

Survivable optical grid dimensioning: anycast routing with server and network failure protection

Chris Develder*, Jens Buysse*, Ali Shaikh†, Brigitte Jaumard†, Marc De Leenheer*, and Bart Dhoedt*

* Ghent University – IBBT, Dept. of Information Technology – IBCN, Ghent, Belgium

† CSE Dpt, Concordia University, Montreal (Qc) Canada

Email: chris.develder@intec.ugent.be, bjaumard@ciise.concordia.ca

Abstract—Grids can efficiently deal with challenging computational and data processing tasks which cutting edge science is generating today. So-called e-Science grids cope with these complex tasks by deploying geographically distributed server infrastructure, interconnected by high speed networks. The latter benefit from optical technology, offering low latencies and high bandwidths, thus giving rise to so-called optical grids or lambda grids.

In this paper, we address the dimensioning problem of such grids: how to decide how much server infrastructure to deploy, at which locations in a given topology, the amount of network capacity to provide and which routes to follow along them. Compared to earlier work, we propose an integrated solution solving these questions in an integrated way, i.e., we jointly optimize network and server capacity, and incorporate resiliency against both network and server failures. Assuming we are given the amount of resource reservation requests arriving at each network node (where a resource reservation implies to reserve both processing capacity at a server site, and a network connection towards it), we solve the problem of first choosing a predetermined number of server locations to use, and subsequently determine the routes to follow while minimizing resource requirements. In a case study on a meshed European network comprising 28 nodes and 41 links, we show that compared to classical (i.e. without relocation) shared path protection against link failures only, we can offer resiliency against both single link and network failures by adding about 55% extra server capacity, and 26% extra wavelengths.

I. INTRODUCTION

Originating from so-called eScience applications (stemming from various domains, such as astrophysics, climate modeling, and particle physics), Grids were envisioned: heterogeneous resources (computational, storage and networking) are geographically spread (possibly over various administrative domains, implying that resource coordination is not subject to centralized control) to jointly provide the required computational and storage capabilities. Similar ideas are applied in cloud computing and virtualisation. These technologies make network dimensioning a complex problem, especially for providers needing to plan and deploy both network and IT resources (i.e., servers, both for computing and storage). In particular, since users typically do no longer care where their workload is processed (“in the cloud”), freedom arises as to where to install e.g., server farms. Thus, a (source,destination)-based traffic matrix, as assumed in traditional (optical) network dimensioning problems, including many routing and wavelength assignment (RWA) approaches, is not a priori available.

In the current work, we assume the network interconnecting

the Grid server sites to be optical circuit-switched (such as an ASON), based on Wavelength Division Multiplexing (WDM). To deal with potential network failures, various network resiliency strategies for WDM networks have been devised (for an extensive overview, see [1], [2]). A well-known classical shared path protection scheme protects against single link failures: a primary path from source to destination is protected by a link-disjoint backup path which is used in case of a failing link (this link diversity guarantees that the primary and backup paths will never fail simultaneously for any single link failure). In a grid-like scenario however, we proposed the idea of exploiting relocation [3], which is possible due to the anycast routing principle. Since a user generally does not care about the exact location where his workload is being processed, it could be better to relocate the job to another resource (different from the one chosen under failure-free conditions).

In this paper, we expand on the relocation idea to judge the resource requirements to also cater for server site failures. Without considering relocating jobs to other locations, providing resiliency against server site failures would imply doubling server capacity (if no service degradation is allowed) at each location. However, by relocating to other sites (as in [4] for link failure protection) we may be able to reduce the amount of overall backup server capacity. In this paper, we assess this reduction quantitatively by dimensioning the network to survive both single link and server site failures. As a matter of fact, the dimensioning model presented further is generic and can address any failure scenario that can be expressed as a set of jointly failing link resources, i.e., a so-called shared risk link group (SRLG; e.g., to model failures such as fibre duct cuts [5]). Note that node failure can be represented as joint failure of its incident links (hence leading to an SRLG comprising all of them).

