

Continuity-based Resilient Communication

Piotr Cholda*, Anders Mykkeltveit†, Bjarne E. Helvik†, and Andrzej Jajszczyk*

*Department of Telecommunications
AGH University of Science and Technology
Kraków, Poland
E-mail: {cholda,jajszczyk}@kt.agh.edu.pl

†Centre for Quantifiable Quality of Service in Communication Systems
Norwegian University of Science and Technology (NTNU)
Trondheim, Norway
E-mail: {mykkeltv,bjarne}@q2s.ntnu.no

Abstract—The paper advocates that in some communication service providing settings, it is more appropriate to focus on the continuity of a connection as the prime reliability attribute for defining requirements and establishing Service Level Agreements, rather than the traditionally used availability. The justification for this approach and types of services where continuity is relevant are given.

For illustration, a transport network with an optical control plane for a utility grid along with some theoretical background is studied.

I. INTRODUCTION

Digital communication networks form a part of the critical infrastructure of the today's society. Thus, provisioning of reliability, i.e., the proper behavior from the viewpoint of faults, is very important for this environment, also in the context of different recovery classes. Establishment of connections in a network is on one hand based on the capacity usage optimization techniques, and on the other hand on modeling of some survivability-related measures. There is a common practice to base the design on availability requirements, that is the sufficiently large probability that a service is provided when needed. While this approach is appropriate for long-lasting connections traditionally seen in the networking area (e.g., for optical lightpaths [1]), attention must be paid to the fact that this is not always the case. From the service provisioning, operational and commercial viewpoint, the continuity, i.e., the probability of obtaining an uninterrupted service, may in many settings be a better measure. The objective of this paper is to identify the settings where the continuity is more appropriate and to illustrate the feasibility of this measure.

An important issue for the coming generations of communication systems is to be able to tailor the reliability of the provided services to the users needs and the characteristics of the service itself. The authors of this paper realized that although there have been many ideas on recovery-related differentiation and modeling of communication connections, none of them thoroughly analyzed the continuity criterion, related to the length of the uninterrupted working time [2]. However, as it will be discussed in Section II, the continuity-based specification and differentiation may be most adequate sometimes. This is typically the case when the service delivery cannot be interrupted or the duration of the Service Level Agreement (SLA) is short. In these cases, the continuity expressed as the *Mean Time to Failure*, $MTTF$, or a value of the *reliability function*, $R(t)$, serve as the prime reliability requirements.

In recent years, many recovery methods have been invented, and the researchers have focused on the selection of a method for particular requirements [2], [3]. It can be stated that although the continuity-based parameters are hardly ever used in the communications community, the computing field has paid a considerable attention to this mission-oriented aspect, cf. for instance the recent analyzes on crash-recovery failure detectors [4]. For the computing reliability many analytical models of the continuity can be found. As a classical text, Beaudry's seminal work [5] should be mentioned.

The paper is organized as follows. Section II discusses the rationale behind introduction of the continuity-based approach and potential pitfalls related to the usage of the availability-based one. Next, Section III presents how this approach may be used in a grid environment. A simple modeling approach only to present the usage of the method has been adopted. Then, in Section IV a simulation study to show the usefulness of the operation is presented.

II. NECESSITY FOR THE CONTINUITY-BASED ASSESSMENT

The current practice is to evaluate the survivability of connections on the basis of the steady-state availability. In addition, an upper bound on the downtime (or recovery time) is defined. Regarding the well known relation between fault-related times in a system given in Fig. 1, see Table I for notation, the availability is given as follows:

$$A \equiv \frac{MUT}{MDT + MUT} \simeq \frac{MTTF}{MTTR + MTTF} \quad (1)$$

With typical values for a lightpath (a long-haul link) $MUT = 10,000$ hours, approximately one fibre cut per year; whereas $MDT = 10$ hours: an average time necessary to repair a cable. Then, $A \approx 0.999$, the so-called three nines' level of the availability. Assuming Poisson failure process, $MTTF = 10,000$ hours and $R(1 \text{ hour}) \approx 0.9999$ which is satisfactory for most services. However, with other systems it is possible to have MUT as low as 100 hours (some five-six days) and an automatic fault handling that yields $MDT = 0.1$ hour. In this case A does not change, but $R(1 \text{ hour}) \approx 0.99$ may become unacceptably low for many applications resulting in frequent interruptions—this will be harmful for applications that cannot accept discontinuity longer than 0.1 hour.

This phenomenon, though in different scale, can be observed in multilayer networks. Therein, the higher layer is

TABLE I
NOTATION USED IN THE PAPER

Symbol	Description
T_{dur}	Predicted duration of a requested connection
T_U	Up Time (inter-event measure)
MUT	Mean Up Time: $MUT \equiv E[T_U]$
T_{TF}	Time to failure (forward recurrence measure)
$R(t)$	Reliability function: $R(t) \equiv \Pr\{T_{TF} > t\}$
$MTTF$	Mean Time to Failure: $MTTF \equiv E[T_{TF}] = \int_0^{\infty} R(t)dt$
$MTTF_{req}$	Value of $MTTF$ demanded for the request
λ	Failure rate
T_D	DownTime (inter-event measure)
MDT	Mean DownTime: $MDT \equiv E[T_D]$
T_{TR}	Time to repair (forward recurrence measure ^a)
$MTTR$	Mean Time to Repair: $MTTR \equiv E[T_{TR}]$
A	Asymptotic (steady-state) availability
t_p	Maximum value of time for which $R(t) \geq p$
ser	Parameter related to the serial reliability scheme
par	Parameter related to the parallel reliability scheme
WP	Parameter related to the working path
BS	Parameter related to the backup path and paths that compete for resources with the working path

^aSee the description of Fig. 1 in Subsection III-B for explanation of the difference between the inter-event and forward recurrence measures.




