On the Trade-off between Cost and Availability of Virtual Networks

Sandra Herker, Wolfgang Kiess, Xueli An
DOCOMO Communications Laboratories Europe GmbH
Munich, Germany
{herker, kiess, an_de_luca}@docomolab-euro.com

Andreas Kirstädter
Institute of Communication Networks and Computer Engineering (IKR), Universität Stuttgart, Germany
andreas.kirstaedter@ikr.uni-stuttgart.de

***Abstract*—To minimize cost, Virtual Network Operators (VNOs) need to consider the required network availability already at the network design stage. One generic approach to reach the availability target is to select only high-quality physical network elements that offer high availability and consequently demand high expenses per element. The other generic approach to achieve high availability is to add protection capacity on the level of the virtual network based on lower cost components. In this paper, we analyze both alternatives with a simulation tool to answer the fundamental question how quality can be traded against capacity. For this purpose, we consider different network topologies and the influence of different parameters and provide a framework to find an optimal strategy between Mean Time Between Failures (MTBF) targets for the physical infrastructure and the usage of additional backup paths on the virtual network level.**

I. Introduction

Besides connectivity and capacity, any carrier-grade virtual network has also to comply to availability targets at coping with fiber cuts and other failures. Constantly trying to minimize cost at renting (virtualized) links and nodes from Physical Infrastructure Providers (PIP) [1], the Virtual Network Operator (VNO) therefore faces a basic choice: it may build upon highly reliable network elements (nodes and links) or apply protection and restoration mechanisms on the basis of a larger number of network elements with lower availability figures (and thus lower cost).

The first option – the 'high cost physical network' approach – uses direct, shortest paths with high availability basing on high cost links, i.e. the operator will invest in the infrastructure. The second option – the 'low cost physical network' approach – realizes the necessary network availability in the virtual network domain by combining several parallel paths with lower availability and lower individual cost. Here, the individual physical network elements can be kept cheap, while a larger number of them are required to realize the parallel paths. Thus, a trade-off can be expected between link quality (small number of expensive paths) and capacity (combination of multiple cheap paths).

In this paper, we examine this trade-off between the two choices of a 'high cost physical network' approach and a 'low cost physical network' approach focusing on the optical links. Our contribution is a framework to identify the cost optimal values of the Mean Time Between Failures (MTBF) parameters of the physical links considering the influence of different parameters and network topologies. Further we analyze the fundamental inter-dependencies and give recommendations. To the best of our knowledge this is the first paper to consider how changes in the underlying physical infrastructure (higher MTBF values or usage of several backup paths) will interact to achieve the desired link availability of a Virtual Network Request (VNR) - the embedding with lowest cost.

In the following sections, we first consider the (fiber) availability and how to determine it from measured MTBF values. Further, we model the relationship between MTBF and cost derived from existing values. Next, we explain our virtual network embedding algorithm that is mapping a VNR with desired availability onto the physical network while mapping the virtual nodes to the matching physical nodes and the virtual link to a single or several parallel physical (backup) paths. This algorithm is then used to investigate the physical infrastructure deployment strategy considering the MTBF to achieve cost reduction. Several cost-model parameters, different network topologies and sizes, and requested link availability values will be considered. We show that the network topology (nodal degree) and size is strongly influencing the Virtual Network Embedding (VNE) deployment strategy to achieve minimum cost. In the last section we give some recommendations for network operators derived from our results.

II. Related work

In this section, we discuss work related to our problem in the areas of network availability calculation, Virtual Network Embedding (VNE), and cost modeling.

A. Related Work on Network Availability Calculation

Several authors describe the computation of the network availability in general and in the area of optical networks: [2] and [3] describe how the basic availability of a network can be calculated and providing exact calculation [3] or analytical expressions for several different network topologies like star or crown [2]. Other works focus on end-to-end connection availability in optical transport networks. The connection availability for different resilience mechanisms like unprotected, dedicated and shared path protection and path restoration in [4], in [5] and [6] dedicated and shared path protection and

ISBN 978-3-901882-58-6 © 2014 IFIP

in [7] only dedicated protection is compared. Additionally in [4], the authors list general availability numbers for several network equipment types.

In [8], the authors examine the relation between the path availability (the product of the availabilities of the components – nodes and links – that belong to the path) and the restorability of a network to dual failures (i.e., two failures present at a given time).

These works compute or analyses the availability of an existing network with or without protection. However, they do not consider availability at the planning stage.

B. Related Work on Virtual Network Embedding (VNE)

The objective of VNE is to find an effective and efficient mapping for the VNR to the physical network using a cost function. The VNE can be considered as a process with two stages: virtual node mapping and virtual link mapping. In the first stage, virtual nodes are mapped to resource nodes in the physical network. In the link mapping, for each virtual link a feasible path between the corresponding physical nodes that host the virtual nodes of that virtual link is calculated.

Many different basic solutions for embedding VNs exist [9], [10], [11], [12], however, the issue of availability in VNE is not considered therein. Other VNE algorithms consider survivability methods like protection [13], [14], [15] and restoration [16]. The presented algorithms assume simple cost functions at minimizing resource consumption (especially bandwidth) or maximizing the revenue and acceptance ratio of VNRs. In [17], a cost function is used for the VNE where the cost increases exponentially as the link traffic increases.

C. Related Work on Cost Modeling

The authors of [18] argue that reductions in the physical Mean Time To Repair (MTTR) can also enhance availability at full or partial dual-failures - besides adding protection capacity. They show that an economic strategy exists for balancing the trade-off between capacity investment and MTTR reduction efforts to achieving high availability in networks designed to be 100% restorable against single failures. They model the cost functions for maintenance expenditures (considering the required repair time) and also the spent protection capacity and survivability mechanism. As reported in [18], with a reasonable approximation, the value of MTTR can be shown to be directly proportional to the physical unavailability of each span.

Existing work provides availability calculation and analysis in networks, however, not in virtual networks. Changing the underlying physical infrastructure to influence availability is not considered at network planning stage. Especially, the MTBF for the fiber is considered to be constant.

