Guaranteeing service availability in optical network design

Dominic A. Schupke

Siemens, Corporate Technology
Information and Communications
Optical Networks and Transmission
Otto-Hahn-Ring 6, 81730 Munich, Germany
Tel: +49 89 636-43518, Fax: -51115
Email: dominic.schupke@siemens.com

ABSTRACT

We show the principal methodology for providing connections with minimum cost and subject to target availability constraints, that can be used to guarantee service availability in a service level agreement (SLA). For connections without protection and for connections with dedicated protection, we develop path-based design models and evaluate the gain of availability-based provisioning.

Keywords: Availability, Network Design, Optical Networks, SLA.

1. INTRODUCTION

Backbone networks provide services with high data rates. Because of the geographical size alone, these networks are subject to frequent failure occurrence and long repair durations. Therefore, outages in backbone networks can have a critical impact on the provided services. As a result, there is a need for network survivability, demanding on the one hand fast recovery and high availability, and on the other hand resource efficiency (i.e., low cost). In this paper, we develop methods for providing services with minimum cost and subject to target availability constraints.

We can obtain the target availabilities from service level agreements (SLAs). A SLA is a contract between a customer and a network operator that specifies an availability performance, e.g., by

$$a_{SLA} = 1 - \frac{\sum \text{ outage times in } \Delta t}{\Delta t}. \tag{1}$$

This definition is based on a measurement period of time Δt (e.g., $\Delta t = 1$ year). The target availabilities for the constraints can be lower than the SLA availabilities, e.g., if we include margins between the period-based availability definition of Equation (1) and the steady-state availability definition used in calculations.

Technologically, we mainly consider optical networks such as wavelength division multiplexing (WDM) networks. We assume circuit-switched connections (the almost only used switching type in optical networks) of service-types unprotected connection and dedicated path protection (e.g., 1+1 path protection) connection, along with an outlook on the service-types shared (backup) path protection connection and path-restorable connection.

This paper is structured as follows. Section 2 forms the basis to set up models for availability-based network design. Section 3 describes how the selection of failure scenarios in the design impacts network performance. Section 4 describes provisioning design models. In Section 5, a network case study shows the benefit we have from guaranteeing service availability within the design. In Section 6 we draw conclusions and provide an outlook.

2. MODELING BASIS

Generally, we model a system composed of a set of items, where an item can be a device or service, or again a system. For each item we can define a failure rate, which we assume to be constant over time. The failure rate is denoted λ and often measured in units of "failure in time (FIT)" with 1 FIT = 1 failure in 10^9 hours. A constant failure rate results in an exponentially distribution for the time to failure, with a mean time to failure (MTTF) obeying $MTTF = \frac{1}{\lambda}$.

Optical networks are repairable systems. For the items we assume that the time to repair is also exponentially distributed. The mean time to repair (MTTR) obeys $MTTR = \frac{1}{\mu}$, where μ denotes the repair rate.

We define the instantaneous availability $a(t)$ as the probability that an item will perform its required function under given conditions at a given *instance* of time. Availability is only defined for a given usage time (often also called life time or service time). We assume a two state model for repairable systems consisting of the operating (up) state and failed (down) state. The item is up at $t = 0$ (when new) and "as new" after repairs. No further failures occur during repairs.

From this, we define the steady-state availability as

$$a = \lim_{t \to \infty} a(t) = \frac{\mu}{\lambda + \mu} = \frac{MTTF}{MTTF + MTTR}. \tag{2}$$

When we write "availability" we mean the steady-state availability. The unavailability is defined as

$$u = 1 - a = \frac{MTTR}{MTTF + MTTR}. \tag{3}$$

As the average repair time is much shorter than the average time to failure (thus $\mu \gg \lambda$), the mean time between failures (MTBF) approximates MTTF:

$$MTBF = MTTF + MTTR \simeq MTTF. \tag{4}$$

Note that reliability is sometimes employed in the sense of availability in literature. However, reliability differs from availability, since reliability is defined as the probability that an item will perform its required function under given conditions for a given *period* of time.

