Les Cahiers du GERAD ISSN: 0711–2440

Dimensioning Resilient Optical Cloud Networks

C. Develder
B. Jaumard

G–2013–39

June 2013

Dimensioning Resilient Optical Cloud Networks

Chris Develder

Department of Information Technology (INTEC) - IBCN
Ghent University - iMinds
Ghent, Belgium

chris.develder@intec.ugent.be

Brigitte Jaumard

GERAD & CSE Department
Concordia University
Montreal (Quebec) Canada

bjaumard@cse.concordia.ca

June 2013

Les Cahiers du GERAD

G–2013–39

Abstract: Optical networks, given their excellent characteristics in terms of high bandwidth and low latency, play a crucial role in enabling today's applications that heavily rely on network connectivity. This is exemplified by the evolution towards grid and cloud computing, as observed during the last decennium. Such cloud computing based applications call for a tighter integration between the control and management of data centers, and the (optical) wide area network connectivity they require. The advent of the resulting "optical clouds" also demands an appropriate network (and data center) planning methodology, which exhibits significant differences compared to, e.g., the classical routing and wavelength assignment (RWA) methods. In this paper, we explain fundamental concepts that cloud computing is based on (i.e., anycast routing and virtualisation), and in particular outline what challenges they bring about in network dimensioning. We survey possible solutions, and specifically argue how large scale optimization (column generation) can be adopted to dimension resilient optical networks for cloud computing. Sample case studies illustrate potential capacity savings enabled by exploiting the particular cloud concepts.

Key Words: Cloud Computing, Grid Computing, Network Planning, ILP, Column Generation, Virtualisation, Anycast routing.

Acknowledgments: B. Jaumard was supported by NSERC (Natural Sciences and Engineering Research Council of Canada) and by a Concordia University Research Chair (Tier I) on the Optimization of Communication Networks.

1 Introduction

In the last decennium, we have observed growing importance of network-based applications. The evolution towards grid computing mainly in scientific applications, as well as towards more user- and business-friendly cloud computing paradigms incontestably is based on advanced and cost-effective network technologies. Especially the high bandwidth and low latencies offered by the underlying optical wide area networks today make it possible to remotely and interactively run applications in data center infrastructures, thus also alleviating heavy processing on the user's own device (which therefore can be relatively lightweight, as in thin client systems). Clearly, the other pillar (besides networking) that enabled this evolution, is formed by recent advances in data center infrastructures and their associated control and management software.

It is such relatively easy access to powerful, distributed software and hardware resources that formed the foundation of scientific applications, cf. the quintessential European Grid Infrastructure (EGI) that originated as computing infrastructure to serve the needs of efficiently distributing and processing the vast amount of physics experimental data produced by the Large Hadron Collider (LHC) at CERN. Yet, also more mundane business applications benefit from similarly distributed (cloud) computing infrastructure (e.g., Amazon EC2): this ranges from transactional systems over collaboration tools and multimedia processing to data mining. From a consumer perspective, managing and manipulating personal content services (mail, images, video, etc.) nowadays are virtually unthinkable without relying on any data center based application. The role of the optical core network is to interconnect all these entities, in terms of data sources (such as sensors, experimental facilities, etc.), users, and distributed computing infrastructure, as illustrated in Figure 1. For a more elaborate discussion of these applications and their requirements, we refer to [1].