Our previous work [19] focuses on achieving the requested link availability for path protection for virtual network embedding. The optimization goal for the algorithm is to minimize the bandwidth consumption. The link availability is not calculated from the link length, instead it is distributed randomly. In this work, we first need to calculate the minimum overall cost for the embedding for the requested availability before we can examine the trade-off between investment in physical infrastructure or backup paths. In order to calculate this minimum overall embedding cost, we extend the algorithm from [19] accordingly such that it can be used as a tool for solving the trade-off problem.

TABLE I
MTBF VALUES FROM DIFFERENT PAPERS FOR FIBER CABLE

Reference	MTBF(hours) for 1 km	FIT (/km)	MTTR (hours)
Fiber, aerial [4]	1.75×10^5		6
Fiber, buried [opti] [4]	5.5×10^6		9
Fiber, buried [nomi] [4]	2.63×10^6		12
Fiber, buried [cons] [4]	2.41×10^6		24
Optical fiber (PON) [21]	5×10^6	200	14
Fiber [22]	1.75×10^6	570	24
Fiber and inline amplifier [7]	3.23×10^6	310	12
Fiber [23]	8.773×10^6		-
Fiber [23]	4.99×10^6		-
Fiber [23]	3.21×10^6		-

III. Technical and Practical Considerations

Here, we will provide background information on MTBF and fiber cost modeling to prepare the next sections.

A. Availability and MTBF

The availability of any component is defined as the percentage of time when the component is operational and fulfilling its requirements, i.e. as the relation between its uptime and the sum of its uptime and downtime [20]. Using common parameters as the Mean Time Between Failures (MTBF, i.e. the expected time between two failures of the component) and the Mean Time To Repair (MTTR, time required to repair or replace the component) [4], this can be expressed by

$$A = (MTBF - MTTR) \, / \, MTBF \tag{1}$$

The availability of a simple path through a network can be determined as the product of the availabilities of the components (nodes and links) that belongs to the path [2]. Network availability is then the minimum path availability over all shortest paths between two distinct node pairs.

As we consider wide area transport network scenarios, we need to determine the availability of each individual optical fiber cable. It is commonly calculated [4] by considering measured values for the average cable length CC (in kilometer) that experiences one cable cut within an one year period. The MTBF value (in hours) of optical fiber cable, from which the availability can be derived using the MTTR, can be calculated as:

$$MTBF\ (hours) = \frac{CC(km) \times 365 \times 24}{total\ cable\ length\ (km)} \tag{2}$$

Table I shows MTBF and MTTR values for fiber optical links taken from the literature for fiber cable. It can be seen that realistic MTBF values range between 1.5 and 5.5×10^6 hours for 1 km fiber.

As we want to calculate the End-to-End (E2E) physical path availability, we also have to consider the nodes along the path. The availability of an optical cross-connect (OXC) can also be derived from the data in [4]: With an MTBF of 1×10^5 hours and an MTTR of 6 hours we have an availability of 0.99994 per OXC that is used in further calculations.

B. Fiber Deployment and Leasing Costs

As we are interested in the relationship between availability and cost in this paper, we first consider the cost of fiber deployment and leasing.

The cost for deploying buried fiber ranges between US$ 10 000 and 100 000 per kilometer [24], [25], [26], [27] being smaller in rural areas than in cities [26]. The cost depends on the type of ground, e.g. deploying in hard rock is more expensive than in sand ground [24].

Compared to this, aerial fiber deployment cost is only between US$ 2 000 and 10 000 per kilometer [24], [25], [26]. However, the maintenance cost for the areal fiber will be much higher [24]. In rural areas the cost of setting up towers could be about 30-40% higher than in urban areas, as there the towers need to be ground-based and consume more material [24].

The cost for fiber leasing varies strongly between different providers as data from the Web indicate. Therefore, we assume in our calculations that the leasing of one km of fiber for one month is 0.1% of the deployment cost. [1]

These fiber cost value build the basis for constructing our realistic cost model in the next section.

IV. Solving of the Problem

This section provides a model for the relation between fiber MTBF and cost, followed by an explanation how our algorithm together with the cost model is able to identify optimum MTBF values for a deployment.

A. Modeling the Relation between fiber MTBF and Cost

For the further analysis of our problem, the relation between the MTBF and availability of the underlying physical infrastructure and the involved cost has to be modeled in form of a cost function. This cost function will then serve as an input for the algorithm.

1) Basic relation between MTBF and cost: Identifying the dependency between fiber-link MTBF and the associated cost is challenging. But also in this case the typical relationship (well-known from the field of micro-economics) between the value of an output result and the amount of an input factor applied for its increase can be observed. With increasing deployment, each additional unit of the input factor will only lead to a smaller increase of the output than the deployment of the previous unit. That is, the benefit of each additional input factor unit is diminishing - more and more units of the input factor have to be spent to achieve a certain additional rise of the output. A classical example is the use of fertilizer to increase the amount of crop that can be harvested.

[1]Note that during deployment, normally multiple strands of fiber are deployed, thus the deployment costs can be shared among these strands. This is also reflected in current leasing fiber prices found on the Web.