Complex systems are (recursively) partitionable in subsystems/items which are connected in series or parallel. This is under the assumption of independent failure processes (holds typically for optical networks) and independent repair processes (holds, e.g., in presence of several repair teams).

Figure 1 shows an example for a WDM network composed of several network elements (items). The figure also shows a connection configured between nodes S and T. In Figure 2, the WDM network is mapped onto a graph that consists of nodes and edges representing the network elements. The connection passes through nodes and edges and represents further network elements such as transponders at the end-nodes.

For a notion of typical FIT values, we list several published values. Table 1 shows FIT values for network elements as published in literature.[1] For fiber ducts we can state these typical FIT values[2]:

1. COST239 values[3]: 114 FIT per kilometer (MTBF = 1000 years)
2. Average of European network operators in FIRST project[4]: 200 FIT per kilometer (MTBF = 570 years)
3. Bellcore rates[5]: 4.39 per year per 1000 miles $\simeq$ 311 FIT per kilometer (MTBF = 367 years)

MTTR values depend on the system type. Typically, it is assumed that components have a MTTR ranging from 4 to 6 hours and fiber links have a MTTR ranging from 12 to 24 hours. This reflects that components are easier to access (e.g., locally and in-house), whereas fiber links are harder to access (e.g., far away and buried underground).

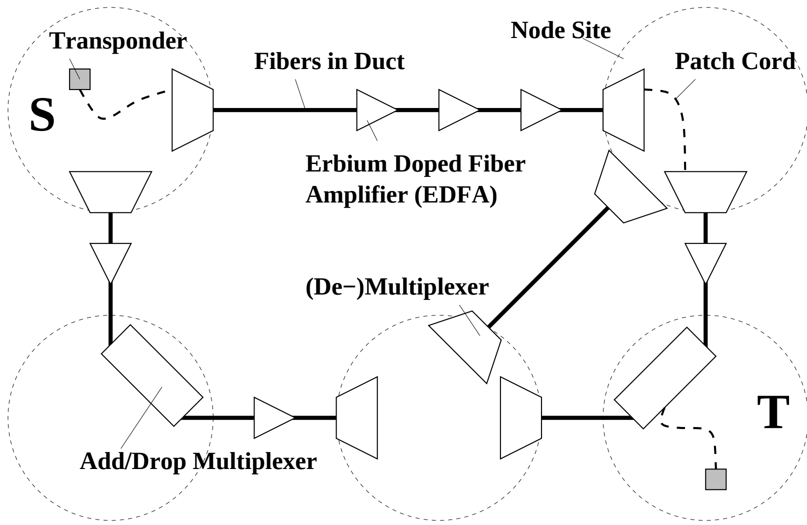

Figure 1. WDM network example.

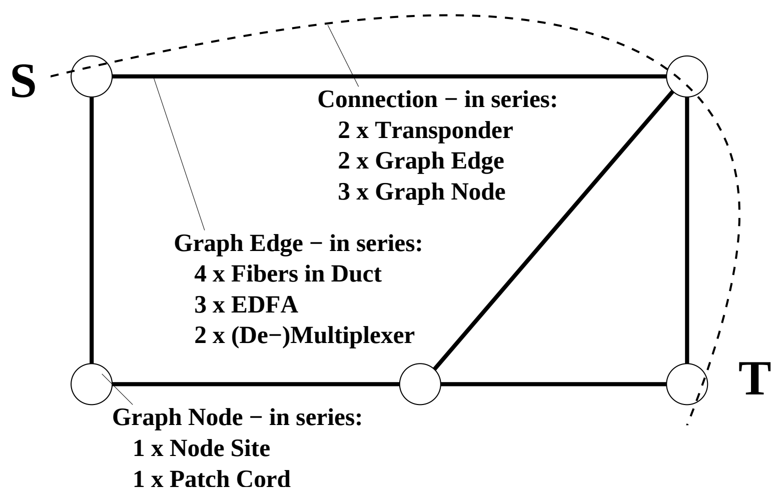

Figure 2. Example network mapped to a graph.