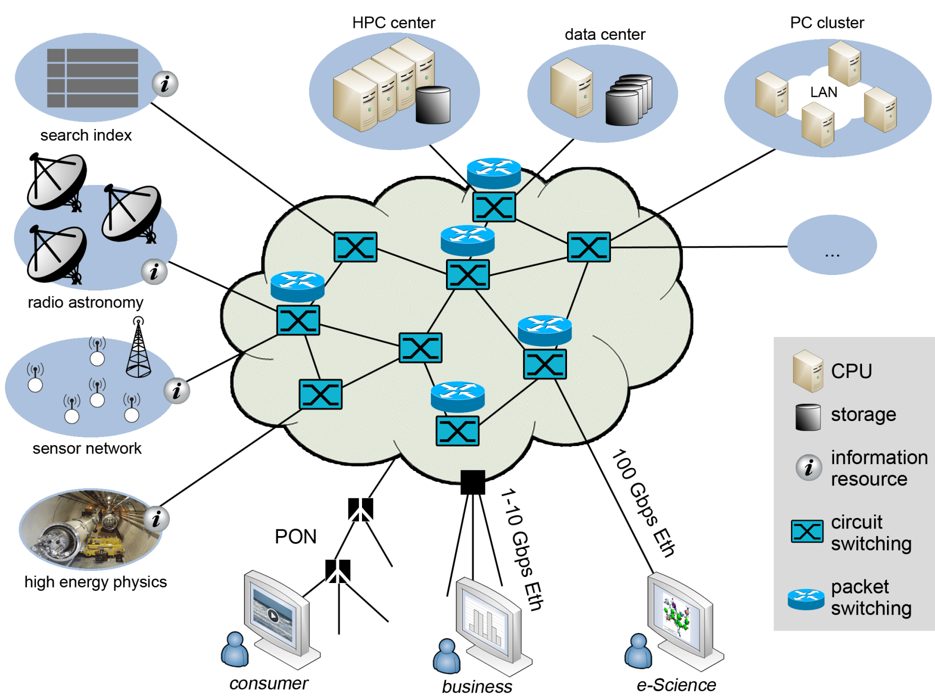

Figure 1: Optical grids/clouds: an optical core network interconnects various users (bottom: consumers, businesses, scientists) to distributed server infrastructure (top) for computing and storage. (Adapted from [1].)

1.1 Cloud and grid computing fundamentals

Grid computing mainly originated from so-called e-Science applications, and aimed to interconnect mainly high performance computing (HPC) facilities to offer access to scientists at various locations. The term "grid" relates to the power grid, where users have access to energy via a simple, standardised interface (i.e., the power plug): conceptually, grid computing should offer to scientists the same simple access to computing resources. In a typical scenario, scientific experiments to run on a grid are submitted as so-called "jobs", through specific job submission interfaces and schedulers with appropriate queueing policies to decide when and at what machine to run them. Since the grid users are typically geographically dispersed, the core network needs to be planned to have appropriate connectivity and sufficient bandwidth capacity between the locations of the scientists (the users) and, e.g., the HPC resources (a server facility). Note that the latter HPC center also needs to be designed to have sufficient capacity to deal with the volume of, e.g., scientific computing tasks delegated to it: clearly, the server and network capacities need to be aligned. Here, it is important to realize that in practice, the user does not greatly care at what exact location the computing tasks are served: this amounts to the anycast routing principle. *Anycast* means that the destination of traffic is not fixed a priori, but rather can be freely chosen by the system from a given set of locations. This can be exploited when dimensioning the network, as we will discuss in further sections.

Cloud computing manifests itself in more commercially oriented applications (versus the typically more academic nature of grids), and their applications are served in data centers. This can be seen as a realisation of the seminal idea of "utility computing" launched back in 1961 by John McCarthy, who envisioned that computing facilities could be offered on-demand, as traditional utilities. From an architectural perspective, clouds can be cast in a layered "everything as a service" (XaaS) taxonomy. In this paper, we will mainly focus on the so-called Infrastructure-as-a-Service (IaaS) paradigm: an IaaS cloud offers resources for storage and computing, where "resources" can refer to physical resources (i.e., servers), but these are typically virtualised. This *virtualisation* is a fundamental concept of clouds, and refers to either partitioning a single physical resource in multiple virtual resources (1:N) or aggregation of multiple physical resources in a single virtual resource (N:1). A concrete example is where a single physical server is used to host multiple so-called virtual machines (VMs), each running their own (and not necessarily the same) operating system, etc. Thus, a single physical server can be shared by multiple users, each running their own application and not interfering with one another. This abstraction mechanism thus allows to exploit statistical multiplexing: it can be assumed not all users will simultaneously use the full physical server capacity, and by collocating their VMs, the physical resource capacity can be better utilised.

The same virtualisation concept is also applied in a network context: the network can be logically partitioned into distinct "slices", each comprising a subset of the overall network, where partitioning can be realised at the optical equipment level (e.g., certain ports, wavelengths, etc. are assigned to a particular slice and cannot be used for another). This can be enabled by technologies such as OpenFlow, and implies that the network control plane running in each of the slices could even be diversified with, e.g., specific policies, routing algorithms, etc. It is notable that recently, combined virtualisation of both networking and server (i.e., IT) resources has gained widespread attention (see, e.g., [2]), mainly due to the success of grid/cloud computing concepts.

2 Network dimensioning for anycast routing

The anycast routing problem, e.g., in an IP context, typically consists of finding for each source node a path towards one of a set of candidate destinations, such that a particular cost (e.g., bandwidth consumed, delay) is minimised. Since this is an NP-hard problem, typically heuristic approaches have been proposed. Here, we focus on optical core networks, and in particular circuit switching (OCS). These networks adopt wavelength division multiplexing (WDM): on an optical fiber multiple wavelengths are concurrently transmitted. Thus, a classical optical network planning problem is that of routing and wavelength allocation (RWA): a given set of (source, destination)-pairs is given, each with the required capacity (expressed in number of wavelengths) to be transmitted between them. In the case of optical grids/clouds, where anycast applies, the specific destination is not given a priori, and thus we cannot easily specify the (source, destination)-based traffic matrix assumed in traditional optical network design. Thus, for grids/clouds the RWA problem needs to be extended to the anycast routing and wavelength assignment (ARWA) problem: we need to decide what destination to choose for each demand originating from particular source, as well as what route (and wavelength) to take to reach it (see Figure 2).