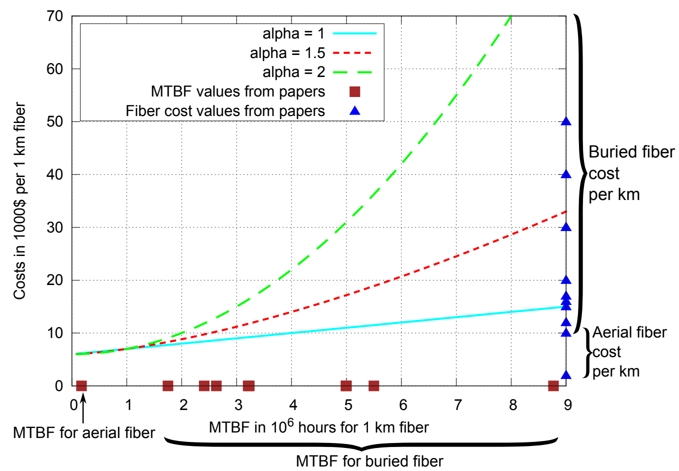

Fig. 1. Different cost models: $y = x^\alpha + 6$

Accordingly, for each additional increase of the reliability of a component (the 'output') more and more effort and cost have to be spent (the 'input'). We therefore use an exponential behavior to model the cost of a fiber link depending on its MTBF.

$$y = x^\alpha + \beta \tag{3}$$

Here, x is the MTBF value, y is the cost, α is a scaling parameter and β is used to adjust the cost curve. The scaling parameter α reflects the relative growth of the cost according to increasing the MTBF. The β value marks the starting cost, since even to deploy fiber with the lowest MTBF value, a certain amount of money needs to be spend.

As to the best of our knowledge, no such model exist up to now, we use the values for fiber deployment costs from Section III-B and the MTBF for different fiber types from Table I to calibrate the model. Combing these, we can adjust our curves like in Figure 1 that it fits the MTBF values to realistic cost per kilometer of deployed fiber. We consider three cases: first we assume the cost rises linear, second quadratic and the last is steeply with the higher MTBF values.

2) Cost function: In our cost function, the optimization objective of our embedding is to minimize the overall link cost for the embedding. The overall link cost is defined as the sum of the physical links that are used for the embedding of all virtual links in the virtual network request. Each physical link has a cost which depends on the length of the physical link of the underlying infrastructure and its MTBF value.

Objective:

$$\text{minimize} \sum_{e_v \in E_V} \sum_{e_s \in E_S} cost(e_s) x_{e_v e_s} \tag{4}$$

Bandwidth Constraints:

$$\sum_{e_v \in E_V} BW(e_v) x_{e_v e_s} \leq BW^R(e_s), \forall e_s \in E_S \tag{5}$$

$$x_{e_v e_s} \in \{0, 1\}, \forall e_v \in E_V, \forall e_s \in E_S \tag{6}$$

Path Availability Constraints:

$$A(e_v) \leq \prod_{e_s \in E_S} A(e_s) x_{e_v e_s} \prod_{v_s \in V_P} A(v_s), \forall e_v \in E_V \quad (7)$$

e_v is a virtual link of the virtual network request. e_s is a physical link of the physical network. v_s is a physical node of the physical network. $x_{e_v e_s}$ is a binary variable as indicated by (6) denoting, whether physical link e_s is part of the mapping of virtual link e_v: 1 if true, otherwise 0.

The physical link cost $cost(e_s)$ is defined as: $cost(e_s) := MTBFcost/1000 \times distance(e_s)$. The $MTBFcost$ is the cost for deploying one kilometer of fiber with the selected MTBF value using the selected cost models. For the embedding of a virtual network request only a fraction of the fiber is needed, i.e. virtual network embedding is like leasing fiber. $distance(e_s)$ is the distance (in km) of the e_s from its start node to the end node. Our assumption is that the relation of distance-to-cost is linear, because the VNO does not multiplex on its own (would also need to rent the multiplexing equipment), instead only the wavelength is assigned and the distance has to be considered for each wavelength usage on the link individually.

Equation (5) represents the bandwidth constraint that the total bandwidth $BW(e_v)$ of all the virtual links on the physical link e_s is limited by its bandwidth constraint $BW^R(e_s)$. Equation (7) represents the path availability constraint. It calculates the path availability out of the availabilities $A(e_s)$ of physical links and $A(v_s)$ of the physical nodes along the path and ensures that the path for the virtual link e_v has equal or higher availability than the requested link availability $A(e_v)$ of e_v.

After defining our cost models and cost function, we need an embedding algorithm to solve our object function which uses our cost models as input to map the VNR to the physical network.

B. The Heuristic as Tool to solve the Problem

For solving the problem of getting lowest embedding cost while examining the trade-off between high cost direct paths (high MTBF) and a primary path with backup path(s) to achieve the desired link availability, we extent the algorithm of our previous work [19].

The original algorithm embeds virtual networks with path protection in a bandwidth efficient way while achieving the requested link availability. As the physical network links cannot always provide the requested availability, several independent parallel links or paths are combined to achieve the availability. The idea of the algorithm is to calculate the primary paths and if needed one to several backup paths which together have the requested availability.

The main method of the algorithm stays the same: the node mapping is based on the geographic constraints and node constraints. Note that we assume virtual nodes to be in different geographical locations and that they must not be mapped to the same physical node, i.e. each virtual node is mapped to a separate physical node. The mapping of virtual links to physical paths is first determined by a graph search algorithm. The End-to-End (E2E) paths are calculated using a Constrained-based Shortest Path (CSPF) algorithm on bandwidth. Afterwards the link availability constraint is checked for the paths. For each virtual link we compute k candidate paths. After checking the availability constraint on the candidate paths and calculating required backup path(s), the most cost efficient combination of these candidates for the complete VNR will be calculated using Integer Linear Programming (ILP). With increasing k, the bandwidth consumption/embedding cost is getting closer to the optimal value due to selecting the best paths out of more possible candidates. In our previous work, we showed that a value of 10 for k can achieve close to an optimum cost. Therefore, we chose $k = 10$ in our following simulation.

We modify the algorithm from [19] with the following extensions. First for the physical network the link availability is calculated for all links and the MTBF-to-cost model is applied to the links (link costs) which are needed as input for the algorithm. The link availability of the physical link is calculated with the Equation (1) and (2) using the distance between source and destination node. For the E2E path availability calculation the node availability is considered using the OXC availability from section III-A. For a path consisting of x physical links and z physical nodes, the availability is calculated as $A_{path} = \prod_{i=1}^{x} LA_i \prod_{j=1}^{z} NA_j$, where LA_i is the link availability of the link i and NA_j is the node availability of node j. We consider that every physical node in the network has the same node availability value (OXC availability) of 0.99994. The link mapping is modified by finding fully-disjoint backup path if the primary path availability is lower than the requested link availability of e_v. After finding candidate path pairs (primary and backup path(s)) for each virtual link in the VNR, with the ILP the suitable candidate for each connection using our cost function from above IV-A2 is selected and resulting in the minimal cost for each embedding.