A connection is provisioned by one or more paths in the network. Two paths are *edge-disjoint* if they do not share common edges. A path is *node-disjoint* from another path if they do not have common nodes, except for possibly the initial and final nodes. Node-disjointness includes edge-disjointness.

Unprotected connections need only one path. Dedicated path protection (e.g., 1+1 path protection) connections need two paths being edge-disjoint or node-disjoint. Edge-disjoint paths are chosen to survive single link failures and node-disjoint paths are chosen to survive single failures of both links and the paths' intermediate nodes.

3. FAILURE SCENARIOS

In optical network design, the choice and dimensioning of a recovery is today mostly based on failure scenarios. Failure scenarios are sorted in a "priority list" according to the probability of occurrence and the involved impact.

Covering more failure scenarios comes with a higher cost, thus design considers only the top entries of the priority list.

An example of a priority list is as follows:

1. Single link failures
2. Dual link failures (second link failed before first link repaired)
3. Node failures (affects traffic through node)
4. Joint single link and single node failures
5. Dual node failures
6. Higher order failures

In most literature, we see that design is done assuming single link failures, i.e., there is only at most one link failed at any time. Several publications,[2] however, have shown the significance of failures higher in order than single failures. This holds especially for dual link failures, i.e., there are at most two links simultaneously failed. Figure 3 shows the presence of three link failure types in the COST 239 pan-European network[3] for a fiber duct failure rate of 200 FIT/km. The figure shows the per-year presence of multiple failures, of two failures on a ring (for a ring design in[2]), and of a pair of disjoint paths (lower bound). For a MTTR of 24 hours, we have to cope with higher order failures in the range of minutes to hours per year.

The selection of failure scenarios facilitates network design, since the scenarios are describable as a direct parameter input into the model. However, whenever we assume a selection level for the design, we may still miss the service availability targets (e.g., not enough failure scenarios covered) or we may excessively overprovision capacities (e.g., too many failure scenarios covered).

The latter aspect is particularly apparent in dual failure designs. For instance, dedicated path protection involves three disjoint paths to survive dual failures, thus, we need at least to triple the capacity of a single path. If we adopt that the design covers dual failures, every path needs this capacity and, as a result, we need very high capacities in the network. However, several paths may require less than three paths to meet the service availability targets, which means lower overall network capacity requirements.

Here, the paradigm of availability-based provisioning comes into context. Under this paradigm, we configure paths in the network under constraints ensuring that availability requirements are met. The design does not

Table 1. Typical FIT values for network elements. Variable W denotes the number of wavelengths and variable N the number of incoming fibers. All components are unidirectional.

Component	**Failure Rate in FIT**
(De-)Multiplexer	$25 \times W$
EDFA	2850
Optical Switch 1	$21W \times W/4$
Optical Switch 2	$21 \times 2 \times 2 \times N$
Coupler 1	25×2
Coupler 2	$25 \times W/4$
Coupler 3	$25 \times (N-1)$
Tunable Transmitter	745
Fix Transmitter	186
Tunable Receiver	470
Fix Receiver	70
Digital Switch 1	$875 \times W$
Digital Switch 2	$875 \times W \times N$
Wavelength Blocker	$50 \times W$

