Aalborg Universitet

AALBORG UNIVERSITY
DENMARK

Converged Optical Network and Data Center Virtual Infrastructure Planning

Georgakilas, Konstantinos; Tzanakaki, Anna ; Anastasopoulos, Markos; Pedersen, Jens Myrup

Published in:
Journal of Optical Communications and Networking

DOI (link to publication from Publisher):
[10.1364/JOCN.4.000681](https://doi.org/10.1364/JOCN.4.000681)

Publication date:
2012

Document Version
Accepted author manuscript, peer reviewed version

[Link to publication from Aalborg University](#)

Citation for published version (APA):
Georgakilas, K., Tzanakaki, A., Anastasopoulos, M., & Pedersen, J. M. (2012). Converged Optical Network and Data Center Virtual Infrastructure Planning. *Journal of Optical Communications and Networking*, 4 (9), 681-691. <https://doi.org/10.1364/JOCN.4.000681>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal -

Take down policy

If you believe that this document breaches copyright please contact us at vbn@aub.aau.dk providing details, and we will remove access to the work immediately and investigate your claim.

Converged Optical Network and Data Center Virtual Infrastructure Planning

Konstantinos N. Georgakilas, Anna Tzanakaki, Markos Anastasopoulos, and Jens Myrup Pedersen

Abstract—This paper presents a detailed study of planning virtual infrastructures (VIs) over a physical infrastructure comprising integrated optical network and data center resources with the aim of enabling sharing of physical resources among several virtual operators and services. Through the planning process, the VI topology and virtual resources are identified and mapped to the physical resources. Our study assumes a practical VI demand model without any in advance global knowledge of the VI requests that are handled sequentially. Through detailed integer linear program modeling, two objective functions—one that minimizes the overall power consumption of the infrastructure and one that minimizes the wavelength utilization—are compared. Both are evaluated for the virtual wavelength path and wavelength path optical network architectures. The first objective results in power consumption savings and the two optical network architectures provide similar performances. However, the trend changes for higher load values, due to the inefficient wavelength utilization that the first objective leads to. Finally, we compare the virtual infrastructures created by the two objectives through online traffic provisioning simulations. The objective minimizing wavelength utilization results in VIs suffering higher request blocking compared to the VIs created by the objective minimizing the overall power consumption.

Index Terms—Data center networks; Optical WDM networks; Power consumption; Virtual infrastructure planning.

I. INTRODUCTION

Cloud-based services are being increasingly deployed, taking advantage of the continuous advancements of data centers (DCs), while large-scale service providers like Amazon and Google are increasingly deploying geographically dispersed DCs [1] to satisfy the requirements of the offered services. These services include storage, processing, e-mail, Web services and gaming, whereas private enterprise DCs are also used for data-intensive tasks like Web page indexing and large data-set analysis [2–4]. On-demand self-service, location independence, rapid elasticity, reliability and disaster recovery are some key requirements [5] that the cloud infrastructure

needs to satisfy. It is true to say that cloud infrastructures have emerged as an evolutionary step from computing grid infrastructures, adopting some of the main technologies and approaches used to serve these requirements. These include the most promising optical networking architectures [4] in terms of technology, as well as advanced routing, virtualization, control [4] and joint consideration of network and computing resources [6].

In this context, the concept of virtualization applied on an infrastructure comprising DCs interconnected through a wavelength division multiplexing (WDM) optical network supporting cloud services can offer performance advantages [7] and facilitate sharing of physical resources. This enables the introduction of new business models [8] that suit well the nature and characteristics of the future Internet and enable new exploitation opportunities for the underlying physical infrastructures (PIs). In this environment, service providers are able to establish their own virtual infrastructures (VIs) over the underlying physical infrastructure.

In the converged infrastructure described above there are two kinds of resources: a) network resources including fiber links and nodes and b) DC resources comprising storage, processing cores and memory. The infrastructure planning process is usually formulated as an optimization problem with common objectives dealing with resource utilization. Efficient resource sharing, minimum resource allocation and load balancing are variants of this objective class. Previous work [3,6], mainly addressing grid computing solutions, has already identified the joint consideration of these two kinds of resources as the most effective route towards efficient utilization of the infrastructure. However, in cloud computing, fault management and load balancing usually require all available DCs to be active and the optimization related to resource utilization to take place internally within the DCs.

As information and communication technology (ICT) is estimated to be responsible for about 4% of the primary energy worldwide, a percentage expected to double by 2020 [9], a lot of attention has been recently paid to the energy efficiency of such converged infrastructures [3,8]. Optical networking is an energy-efficient technology that can be further optimized with regard to energy consumption through power-aware network design and protocol implementations [10]. However, the operation of DC resources requires very high levels of power and their conventional operating window is commonly not optimized for energy efficiency. Allocating IT processing jobs in an energy-aware manner through using an optical network with relatively low energy consumption and switching off unused IT resources can offer significant energy savings.

Manuscript received February 23, 2012; revised July 16, 2012; accepted July 23, 2012; published August 20, 2012 (Doc. ID 163623).

Konstantinos N. Georgakilas (e-mail: kg@es.aau.dk, koge@ait.edu.gr) is with the Network Design and Services Group at Athens Information Technology, Peania, Greece, and is also with Aalborg University, Denmark.

Anna Tzanakaki is with the Network Design and Services Group at Athens Information Technology, Peania, Greece.

Markos Anastasopoulos is with the Network Design and Services Group at Athens Information Technology, Athens, Greece.