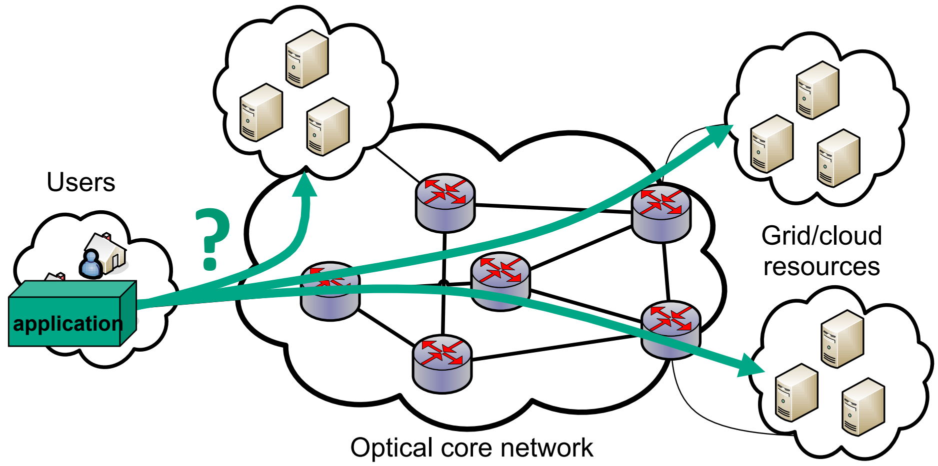

Figure 2: The anycast routing principle: grid/cloud users typically don't care about the exact location their application (be it a VM in the IaaS cloud computing paradigm, or a scientific computation job in an e-science grid context) is being served. Thus, anycast routing needs to decide both on the eventual destination and the route towards it.

Here we focus on the *offline* variant of ARWA: we look for routes and wavelength assignments for a static set of requests (as opposed to the dynamic light path establishment at runtime, in an online scenario). Stated more formally, a first grid/cloud planning problem to address can be summarised as follows:

Given

- *Topology* comprising the sites where grid/cloud requests originate, as well as the optical fiber network interconnecting them;
- *Demand* stating the amount of requests originating at each of the source sites; and
- K *server site locations*, where server infrastructure is provided to cater for the grid/cloud demands;

Find

- *Destination sites and routes* to follow for each grid/cloud request (where each destination should be one of the K server locations);
- *Network capacity* to provide on each optical fiber link;

Such that the total network capacity is minimised.

Traditional RWA problems are typically solved using either integer linear programming (ILP) approaches for medium size data instances [3] or heuristics for large instances (because of ILP scalability issues). Clearly, by adding the anycast principle, which introduces an extra degree of freedom (i.e., the destination of traffic), the scalability issue occurs earlier. Hence, we note several heuristic approaches to the offline ARWA problem (see references in [4]).

As opposed to these heuristic approaches, in several optical cloud/grid dimensioning problems, we have successfully developed a column generation model with an adapted solution scheme. Its basic idea is to formulate the optimisation problem as one that has to decide on which (and how many times) "configurations" to use. A concrete example of a "configuration" could be a set of a two paths, one to use under failure-free conditions, and one to use in case the first is affected by a failure (see further, Section 4). Then, instead of listing all possible configurations, column generation initially only considers a limited subset of them. The best combination (e.g., minimising total network capacity to satisfy a given demand) of configurations is then found by solving a so-called Restricted Master Problem (RMP). Subsequently, a Pricing Problem is derived from the current RMP solution to find possibly new (cost-reducing) configurations, which are then added to the RMP. Thus, RMP and PP are solved alternately until no further cost-reducing configurations are found. A particular column generation + ILP scheme is illustrated in Figure 3; for more details on column generation, see, e.g., [5].