We here consider individual connections between a source site generating tasks to execute, and a server site processing them. Thus, we do not consider provisioning virtual networks interconnecting multiple sites. For an overview of such work, see [6].

The remainder of this paper addresses the offline grid dimensioning problem, first stated in [7] for the case without resiliency against failures. Whereas there we proposed a phased approach, we now (i) solve the sub-problems of establishing server and network capacity in an integrated way, and (ii) additionally provide resiliency against both network

and server failures. To this end, we propose an integer linear programming (ILP) formulation using a column generation approach, similar to [4]. Compared to the latter work, we now (i) consider not just network but also server site failures to protect against, (ii) simultaneously optimize the amount of server resources (instead of just network), and (iii) do not fix the destination server site (under failure free conditions) a priori.

The remainder of this paper is structured as follows. In Section II, we outline the problem statement addressed in this paper. Our solution method and the associated ILP models are detailed in Section III. Subsequently, we discuss a quantitative case study in Section IV. Section V summarizes our conclusions and future work.

II. DIMENSIONING RESILIENT OPTICAL GRIDS

The dimensioning problem we address is the following:

Given

- *Topology* comprising the sites where jobs originate, as well as the optical network interconnecting them;
- *Demand* stating the amount of jobs arriving at each of the sites; and
- *Survivability requirements* specifying the failures that should be protected against,

Find

- K *server locations*, chosen out of the given grid sites, where server infrastructure should be provided;
- *Destination sites and routes* to follow for all grid traffic, originating with given intensity at the various grid sites (where each destination should be one of the K server locations);
- *Network and server capacity* to provide on each of the links and server sites;

Such that the latter resource capacity (server and network resources) is minimized.

To achieve the latter objective, we will allow sharing capacity (wavelengths, servers) for the backup of connections whose resources under working conditions are disjoint. In particular, we will adopt a shared path protection concept. Similarly, at each server site, we will install the minimum capacity required to cope with each one of the considered failure scenarios (as well as the failure-free case, obviously). Thus, we will allow reclaiming of server and network resources for backup purposes, if they are no longer used as primary under failure conditions.

We will express the *survivability requirement* through the concept of a shared risk link group (SRLG): a set of resources (links) that may fail jointly, because of shared dependencies (e.g. fibre ducts [5]). Thus, to protect against failure of an SRLG, the backup resources should not include any of the SRLG elements. In our case studies, we will protect against single server site or single network link failures (whereas in earlier work, we only considered network failures [3], [4]). Note that server site failures will be modeled as a failure of the access link towards the server resources.




Fig. 1. Input data: (i) network links and nodes (OXCs labeled A-D), (ii) source nodes with job arrival intensity (represented as users), and (iii) candidate server sites.

III. NETWORK MODEL AND ILP SOLUTIONS

Solving the stated dimensioning problem in a single integrated way is quite hard. We previously we developed an iterative approach, comprising four phases [7], [8]: (S1) Find the K best locations to install servers; (S2) Determine the total number of server CPUs and distribute them heuristically over the K locations; (S3) Determine the number of jobs sent to each site, thus finding a traffic matrix; and (S4) Use classical traffic-matrix based dimensioning techniques to find the required routes and network resources. (Note that the time consuming simulations in step S1 can be replaced with much faster analytical modeling techniques, as in [8].) In the current paper however, we will use a single integrated ILP-based solution of steps S2–S4. We will assume WDM networks and consider the following network model, illustrated in Fig. 1:

- $G = (V, L)$, directed graph representing an optical grid, where V is the node set and L is the set of (directed) links, where every link has the same unlimited transport capacity.
- $V = V_{\text{SRC}} \cup V_{\text{NET}} \cup V_{\text{DST}}$, the set of all nodes indexed by v , comprising pure OXCs (V_{NET}), grid server sites V_{DST} (with $|V_{\text{DST}}| = K$), and explicitly modeled sources V_{SRC} .
- $L = L_{\text{SRC}} \cup L_{\text{NET}} \cup L_{\text{DST}}$, the set comprising all directed network links, indexed by ℓ , again split into the core network links L_{NET} interconnecting OXCs, and the modeled access links L_{SRC} from job sources and those towards the server sites L_{DST} .
- Δ_v the number of unit demand requests, originating from a source node $v \in V_{\text{SRC}}$.
- $\text{SRGL}(\ell)$ The SRLG comprising all links that can fail simultaneously with $\ell \in L$.

We also will use the following notations:

- $\omega^+(v)$ represents the set of v 's incoming links.
- $\omega^-(v)$ represents the set of v 's outgoing links.

Note that we assume the server site locations to be given (through specification of V_{DST}), but we are still free to choose any of them to route the demands Δ_v to (as opposed to [4]). To find K server site locations prior to solving the dimensioning problem as discussed next, we use the ILP presented in [7].

To jointly solve steps S2–S4, we use a column generation (CG) approach to find so-called configurations and the number

of times to use them. A configuration will be associated with a particular source-site $v \in V_{\text{SRC}}$, and comprise a pair of working and backup paths, both originating from v and terminating in one of the server sites in V_{DST} (possibly different in case of relocation). This involves solving a so-called Restricted Master Problem (RMP) and a Pricing Problem (PP) iteratively. We will detail the constituent phases of the following scheme next:

- 1) Find a set of initial configurations C ;
- 2) Solve the LP relaxation of RMP, minimizing required network and server resources,
- 3) Solve the PP to try and find a new configuration c for a source node $v \in V_{\text{SRC}}$, with negative reduced cost. If successful, add c to the set of configurations C .
- 4) Repeat steps 2–3 until no new configurations (with negative reduced cost) can be found.
- 5) Finally solve the RMP as ILP, to find an integer solution.

In each iteration of Step 3, source nodes V_{SRC} are considered in a round-robin fashion. Step 2 is performed every time a new configuration was added in Step 3. (For the CG methodology, see also [4].)

A. Initial configurations

To find initial configurations, we used a heuristic inspired by [9], and detailed in Algorithm 1. We introduce the set of candidate server locations V_{LOC} . For the case without relocation, $V_{\text{LOC}} = V_{\text{DST}}$. Yet, for the case with relocation, we add a (virtual) node σ to the node set V of the graph, and introduce additional links $(v, \sigma), \forall v \in V_{\text{DST}}$ and set $V_{\text{LOC}} = \{\sigma\}$. Then, for each source site $v \in V_{\text{SRC}}$, we find initial configurations by finding the shortest pair of disjoint paths to each candidate server site in V_{LOC} . For this, we use the algorithm originally developed by Suurballe [10]. In a subsequent step, we find additional configurations by trying to find alternate backup paths that share links with other configuration's backup.

B. Restricted master problem (RMP)

The parameters of our column generation ILP are:

- c A configuration, defined for a given source node $v \in V_{\text{SRC}}$.
- C_v The set of configurations associated with a source node $v \in V_{\text{SRC}}$.
- $C = \bigcup_{v \in V_{\text{SRC}}} C_v$
- S The set of SRLGs, indexed by s .
- z_c Integer decision variable, counting the number of times configuration c is used.
- p_{cl}^W Binary, equaling 1 if and only if link ℓ is used on the *working* path in configuration c .
- p_{cl}^B Binary, equaling 1 if and only if link ℓ is used on the *backup* path in configuration c .
- w_ℓ Auxiliary integer variable, counting the number of wavelengths used on link ℓ .