Jens Myrup Pedersen is with the Networking and Security Section, Aalborg University, Denmark.

Digital Object Identifier 10.1364/JOCN.4.000681

In this paper we focus on the design of virtual infrastructures (VIs) over a physical infrastructure (PI) taking into consideration jointly the network and DC resources. The VIs are slices of the PI comprising subsets of the optical WDM network and DC resources enabling sharing of the available physical resources among several virtual network operators and services [11]. Through the design process, both the topology and the required virtual resources are identified and mapped to the physical resources and the associated operating parameters. In this context, we compare two objectives—one minimizing the joint power consumption of network and DC resources and one minimizing the network resources used. The goal of this comparison is to identify suitable design objectives, tradeoffs and trends for realistic VI request scenarios and a variety of traffic loading conditions. Moreover, we study the impact of the design objectives on the resulting virtual topologies and their performance under dynamic traffic. We make for the first time, to the best of our knowledge, the realistic assumption that there is no global knowledge of the requests for all the VIs. Therefore, we perform the planning of each VI in sequence according to the arrival order of the VI requests over the underlying PI that is already supporting previously established VIs. This work investigates the impact of establishing a set of VIs, over a PI, on the overall network and DC resource utilization, the power consumption and the blocking performance. Power consumption has been chosen to be used as it is of major importance in the ICT, it is relevant to both network and IT resources and it reflects to a large extent the operational costs of running such converged infrastructures.

The planning problems are formulated for the two aforementioned objectives, namely MinJointPower (MJP) and MinNetRes (MNR). Both objectives are evaluated over two network architectures: virtual wavelength path (VWP), where full wavelength conversion is available across all network nodes, and wavelength path (WP) [12], where wavelength continuity is a strict constraint. Finally, the respective VIs generated solving these planning problems are evaluated through online traffic provisioning simulations.

Our results demonstrate that although the MJP objective achieves lower power consumption compared to the MNP as expected, the benefit decreases as the number of established VIs and the volume of demands supported increases. The performance comparison of the different planned VIs shows that the gain in the power consumption offered by the MNR objective introduces a penalty in the blocking performance. The presence of wavelength conversion in the network increases the overall power consumption but improves the blocking performance.

The rest of the paper is structured as follows. Section II provides a review of previous work related to this paper. Section III presents the system adopted for the modeling of the converged network and DC infrastructure and its power consumption. Section IV provides the detailed problem formulations for the infrastructure planning and the online traffic simulation. Section V presents the evaluation results and discussion, and Section VI concludes the paper.

II. RELATED WORK

In [3] the authors provide a comprehensive study of the major approaches for achieving energy efficiency in optical networks. While they present the respective technologies and techniques such as green routing and energy-efficient design in all three network segments (core, metro and access), they also focus on important services running over optical networks in grid and data center networks, identifying the most important approaches for achieving energy efficiency like network connectivity proxying, green TCP/IP protocol design and green grid computing. In the works mentioned, the concepts of anycast routing, equipment switching off and efficient job scheduling are widely used. In [6], the authors focus on grid applications and the joint scheduling of computing and network resources, minimizing the completion time of a job. In [4] the authors present the most important applications that drive the evolution from grid to cloud, like business, scientific, consumer and gaming applications. They also provide their basic characteristics and requirements such as on-demand setup, scalability and elasticity. Then, they present the details of grid and especially cloud computing and finally explain how optical networking characteristics related to technology, routing, virtualization and control are able to serve the requirements of these demanding paradigms and applications. In [13] the authors demonstrate the efficiency of joint network and IT consideration in terms of power consumption over the physical infrastructure and demonstrate a benefit of 3%–55%. This benefit depends on the ability of a data center to switch on/off servers and is compared to the case where only network power consumption is minimized. In [8], a model describing the concept of planning virtual infrastructures over a converged network and IT infrastructure is presented and an energy-aware VI planning problem is formulated and compared to an approach that allocates the IT services to the server located closest to the source node, providing savings of the order of 30% for a single VI establishment. Moreover, energy-aware offline service provisioning for the VWP case on top of the planned VI is presented and compared to three other approaches, achieving minimum overall power consumption. Finally, in [14] the authors propose multiple VI energy-aware planning, assuming global knowledge of the VI requests, and provide the optimal solution for the establishment of all VIs concurrently.

III. THE SYSTEM MODEL

A. The Network Model

The network is modeled as a graph, comprising a set of nodes N interconnected by a set of unidirectional links L . The nodes are optical cross-connects (OXCs) based on a photonic switching matrix that is realized by 3D microelectromechanical systems (MEMS) [15].

Each node supports M input and M output fibers, each employing a maximum number of wavelengths W . Apart from the passive elements, which are the multiplexers (MUX) and de-multiplexers (DEMUX), Fig. 1 illustrates the active

Fig. 1. Optical cross-connect architecture.

elements of the OXC: the switch matrix, one erbium-doped fiber amplifier (EDFA) per input/output fiber port, one optical–electrical–optical (OEO) transponder per output wavelength port and one transmitter (Tx)–receiver (Rx) pair per lightpath. The OEO transponders are used to support the wavelength conversion functionality in the case of the VWP network. The number of through (express) ports is calculated as the number of input fibers M times the fiber wavelength capacity W . It is further assumed that the add/drop capability of the node is 50% of the through traffic. No OEO converters are included in the WP network architecture.