Fig. 1. Renewal process of a single item with different relevant times [9]. F: failure, R: repair. For further description, see Table I

sensitive to the recovery time of the lower layer: if it is short enough, the application in question does not even notice the recovered failure. However, if a slow recovery method is used, the application operation can be interrupted. Then, from the viewpoint of the client, the fact that there is some recovery is not important any more [6]. On the other hand, the character of failures reported in [7] suggests that they are not memoryless (e.g., they can be better fitted with the usage of the Weibull reliability function based modeling), and in fact there are quite many of them from the viewpoint of the whole network: up to 50% of consecutive failures of routers (not necessarily the same) happen up to thousands of seconds (tens of minutes) while for optical links it is tens of thousands of seconds (single hours). As can be seen from the failures duration logs published by UNINETT network [8], links have generally high availability, but this is usually related to the fact that the downtimes are quite short (up to 10 minutes) but can be quite frequent (up to tens in one year).

Hence, it is necessary to control the continuity of the provided service. The remaining part of this section discusses four groups of applications where continuity should serve as

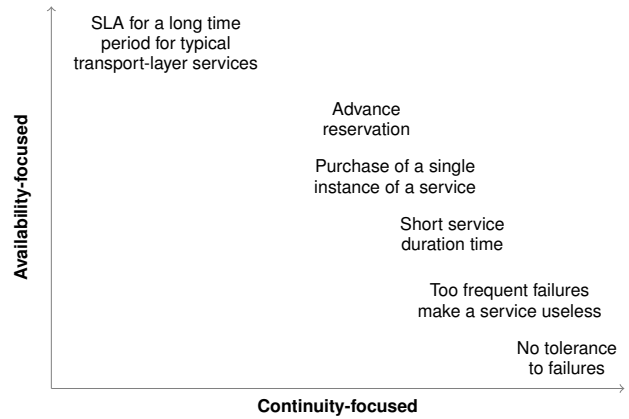


Fig. 2. Sensitivity to the availability- and continuity-based provisioning

the prime reliability measure and the subject for specification and differentiation. Fig. 2 illustrates different types of relations between reliability needs.

A. Failures May be Catastrophic

There are some services where interruptions may cause *catastrophic failures*, e.g., remote control of critical tasks, like offshore/subsea maintenance operations handled from on-shore supporting centers, remote surgery, or military applications. In this group, there are services that need a *very high level of QoS*, not accepting discontinuities, like real-time control applications. For them, virtually failure-free connections are necessary, as even short recovery time may harm jitter requirements. In other cases, *interruptions may be highly undesirable and incur losses to the user far beyond the cost of the communication*, like for online bidding processes and for loss of results in a compound real-time grid operations.

B. Frequent Interruptions Reduce the Quality of Experience Significantly

The availability, and hard or soft bounds on incurred downtimes, are insufficient parameters to describe the quality of some services. For instance, the availability may be acceptable but the *rate of interruptions render the service useless* cf. a scaling example above. Hence, the interruption intensity or a similar parameter related to $MTTF$ should be added in SLAs. Applications with the so-called *advance reservation of resources*, like grids or burst switched systems can be given as examples here. For them, a relatively non-breakable data transfer is needed, as a determined time for the transport is defined. Otherwise, the discontinuity may cause the resources seizure by some other processes awaiting them to optimize the usage. Such a type of reservation is related to systems with a high level of dynamics and many users, especially when the convergence between networking and computing areas takes place.

C. Short Duration of Service Delivery

Sometimes the *service provisioning period is so short that the availability becomes useless* as a contractual and opera-

tional measure. SLA with an acceptable risk for both the user and the service provider cannot be based on the availability when $MTTF$ is of the same order or less than the duration of the contract (connection). For a very short service duration that is shorter than the actual repair or service recovery time, e.g., as that found in computing applications, it is evident that the availability is not meaningful as a measure since it includes MDT , which is irrelevant for the quality of the provided service.

D. Single Instance of Service Delivery

When users of a service do not subscribe to it, but *buy a single instance* of it when needed, they are interested in getting delivered what they have purchased, i.e., the continuity during service provisioning. In this commercial context with *SLAs signed for a short time*, it is of course in the interest of the provider that the service has a high availability, so the users do not get the same service from a competitor. Nevertheless, this is not an issue for the contract, and the continuity is what matters most. As an example, entertainment services can be given, like transmission of the opening of the Olympic Games, where a TV station hires the transfer service from some communications carrier.