V. PARAMETER STUDY OF THE AVAILABILITY PROBLEM

In this section, we study the influence of different parameters of the availability problem on the embedding cost. One of the main questions is: What is the influence on the optimum MTBF for certain scenarios?

A. Simulation Setup

Our framework for analyzing the trade-off between a high-cost, high-availability infrastructure and a lower-cost, redundant infrastructure is custom-built and written in Java.

We simulate the arrival of VNRs (Virtual Network Requests) as discrete events. The input parameters for our algorithm are the MTBF values between 0.001 and 1000×10^6 hours and three different α values. As we know from Section III-A, current realistic values for MTBF are between 0.1 and 9×10^6 hours. However, we want also examine the effects of very small and very large MTBF values to check if there exists any abnormality and if using these values could result in lower

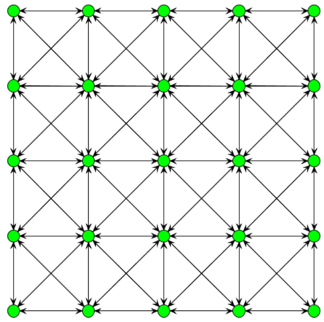

Fig. 2. Example grid network with 25 nodes

physical deployment costs for the operators. The MTTR value is constant during the whole simulations and has a value of 12 hours. The α value for the cost model is chosen between 1 (linear growth) and 2 (quadratic growth).

For each MTBF and α value the simulation is run 100 times and the average embedding cost is calculated. Different physical and virtual networks are created for each simulation run. The VNR is a connected graph with 5 nodes. For the VNRs, different values for the requested link availability are examined. The values are between 0.999 and 0.999999.

Different physical topologies are examined: a grid network structure like in Figure 2 and two real-world networks of different size. The area across which the nodes are distributed in the grid network is varied in size from ten km to several thousand km. We assume that each physical node and link has sufficient capacity.

B. Influence of Different Parameters

For these physical topologies, we examine the influence of different parameters like different cost scaling factors (α values), different extensions and different requested link availabilities by running simulations with our algorithm. From the simulations we determine the total lowest cost for embedding and the corresponding MTBF value and how this cost minimum changes with the MTBF. The resulting MTBF and cost values are mean values. These values are then compared in the following sections and the results are described.

The physical network we use as basis for our simulation is a 5 × 5 grid network as in Figure 2. We use the grid with a high average nodal degree to achieve a high acceptance ratio for the embedding and to examine the general behavior at the parameters.

1) Influence of cost parameter α: First the general behavior of the resulting curve is examined. We run the algorithm with the different scaling factors (α values) of the cost model each with the different MTBF values and always plot the lowest embedding cost. Intuitively, we would assume a growth of cost in relation to the cost model Figure 1 and that minimum embedding cost is reached at the lowest MTBF value. An example result curve can be seen in Figure 3.

The cost in relation to the MTBF value is shown for the embedding using different α values. The acceptance ratio of the embedding is also plotted in the picture. The curve can be divided in two parts: the first part shows a decrease in the costs until the minimum. The second part is the increase of the costs. Between these two parts, there is a turning point which has the lowest cost for the embedding. With increasing MTBF, the curves first show a decrease of the embedding cost until a minimum is reached. The slope of the rise beyond this turning point is strongly reflecting the parameter α. For α = 1 (linear increase of the cost) the increase is much slower than for α = 2. If the physical network has very low availability (MTBF values close to zero) low acceptance ratios (i.e. the percentage of successful embedding) of the embedding are the consequence. Here, embedding cost often are very low - however, this only shows that the less complex VNRs were embedded which themselves lead to low cost. Thus, no real conclusion can be drawn at these low acceptance ratios. Figure 4 shows the zoomed in on Figure 3 where it can be seen that MTBF values lower than 0.07×10^6 hours cannot achieve any successful embedding. Even at 0.07×10^6 there is a local minimum of the cost which is not utilizable due to low acceptance ratio.

The observed turning points are generated by the number of physical paths for embedding a virtual link: For MTBF values towards zero, several backup paths are needed to achieve a successful embedding for the requested virtual link availability which results in high cost. This can be seen in Figure 5, where up to five physical paths are needed for embedding one virtual link. For higher MTBF values less backup paths are required and the cost decrease. After the turning point, the cost is dominated by the MTBF and rise again. In the figure, we find the turning point at a place where the MTBF values are large enough to allow the usage of only two paths (primary and backup path). We also see that while the resulting embedding cost depend on α this is not the case for the embedding decision itself (selection of the individual physical paths): For different values of α, the VNRs are embedded in the same way (same paths are selected). Figure 6 shows the ratio of virtual network links that can be embedded with only a primary path. We see that until a value of about 8.5×10^6 always a backup path is needed. Compared to Figure 5, we can recognize the same behavior of the ratio of using primary path in relation to the MTBF values.

In this section, we have seen that the minimum cost is not achieved at the lowest MTBF value, but rather that there is a turning point delivering the minimum. In the next section we examine some of the other influence factors. The global minimum (turning point) is determined by simulations by changing some of the parameters. We will see that the position of the turning point depends on the topology, size and network extension of the network and the requested link availability values.