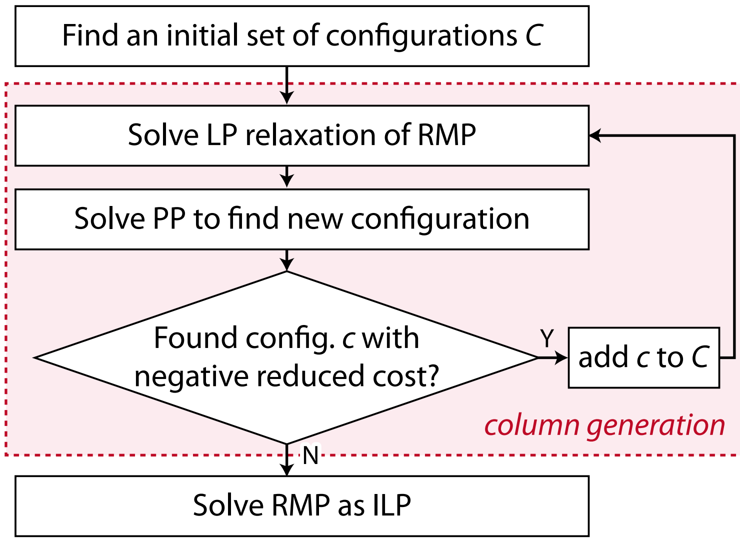

Figure 3: A column generation + ILP approach to solve large-scale integer linear programs. (RMP: restricted master problem; PP: pricing problem; LP: linear program; ILP: integer linear program.)

3 Integrating server dimensioning

In the simplest optical grid/cloud dimensioning problem as stated before, it should be noted that the locations of the server infrastructures (i.e., the K data centers) are assumed to be given at the outset. Yet, as pointed out earlier, the resulting network dimensions may be highly dependent on both the number and the chosen location of the data centers. Thus, a more general grid/cloud dimensioning problem would rather optimise the data center locations, jointly with the network dimensions and associated RWA. Clearly this highly increases the complexity, and hence an obvious approach would be to separate the planning problem in two distinct steps: (S1) first determine the K data center locations, and then (S2) solve the ARWA problem to derive the network dimensions.

The first step (S1) resembles the classical facility location problem [6], and in case we aim to minimize the network resources, in particular the K-median (also called *minisum*) problem. The latter finds the location of K facilities with the objective of minimising the overall distance between facilities and demand sources. Therefore, an obvious choice to solve step (S1) would be to use a K-means[1] clustering variant. However, as shown in [7], we can easily define a straightforward ILP formulation that can be solved quickly (and can be extended to also cover the resilient cloud dimensioning problem as introduced further in Section 4, see [4]).

The second step (S2) can be extended to also determine the required server capacity at each of the chosen data center locations, using a single ILP formulation (or a column generation version thereof). To this end, for the dimensioning of the optical core network as well as the server capacity at data center locations, a graph model $G = (V, L)$ can be constructed as illustrated in Figure 4: the vertices V are partitioned into sources of the traffic ("users"), the optical cross-connects (OXCs), and the (candidate) data center locations. The capacity on the links L_{NET} between the OXCs represents the optical fiber capacity (e.g., expressed in number of wavelengths of a given bitrate), while the capacity of links L_{DST} towards the data centers models the server resource requirements.

[1]Technically, we need K-medoid clustering, in case candidate locations (e.g., nodes in the given topology) are given a priori, since the locations then have to be an existing node and cannot be chosen completely arbitrarily.

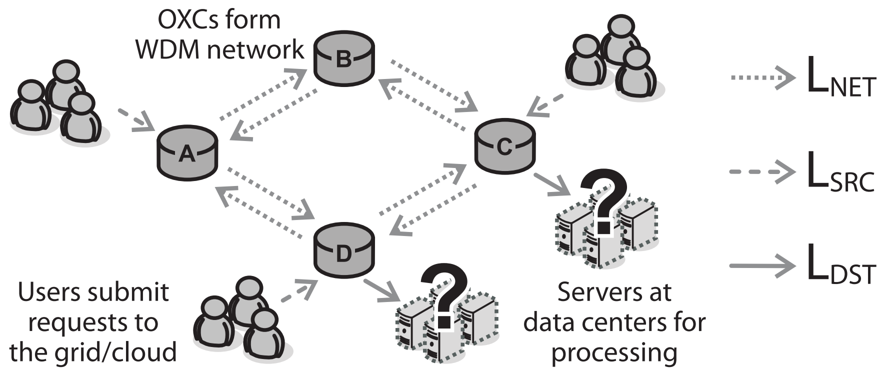

Figure 4: A graph model $G = (V, L)$ for dimensioning both the optical core network and the server capacity in data center locations. The vertices V and directed edges L in the graph are partitioned among (i) the users that request services from the grid/cloud and hence are traffic sources, (ii) the optical core network comprising the OXCs and optical fiber links L_{NET}, and (iii) the data centers whose capacity will be modeled as link capacities of $\ell \in L_{\text{DST}}$.

Recent works also have solved the steps (S1) and (S2) in a unified, single-step approach, thus jointly deciding on the locations of data centers as well as the optical core network dimensions. An illustrative example of such an integrated network and data center dimensioning approach is [8], which actually addresses resilience as well.