E. Comparison of the Continuity- and Availability-based Approaches

The *availability-based approach* is related to the decision of a reliability-aware customer who wants to purchase a service and who needs this service for a long period (say, in the order of months to years) for a more or less typical usage, e.g. leasing an optical lightpath for providing “LAN services” in a company with many sites. The lightpath may fail, but should be repaired effectively and fast. However, the customer is aware that a few relatively short repair periods of the service across the connection lifetime will be necessary and arranges his business accordingly. Depending on the customer’s trade-off between cost of downtime and the cost of a high availability, the service provider finds an offer for them, i.e., a connection with an appropriate availability. Thus, the availability must be defined by the customer at start. Additionally, if such a service is purchased, it is often not utterly important that the service is working when the service period commences. Primarily, it is the average state of the connection that matters.

On the other hand, the *continuity-based approach* is relevant for a situation when a customer would like to hire some amount of bandwidth for a short period of time (a week, a few hours). Obviously the connection must be reliable, as a broken service due to repairs cannot be accepted; the service is needed now and after a repair period the communication need is gone, cf. the example with opening of the Olympic Games in II-D. The user’s interest in purchasing the service depends on the (probabilistic) guarantee that it will not fail. The guarantee is given along with its cost. In this case, what primarily matters is the continuity—the asymptotic availability is not a significant attribute. A second important issue is that

the service must work properly when its use starts, and hence, knowledge about the current state of the network resources is very important. Additionally, recovery is acceptable only if the recovery time does not disturb the application, and the modeling of a recovery operation must be distinct from that of a repair.

The above engineering intuitions are supported by theoretical modeling. For instance, Grottke et al. [10] states that sometimes the availability- and continuity-based approaches are mixed. They show with a theoretical model that when the user’s perception of up- and downtimes is crucial, the manipulation of the mutual relation between MUT and MDT , not changing the availability, may have catastrophic consequences. However, it should also be remembered that, even in steady-state situations, the failure frequency of the system, that is a metric of a continuity-related character, is another helpful performance measure that can be actively used jointly with the availability when services are dimensioned [11]. Using a more realistic Weibull-like failure process, introduces the dynamics into account, adding some new information in comparison to the usually used steady-state failure intensity obtained on the basis of the availability calculations.

III. EXAMPLE APPLICATION—A UTILITY GRID

A. Utility Grid and Its Resilience

An application area well suited for the continuity-based service provisioning is transport network services for the workflow support in a utility grid as this type of grids has strict transport service requirements [12]. Relations between networking and grids, as well as their transport requirements, also in the reliability context, have been thoroughly discussed in the literature [13]. Moreover, some grid-related documents mention the necessity to differentiate transport services from the viewpoint of the resilience [14]. Nevertheless, it is very hard to give general requirements for reliability in grids [12].

Four potential fields of applications for a proposed methodology were presented in Section II. Three of them can be found in the grid environment. First, there are many connections with *very high level of QoS requirements*, preventing even short recovery, not mentioning a repair, as some grid-based applications in particle physics or radio astronomy have so high QoS requirements that they seem not to accept even very limited packet loss [15]. Second, some of those demanding connections have a *very short duration*, starting from a few microseconds [13]. And third, since grids involve many shared computing resources, an *advance reservation* is frequently appropriate: very expensive computing or visualization resources [16] must be reserved this way for well defined starting times and durations [17]. When a service based on such type of reservation is broken, then resources will be released and cannot be used by the service again.

The transport model (Fig. 3) for grids is assumed to be based on the layering structure, where as a physical, bearer layer an optical network is used, whereas the higher layer is organized as an IP/MPLS network. This model has been extensively discussed as appropriate for grids [13].

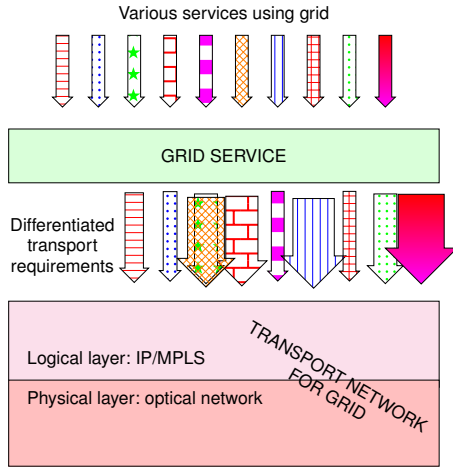


Fig. 3. Model of the functional cooperation between a grid service and its transport network

B. Provisioning Model

The provisioning of connections with a specified continuity, based on the proposed approach, goes as follows (see Fig. 4).

A connection is requested with determined duration T_{dur} and continuity parameters, either hard (for notations, see Table I):

$$R(T_{dur}) \geq p$$

meaning explicitly that:

$$\Pr\{T_{TF} > T_{dur}\} \geq p,$$

with the values of p at the level of 0.95, 0.999, etc., or soft:

$$MTTF_{req} \gg T_{dur}$$

A connection may be established as **Class 1** where failures are not recovered, i.e., yielding a serial reliability system (see Subsection III-B1); or if necessary, at the expense of an increasing capacity usage, as **Class 2** for a protected connection (see Subsection III-B2). With dedicated protection a parallel reliability system is yielded. However, the case with a shared protection, which requires a more extensive analysis, is studied in Subsection III-B3. If a single protection path is not sufficient, the level of protection within **Class 2** could be increased, e.g., to 1 + 2 dedicated protection, but such an extension is not dealt with here.