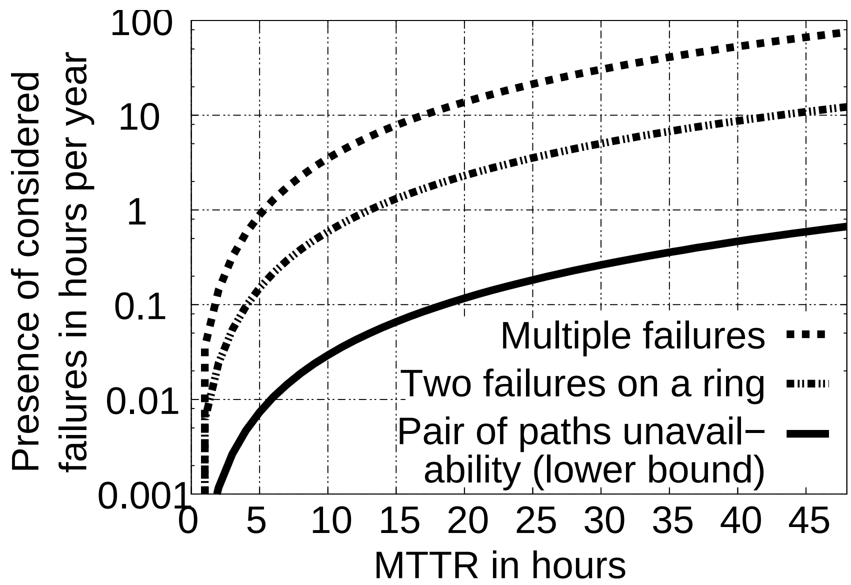

Figure 3. Per-year presence of higher order failures in the COST 239 network.

rely on a selection of covered network-wide failure scenarios, rather it chooses the protection level to fulfill the availability requirements for each connection individually.

4. NETWORK AND PATH DESIGN

In the traditional two-step provisioning process, we firstly design the network (or an individual path) such that it recovers given failure scenarios. Secondly, we calculate the availability performance (a_c). If we miss the availability target ($a_c < a_t$), we can choose from the following means: (i) Include more failure scenarios and restart. (ii) Add redundant resources (e.g., redundant sender/receiver) and recalculate the availability performance. (ii) Consider monetary penalties (which are also specified in SLAs).

In the availability-based provisioning process, we take service availability directly into account and aim at providing minimum resources. The recovery mechanism is either given or selectable. In the former case, we optimize subject to $a_c \geq a_t$, otherwise we aim to get a best-effort value $\max a_c$. In the latter case, we find the optimal recovery mechanism meeting $a_c \geq a_t$ (recovery type differentiation).

To design paths, we can choose from several alternatives[6]: Integer linear programs (ILPs) with optimization using a link-flow model (LFM) or a path-flow model (PFM), as well as direct routing algorithms on a graph model. In the following, we elaborate on PFMs. While we consider only one demanded connection, the PFMs are simple to extend for multiple demands (e.g., for static network design).

We model the network as undirected and simple graph $G(V, E)$, with the set of nodes V and the set of edges E. Per edge $e \in E$, we define a cost parameter γ_e. The end nodes of the connection are specified by o and d. The indicator $\delta_{e,n}$ is 1, if node $n \in V$ is one of the end nodes of $e \in E$, 0 otherwise. In the PFM we precompute a set of m candidate paths $P_p \subset E, \forall p \in \{1, \ldots, m\}$, using, e.g., a standard m-shortest path algorithm.

In a simple form, we can state the general path-flow model PFM_{1,a_t} for a single path as follows:

$$\min \sum_{e \in E} \gamma_e \sum_{p \in \{1,\ldots,m\}: e \in P_p} z_p \tag{5}$$

$$\sum_{p \in \{1,\ldots,m\}} z_p = 1 \tag{6}$$

$$z_p \in \{0, 1\}, \forall p \in \{1, \ldots, m\} \tag{7}$$

The Objective (5) minimizes the path's overall edge costs. Constraint (6) ensures that the demand between o and d is satisfied. The integer flow variables z_p are defined in Expression (7), meaning that z_p is one, if the path $p \in \{1, \ldots, m\}$ is selected, zero otherwise. Note that the PFM can be extended by further constraints (e.g., capacity constraints). To keep things simple, we concentrate only on the necessary model contributions.

We can introduce the availability requirements by admitting only those paths $p \in \{1, \ldots, m\}$ meeting

$$a_p = \prod_{e \in P_p} a_e \prod_{n \in V : \exists e \in P_p : \delta_{e,n}=1} a_n \geq a_t \tag{8}$$

where a_p, a_e, and a_n are the availabilities for a path, an edge, and a node.