4 Resilient cloud dimensioning

Given the increasing dependence of various businesses and users on distributed (cloud) computing facilities, there is a need for resilient provisioning. Thus, at the planning stage, the network and server resources should be dimensioned in a way that they are sufficient to cater for counter-measures against failures. In general, resilience strategies can be classified in either protection or restoration. Restoration amounts to finding backup resources at the time a failure occurs, whereas protection generally refers to the case where remedying actions are pre-planned proactively (e.g., alternate backup routes are already pre-provisioned).

For optical WDM networks, multiple resilience approaches have been proposed [9]. A well-known example is shared path protection: a primary path from source to destination is protected by a failure-disjoint backup path. Here, the term *shared* refers to the fact that the same resources (in casu, link capacity, i.e., wavelengths) can be reused among different backup paths under different failure scenarios – this avoids to excessively set aside capacity for (hopefully quite infrequent) specific failures. Further, *failure-disjoint* means that primary and backup paths do not share any resources (i.e., nodes and/or links) that may jointly fail. A basic example is protection against single link failures: then primary and backup will need to be link-disjoint. In general, it is often convenient to model multiple failure scenarios by so-called shared risk groups (SRGs); typically these are shared risk link groups (SRLGs): sets of links that can be affected by a particular common root cause (e.g., fibers that are bundled in the same duct). This general SRG model can also be used to model the impact of, e.g., natural disasters.

One could be tempted to believe that traditional optical network resiliency approaches, and associated dimensioning methodology, can be straightforwardly applied in the case of optical grids/clouds. Yet, the anycast routing principle needs to be catered for, as explained before. In case of failure protection, we can advantageously exploit anycast to limit the amount of network resources: since end users are not very much concerned where their applications are physically hosted, we can wisely choose this location as to minimize the required resources for ensuring resiliency. This amounts to the idea of *relocation*: we can use an alternate destination under failure conditions, compared to the failure-free scenario, as illustrated in Figure 5(a).

An extensive discussion of joint dimensioning of optical core network and datacenter capacities is presented in [4], of which Figure 5(b) presents a sample result. Here, we compare the naive approach without relocating

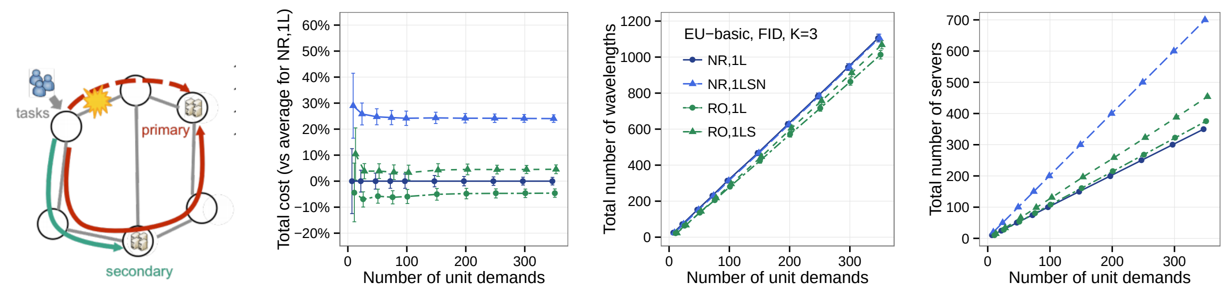

(a) (b)

Figure 5: (a) The idea of relocation: instead of provisioning a backup path to the same destination as the primary, we can choose an alternate destination under failure conditions. (b) Illustrative result showing the effect of adopting relocation: network resource requirements can be reduced, at the slight expense of some extra server capacity, which typically still amounts to a net cost benefit.

(NR) with optional relocation (RO), where we choose to relocate under failure conditions if this leads to cost savings. Furthermore, results show the required capacities required to protect against single link failures (1L) or single link or server failures (1LS). We adopted a column generation approach on a 28-node, 41-link European network topology, for increasing amount of unit demands (see [4] for full details and assumptions). The leftmost graph shows the relative difference in total cost (summing both network and server capacities) compared to the baseline of single link protection without relocation (NR, 1L). The middle graph shows the total network cost, expressed as the number of wavelengths summed over all optical fibers, whereas the rightmost graph plots the total server capacity. For the specific case at hand, with $K = 3$ data center locations, we observe a network capacity reduction of around 9% when protecting against single link (1L) failures using relocation (RO) compared to the no-relocation (NR) baseline. Yet, since we relocate, it means a slight increase in server capacity at the relocation sites, reflected in the total server cost graph. Nonetheless, we still achieve a net cost benefit as shown in the total cost plot. The additional server capacity can be reused to protect also against server failures, and therefore the relocation case only requires a modest total cost increase to cater for both single link or server failures (RO, 1LSN). Concluding this discussion, we see a clear infrastructure cost saving potential of relocation in cloud/grid scenarios — an advantage which even becomes more pronounced in case the number of data center locations increases [4].