Figure 2 illustrates the link architecture [16] employed for the interconnection of the OXCs. It is modeled as a sequence of alternating single-mode fiber (SMF) and dispersion compensation fiber (DCF) spans, to address fiber dispersion effects including pre-compensation and post-compensation DCF spans at the beginning and the end of each link, respectively. To compensate for the insertion loss of the fiber spans, optical amplifiers based on EDFA technology are allocated at the end of each transmission span.

B. The Data Center Model

The main building block of the data center model is based on [17]. It is a full rack implementation of a hardware platform that is used in real DCs and its main characteristics are summarized in Table I. For the data center throughput, we have assumed it to be equal to four times the 75 GB/s per rack uncompressed I/O bandwidth reported in [18].

C. Power Consumption Models

1) *The Network*: The total network power consumption is determined by the power consumption of OXCs and fiber links. Figure 1 illustrates the OXC architecture and its

Fig. 2. Fiber link model.

TABLE I
DATA CENTER BUILDING BLOCK CHARACTERISTICS

Resource	Type	Capacity
CPU	30 servers/360 cores	2.93 GHz Intel Xeon six-core processor
Storage		2.9 TB
T_{DC}		2.4 Tbps
Memory	1333 MHz	40 TB

TABLE II
NETWORK EQUIPMENT POWER CONSUMPTION FIGURES

Symbol	DESCRIPTION	Power (W)
P_{port_pair}	Input/output port pair of the switch fabric	0.107
$P_{Transponder}$	O/E/O: Line-side WDM Transponder (10G)	6
$P_{Tx/Rx}$	E/O, O/E: Transmitter or receiver	3.5
P_{EDFA}	EDFA	13

power-dissipating elements with gray shading. The node power consumption (P_{OCX}) depends on four factors: (a) the switch fabric (P_{SF}), the transponders for (b) transmission (P_{Transm}) and (c) wavelength conversion (P_{Conv}) and (d) the optical amplifiers (P_{Amp}). Equations (1)–(4) demonstrate the computation of these power consumption values, whereas Table II provides a short description and typical power consumption values for the required equipment [19,20].

$$P_{SF} = \text{ports}_{total} \times P_{port_pair} = (\text{ports}_{th} + \text{ports}_{a/d}) \times P_{port_pair}, \quad (1)$$

$$P_{Transm} = \text{ports}_{a/d} \times P_{Tx/Rx}, \quad (2)$$

$$P_{Conv} = \text{ports}_{th} \times P_{Transponder}, \quad (3)$$

$$P_{Amp} = (f_{in} + f_{out}) \times P_{EDFA}. \quad (4)$$

In our case, we assume a symmetric switch, $M \times M$, whereas M is not identical across the OXCs, but computed after the planning process based on the traffic volume supported.

The only power-consuming elements included within the optical network links in Fig. 2 are the optical amplifiers installed per span. The amplifier span length (span) is assumed to be 80 km. Thus the power consumption P_l of a fiber link l is length dependent and is calculated as depicted in Eq. (5):

$$P_l = \left\lceil \frac{\text{length}(l)}{\text{span}} \right\rceil P_{EDFA}. \quad (5)$$

The total network power consumption of the physical infrastructure is computed by using Eq. (6):

$$P_{Net} = \sum_{n \in N} P_{OCX_n} + \sum_{l \in L} P_l. \quad (6)$$

TABLE III
DATA CENTER POWER CONSUMPTION VALUES

Symbol	Description	Value (kW)
P_{busy}	Power consumption under full utilization	17.5
P_{idle}	Power consumption under idle state	8.75

2) *The Data Center*: The power consumption of the data center is based on typical power consumption values taken from [17] and on the simple linear model illustrated in Eq. (7):

$$P_{\text{DC}} = P_{\text{idle}} + (P_{\text{busy}} - P_{\text{idle}})u_s. \quad (7)$$

P_{idle} is the power consumption of the DC in the idle state, which is when no server is utilized. Its value is considered to be 50% of the DC power consumption under full utilization, P_{busy} . In our model, this value is assumed to be the maximum power consumption value from [17], as shown in Table III. Finally, u_s represents the utilization of a data center s that is defined as the sum of lightpath requests currently served by the data center s over the DC throughput expressed in wavelengths. Equation (8) illustrates the computation of the DC utilization:

$$u_s = \frac{\text{number of lightpaths arriving at } s}{T_{\text{DC}}}. \quad (8)$$

Table III summarizes the power consumption values used for each DC, assuming the same configuration and capacity for all DCs employed in the converged infrastructure.

IV. PROBLEM FORMULATION

A. Virtual Infrastructure Planning

The virtual infrastructure planning problem is formulated as an integer linear program (ILP) and is based on the well-studied routing and wavelength assignment (RWA) [12] problem. As also defined in the introduction section, the term PI refers to the physical infrastructure, which is the set of DCs and WDM nodes and links that interconnect them. The term VI refers to the virtual infrastructure, which is a slice of the PI comprising a set of DC and network resources in terms of DC throughput and wavelengths, respectively. In this work we concentrate on optimal planning of virtual infrastructures with respect to specific objectives, while further virtualization implementation details are not taken into account. More specifically we provide two sets of formulations using path and flow variables, one for the virtual wavelength path and one for the wavelength path case.