2) Different network size extension: The influence of physical network extension is investigated. The physical network extensions are between a square area of 10 × 10 kilometers up to 4000 × 4000 kilometers. For each extension the minimum cost with its related MTBF value is extracted from the simulation results. For the example of Figure 3 and 4, this is a MTBF

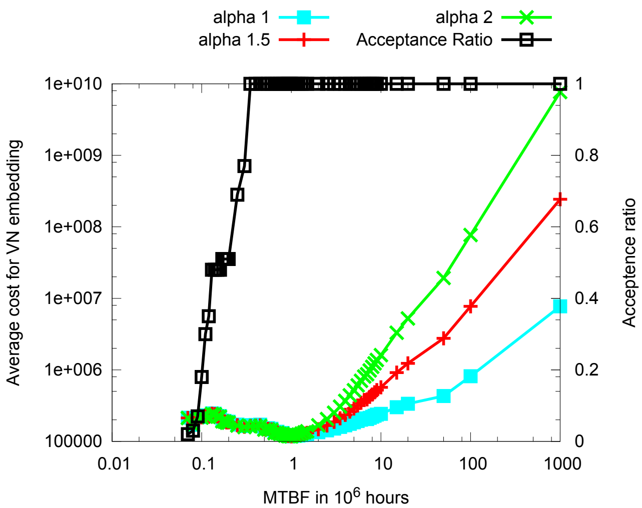

Fig. 3. Example: Results of the embedding cost using different α values, 25 nodes, area 2000km $\times$ 2000km, VNR with 5 nodes and availability 0.999

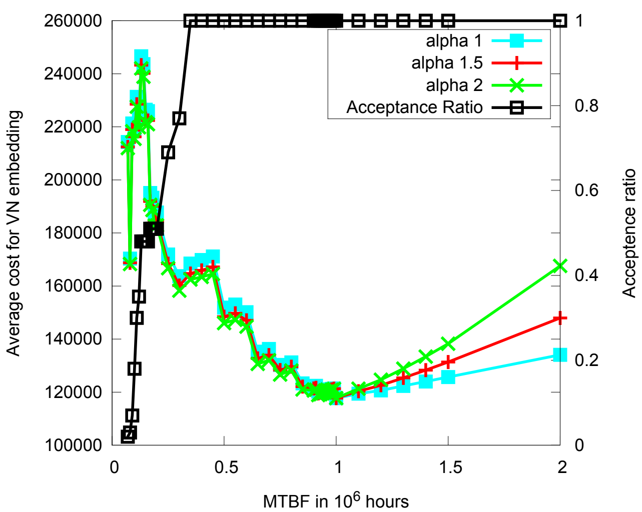

Fig. 4. Results of the embedding cost using different α values (zoomed in on Figure 3)

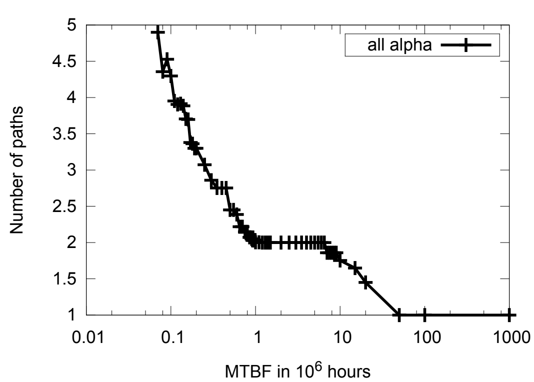

Fig. 5. Number of (parallel) paths needed for successfully embedding of one virtual link

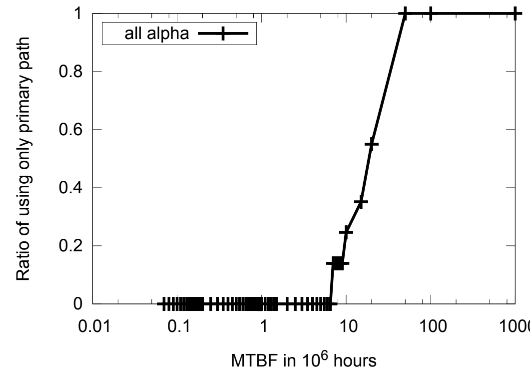

Fig. 6. Ratio of using only primary path

value of 1.0×10^6. The effect of different physical network extensions is shown in Figure 7 for a physical network of 25 nodes and virtual networks of 5 nodes.

The result curves show following behavior: For very small network extensions, a low MTBF is sufficient to fulfill the required availability. Here, only a primary path is needed and the MTBF value with lowest cost increases until it is not anymore economic enough. After that region there is a drop (depending on the α value at different extension values). This is the point from which onwards it is not economic to use only a primary path since the costs increase strongly for larger MTBF value. After this point, it is cheaper to have backup paths than a single primary path with a high MTBF value. For the different values some different behavior can be seen (especially for α = 1, linear), because in the linear cost model, the cost increase much slower and single paths with higher MTBF values can achieve lower cost than several backup paths. We can observe that this behavior is repeating for larger network extension.

Beside the dependency on the extension and on the α value, there can be seen a relation also to the number of nodes in the physical network in Figure 8 which has an influence in the region of 50 and 300 km length of the square. Compared to the results in Figure 7, the MTBF values with lowest cost only goes up until 3.5×10^6 and a slightly different behavior is seen for the α value.

3) Different requested link availabilities: Now the effect of different requested link availabilities is investigated. The other parameters are kept the same and again the MTBF value with the minimum cost is derived from the simulation results. The effect of higher requested link availabilities is shown in Figure 9. The requested link availability values are 0.999, 0.9993, 0.9996, 0.9999, 0.99999 and 0.999999.