Work on resilient network dimensioning has also distinguished between failure-independent (FID) versus failure-dependent (FD) rerouting. In failure-independent rerouting, the same backup solution is chosen for any of the possible failures affecting the primary paths. With a protection approach, the advantage is that per primary path (as well as, in a cloud context, the accordingly chosen data center destination) only a single backup path needs to be provisioned. Also, the protective action (namely, switching to the backup) is independent of the exact cause of the primary's failure. This clearly has its advantages in terms of control and management complexity, including maintenance of associated state for the backup. Yet, when adopting a shared protection approach, flexibility in choosing a different backup solution that depends on the failure (i.e., FD) might at least intuitively allow better sharing the available capacity and hence in a dimensioning algorithm allow for lower infrastructure cost. Still, earlier work on unicast problems — i.e., not in a cloud context with the opportunity to exploit relocation and hence potentially increasing capacity sharing — reported mixed advantages in terms of network capacity [10,11]. For the current case of grid/cloud dimensioning, we noted that infrastructure cost savings can be more substantial, at least for the case with a sufficiently large number of data centers (e.g., [4] showed cost reductions of below 2% for $K = 3$ data center locations, but around 6% for the case of $K = 7$ data centers, in the 28-node EU network mentioned before).

4.1 Disaster resilience

Note that so far, we discussed the generic grid/cloud resiliency. Whereas traditionally, failure scenarios typically considered equipment/component failures because of malfunctions or, e.g., infrastructure works (such as construction workers accidentally cutting optical fibers), more recently also more larger scale failures in the context of disasters are being studied. These disasters can include natural disasters, but also human-induced ones (e.g., terrorist attacks). From a network modelling perspective, the aforementioned SRGs can be used to represent so-called "disaster zones", as, e.g., in [12]. The latter work addresses a cloud context, and proposes replication of services (or content) in distinct locations. Thus, relocation under failure conditions is enforced (whereas, e.g., the results discussed earlier and presented in Figure 5(b) had *optional* relocation that was only adopted when it induced cost savings). Also, from the data server perspective, this amounts to dedicated protection, and [12] even considers multiple replicas (i.e., up to K locations where a particular service/content is hosted, with $K \geq 2$). Since this can be conceptually perceived as a variant of the relocation idea, [12] finds similar network cost savings compared to traditional single protection. (In terms of methodology, a single (unscalable) ILP formulation is presented to decide jointly on location choices and routings, and subsequently, either a 2-step approach (first deciding on the locations, then on the routes) as well as a combination of LP relaxation and heuristics is applied.)

4.2 Virtual optical network planning

As outlined in the introduction, virtualisation is one of the key concepts behind cloud computing, and this idea is also applied on the network level. From a business case perspective, this introduces different roles that can be played by telecom/cloud actors: we can distinguish Physical Infrastructure Providers (PIP) and Virtual Infrastructure Operators (VIO). The physical infrastructure providers offer the physical resources, i.e., the actual hardware in terms of (core) network and server (data center) equipment. Using virtualisation, these physical resources can be logically partitioned into different "slices" where each slice thus represents a virtual network/cloud instance, and can be operated by a VIO. Thus, multiple network/cloud service providers as VIOs can share the same underlying infrastructure, but operate their share of it independently of the other. As an example, to technically implement this, the Geysers project [2] proposes a so-called Logical Infrastructure Composition Layer (LICL) to maintain the mapping of logical virtual networks of the VIOs to the physical infrastructure of the PIPs, while within a VIO's (logical) cloud network an advanced control plane (the so-called NCP+) is adopted to setup the network connections (including choosing the appropriate data center location via anycast routing).

From a resilience perspective, the question now arises how to ensure connectivity to a data center as to not disrupt services offered under physical resource(s) failures. There are two fundamental approaches to this [13], as illustrated in Figure 6 for network mappings only: provide resilience directly in the physical layer, i.e., under responsibility of the PIP (hence coined PIP-resilience), or rather defer it to the VIO (thus,