1) *The Virtual Wavelength Path*: Let $G = (N, L)$ be the directed graph that represents the network topology, where N is the set of network nodes (OXCs) and L is the set of directed links interconnecting the nodes. Let also S be the subset of N that represents the nodes where data centers are attached and D be the set of requests for each virtual infrastructure. Multiple D sets, each one corresponding to a VI

request, are treated sequentially in the order of arrival of the VI requests without any prior knowledge of the VI requests. We solve one optimization problem for each virtual infrastructure that needs to be formed and established over the converged physical infrastructure. The solution of each problem updates the network and DC capacity and gives a new instance of the infrastructure as an input to the next problem. The solution of the problem provides an optimal mapping of the virtual infrastructure to the physical infrastructure, in terms of both network topology and infrastructure resources, according to the objective function.

The planning problem considered in this work is treated as a capacitated problem. We start with a given DC capacity offered over the entire infrastructure and a relative network capacity that is sufficient to accommodate the traffic requests aiming to utilize the DC resources. The definition of the demand volume needs to include the amount of resources requested, in terms of both the network and the data center. The characteristics that enable the definition of an exact relationship between these two kinds of resources are the DC throughput T_{DC} and the wavelength bit rate R_w . Equation (9) gives the number of wavelengths k that are needed to satisfy the DC throughput:

$$k = \frac{T_{\text{DC}}}{R_w}. \quad (9)$$

Each demand d of the set D is described by the source node and by the constant h_d , which is the demand volume assuming wavelength level granularity. For each virtual infrastructure problem, the demand sources are randomly selected among all the network nodes apart from the ones that are directly connected to DCs. Each problem is solved for a set of total lightpath requests that span from 30 to 150, whereas for each such value the reported results correspond to averages over a number of repetitions.

We refer to physical topology links (or physical links) by the index l ($l = 1, \dots, L$) and to virtual topology links (or virtual links) by the index e ($e = 1, \dots, E$). The physical link set L is straightforward to obtain since it represents all the directed links of the graph G . The virtual link set E is assumed to be the set of links in a full mesh graph with the same set of nodes N . Moreover, network nodes are indexed by n over the set N ($n = 1, \dots, N$) and data centers are indexed by $s = 1, \dots, S$. Demands are also indexed by $d = 1, \dots, D$. Finally, physical and virtual candidate path lists are indexed by q ($q = 1, \dots, Q_e$) and p ($p = 1, \dots, P_d$), respectively.

Besides the demand volume h_d , we use two more constants in the problem formulation. As already mentioned, the destination is not part of the demand description, since all DCs with available capacity are candidate destinations for every request. Thus for every demand and every candidate destination, a candidate path list is generated on the basis of k -shortest-paths routing. We refer to this as the virtual candidate path list and use the parameter δ_{edp} (Eq. (10)) to represent the use of candidate virtual path p in the realization of virtual link e for demand d :

$$\delta_{edp} = \begin{cases} 1, & \text{if candidate virtual path } p \text{ of demand } d \\ & \text{is used to realize virtual link } e \\ 0, & \text{otherwise.} \end{cases} \quad (10)$$

TABLE IV
PROBLEM VARIABLES

Variable	Domain	Description
x_{dsp}	Z_+	Flow realizing demand d towards CD s on candidate virtual path p (number of lightpaths)
z_{eq}		Flow realizing virtual link e on candidate physical path q
w_e	Z_+	Number of wavelengths utilized on virtual link e
y_l	Z_+	Number of wavelengths utilized on physical link e
f_s	$\{0, 1\}$	$\begin{cases} 1, & \text{if DC } s \text{ is used} \\ 0, & \text{otherwise} \end{cases}$

Since we need to map the virtual to the physical network resources, we use one more constant to illustrate the realization of virtual link e by the corresponding physical path q , illustrated in Eq. (11):

$$\gamma_{leq} = \begin{cases} 1, & \text{if physical link } l \text{ is part of the candidate} \\ & \text{physical path } q \text{ of the virtual link } e \\ 0, & \text{otherwise.} \end{cases} \quad (11)$$

For both candidate path lists, we use a length-based Yen's k -shortest-paths algorithm [21] and retrieve the two shortest paths from every source to all candidate destinations.

Table IV summarizes the variables whose optimal values form the solution of the ILP and represent the resource assignment that accommodates the requested traffic and provides the mapping between physical and virtual resources.

Two different optimization problems are formulated by using two objective functions, namely, "NetRes" (Eq. (12)) and "JointPower" (Eq. (13)). As the names indicate, the objectives minimize the total number of wavelengths (representing the network resources) used and the total power consumption of both network and DC resources, respectively. On the basis of the problem variables and the power consumption models presented in Subsection III.C, the two objectives are

$$\text{NetRes} = \sum_l w_l \text{length}(l), \quad (12)$$

$$\begin{aligned} \text{JointPower} = & \sum_s [(1 - f_s)P_{\text{idle}} + (P_{\text{busy}} - P_{\text{idle}})u_s] \\ & + \sum_n \sum_{l \in \omega^-(n)} \left[(1 - \alpha_l)P_{\text{link}} + \alpha_l P_{\text{link}} \left(\frac{w_l}{\sum_l \Lambda_l} + 1 \right) \right] \\ & + \sum_n \sum_{l \in \omega(n)} w_l \frac{1}{2} P_{\text{Transponder}} + \sum_n \sum_{l \in \omega^-(n)} w_l P_{\text{Transponder}}, \end{aligned} \quad (13)$$

where

$$P_{\text{link}} = \sum_n \sum_{l \in \omega^-(n)} \frac{1}{W} w_l \left(\left\lfloor \frac{\text{length}(l)}{\text{span}} \right\rfloor + 2 \right) P_{\text{EDFA}}, \quad (14)$$

$$f_s = \begin{cases} 1, & \text{if DC } s \text{ is already used by another VI} \\ 0, & \text{otherwise} \end{cases}, \quad (15)$$

and

$$\alpha_l = \begin{cases} 1, & \text{if link } l \text{ is already used by another VI} \\ 0, & \text{otherwise.} \end{cases} \quad (16)$$