We can generalize this ILP for $k \in \{1, 2, \ldots\}$ disjoint paths, denoted PFM_{k,a_t}. We describe two alternative methods, applicable in different modeling environments.

1. Method: We use candidate path tuples $(p_1, p_2, \ldots, p_k)$ fulfilling

i) link-disjointness:

$$P_{p_i} \cap P_{p_j} = \emptyset, \forall i, j \in \{1, \ldots, k\}, i > j \tag{9}$$

ii) node-disjointness (optional):

$$\sum_{i \in \{1,\ldots,k\}, e \in P_{p_i}} \delta_{e,n} = 2, \forall n \in V / \{o, d\} \tag{10}$$

To meet the availability requirements, we admit only feasible tuples:

$$1 - \prod_{i \in \{1,\ldots,k\}} (1 - \frac{1}{a_o a_d} a_{p_i}) \geq \frac{a_t}{a_o a_d} \tag{11}$$

If the paths are not node-disjoint, Condition (11) holds only for $a_n = 1, \forall n \in V$. As the final step, we employ the PFM, however, with path tuples instead of single paths

$$P_1^* = \cup_{p \in \{1,\ldots,k\}} P_{p_i}, P_2^* = \ldots \tag{12}$$

2. Method: We make use of replication, where the variables and constraints of the PFM_{1,a_t} are replicated k times. The objective sums the with cost weighted variables over all k disjoint paths. Similar to Equations (9)-(11), further constraints force variables to zero (i.e., $\sum_{i \in \{1,\ldots,k\}} z_{i,p}$=0) which fail in either of

i) link-disjointness

ii) node-disjointness (optional)

iii) availability met.

5. NETWORK CASE STUDY

To gain insight of the benefit of availability-based connection provisioning, we consider a case study using the COST 239 network.[3] We assume edge failures only and we determine the availability of an edge by its length, using the failure rate of 200 FIT/km and a MTTR of 12 hours. The network is three-edge-connected so that it is topologically feasible to survive dual failures if we provide adequate protection capacity.

The network has eleven nodes, yielding 55 node pairs between which demand can occur. Figure 1 shows the number of node pairs which are connectable by one path, by two disjoint paths, and by three disjoint paths for a minimum target availability value.

Clearly, three edge-disjoint paths always achieve the target availability, however, they need excessive resources. Single path solutions are only feasible for low availability values (at most 99.9%). Two edge-disjoint paths can cover all or a large portion of the node pairs until an availability level of 99.9995%, which is in the region of a five-nines service (i.e., an service with 99.999% SLA availability) when taking node availability and margins into account. At a target availability of 99.9995%, approximately half of the connections can satisfy the availability constraints by two edge-disjoint paths, while the other connections need three edge-disjoint paths. Hence, a corresponding path assignment can significantly spare on network resources.

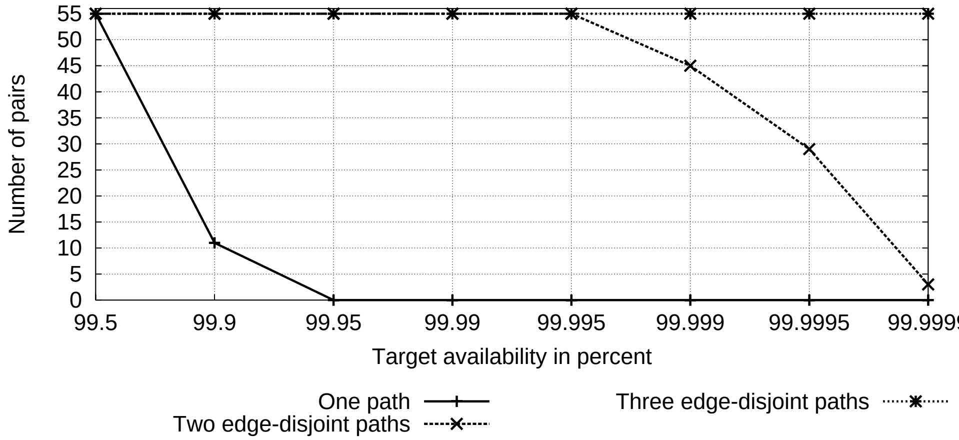

Figure 4. Number of connectable node pairs over target availability.