The master problem will choose the appropriate amount of configurations, using decision variables z_c . The objective

Algorithm 1 Finding an initial set of configurations C

- 1: **for all** $v \in V_{\text{SRC}}$, and $s \in V_{\text{LOC}}$ **do**
 - 2: $c \leftarrow \text{DisjointPaths}(v, s)$
 - 3: Add c to C .
 - 4: **end for**
 - 5: **for** $c \in C$ **do**
 - 6: Construct a copy G' of the graph G .
 - 7: Remove links of working path of c from G' .
 - 8: **for all** $c' \in C$ with $c' \neq c$ **do**
 - 9: **if** working paths of c' and c are disjoint **then**
 - 10: In G' , set weights of backup path links of c' to 0
 - 11: **end if**
 - 12: **end for**
 - 13: Construct a new configuration c' .
 - 14: Set working path of c' to that of c .
 - 15: Set backup path of c' to shortest path in G' between source and backup destination of c .
 - 16: **if** backup path c' shorter than that of c **then**
 - 17: Add c' to C .
 - 18: **end if**
 - 19: **end for**
-

function is given in (1): we minimize the amount of network resources (wavelengths w_ℓ) and the amount of server resources, which we can conveniently count as the number of wavelengths towards server nodes. We introduce a factor α that expresses the cost ratio of the server capacity corresponding to a workload filling a single wavelength with data, compared to the cost of a single wavelength on a single link.

$$\min \left(\sum_{\ell \in L_{\text{NET}}} w_\ell + \alpha \cdot \sum_{\ell \in L_{\text{DST}}} w_\ell \right) \quad (1)$$

The first set of constraints (2) are obviously to meet the requested demands. Next, in constraints (3) we enforce the number of wavelengths to be sufficient to carry all selected configurations under failure-free conditions. For each considered failure case, represented as an SRLG $s \in S$, we have constraints (4), whose right hand side comprises as first term a summation covering all unaffected configurations and secondly the affected ones. Therefore, we define two auxiliary variables (whose values in this RMP are constants, depending on the configuration at hand):

- π_{cls}^W Binary, equaling 1 if the working path of configuration c crosses link ℓ , which remains unaffected by failure of SRLG s .
- π_{cls}^B Binary, equaling 1 if link ℓ is part of the backup path of configuration c , whose working path is affected by SRLG s (that is, $\text{workingPath}(c) \cap s \neq \emptyset$).

(Note that within (4), $\ell \notin s$, and therefore $\pi_{cls}^W = p_{cl}^W$.)

$$\sum_{c \in C_v} z_c \geq \Delta_v \quad \forall v \in V_{\text{SRC}} \quad (2)$$

$$w_\ell \geq \sum_{c \in C} p_{cl}^W \cdot z_c \quad \forall \ell \in L \quad (3)$$

$$w_\ell \geq \sum_{c \in C} \pi_{c\ell s}^W \cdot z_c + \sum_{c \in C} \pi_{c\ell s}^B \cdot z_c \quad \forall s \in S, \forall \ell \notin s \quad (4)$$

C. Pricing Problem (PP)

In order to answer the dimensioning question, the master problem should include all possible configurations to find the overall optimum. However, in the column generation approach, we start from an initial limited set of promising configurations: this is the Restricted Master Problem (RMP). Based on the solution of the RMP, we subsequently add new configurations c by solving the pricing problem (PP) if it finds such c that is able to reduce the RMP objective value. In our case, a PP is associated with a given source node $v_{\text{SRC}} \in V_{\text{SRC}}$. The PP uses the values (as found by the RMP, relaxed as linear program) of dual variables corresponding to constraints of the RMP:

- u_v^1 value of RMP dual variable corresponding to (2).
- u_ℓ^2 value of RMP dual variable corresponding to (3).
- $u_{\ell s}^3$ value of RMP dual variable corresponding to (4).