The indices $\omega^+(n)$ and $\omega^-(n)$ represent the outgoing and incoming links of node n , respectively. Λ_l is the number of wavelengths on link l already established by previous VIs. As indicated by Eq. (13), the power consumption of the optical links is calculated as follows: when a VI is the first to utilize a link, the total power consumption of the link is assigned to this VI. When more than one VI uses the same link, each VI is assigned proportionally a power consumption level reflecting the utilization of the corresponding link resources (wavelengths).

The constraints that follow complete the ILP formulation and ensure that network and DC resources are optimally assigned, following flow conservation and capacity rules. The constraint in Eq. (17) ensures that all demands d are served:

$$\sum_s \sum_p x_{dsp} = h_d, \quad d = 1, \dots, D. \quad (17)$$

The next three constraints (18)–(20) deal with capacity bounds and require that the virtual and physical link capacities, w_e and y_l , respectively, are enough to accommodate all lightpaths, noting that these two variables are upper bounded by the fiber capacity and that, based on commercially available WDM products, is assumed to be 96 wavelengths. Moreover, they provide the mapping between physical and virtual links. The mapping is performed through constraints (18) and (19), which map the virtual link capacities to the physical link capacities through the flow variable z_{eq} .

$$\sum_d \sum_s \sum_p \delta_{edp} x_{dsp} \leq w_e, \quad e = 1, \dots, E, \quad (18)$$

$$\sum_q z_{eq} \leq w_e, \quad e = 1, \dots, E, \quad (19)$$

$$\sum_e \sum_q \gamma_{leq} z_{eq} \leq y_l, \quad l = 1, \dots, L. \quad (20)$$

Finally, two more constraints ensure the correct assignment of the binary variable f_s that represents whether a data center is used or not:

$$f_s \leq u_s, \quad s = 1, \dots, S, \quad (21)$$

$$f_s T_{\text{DC}} \geq u_s, \quad s = 1, \dots, S. \quad (22)$$

The utilization of the DC u_s is defined as

$$u_s = \sum_d \sum_p x_{dsp}, \quad s = 1, \dots, S. \quad (23)$$

2) *The Wavelength Path*: The path formulation for the WP case follows the same principles and notation as the VWP case described above with the addition of the index c that represents the distinct wavelengths of a fiber link. The integer flow variable of the problem that indicates the number of lightpaths using the virtual path p to support demand d that will be serviced by the corresponding DC s is now updated to

x_{dspc} and indicates the assignment of a specific wavelength c across the path. Accordingly, the flow variable z_{eq} is updated to z_{eqc} . The rest of the variables, indices and constants remain the same.

The constraints of the problem are updated to ensure the correct assignment of capacity across the paths for the establishment of all the demands: constraint (24) ensures the establishment of all the demands, whereas constraints (25)–(27) assign to virtual and physical links the capacity required for accommodating the traffic flows:

$$\sum_s \sum_p \sum_c x_{dspc} = h_d, \quad d = 1, \dots, D, \quad (24)$$

$$\sum_d \sum_s \sum_p \sum_c \delta_{edp} x_{dspc} \leq w_e, \quad e = 1, \dots, E, \quad (25)$$

$$\sum_q \sum_c z_{eqc} \leq w_e \quad e = 1, \dots, E, \quad (26)$$

$$\sum_e \sum_q \sum_c \gamma_{leq} z_{eqc} \leq y_l, \quad l = 1, \dots, L. \quad (27)$$

Constraint (28) ensures that each distinct wavelength on a physical path q realizing a virtual link e is assigned only to one flow:

$$\sum_q \gamma_{leq} z_{eqc} \leq 1, \quad e = 1, \dots, E, l = 1, \dots, L, c = 1, \dots, C. \quad (28)$$

Constraints (21) and (22) of the VWP formulation remain the same. Finally, the data center utilization (Eq. (29)) is expressed as the sum of lightpaths arriving at DC s :

$$u_s = \sum_d \sum_p \sum_c x_{dspc}, \quad s = 1, \dots, S. \quad (29)$$

The optimization objectives remain as defined for the VWP case.

B. Online Traffic Provisioning

The virtual infrastructures planned through the integer programs based on the two objective functions presented previously are further evaluated under an online traffic provisioning scenario. The simulations are based on a custom Matlab tool [22] and traffic is modeled as a Poisson arrival process with exponential service time distribution. Requests are randomly generated from all sources, apart from the nodes directly connected to DCs. The traffic load spans from 30 to 80 erlangs.

The goal of these simulations is to evaluate the performance of the different VIs in terms of request blocking. The request destination is a decision based on the closest-DC scheme, where for every request we choose as destination the DC that can be reached through a length-based shortest path with available capacity. The granularity of the requests is one wavelength. The performance evaluation takes place for virtual infrastructures that correspond to both VWP and WP planning problem solutions and both optimization objectives, as presented in Section IV. For the WP case, the “First Fit” wavelength assignment algorithm is used.