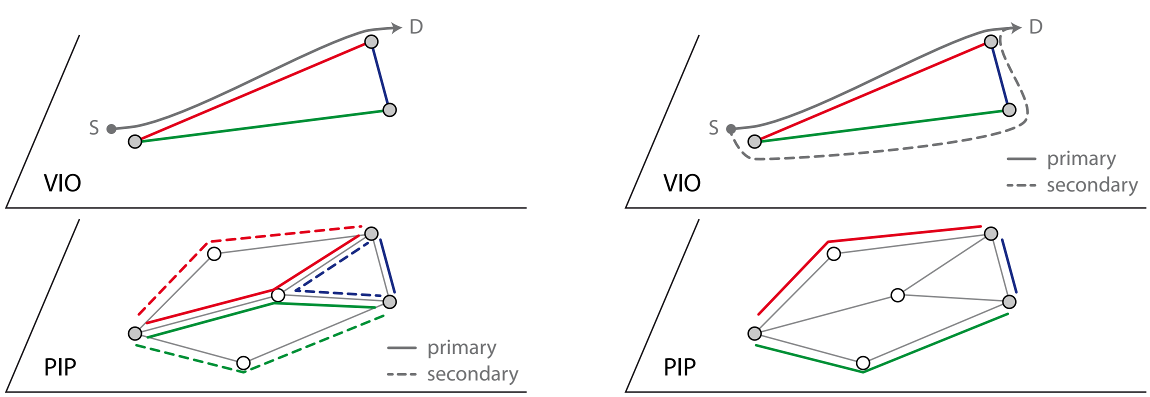

(a) PIP-resilience (b) VIO-resilience

Figure 6: Resilience provided by the physical infrastructure provider (PIP) vs the virtual infrastructure operator (VIO), for single physical link failures, and an illustrative connection (S, D) provisioned in the VIO network.

VIO-resilience). PIP-resilience can be realised on the network level by mapping a logical VIO link not to just a single connection in the physical network, but rather adopt protection in the PIP: for instance to realize 1+1 protection, two failure-disjoint connections can be set-up between each pair of end-points corresponding to a logical link of the VIO's network. VIO-resilience on the other hand requires the VIO to, e.g., reroute a connection in its logical network via an alternate backup path, such that the underlying physical resources of primary and backup paths are failure-disjoint. Clearly, the latter requires awareness of the VIO of the possible physical resource sharing of his mapped logical links (which, from a business confidentiality perspective, the PIP may want at least to abstract in some way). Yet, as shown in [14], VIO-resilience might exhibit total network capacity savings compared to PIP-resilience.

5 Energy efficiency

Given increasing environmental concerns, and global commitments to limiting if not reducing our overall carbon footprint, recent research efforts focused on minimizing the energy consumption of network solutions. Optical networks in this perspective offer some inherent advantages, given their low power consumption compared to solely electronic technologies. The evaluation of cloud solutions in terms of energy use needs careful balancing of both network and server (for storage and processing) energy. In the overall picture, typically the power consumption of data centers is dominant compared to optical core networks, yet also in the latter domain, e.g., energy-aware routing approaches have been proposed. The virtualisation principle, given its underlying statistical multiplexing idea, promises accurate matching of on-demand provisioned resources to the instantaneous demand, where currently unused resources can be powered off or at least put in a low-power sleep state. More progressive proposals include the so-called "follow the sun / follow the wind" idea to locate data centers close to renewable energy sources (e.g., solar plants, wind farms) and moving the catered services (e.g., VMs in a cloud context, computation jobs in a grid setting) from one data center to another, depending on their associated renewable energy availability. Obviously, such relocation to data centers with lower CO2 emissions can come at the cost of higher network capacity utilisation (which slightly reduces the net advantage of relocating). Moreover, one should be aware that in order to easily move services from one location to the other, sufficient data center capacity needs to be installed, which entails over-provisioning. In the overall picture, the latter also implies an energy cost (even if the installed resources are not turned on), since for the CO2 footprint assessment full lifecycle assessment is required. As shown by [15], for a given required total operational server capacity, and given probability of renewable energy availability, an optimum number of locations (and hence associated over-provisioning factor) can be found to achieve maximal CO2 footprint reduction.

6 Conclusion

The fundamental principles that optical clouds/grids are built on include anycast and (server as well as network) virtualisation, which introduce both challenges and opportunities from a capacity planning perspective. We have discussed how the anycast principle, relying on the fact that cloud/grid users basically do not care where their applications are running, can be exploited to our advantage for both resilience purposes (i.e., route requests to an alternate destination) and energy concerns (i.e., route to the site with the lowest carbon footprint). Yet, from a dimensioning problem perspective, this implies the lack of a fully specified (source, destination)-based traffic matrix: only the sources are completely known in advance. We argue that approaches based on column generation can prove to be scalable solution methods to solve the increasingly complex (because of the extra degrees of freedom that we need to decide on) problems that result. Moreover, we have also shown illustrative results of dimensioning approaches that allow to jointly dimension both network (wavelength) and data center (server) capacities. Overall cost savings are possible by carefully deciding on the data center locations. Also the virtualisation principle has its impact: it allows for statistical multiplexing, and thus potential cost savings stemming from reduced physical infrastructure requirements. We further explained how resilience can be provided either by the physical infrastructure provider (PIP), or rather by the virtual infrastructure operator (VIO).