For higher link availability values (larger than 0.999) the MTBF values with lowest embedding cost increase. At a certain point, higher MTBF values result in higher costs than an additional backup path. Therefore, we detect a drop at high requested availability values. This can be seen in Figure 9 for a value of 0.99999 where there is a drop and for the embedding an additional backup path is used. All α values show nearly

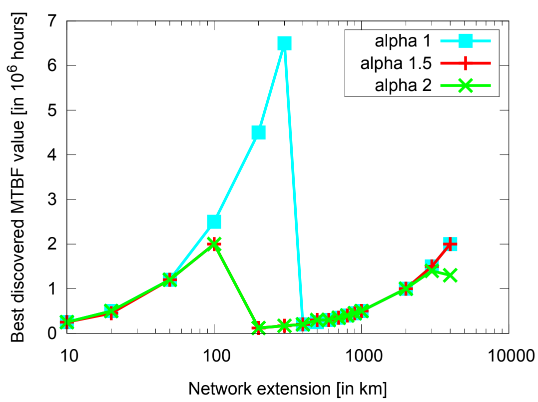

Fig. 7. Result for different physical network extensions for 5 × 5 grid network (25 nodes) and VN with 5 nodes and a requested availability of 0.999

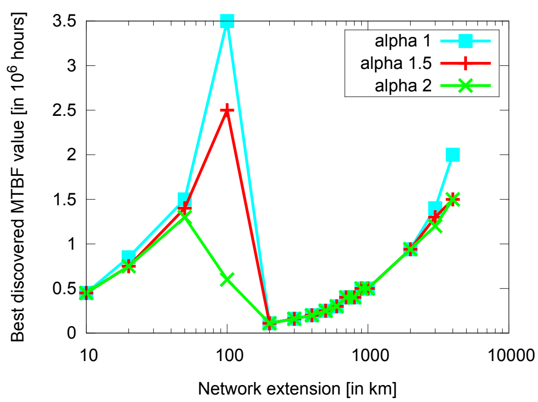

Fig. 8. Result for different physical network extensions for 10 × 10 grid network (100 nodes) and VN with 5 nodes and a requested availability of 0.999

the same results, due to the reason of the same embedding behavior and the similar cost in the resulting MTBF range. Also a larger physical network with 100 nodes (see Figure 10) shows no significant change in the behavior compared to Figure 9, because the embedding is done similarly and identical numbers of physical paths are required for embedding the virtual links.

As a summary, we see that for low requested virtual link availabilities it may be cheaper to use a physical network with high MTBF values than to use backup paths. However, if higher virtual link availability values are requested, they can only be achieved using backup paths and even with higher MTBF values.

C. Real-world Network Topologies

In this section we examine the influence of the requested link availability on the basis of two different real-world networks from Germany and North America which are publicly available network topologies from the Internet Topology Zoo [28]. In these real-world networks, the cities are mapped to the nodes and their interconnections to the links of the graph.

1) German Network: The German network, denoted DFN, is the national IP-based research backbone network in Germany. It consists of 47 nodes and 72 fiber links which results in an average nodal degree of 3, see Figure 11. The average length of a link is 116 km.

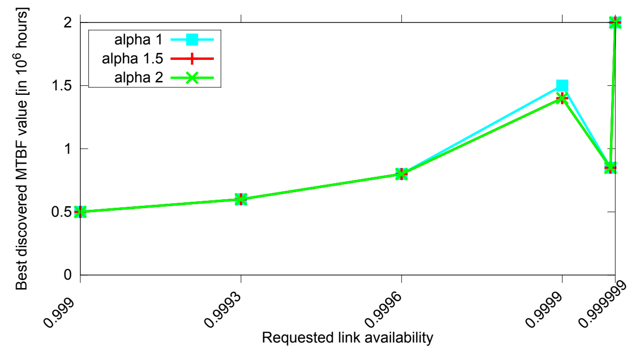

Fig. 9. Result for different requested link availabilities using a 5 × 5 physical grid network (25 nodes) of extension 1000km × 1000km

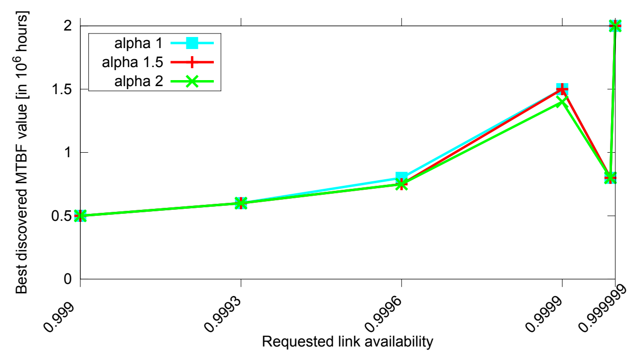

Fig. 10. Result for different requested link availabilities using a 10 × 10 physical grid network (100 nodes) of extension 1000km × 1000km

Figure 12 shows the behavior for requested link availabilities in the range between 0.999 and 0.999999. First, the MTBF value providing lowest cost is small and slowly increasing before rising steeply. The requests can be satisfied with one primary path and a backup path with higher availabilities requiring higher MTBF values. For very high link availability values (0.999999), we again find a drop resulting from the fact that additional backup paths allow lower MTBF values (resulting to three paths on average). The acceptance ratio of the embedding is getting lower as it is getting harder to identify three paths to satisfy the required embedding constraints. Therefore, we only find a quite low acceptance ratio of about 20% trying to achieve availability values of 0.999999 and identifying MTBF value with the lowest. High acceptance ratios (above 60%) are only reached at MTBF values higher than 10×10^6 hours when embedding can be done with one primary and one backup path, which however lead to very high cost.

Using the more complex network of Germany instead of the regular grid in the previous section, we see that the minimum cost can be achieved with low MTBF values (between 0.25 and 1×10^6 hours) already. However, increasing availability requirements enforces an investment in the physical infrastructure resulting in higher MTBF values.

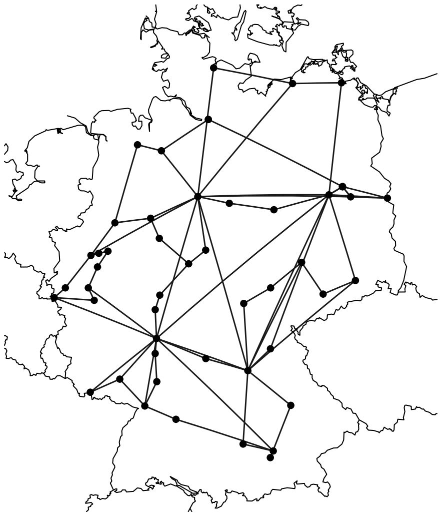

Fig. 11. Topology of the German network

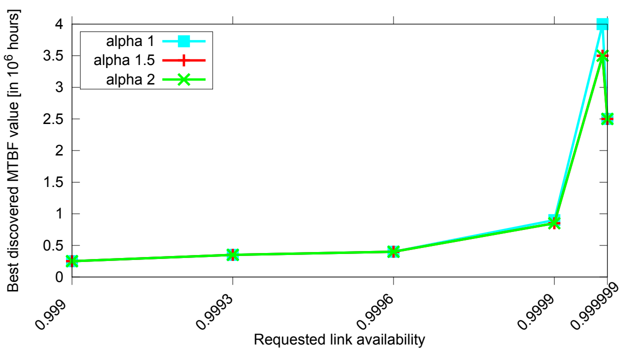

Fig. 12. Result for different requested link availabilities with German network and VNs with 5 nodes