Connections which commence at a random time during the operation of a provisioning system are regarded. Thus, it is assumed that a steady operational state is achieved, and it is natural to base the reliability function on the time to failure, T_{TF} , cf. Fig. 1 and Table I. Note that if the failure repair process is regarded as an alternation point process, T_{TF} and T_{TR} are forward recurrence times, as they are measured from some steady state time points (exemplary t_2 and t_1 , respectively, in Fig. 1) to event occurrence. Contrary to them, T_U and T_D are inter-event times measured between consecutive events. Hence, unless the failure process is of

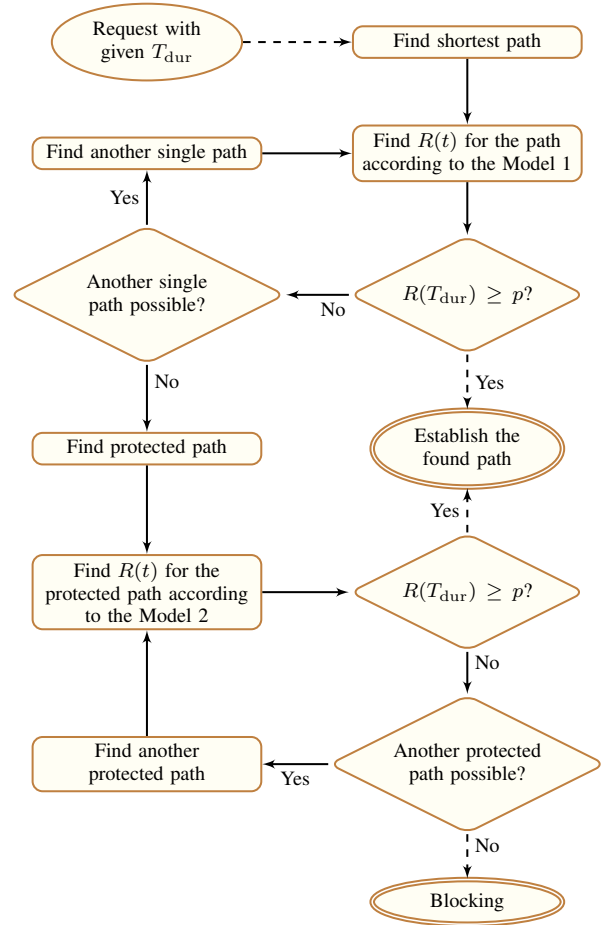


Fig. 4. Algorithm using the derived models (example for the hard requirement)

Poissonian character (see below) $MTTF \neq MUT$. Similarly for the repair process. This is an important fact related to the limited generality of the approximation given in (1), for instance when the underlying processes are of a Weibull-like character.

The network elements are assumed to be related to Poisson failure processes and hence, negative exponential distributions (n.e.d.) of failure times with the reliability functions of the following form:

$$R(t) = e^{-\lambda_i t} \text{ and } \lambda_i = \frac{1}{MTTF_i} \quad (2)$$

This distribution is used here only for exemplification (and availability of network failure statistics based on this assumption) and is not an assumption behind our proposal.

In addition to being a reasonable and common assumption, this enables simple analytical models, which may be embedded in the software supporting the connection provisioning system at the level of an automatic control plane. This is important as there are expectations that some connections, for instance in grids, are going to be highly dynamic [13] and not as static as

in the case of lightpaths, and therefore requiring a very short provisioning time [12].

1) *Class 1, Model 1: Non-protected Connections:* For **Class 1**, that is a non-protected, single connection, the serial reliability model is used, i.e., the connection fails if at least one of its components fails. This model is quite trivial for exponential failure times of elements described by (2). The reliability function for the whole serial connection can be expressed as follows:

$$R_{\text{ser}}(t) = \prod_i R_i(t) = \prod_i e^{-\lambda_i t} = e^{-\lambda_{\text{ser}} t},$$

where $\lambda_{\text{ser}} = \sum_i \lambda_i$ is the failure rate of the connection. Due to Poissonian failure assumption $\lambda_i = 1/MTTF_i$ and consequently $MTTF_{\text{ser}} = 1/\lambda_{\text{ser}}$. Hence, a connection may be established on the basis of the non-protected path, i.e., on the basis of **Class 1**, if only:

$$t_p = -\frac{\ln p}{\lambda_{\text{ser}}} \geq T_{\text{dur}}$$

2) *Class 2, Model 2: Connections Recovered by Protection Schemes:* For **Class 2**, a connection is protected either by a higher or lower layer protection method. A prerequisite for using this method is that the involved recovery time is not perceivable by the application, i.e., it does not disturb the required quality. This may be achieved by using a dedicated disjoint backup path which may be activated involving short recovery times (1 + 1 protection). From a modeling standpoint, this will be a simpler special case of a shared protection scheme which is studied in Subsection III-B3. A recovery scheme that is sufficiently fast for making failures not detectable at the level of the working path for the application is assumed. Furthermore, there must exist a disjoint backup path with the sufficient capacity, where links of this path may also serve as the backup for other connections. Seen from the viewpoint of the connection, the backup path fails with rate λ_{BS} which either may be due to failures of the backup path itself or one or more of its links serving as a backup for another connection, cf. Subsection III-B3. Being conservative and assuming that no repair is completed during the operation of the considered connection, the system may be regarded as a parallel reliability structure, i.e., the connection fails if both of its component branches fail, yielding the following reliability function for **Class 2**:

$$R_{\text{par}}(t) = 1 - [1 - R_{WP}(t)][1 - R_{BS}(t)] \\ = e^{-\lambda_{WP} t} + e^{-\lambda_{BS} t} - e^{-(\lambda_{WP} + \lambda_{BS})t} \quad (3)$$

Our model is of a system where both the working and backup paths are established at the opening of the connection, and the backup path may fail ahead of the working path without being replaced due to the short duration of the connection. Other formulas may be used for other modes of operation, but it is problematic if other situations (e.g., fault free backup path in the moment of the failure in all cases) are realistic in the settings we are interested here.

Mean Time to Failure is obtained from (3) as

$$MTTF_{\text{par}} = \frac{1}{\lambda_{WP}} + \frac{1}{\lambda_{BS}} - \frac{1}{(\lambda_{WP} + \lambda_{BS})}$$

To obtain t_p , equation $R_{\text{par}}(t_p) = p$ can be solved algebraically using a series expansion of the exponent [18]:

$$e^{-Z} = \sum_{i=0}^{\infty} (-1)^i \frac{Z^i}{i!} = 1 - Z + \frac{Z^2}{2!} - \frac{Z^3}{3!} + \dots,$$

with truncation satisfactory for interval of t , where $e^{-Z(t)} \geq p$.

Similarly as for **Class 1**, if only $t_p \geq T_{\text{dur}}$, then the connection can be established on the basis of the path protected by this shared scheme.

3) *Estimates of Shared Backup Path Failure Rate:* Having shared protection, it is in general very difficult to precisely determine λ_{BS} , the failure rate of the backup path. When repairs are not included in the estimation model, it is feasible to estimate failure rates of shared backup paths by adapting models used for estimating the availability of such paths. One of the most accurate models is proposed by Zhang et al. [19]. However, methods of a limited complexity are likely to be the only feasible approach when implementation in a large network is considered. Mykkeltveit and Helvik [20] proposed two relatively simple, conservative methods to be used for provisioning of connections with availability guarantees. The methods may easily be used for estimating the failure rate. The discussion here is limited to link failures only, but it is straightforward to extend both methods to include node failures as well.

The first *ultra-conservative* method assumes a worst-case scenario where any failure on other links in the network, except on the working path, makes the backup path inaccessible. Let the failure rate of link l_i be denoted as λ_i . Then:

$$\lambda_{BS} \leq \sum_{\forall i | l_i \notin WP} \lambda_i \quad (4)$$

The second *preemptive sharing* method is based on temporal priority and preemption. Access to a shared backup path is given a priority similar to the order of assignments. Hence, a currently active backup path may be preempted by the failure of a primary path having assigned one or more links on the shared backup path before the current one. The failure rate of the backup path is increased with the failure rate of the primary paths of the connections having a higher priority (i.e., established earlier) than ours. Focusing on our connection c , let $c^*(i)$ denote the set of connections that share backup link l_i with c and were established before c . Then:

$$\lambda_{BS} \leq \sum_{\forall i | l_i \in BS} \left(\lambda_i + \sum_{\forall j | c_j \in c^*(i)} \lambda_{WP}^{c_j} \right), \quad (5)$$

where $\lambda_{WP}^{c_j}$ denotes the failure rate of the working path of connection c_j . The method is also conservative, since it assumes that none of the connections in $c^*(i)$ terminates before c and since it pessimistically assumes that c does not share wavelengths on more than one link with any other connection.

Note, that this estimate can in theory be more pessimistic than the ultra-conservative one, but this will not be the case in practice.

IV. NUMERICAL RESULTS

A. Main Data on Simulations

The selected scenario considered here is based on a scientific computational grid used for astronomy-related calculations purposes described in the grid literature [21, Section 2.7]. A large amount of data is received by five telescope centers and synchronously transmitted to a central computing node.

The COST 266 Pan-European network [22] with 37 nodes is used as a transport network topology. We assume that telescope centers are located in Birmingham, Brussels, Milan, Krakow, and Stockholm. They send data to a center located in Amsterdam.

Simulations were carried out using an extended version of a purpose-built Java-based simulator [20]. A series of grid sessions is simulated. It is assumed that the session duration may vary from one hour up to a week. The network has also a steady-state background traffic which leaves some wavelengths open while others may be re-used as shared backup path components. Background traffic connections requests are uniformly distributed among all pairs of nodes. Each connection requires a single wavelength. One hundred connection requests arrives according to a Poisson process between each grid establishment attempt. The holding times are n.e.d. distributed with expectation of one inter-grid session interval. Hence, the number of background connections existing when a grid session is scheduled is generated due to Poisson distribution with an expectation of 100.