Note that in our discussion, we only focused on offline planning problems: deciding on the infrastructure dimensions, for a given (static) demand. Clearly, similar approaches can be applied in an online context (as part of the online provisioning process, as in a path computation engine (PCE) in a GMPLS context — for which the aforementioned Geysers project also proposed control plane extensions). Also the mapping of a VIO's virtual infrastructures to a PIP's physical resources clearly needs to be performed in an online context, and may involve re-planning (thus revising earlier mapping decisions) to obtain, e.g., better resource utilisation. Apart from the planning algorithms, also the control plane operations need to be (re)designed, and clearly not solely on the network part, but they need also to integrate with the management procedures of data centers (to realize e.g., the relocation of VMs).

Furthermore, we mainly debated optimisation in terms of infrastructure cost (related to their capacities in terms of bandwidth, processing power, etc.). Past and ongoing research efforts also look into taking account additional constraints (such as impairments in transparent optical networks), and more advanced technologies beyond classical WDM (such as optical flexible grids). These imply associated challenges to exploit/integrate them efficiently in the present and future cloud/grid context.

References

[1] C. Develder, M. De Leenheer, B. Dhoedt, M. Pickavet, D. Colle, F. De Turck, and P. Demeester, "Optical networks for grid and cloud computing applications," *Proc. IEEE*, vol. 100, no. 5, pp. 1149–1167, May 2012.

[2] E. Escalona et al., "Geysers: A novel architecture for virtualization and co-provisioning of dynamic optical networks and it services," in *Proc. Future Netw. Mobile Summit (FNMS 2011)*, Warsaw, Poland, 15–17 Jun. 2011, pp. 1–8.

[3] B. Jaumard, C. Meyer, and B. Thiongane, "Comparison of ILP formulations for the RWA problem," *Optical Switch. Netw.*, vol. 4, no. 3-4, pp. 157–172, November 2007.

[4] C. Develder, J. Buysse, B. Dhoedt, and B. Jaumard, "Joint dimensioning of server and network infrastructure for resilient optical grids/clouds," *IEEE/ACM Trans. Netw.*, 2013, to appear.

[5] V. Chvatal, *Linear Programming.* Freeman, 1983.

[6] M. Melo, S. Nickel, and F S. da Gama, "Facility location and supply chain management – a review," *Eur. J. Oper. Res.*, vol. 196, no. 2, pp. 401–412, 2009.

[7] C. Develder, 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. Jaumard, A. Shaikh, and C. Develder, "Selecting the best locations for data centers in resilient optical grid/cloud dimensioning (invited paper)," in *Proc. 14th Int. Conf. Transparent Optical Netw. (ICTON 2012)*, Coventry, UK, 2–5 Jul. 2012.

[9] S. Ramamurthy, L. Sahasrabuddhe, and B. Mukherjee, "Survivable WDM mesh networks," *IEEE J. Lightwave Technol.*, vol. 21, no. 4, p. 870, Apr. 2003.

[10] Y. Xiong and L. Mason, "Comparison of two path restoration schemes in self-healing networks," *Comput. Netw.*, vol. 38, no. 5, pp. 663–674, Apr. 2002.

[11] H. Hoang and B. Jaumard, "A new flow formulation for FIPP p-cycle protection subject to multiple link failures," in *Proc. 3rd Int. Workshop Reliable Networks Design and Modeling (RNDM 2013)*, Warsaw, Poland, 5–7 Oct. 2011, pp. 1–7.

[12] M.F. Habib, M. Tornatore, M. De Leenheer, F. Dikbiyik, and B. Mukherjee, "Design of disaster-resilient optical datacenter networks," *IEEE/OSA J. Lightwave Techn.*, vol. 30, no. 16, pp. 2563–2573, Aug. 2012.

[13] I.B. Barla, D.A. ScSchupke, M. Hoffmann, and G. Carle, "Optimal design of virtual networks for resilient cloud services," in *Proc. 9th Int. Conf. Design of Reliable Commun. Netw. (DRCN 2013)*, Budapest, Hungary, 4–7 Mar. 2013, pp. 218–225.

[14] M. Bui, B. Jaumard, and C. Develder, "Anycast end-to-end resilience for cloud services over virtual optical networks (invited)," in *Proc. 15th Int. Conf. Transparent Optical Netw. (ICTON 2013)*, Cartagena, Spain, 23–27 Jun. 2013.

[15] W.V. Heddeghem, W. Vereecken, D. Colle, M. Pickavet, and P. Demeester, "Distributed computing for carbon footprint reduction by exploiting low-footprint energy availability," *Futur. Gener. Comp. Syst.*, vol. 28, no. 2, pp. 405–414, Feb. 2012.