The objective function (5) of the PP will be to minimize the reduced cost. (The first explicit 0 term is the coefficient of z_c in the RMP objective.) The PP's decision variables p and its auxiliary variables π have the same definitions as before, but we drop the c index.

$$\min \quad \overline{\text{COST}}(p, \pi) = 0 - u_{v_{\text{SRC}}}^1 + \sum_{\ell \in L} u_\ell^2 \cdot p_\ell^W + \sum_{s \in S} \sum_{\ell \notin s} u_{\ell s}^3 \cdot (\pi_{\ell s}^W + \pi_{\ell s}^B) \quad (5)$$

The first set of equations (6) represent the flow conservation equations, expressing that the net flow going into a node should be either -1 (for the source node), $+1$ (for a destination node) or 0 otherwise.

$$\sum_{\ell \in \omega^+(v)} p_\ell^\star - \sum_{\ell \in \omega^-(v)} p_\ell^\star = \begin{cases} -1 & \text{if } v = v_{\text{SRC}} \\ +p_{\ell'}^\star & \forall \ell' \in L_{\text{DST}} \\ 0 & \text{otherwise} \end{cases} \quad (6) \quad \forall v \in V \text{ and } \star = W, B$$

Next, constraints (7) assure that there will be no loops, and exactly 1 working and backup path will be constructed. Additionally (8) enforces that a single working and backup destination will be chosen. Finally, working and backup paths obviously need to be disjoint (9) with respect to an SRLG $s \in S$.

$$\sum_{\ell \in \omega^+(v)} p_\ell^\star \leq 1, \quad \sum_{\ell \in \omega^-(v)} p_\ell^\star \leq 1, \quad \forall v \in V, \star = W, B \quad (7)$$

$$\sum_{\ell \in L_{\text{DST}}} p_\ell^\star = 1, \quad \text{for } \star = W, B \quad (8)$$

$$p_\ell^W + p_{\ell'}^B \leq 1, \quad \forall \ell \in \bigcup_{s \in S} s, \forall \ell' \in \text{SRGL}(\ell) \quad (9)$$

It remains to define constraints so that the definitions of π^W and π^B apply as before. For this, we define additional auxiliary variables a_s , each associated with an SRLG $s \in S$:

- a_s Binary variable, equaling 1 if any of the links $\ell' \in s$ is used as working link (i.e. $p_{\ell'}^W = 1$), hence if the chosen working path is *affected* by the SRLG.

Looking at (5), we note that the last summation only contains links not in the SRLG at hand, hence (10) applies. Constraints (11) ensure the logical relation $\pi_{\ell s}^B \equiv p_\ell^B \wedge a_s$ that expresses the definition given before. The above definition of a_s translates to the logical relation $a_s \equiv \bigvee_{\ell' \in s} p_{\ell'}^W$, or thus (12).

$$\pi_{\ell s}^W = p_\ell^W \quad \forall s \in S, \forall \ell \in L \quad (10)$$

$$\left. \begin{aligned} \pi_{\ell s}^B &\geq p_\ell^B + a_s - 1 \\ \pi_{\ell s}^B &\leq p_\ell^B \\ \pi_{\ell s}^B &\leq a_s \end{aligned} \right\} \quad \forall s \in S, \forall \ell \notin s \quad (11)$$

$$M \cdot a_s \geq \sum_{\ell' \in s} p_{\ell'}^W \quad \forall s \in S, M = |s| \quad (12)$$

The above constraints all apply regardless whether we consider relocation or not. Yet, if we do not want to relocate, we need to enforce one additional constraint (13), stating that working and backup destination need to be the same:

$$p_\ell^W = p_\ell^B, \quad \forall \ell \in L_{\text{DST}} \quad (13)$$

IV. CASE STUDY

We evaluated the above methodology on a European network topology comprising 28 nodes and 41 links (the same as in [4]). To assess the cost of resilience against both link and server failures, as well as the benefit of relocation, we compared the following three scenarios:

- *IL, No Reloc.*: Single link failures only, no relocation;
- *IL, Reloc.*: Single link failures only, with relocation;
- *ILS, Reloc.*: Single link or server site failures, with relocation.