Fig. 3. COST 239-based converged network and data center infrastructure.

V. EVALUATION RESULTS

To evaluate the different VI planning approaches described above, we assume the converged network and DC infrastructure presented in Fig. 3. The reference network topology used is based on the COST 239 pan-European network [23] interconnecting four data centers. One fiber link per direction ensures bi-directional communication with a fiber capacity of 96 wavelengths. The data center throughput in terms of wavelengths is $\frac{2.4 \text{ Tbps}}{10 \text{ Gbps}} = 240$ wavelengths. The mapping of the data center and network resources requires that the four DCs' total capacity is $240 \times 4 = 960$ lightpaths and our traffic demands reach the level of 150 lightpath requests per VI, and thus 600 lightpaths for the case where all four VIs are present. We assume four virtual infrastructure request sets, each one supporting a total number of lightpath requests ranging from 30 to 150. Moreover, we assume that only active (utilized) network and data center resources consume power, whereas the respective granularity is the fiber link and the DC as a whole.

To address the statistical uncertainty associated with the input traffic, we report confidence interval (CI) limits of the mean network utilization for a confidence level of 95%. Each value corresponds to a specific VI and lightpath load combination and is computed as the lower and upper ($\pm$) CI limit over the mean network utilization for 50 repetitions. For each repetition, we generate a set of lightpath requests that sum up to the corresponding loading value and the source nodes are randomly selected according to a uniform distribution. Table V provides the corresponding values for the case of wavelength path formulation and the MNR objective. All other formulations result in similar values that always remain in the range of 3%–15%.

A. Power Consumption

The planning problems presented in this paper model two different objective functions (MJP and MNR) under two

TABLE V
CI LIMITS OF MEAN POWER CONSUMPTION FOR 95%
CONFIDENCE LEVEL

VI	Lightpaths			
	30	60	90	120
1	0.1026	0.0282	0.0129	0.0147
2	0.0335	0.0279	0.0129	0.0148
3	0.0335	0.0280	0.0128	0.0145
4	0.0334	0.0278	0.0128	0.0143

network technologies (VWP and WP). This study aims at investigating: a) the impact of the VI model and request model on the power consumption of the infrastructure and b) the impact of the objective function on the power consumption when VIs are sequentially planned for both VWP and WP networks. These issues are illustrated in Figs. 4–14 and are analyzed below.

a) Figures 4 and 7 demonstrate two result sets for the VWP and WP cases, respectively, and concentrate on the total power consumption corresponding to the planned VIs for each objective and for all four VIs and the breakdown of power consumption on DC and network parts. We first observe that the MJP objective achieves lower power consumption values across lower traffic volumes, as expected, whereas a change of this trend is observed for high numbers of VIs and load values.

Focusing on the VWP case, it is observed (Fig. 4) that the MJP objective achieves significantly lower power consumption for 120 and 240 lightpaths (corresponding to 30 and 60 lightpaths per VI), very similar values with the MNR objective for 360 and 480 lightpaths, and finally higher power consumption for 600 lightpaths. The breakdown of network and DC power consumption provides a more detailed understanding of how network and DC resources need to be powered up based on the output of the two objectives and explains why the MJP objective does not lead to the optimal power consumption across all load values and VI numbers. In this context, there are two main observations:

1) The MJP demonstrates almost constant network power consumption across all load values (small variations are observed due to statistical error), which is also much higher than the respective network power consumption when applying the MNR objective. This is verified in Fig. 5, which demonstrates the average lightpath length for the two objectives.

2) At the same time, the DC power consumption of the MJP objective is much lower than that of MNR for lower load values and grows to the same level for higher load values. This is clearly verified in Fig. 6, where the number of powered up DCs versus the number of requests is illustrated. We observe that the MJP objective causes powering up of DCs only when the already powered up DCs are not sufficient to accommodate the new requests. Since the DCs are the most energy-consuming elements of the infrastructure, the decision to power up DCs plays an important role in the optimization process and leads to higher average lightpath length.

These results have demonstrated that the MJP objective gradually causes powering up of DCs and introduces longer

Fig. 4. Power consumption contributions—VWP.

Fig. 5. Average lightpath length—VWP.

Fig. 6. Number of data centers powered up.

paths to reach them. Therefore, the objective of achieving lower power consumption leads to over-utilization of the network resources. On the other hand, the MNR objective achieves better network utilization, as it aims at minimizing the total path length. For high load values, all available DCs need to be powered up for both objectives. However, in this case, applying the MNR achieves lower network resource utilization.

In Fig. 7 we demonstrate the same set of results for a WDM network without wavelength conversion capability. The results acquired applying the two different objectives are similar to those for the VWP case, demonstrating a similar trend, where the MJP achieves lower power consumption across most load values and the trend is changing only for the highest load value of 600 lightpaths. The difference of the absolute values

Fig. 7. Power consumption contributions—WP.

Fig. 8. Total power consumption versus number of VIs—VWP—30–60 lightpaths.

is attributed to the absence of wavelength converters, which significantly reduces the network power consumption.

b) Aiming at a more detailed evaluation of the objective impact on the power consumption of the converged infrastructure, we provide a set of results that represent the total power consumption of the infrastructure across the VIs established over the physical infrastructure and for three different lightpath load volumes. Figures 8 and 9 present these results for 30–60 and 120–150 lightpaths, respectively, for the virtual wavelength path case.