6. CONCLUSIONS AND OUTLOOK

In optical network design, we can guarantee service availability by direct inclusion of availability constraints. This paper has exemplified the methodology for unprotected and protected connections, showing sparings in network resources. In contrast to network design based on failure scenarios, we are able to assign as much network resources as needed for the individual connections.

The approach is not restricted to unprotected connections and dedicated path protection connections, as presented in this paper. Recent literature has shown that it can also be used for shared (backup) path protection connections[1, 7–9] and path-restorable connections.[6] From the availability point of view, e.g., shared path protection is like dedicated path protection with reduced availability of the backup path, because of sharing with other connections. Different approaches can approximate this reduction.

ACKNOWLEDGMENTS

The author thanks Franz Rambach for his comments on the paper.

REFERENCES

1. D. Mello, D. Schupke, M. Scheffel, and H. Waldman, "Availability Maps for Connections in WDM Optical Networks," in *Proceedings of the International Workshop on Design of Reliable Communication Networks (DRCN)*, (Island of Ischia (Naples), Italy), Oct. 2005.
2. D. A. Schupke, A. Autenrieth, and T. Fischer, "Survivability of Multiple Fiber Duct Failures," in *Proceedings of the International Workshop on Design of Reliable Communication Networks (DRCN)*, (Budapest, Hungary), Oct. 2001.
3. P. Batchelor, B. van Caenegem, B. Daino, P. Heinzmann, D. R. Hjelme, R. Inkret, H. A. Jäger, M. Joindot, A. Kuchar, E. Le Coquil, P. Leuthold, G. de Marchis, F. Matera, B. Mikac, H. P. Nolting, J. Späth, F. Tillerot, N. Wauters, and C. Weinert, *Ultra High Capacity Optical Transmission Networks*, Faculty of Electrical Engineering and Computing, University of Zagreb, Zagreb, 1999. Final Report of COST Action 239.
4. D. Gardan, L. Haagh, L. Larson, R. Margarito, M. Redstall, A. Socard, H. Steffensen, and J. Stroetzel, "Availability analysis of the optic local loop," in *EFOC & N*, 1994.
5. M. To and P. Neusy, "Unavailability analysis of long-haul networks," *IEEE Journal on Selected Areas in Communications* **12**, Jan. 1994.
6. D. A. Schupke, "Availability-Based Service Provisioning in Optical Networks," in *IEEE Global Telecommunications Conference (GLOBECOM)*, (St. Louis, Missouri, USA), Nov. 2005. Tutorial.

7. J. Zhang, K. Zhu, B. Mukherjee, and H. Zang, "Service provisioning to provide per-connection-based availability guarantee in WDM mesh networks," in *Proceedings of the IEEE/OSA Optical Fiber Communication Conference (OFC)*, (Atlanta, GA, USA), Mar. 2003.
8. J. Zhang, K. Zhu, H. Zang, and B. Mukherjee, "A new provisioning framework to provide availability-guaranteed service in WDM mesh networks," in *Proceedings of the IEEE International Conference on Communications (ICC)*, (Anchorage, AK, USA), May 2003.
9. D. A. A. Mello, J. U. Pelegrini, R. P. Ribeiro, D. A. Schupke, and H. Waldman, "Dynamic Provisioning of Shared-Backup Path Protected Connections with Guaranteed Availability Requirements," in *Proceedings of the IEEE/CreateNet International Workshop on Guaranteed Optical Service Provisioning (GOSP)*, (Boston, MA, USA), Oct. 2005.