To model the failures, the SRLGs were constructed as follows: (i) for single link failures (*IL*), S has elements s each containing a pair formed by a link $\ell \in L_{\text{NET}}$ and its reverse; (ii) for single link or server failures (*ILS*), we additionally include in S the singletons formed by each single link $\ell \in L_{\text{DST}}$. In our results below, we consider the case of $K = 3$ server sites, however we did not fix the locations a priori (as opposed to [4]). As pointed out previously, we use the server location ILP of [7] to choose them, based on the demand. For the latter, we varied the total number of unit demands ($D = \sum \Delta_v$) between 10 and 350. For each D , we created 10 random instances. The measures plotted below are averages over those 10 random instances per demand case. The network was dimensioned using the approach outlined in Section III. For the RMP, we set the relative server cost parameter $\alpha = 1$.

Looking at Fig. 2, we confirm our earlier findings [3], [4] that exploiting relocation enables substantial savings in the number of network resources (wavelengths); in this case in the order of 19% (comparing *IL, Reloc.* vs. *ILS, No Reloc.*, averaged over the data points plotted in the figure). Using our model incorporating server capacity dimensioning, we are also able to quantitatively assess that relocation to protect against single link faults incurs an increase in server capacity around 11% (averaged over the [10, 350] demands), as can be derived from Fig. 3. Hence, relocation indeed has a considerable net




Fig. 2. The number of wavelengths required in the EU network for $\alpha = 1$, $K = 3$. Error bars (which are largely hidden behind data point markers) indicate standard deviation among the 10 random instances per data point.

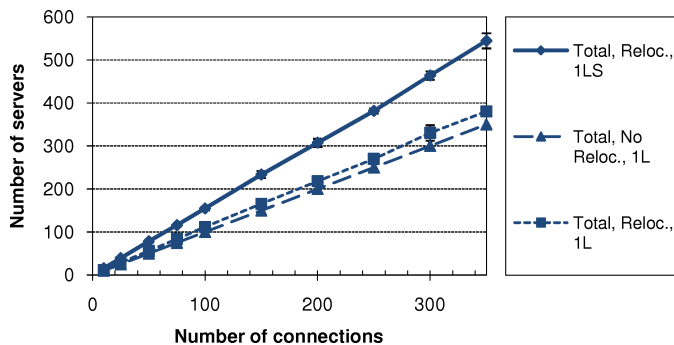


Fig. 3. The number of servers required in the EU network for $\alpha = 1$, $K = 3$. Error bars (which are largely hidden behind data point markers) indicate standard deviation among the 10 random instances per data point.

benefit in resource requirements. (Note that our model jointly minimizes total server and network cost.)

Considering now the protection against both network and server site failures (*1LS, Reloc.*), we note that in terms of server capacity we do better than a 1+1 protection strategy at the server sites, since we only need around 55% extra server capacity (compared to *1L, No Reloc.*, and averaged over the considered 10-350 demands). The additional network capacity amounts to 26% (avg. over the 10-350 demands). Note that the difference in network capacity mainly stems from difference in backup wavelengths; we noted that the maximal deviation in network resources in the relocation cases only amount to around 10% (graphs omitted because of space limitations).

V. CONCLUSION

In this paper, we present an elegant and scalable model to jointly dimension network and server capacity for Grid-like scenarios, where demands for IT infrastructure (servers) and connectivity towards it arise with a freedom in choosing the IT resource location (anycast principle). The model, using column generation, allows providing resilience against both network and server site failures by specifying the appropriate shared risk link groups (SRLGs) comprising possibly jointly failing resources. We have quantitatively evaluated the approach in a

case study on a European network. There, relocation in protection against single link failures achieves network resource savings around 19%, but calls for around 11% extra server capacity. Providing also protection against server site failures incurs 55% extra servers (considerably better than 1+1 server protection), and 26% extra wavelengths (compared to classical shared network protection without relocation). Note that these numbers may depend on the network topology (esp. node degrees, hence meshedness), number of server site locations, and assumed relative server vs. network cost (the α factor). For more server site locations, we expect the relative increase in server and wavelength resources will be lower. Future work will assess those parameters' influence on the differences in server and network resource requirements between traditional shared path protection and relocation cases.