2) North American Network: The North American network, (mostly USA and southern part of Canada) is an IP backbone network. It consists of 42 nodes and 77 links which results in an average nodal degree of 3.66, see Figure 13. The average length of a link is 966 km.

Figure 14 shows the behavior for different requested virtual link availabilities between 0.999 and 0.999999. In contrast to the previous figures, we see an interesting different behavior (no drop in the curve) because of the bigger extension of the network and the large length of the physical links. First, the cost-minimal MTBF values are small and slowly increasing before they rise steeply. The requests can be satisfied with one primary path and a backup path with higher availabilities requiring higher MTBF values. However, for availability requests of 0.99999 and 0.999999 the acceptance ratio of the embedding is getting lower as it is getting difficult to identify two paths to satisfy the required embedding constraint. Very high MTBF values are the consequence.

VI. Recommendations for Network Operators

This section provides some recommendations for operators based on our simulations.

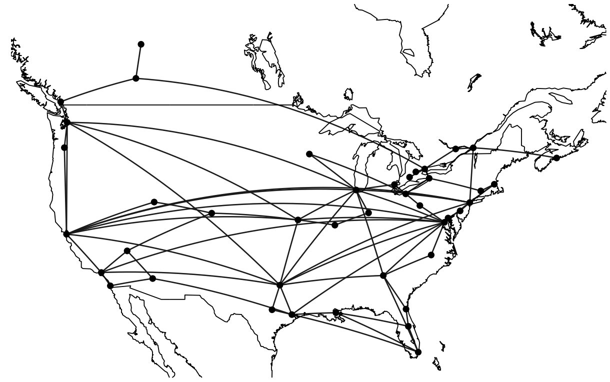

Fig. 13. Topology of the North American network

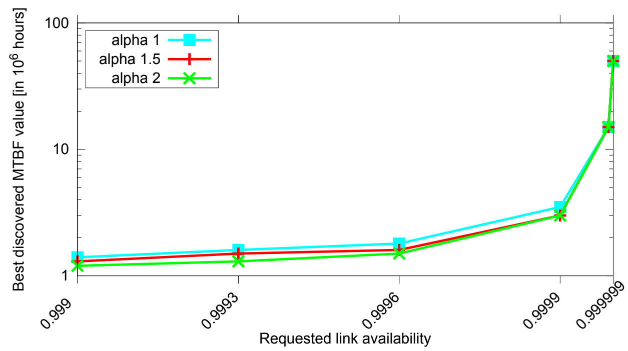

Fig. 14. Result for different requested link availabilities with North American network and VNs with 5 nodes

From the results we can see that the cost-minimum MTBF value depends on the structure of the network as well as its geographical extension. The cost function and especially the scaling factor α play an important role. The lowest-cost MTBF value was typically in the range of 0.1×10^6 to 6×10^6 hours which corresponds to MTBF values found in real fiber networks. MTBF values higher than this result in most cases in enormous costs.

The structure of the results shows that for an operator the turning points of the curves are of most interest - where the cost is the lowest. Whether an investment into higher MTBF values pays off mostly depends on network extension and the requested availability values. Since the increase of the availability for increasing MTBF values is logarithmic, e.g. high MTBF values like 9×10^6 hours can only achieve a 4 nines availability for 1 km fiber length, a lot of money has to be invested to achieve this. In contrast, combining two disjoint paths with low availability (e.g. two parallel paths with MTBF of less than 0.5×10^6 hours achieve already 4 nines for 1 km fiber length) already results in a high path availability. Therefore the turning point can be found at the least MTBF value with the least number of parallel paths. This is for most simulations in the lower range – 0.5 to 2×10^6 hours – for the MTBF values.

Depending on the requested service availabilities, the low cost physical infrastructure approach can be thus cheaper.

However, very high availability values in the region of 6 nines and a moderate acceptance ratio of the embedding can only be achieved by investing in the MTBF of the physical network to get a reliable service. If the physical network of the operators has a low average nodal degree (e.g average nodal degree of 2), in most cases this does not allow to find disjoint paths for a link. Thus it is difficult or even impossible to use the low cost physical infrastructure approach with several parallel paths here.

VII. Conclusion

In this paper, we have examined the trade-off between cost and availability when realizing virtual networks. One end of the design space is marked by the 'low cost physical' approach, where as little money as possible is spent on physical protection. Instead, high availability is realized by combining multiple parallel paths to form one virtual path or link. The other end of the design space can be described as a 'high cost physical' approach. Here, enough money is spent on the physical network to allow already single paths to achieve a requested availability level.

To examine the underlying trade-off between these two philosophies, we have defined a model that sets the network deployment cost in relation to the achieved resiliency. We have determined realistic parameters for this model in an extensive literature study, and then created a cost function to determine the overall cost when realizing a virtual network with a requested availability on a physical network with a different availability. With our algorithm for virtual network embedding, we are then able to determine the minimum embedding cost for such an embedding.

With these instruments, we have examined a number of different networks – both artificial grid topologies and real-world, existing country- and continent-wide networks. The results show that for most configurations, the 'low cost physical' approach with low or medium reliability levels in the physical network results in the lowest cost. This is especially interesting as these reliability levels fit very well to parameters from real fiber deployments: already the lowest availability values for buried fiber found in the literature are sufficient. Therefore, it seems advisable to realize availability in the virtual domain rather than in the physical domain in real networks.

