computer Networks 2001;35(1):5 – 24.

[24] Burchard, L.O., Hovestadt, M., Kao, O., Keller, A., Linnert, B.. The virtual resource manager: an architecture for SLA-aware resource management. In: Cluster Computing and the Grid, 2004. CCGrid 2004. IEEE International Symposium on. 2004, p. 126–133.

[25] Gonzalez, A.J., Helvik, B.E.. System management to comply with SLA availability guarantees in cloud computing. In: Cloud Computing Technology and Science (CloudCom), 2012 IEEE 4th International Conference on. 2012, p. 325–332.

[26] Amazon Web Services. Amacon EC2 service level agreement. 2013. URL: `http://aws.amazon.com/ec2/sla/`.

[27] Microsoft. Windows azure cloud services, virtual machines, and virtual network service level agreement (SLA). 2013. URL: `http://www.microsoft.com/en-us/download/details.aspx?id=38427`.

[28] Nextgen group. Service level agreement (SLA). 2013. URL: `http://www.nextgennetworks.com.au/about/service-management-centre/service-level-agreement/`.

[29] CenturyLink. Savvis SLA attachment. 2013. URL: `http://www.centurylinktechnology.com/legal/sla`.

[30] Verizon. European service level agreement for: Verizon Internet dedicated, Internet DSL office and Internet DSL solo. 2013. URL: `http://www.verizonenterprise.com/terms/emea/at/sla/`.

[31] Takács, L.. On certain sojourn time problems in the theory of stochastic processes. Acta Mathematica Hungarica 1957;8(1):169–191.

[32] Telenor. Prisliste for kapasitetsprodukt. 2014. URL: `https://www.jara.no/produkter/kapasitet/priserogavtaler/`.

[33] Telstra. Standard restoration and SLA premium. 2014. URL: `http://www.telstra.com.au/customer-terms/business-government/other-services/restoration-sla-premium/`.

[34] Barlow, R., Proschan, F.. Statistical theory of reliability and life testing: probability models. Holt, Rinehart and Winston New York; 1975.

[35] Kuusela, P., Norros, I.. On/off process modeling of IP network failures. In: Dependable Systems and Networks (DSN), 2010 IEEE/IFIP International Conference on. 2010, p. 585 –594.

[36] Wolfram Research, I.. Documentation centre. 2011.

[37] Clemente, R., Bartoli, M., Bossi, M., D'Orazio, G., Cosmo, G.. Risk management in availability SLA. In: Design of Reliable Communication Networks, 2005. (DRCN 2005). Proceedings.5th International Workshop on. 2005, p. 8 pp.

[38] Følstad, E., Helvik, B.. Failures and changes in cellular access networks; a study of field data. In: Design of Reliable Communication Networks (DRCN), 2011 8th International Workshop on the. 2011, p. 132 –139.