The simulator uses a WDM network model where each link has sixteen wavelengths. It uses the Disjoint Path-Pair Matrix (DPM) approach [23] to compute a set of candidate working and backup paths which are used in sequence to search for a solution accepting the continuity requirement. The simulator chooses the first feasible path discovered. The simulator estimates link reliability functions on the basis of the geographical length of the links and the reliability-related numerical values reported by Verbrugge et al. [24].

B. Simulation Results for the Continuity-based Approach

The reliability values for different required p levels for the Kraków-Amsterdam connection are given in Table II as an example. The results are based on the first (shortest) pair of paths in the DPM. For the shared protection, the values depend on the background traffic and the connection requirements. For indication and to avoid these dependencies, the ultra-conservative method, cf. (4), is used to get absolute lower bounds irrespective of requirements and the other traffic in the network. It is interesting to see that even though the $MTTF$ s for single connection and shared protection are (accidentally) equal, the value of t_p for $p = 0.9999$ is 17 times higher for the shared protection than for a single connection. This indicates that the soft requirements, $MTTF_{req} \gg T_{dur}$, may not be a good basis for specifying the service continuity.

TABLE II
RELIABILITY OF THE KRAKÓW-AMSTERDAM CONNECTION: VALUES OF THE GUARANTEED PERIOD t_p AND $MTTF$ IN HOURS

Connection type	Reliability requirement, p				$MTTF$
	0.95	0.99	0.999	0.9999	
Single connection	97	19	2	0.2	1890
Dedicated protection	384	159	48.5	15.2	2360
Shared protection ^a	111	37.9	10.6	3.25	1890

^aBased on estimate using the ultra-conservative method in (4).




Fig. 5. Simulation results: blocking probability. Different lines connect points reported for different values of T_{dur}

To evaluate the probability of successful establishment of the grid according to the hard requirements $R(T_{dur}) \geq p$, simulations were run for different values of p and T_{dur} . 40,000 attempts to establish grid connections were simulated to give high confidence in the results (95% confidence level is used). The grid connections may share links in the backup paths with each other or with the background traffic connections. The backup path failure rate, λ_{BS} , was calculated using the preemptive sharing scheme, described by (5). For a successful establishment of the grid, it is required that all the five connections from the telescope stations to the central node can be established meeting the reliability requirement p .

The blocking probability of the grid, that is the ratio of unsuccessful grid setups, is shown in Fig. 5. Blocking may happen for two reasons. First, there may not be sufficient free or shareable capacity on any of the candidate paths to establish all the connections. And second, it may be impossible to find a configuration which is able to meet the requirement on p . Steps in the figure are related to shifts between **Class 1** and **Class 2**.

Fig. 6 shows the average cost of establishing the connections required for the grid in terms of resources usage. The cost is modeled as proportional to units of kilometers of wavelengths used in the working paths and per non-shared wavelength used in the backup paths. The re-use of a wavelength in a backup

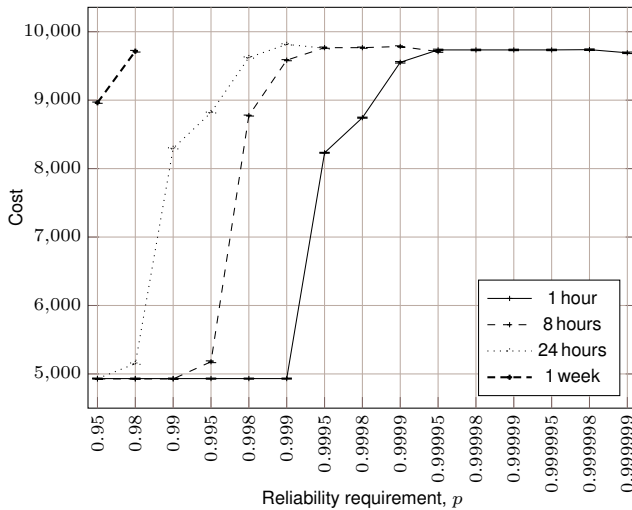


Fig. 6. Simulation results: resources usage. Different lines connect points reported for different values of T_{dur}

TABLE III
RESOURCES USAGE (CONNECTION COST)

Connection type	Resources usage	
	Single connection	Dedicated protection
Birmingham-Amsterdam	520	1633
Brussels-Amsterdam	173	526
Krakow-Amsterdam	1378	3390
Milan-Amsterdam	824	2124
Stockholm-Amsterdam	1494	5058

path has zero cost. For the lower requirements, the cost is of the order of 5000. This corresponds to single connections only. The cost of the order of 10,000 is the result of protected connections, where almost none of the links are re-used, i.e., the arrangement resembles dedicated path protection.

The results show that T_{dur} has a strong impact on the ability to establish the grid connections meeting the requirement. It is only possible to establish connections for all the considered values of p for $T_{dur} = 1$ hour. For this duration, single unprotected connections are sufficient up to $p = 0.999$. On the other hand, for $T_{dur} = 1$ week, it is possible to establish the grid connections only in a few cases and for $p < 0.99$ using dedicated protection.