As a next step we will also investigate in more detail the influence of the network nodes by modeling their cost and availabilities.

References

[1] S. Baucke and C. Görg, "D-3.1. 1 Virtualisation Approach: Concept," FP7 4WARD Project, Tech. Rep., 2009.

[2] W. Zou, M. Janic, R. Kooij, and F. Kuipers, "On the availability of networks," *BroadBand Europe*, pp. 3–6, 2007.

[3] G. Semaan, "Designing Networks with the Optimal Availability," in *Optical Fiber communication/National Fiber Optic Engineers Conference (OFC/NFOEC 2008)*, 2008.

[4] S. Verbrugge, D. Colle, P. Demeester, R. Huelsermann, and M. Jaeger, "General availability model for multilayer transport networks," in *5th International Workshop on Design of Reliable Communication Networks (DRCN 2005)*, 2005.

[5] M. Tornatore, G. Maier, and A. Pattavina, "Availability Design of Optical Transport Networks," *IEEE Journal on Selected Areas in Communications*, vol. 23, no. 8, pp. 1520–1532, 2005.

[6] J. Segovia, E. Calle, P. Vila, J. Marzo, and J. Tapolcai, "Topology-focused availability analysis of basic protection schemes in optical transport networks," *Journal of Optical Networking*, vol. 7, no. 4, 2008.

[7] D. A. A. Mello, D. Schupke, M. Scheffel, and H. Waldman, "Availability maps for connections in WDM optical networks," in *5th International Workshop on Design of Reliable Communication Networks (DRCN 2005)*, 2005.

[8] M. Clouqueur and W. Grover, "Availability analysis of span-restorable mesh networks," *IEEE Journal on Selected Areas in Communications*, vol. 20, no. 4, pp. 810–821, 2002.

[9] M. Yu, Y. Yi, J. Rexford, and M. Chiang, "Rethinking virtual network embedding: substrate support for path splitting and migration," *SIGCOMM Comput. Commun. Rev.*, vol. 38, pp. 17–29, 2008.

[10] Y. Zhu and M. Ammar, "Algorithms for Assigning Substrate Network Resources to Virtual Network Components," in *IEEE INFOCOM*, 2006.

[11] J. Lu and J. Turner, "Efficient Mapping of Virtual Networks onto a Shared Substrate," Washington University in St. Louis, Tech. Rep., 2006. [Online]. Available: http://www.arl.wustl.edu/˜jl1/research/tech_report_2006.pdf

[12] J. Lischka and H. Karl, "A virtual network mapping algorithm based on subgraph isomorphism detection," in *1st ACM workshop VISA*, 2009.

[13] T. Guo, N. Wang, K. Moessner, and R. Tafazolli, "Shared Backup Network Provision for Virtual Network Embedding," in *IEEE International Conference on Communications (ICC 2011)*, 2011.

[14] H. Yu, V. Anand, C. Qiao, and H. Di, "Migration based protection for virtual infrastructure survivability for link failure," in *Optical Fiber communication/National Fiber Optic Engineers Conference (OFC/NFOEC 2011)*, 2011.

[15] J. Shamsi and M. Brockmeyer, "QoSMap: Achieving Quality and Resilience through Overlay Construction," in *4th International Conference on Internet and Web Applications and Services (ICIW 2009)*, 2009.

[16] M. R. Rahman, I. Aib, and R. Boutaba, "Survivable Virtual Network Embedding," in *IFIP/TC6 NETWORKING*, 2010.

[17] H. Kim and S. Lee, "Greedy virtual network embedding under an exponential cost function," in *International Conference on Information Networking (ICOIN 2012)*, 2012.

[18] W. Grover and A. Sack, "High availability survivable networks: When is reducing MTTR better than adding protection capacity?" in *6th International Workshop on Design and Reliable Communication Networks (DRCN 2007)*, 2007.

[19] S. Herker, X. An, W. Kiess, and A. Kirstädter, "Path Protection with Explicit Availability Constraints for Virtual Network Embedding," in *24th Annual IEEE International Symposium on Personal, Indoor, and Mobile Radio Communications (PIMRC 2013)*, London, UK, 2013.

[20] M. L. Shooman, *Reliability of Computer Systems and Networks: Fault Tolerance, Analysis, and Design*. John Wiley & Sons, Inc., 2002.

[21] I. Heard, "Availability and cost estimation of secured FTTH architectures," in *International Conference on Optical Network Design and Modeling (ONDM 2008)*, 2008.

[22] R. Huelsermann, M. Jaeger, A. M. C. A. Koster, S. Orlowski, R. Wessaely, and A. Zymolka, "Availability and Cost Based Evaluation of Demand-wise Shared Protection," in *ITG Symposium on Photonic Networks*, 2006.

[23] D. A. Schupke, A. Autenrieth, and T. Fischer, "Survivability of Multiple Fiber Duct Failures," in *3rd International Workshop on the Design of Reliable Communication Networks (DRCN 2001)*, 2001.

[24] T. Cohen and R. Southwood, "Extending open access to national fibre backbones in developing countries," in *8th Global Symposium for Regulators ITU*, 2008.

[25] G. Taylor, P. Ray, and Memo, "LWA Data Communications - The Fiber Option," LWA, Tech. Rep., 2008.

[26] J. Ellershaw, J. Riding, A. Lee, A. V. Tran, L. J. Guan, R. Tucker, T. Smith, and E. Stumpf, "Deployment costs of rural broadband technologies," *Telecommunications Journal of Australia*, vol. 59, no. 2, 2009.

[27] G. McGuire, "Enhanced cost solutions for buried fiber installation," Tech. Rep., 2009.

[28] S. Knight, H. Nguyen, N. Falkner, R. Bowden, and M. Roughan, "The Internet Topology Zoo," *IEEE Journal on Selected Areas in Communications*, vol. 29, no. 9, 2011. [Online]. Available: http://www.topology-zoo.org/