The MJP objective achieves lower power consumption only for one or two planned VIs, whereas for higher number of VIs, MNR outperforms MJP for both low (30–60) and high (120–150) lightpath requests per VI. The observed change of trend is attributed to the fact that planning of multiple VIs is not based on a global optimization performed having in advance knowledge of the VI requests. Instead a practical planning approach is adopted, according to which the planning procedure takes place for each VI request sequentially, considering that all previous VIs remain established and utilize the already assigned resources. Finally it should be noted that the addition of wavelength converters in the VWP case, assumed to employ OEO transponder technology in this study, significantly increases the overall power consumption of the optical network. This has an impact on the relative proportion of power consumption of the network, as well as the data center resources and the associated tradeoffs.

These observations clearly indicate that in order to maintain the energy efficiency benefit that MJP can offer when planning VIs over a PI, which are dynamically requested in time, there

Fig. 9. Total power consumption versus number of VIs—VWP—120–150 lightpaths.

Fig. 10. Total power consumption versus number of VIs—WP—30–60 lightpaths.

is a need to reconsider the existing resource allocation per VI periodically or following specific triggering events. This can then be followed by suitable reallocation of resources per VI through a VI re-planning phase. If there is a requirement for this reallocation of resources not to be disruptive for services that are already supported, relevant constraints in the VI re-planning process can be applied.

Figures 10 and 11 provide the same set of results for the WP case, where wavelength converters are not present at the network nodes and the corresponding planning problem presented in Subsection IV.A.2 takes into account the fact that the same wavelength has to be assigned across each lightpath established on the physical topology, known as the wavelength continuity constraint. The graphs illustrating the power consumption for the different load values across the VIs are in accordance to the results discussed above and indicate similar performances for the two objectives and the WP case.

B. Blocking Performance

Figure 12 depicts the network resource utilization in terms of the number of link wavelengths utilized by the VIs for the two objectives. The results are presented across all lightpath volumes and after all five virtual infrastructures have been established. For the WP network, MNR achieves link wavelength utilization of the order of 3%–13% compared to 8%–30% for the VWP network. MJP achieves a similar benefit, since no difference in the network resource assignment is observed between the two formulations for the specific problem.

Fig. 11. Total power consumption versus number of VIs—WP—120–150 lightpaths.

Fig. 12. Network capacity utilization.

As already mentioned and observed through the investigation of power consumption contributions from the network and data center resources, the two objectives result in different virtual network infrastructures. The average number of wavelengths per fiber link is the same, but the node degrees and total numbers of links used differ. MNR planned VIs exhibit a uniform node degree of 2 across all load volumes, whereas VIs created by the MJP objective exhibit node degrees of 3 due to the selection of longer paths.

On the basis of the VIs generated as described in detail above, we provide results that evaluate the performance of an online routing algorithm over the established planned VIs. We compare two virtual infrastructures differing in the node degree under both wavelength-converted and wavelength-continuity-constrained optical networks, according to the simulation model described in Subsection IV.B.

Figure 13 depicts the blocking performance of the two VIs in a wavelength-converted network (VWP) and quantifies the performance gain of the MJP category of VIs. The VI with node degree 3 (MJP) reaches a blocking probability value of 5%. The efficiency of the MNR objective in the utilization of network resources results in much higher blocking values that reach the level of 21% for the highest loading value.

In this case where wavelength assignment is applied at each network node, both VI types experience higher demand blocking due to the increased probability for a path to be discarded because of the wavelength continuity constraint. Figure 14 illustrates blocking values of up to 7% for the MJP VI and 23% for the MNR VI. A much greater difference in the blocking probability between VWP and WP networks is observed for lower loading values and for both objectives.

Fig. 13. Blocking probability versus load—VWP.

Fig. 14. Blocking probability versus load—WP.

VI. CONCLUSION

In this paper, a detailed study of planning virtual infrastructures over a physical infrastructure comprising integrated optical network and data center resources was presented. The study assumed a practical VI demand model that did not support any in advance global knowledge of the VI requests and, through detailed ILP modeling, compared two different objective functions—the MJP and MNR—as well as two different optical network architectures—one supporting the VWP and one supporting the WP. The various scenarios under study were compared as regards power consumption, network utilization and blocking performance of the planned VIs. Our results illustrated that although power consumption is an important aspect and underpins an objective function (MJP) that optimizes the energy efficiency of the infrastructure, it may introduce inefficiencies in the utilization of network resources when the number of requests exceeds a certain level. This may in turn compromise the benefit as regards energy efficiency, compared to what is achieved when applying an objective that minimizes the network resource utilization for these high demand levels. To overcome this inefficiency, periodic re-planning of the requests can be applied. Finally, a set of dynamic traffic provisioning results were provided through simulations illustrating that the efficient resource utilization of the second objective (MNR) introduces a penalty on the VIs produced, especially in terms of connectivity, which leads to poor request blocking performance.

ACKNOWLEDGMENTS

This work was carried out with the support of the GEYSERS (FP7-ICT-248657) project funded by the European Commission through the 7th ICT Framework Program. Part of this work was also performed in the context of collaboration between AAU/CTiF and AIT for the joint offering of a Doctoral Program.