C. Comparison with the Availability-based Approach

Here, the above results are compared with the more commonly adopted approach, that is based on the assessment of the asymptotic (steady-state) availability. Two metrics are used: the availability itself and a cost of connections (again, resources usage), without and with dedicated protection. Tables III and IV present the values

First, note that an asymptotic value is obtained and it is impossible to deal with t_p , to take into account how long the service lasts. Thus, independently of the leasing time, the same value of the availability would be required. It can be noted that four of the five connections have the single connection availability at the level of two nines. This is far

TABLE IV
CONNECTIONS AVAILABILITY

Connection type	Availability	
	Single connection	Dedicated protection
Birmingham-Amsterdam	0.997	0.99998
Brussels-Amsterdam	0.9992	0.999998
Krakow-Amsterdam	0.993	0.99994
Milan-Amsterdam	0.996	0.99997
Stockholm-Amsterdam	0.993	0.9998

TABLE V
COMPARISON OF t_p VALUES IN HOURS FOR SINGLE CONNECTIONS

Connection type	p			
	0.95	0.99	0.999	0.9999
Birmingham-Amsterdam	259	50	5	0.5
Brussels-Amsterdam	779	152	15	1.5
Krakow-Amsterdam	97	19	2	0.2
Milan-Amsterdam	163	32	3	0.3
Stockholm-Amsterdam	90	17	1.75	0.2

too low to establish the connection using the availability-based approach, and then the operator would decide to use the protected connection. Even if it is assumed that three nines' availability would be enough¹ for the Brussels-Amsterdam connection, the average cost of using such a layout would be at the level of 13,000 cost units for all connections summed. When these findings are related to the results presented in Fig. 6, it can be seen that protection is not necessary unless connections last for a long time (one week) or a very high continuity assurance is required.

Secondly, it should be kept in mind that three nines in the case of the continuity is definitely not the same as in the case of the asymptotic availability. Look at the results presenting t_p values in Table V. When the continuity-based approach is adopted, values related to the predicted service time are interesting. If it is very short, and the percentile requirement can be met (i.e., t_p is shorter than the service time), then a single connection can be used. For instance, if the transfer is finished in tens of minutes, a guarantee between 0.999–0.9999 that the single connection is up will be sufficient, resulting in large cost savings. A similar argument can be given, but from another point of view, by regarding Table VI presenting the reliability function values for single connections when different mission times are considered. If the time is shorter than one hour, a three nines' reliability guarantee that the single connection is up is sufficient. This is different from the three nines' steady-state availability given in Table IV. It is seen that in fact, it is not possible to infer the continuity attribute from the asymptotic availability attribute. Additionally, even a high availability will not guarantee the satisfactory value of the continuity if failures have high intensity, cf. the arguments in Section II.

V. CONCLUSION

In this paper, it is emphasized that for some types of connections, the continuity criterion, i.e., a delivery of unin-

¹Usually availability requirement is set to much higher value, e.g., five nines or more.

TABLE VI
COMPARISON OF $R(\cdot)$ VALUES FOR SINGLE CONNECTIONS

Connection type	Time			
	1 hour	8 hours	24 hours	1 week
Birmingham-Amsterdam	0.9998	0.9984	0.9952	0.9673
Brussels-Amsterdam	0.9999	0.9994	0.9984	0.989
Krakow-Amsterdam	0.9994	0.9957	0.9874	0.9151
Milan-Amsterdam	0.9996	0.9974	0.9925	0.9486
Stockholm-Amsterdam	0.9994	0.9954	0.9864	0.9089

interrupted service, should be adopted as the prime reliability parameter for SLAs and provisioning. For some types of connections, this is a more appropriate approach than the steady-state availability. A discussion on rationale for such approach is presented and different situations where it is relevant are given: connections where failure consequences can be catastrophic, where frequent failures have a large negative impact on the perceived quality, when a connection lasts for a very short period of time, and when a service is agreed only for a short time period. Then, a comprehensive study with a simple mathematical modeling is presented in a situation well suitable for the continuity-based approach, i.e., to connection provisioning in the utility grid. As a further research, the following two problems are predicted: 1) studies on determining the level of reliability function percentile (t_p) required for different applications, that is, focus on conditions that enable a practical usage of the presented methodology, and, 2) investigations of more precise models, taking into account complex relations between different components of a reliable system, e.g., software reliability of routers or recovery configuration faults.

ACKNOWLEDGEMENT

This work was done within the EU FP7 NoE Euro-NF (<http://www.euronf.org>) framework. The reported work was also supported by the Polish Ministry of Science and Higher Education under grant N517 013 32/2131. Q2S — Centre for Quantifiable Quality of Service in Communication Systems, Centre of Excellence is appointed by The Research Council of Norway and funded by The Research Council, NTNU and UNINETT.