ACKNOWLEDGMENT

Work described in this paper was partly funded by the European Commission through the 7th ICT-Framework Programme (FP7) projects Geysers and BONE (Building the Future Optical Network in Europe, a Network of Excellence). Results were obtained using the Stevin Supercomputer Infrastructure at Ghent University. M. De Leenheer and C. Devellder are supported as post-doctoral fellows of the Research Foundation – Flanders (FWO). J. Buysse is supported by the agency for Innovation by Science and Technology (IWT). B. Jaumard has been supported by a Concordia University Research Chair (Tier I) and by an NSERC (Natural Sciences and Engineering Research Council of Canada) grant.

REFERENCES

- [1] A. Concaro, G. Maier, M. Martinelli, A. Pattavina, and M. Tornatore, "QoS provision in optical networks by shared protection: An exact approach," in *Proc. 2nd Int. Workshop on Quality of Service in Multiservice IP Netw. (QoS-IP 2003)*, Milano, Italy, 24–26 Feb. 2003, pp. 419–432.
- [2] S. Ramamurthy and B. Mukherjee, "Survivable WDM mesh networks, part I – protection," in *Proc. 18th IEEE Conf. Computer Commun. (INFOCOM 1999)*, vol. 2, San Francisco, CA, USA, 29 Mar. – 2 Apr. 1999, pp. 744–751.
- [3] J. Buysse, M. De Leenheer, B. Dhoedt, and C. Devellder, "Exploiting relocation to reduce network dimensions of resilient optical grids," in *Proc. 7th Int. Workshop Design of Reliable Commun. Netw. (DRCN 2009)*, Washington, D.C., USA, 25–28 Oct. 2009.
- [4] B. Jaumard, J. Buysse, A. Shaikj, M. De Leenheer, and C. Devellder, "Column generation for dimensioning resilient optical grid networks with relocation," in *Proc. IEEE Global Telecommun. Conf. (GLOBECOM 2010)*, Miami, FL, USA, 6–10 Dec. 2010.
- [5] H. Zang, C. Ou, and B. Mukherjee, "Path-protection routing and wavelength assignment (RWA) in WDM mesh networks under duct-layer constraints," *IEEE/ACM Trans. Netw.*, vol. 11, no. 2, pp. 248–258, Apr. 2003.
- [6] N. M. K. Chowdhury and R. Boutaba, "A survey of network virtualization," *Comput. Netw.*, vol. 54, no. 5, pp. 862–876, Apr. 2010.
- [7] C. Devellder, B. Mukherjee, B. Dhoedt, and P. Demeester, "On dimensioning optical grids and the impact of scheduling," *Photonic Netw. Commun.*, vol. 17, no. 3, pp. 255–265, Jun. 2009.
- [8] B. Van Houdt, C. Devellder, J. F. Pérez, M. Pickavet, and B. Dhoedt, "Mean field calculation for optical grid dimensioning," *IEEE/OSA J. Opt. Commun. Netw.*, vol. 2, no. 6, pp. 355–367, Jun. 2010.
- [9] J. Buysse, M. De Leenheer, B. Dhoedt, and C. Devellder, "On the impact of relocation on network dimensions in resilient optical grids," in *Proc. Int. Conf. on Optical Network Design and Modelling (ONDM 2010)*, Kyoto, Japan, 31 Jan.–3 Feb. 2010.
- [10] J. W. Suurballe and R. E. Tarjan, "A quick method for finding shortest pairs of disjoint paths," *Networks*, vol. 14, no. 2, pp. 325–336, 1984.