REFERENCES

- [1] L. Song and B. Mukherjee, "On The Study of Multiple Backups and Primary-Backup Link Sharing for Dynamic Service Provisioning in Survivable WDM Mesh Networks," *IEEE J. Select. Areas Commun. Supplement on Optical Communications and Networking*, vol. 26, no. 6, pp. 84–91, Aug. 2008.
- [2] P. Cholda, A. Mykkeltveit, B. E. Helvik, O. J. Wittner, and A. Jajszczyk, "A Survey of Resilience Differentiation Frameworks in Communication Networks," *IEEE Comm. Surv. Tut.*, vol. 9, no. 4, pp. 32–55, 2007.
- [3] W. D. Grover, *Mesh-Based Survivable Networks. Options and Strategies for Optical, MPLS, SONET, and ATM Networks*. Upper Saddle River, NJ: Prentice Hall PTR, 2004.

- [4] T. Ma, J. Hillston, and S. Anderson, "On the Quality of Service of Crash-Recovery Failure Detectors," in *Proc. 37th Annual IEEE/IFIP International Conference on Dependable Systems and Networks DSN 2007*, Edinburgh, UK, Jun. 25–28, 2007.
- [5] M. D. Beaudry, "Performance-Related Reliability Measures for Computing Systems," *IEEE Trans. Rel.*, vol. C-27, no. 6, pp. 540–547, Jun. 1978.
- [6] P. Cholda, J. Tapolcai, T. Cinkler, K. Wajda, and A. Jajszczyk, "Quality of Resilience as a Network Reliability Characterization Tool," *IEEE Network*, vol. 23, no. 2, pp. 11–19, Mar./Apr. 2009.
- [7] 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 Trans. Networking*, vol. 16, no. 4, pp. 749–762, Aug. 2008.
- [8] The Norwegian Research Network UNINETT. (2009) Downtime Statistics. [Online]. Available: <http://drift.uninett.no/downs/>
- [9] B. E. Helvik, *Dependable Computing Systems and Communication Networks. Design and Evaluation*. Trondheim, Norway: Department of Telematics, NTNU Norwegian University of Science and Technology, 2003.
- [10] M. Grottke, H. Sun, R. M. Fricks, and K. S. Trivedi, "Ten Fallacies of Availability and Reliability Analysis," in *Proc. 5th International Service Availability Symposium ISAS 2008*, Tokyo, Japan, May 19–21, 2008.
- [11] J. A. Buzacott, "Markov Approach to Finding Failure Times of Repairable Systems," *IEEE Trans. Rel.*, vol. R-19, no. 4, pp. 128–134, Nov. 1970.
- [12] P. Szegedi, Z. Lakatos, and J. Späth, "Signaling Architectures and Recovery Time Scaling for Grid Applications in IST Project MUPBED," *IEEE Commun. Mag.*, vol. 44, no. 3, pp. 74–82, Mar. 2006.
- [13] A. Jukan and G. Karmous-Edwards, "Optical Control Plane for the Grid Community," *IEEE Comm. Surv. Tut.*, vol. 9, no. 3, pp. 30–43, 2007.
- [14] "Grid Optical Burst Switched Networks (GOBS)," GFD-I.128, OGF Informational Document, Apr. 2008.
- [15] "Optical Network Infrastructure for Grid," GFD-I.036, OGF Informational Document, Aug. 2004.
- [16] "Networking Issues for Grid Infrastructure," GFD-I.037, OGF Informational Document, Nov. 2004.
- [17] K. Munir, S. Javed, and M. Welzl, "A Reliable and Realistic Approach of Advance Network Reservations with Guaranteed Completion Time for Bulk Data Transfers in Grids," in *Proc. 1st International Conference on Networks for Grid Applications GridNets 2007*, Lyon, France, Oct. 17–19, 2007.
- [18] M. L. Shooman, *Reliability of Computer Systems and Networks: Fault Tolerance, Analysis, and Design*. New York, NY: John Wiley & Sons, Inc., 2002.
- [19] J. Zhang, K. Zhu, H. Zang, N. S. Matloff, and B. Mukherjee, "Availability-Aware Provisioning Strategies for Differentiated Protection Services in Wavelength-Convertible WDM Mesh Networks," *IEEE/ACM Trans. Networking*, vol. 15, no. 5, pp. 1177–1190, Oct. 2007.
- [20] A. Mykkeltveit and B. E. Helvik, "On Provision of Availability Guarantees Using Shared Protection," in *Proc. 12th International Conference on Optical Network Design and Modelling ONDM 2008*, Vilanova i la Geltrú, Spain, Mar. 12–14, 2008.
- [21] "Grid Network Services Use Cases from the e-Science Community," GFD-I.122, OGF Informational Document, Dec. 2007.
- [22] R. Wessäly, R. Klähne, and S. Orłowski. (2009) SNDlib: Library of Test Instances for Survivable Fixed Telecommunication Network Design. [Online]. Available: <http://sndlib.zib.de>
- [23] M. Tacca, P. Monti, and A. Fumagalli, "The Disjoint Path-Pair Matrix Approach for Online Routing in Reliable WDM Networks," in *Proc. IEEE International Conference on Communications ICC 2004*, Paris, France, Jun. 20–24, 2004.
- [24] S. Verbrugge *et al.*, "Methodology and Input Availability Parameters for Calculating OpEx and CapEx Costs for Realistic Network Scenarios," *OSA J. Opt. Netw.*, vol. 5, no. 6, pp. 509–520, Jun. 2006.