REFERENCES

- [1] <http://www.google.com/about/datacenters/locations/index.html>.
- [2] T. Benson and A. Akella, "Network traffic characteristics of data centers in the wild," in *IMC*, 2010, pp. 267–280.
- [3] Z. Yi, P. Chowdhury, M. Tornatore, and B. Mukherjee, "Energy efficiency in telecom optical networks," *IEEE Commun. Surv. Tutorials*, vol. 12, no. 4, pp. 441–458, Fourth Quarter 2010.
- [4] 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.
- [5] ADVA Optical Networking, "Virtualized optical networks for sustainable cloud services," White Paper.
- [6] X. Liu, W. Wei, C. Qiao, T. Wang, W. Hu, W. Guo, and M.-Y. Wu, "Task scheduling and lightpath establishment in optical grids," in *27th IEEE Conf. on Computer Communications (INFOCOM)*, 13–18 Apr. 2008, pp. 1966–1974.
- [7] K. J. (Ken) Salchow Jr., "Data center consolidation: Know where you're going and why," F5 White Paper.
- [8] A. Tzanakaki, M. Anastasopoulos, K. Georgakilas, J. Buysse, M. De Leenheer, C. Develder, S. Peng, R. Nejabati, E. Escalona, D. Simeonidou, N. Ciulli, G. Landi, M. Brogle, A. Manfredi, E. Lopez, J. F. Riera, J. A. Garcia-Espin, P. Donadio, G. Parladori, and J. Jimenez, "Energy efficiency in integrated IT and optical network infrastructures: The GEYSERS approach," in *IEEE Conf. on Comp. Communications Workshops (INFOCOM WKSHPS)*, 10–15 Apr. 2011, pp. 343–348.
- [9] M. Pickavet, W. Vereecken, S. Demeyer, P. Audenaert, B. Vermeulen, C. Develder, D. Colle, B. Dhoedt, and P. Demeester, "Worldwide energy needs for ICT: The rise of power-aware networking," in *2nd Int. Symp. on Advanced Networks and Telecommunication Systems (ANTS)*, 15–17 Dec. 2008, pp. 1–3.
- [10] M. Mellia Chiaraviglio and F. Neri, "Reducing power consumption in backbone networks," in *IEEE ICC*, 2009.
- [11] E. Escalona, S. Peng, R. Nejabati, D. Simeonidou, J. A. Garcia-Espin, J. Ferrer, S. Figuerola, G. Landi, N. Ciulli, J. Jimenez, B. Belter, Y. Demechenko, C. de Laat, X. Chen, A. Yukan, S. Soudan, P. Vicat-Blanc, J. Buysse, M. De Leenheer, C. Develder, A. Tzanakaki, P. Robinson, M. Brogle, and T. M. Bohnert, "GEYSERS: A novel architecture for virtualization and co-provisioning of dynamic optical networks and IT services," in *Future Network & Mobile Summit (FutureNetw)*, 15–17 June 2011, pp. 1–8.
- [12] W. D. Grover, *Mesh-Based Survivable Networks: Options and Strategies for Optical, MPLS, SONET and ATM Networking*. Prentice Hall, Englewood Cliffs, NJ, 2003.
- [13] J. Buysse, K. Georgakilas, A. Tzanakaki, M. De Leenheer, B. Dhoedt, C. Develder, and P. Demeester, "Calculating the minimum bounds of energy consumption for cloud networks," in *Proc. of the 20th Int. Conf. on Computer Communications and Networks (ICCCN)*, 31 July 2011–4 Aug. 2011, pp. 1–7.
- [14] A. Tzanakaki, M. P. Anastasopoulos, K. Georgakilas, and D. Simeonidou, "Energy aware planning of multiple virtual infrastructures over converged optical network and IT physical resources," in *ECOC*, 2011, pp. 1–3.
- [15] M. Murakami, "Analyzing power consumption in optical cross-connect equipment for future large-capacity optical networks," *J. Networks*, vol. 5, no. 11, pp. 1254–1259, Nov. 2010.
- [16] K. M. Katrinis and A. Tzanakaki, "On the dimensioning of WDM optical networks with impairment-aware regeneration," *IEEE/ACM Trans. Netw.*, vol. 19, no. 3, pp. 735–746.
- [17] Oracle, *Oracle exalogic elastic cloud X2-2 data sheet*, Mar. 2011 [Online]. Available: <http://www.oracle.com/us/products/middleware/exalogic/index.html>.
- [18] <http://www.oracle.com/technetwork/database/exadata/dbmachine-x2-8-datasheet-173705.pdf>.
- [19] S. Aleksic, "Analysis of power consumption in future high-capacity network nodes," *J. Opt. Commun. Netw.*, vol. 1, no. 3, pp. 245–258, 2009.
- [20] A. Tzanakaki, K. Katrinis, T. Politi, A. Stavdas, M. Pickavet, P. Van Daele, D. Simeonidou, M. O'Mahony, S. Aleksić, L. Wosinska, and P. Monti, "Dimensioning the future pan-European optical network with energy efficiency considerations," *J. Opt. Commun. Netw.*, vol. 3, no. 4, pp. 272–280, Apr. 2011.
- [21] J. Y. Yen, "Another algorithm for finding the K shortest-loopless network paths," *Bull. Oper. Res. Soc. Am.*, vol. 20, p. B/185, 1972.
- [22] K. N. Georgakilas, K. Katrinis, A. Tzanakaki, and O. B. Madsen, "Performance evaluation of impairment-aware routing under single and double-link failures," *J. Opt. Commun. Netw.*, vol. 2, no. 8, pp. 633–641, Aug. 2010.
- [23] P. Batchelor, B. Daino, P. Heinzmann, D. R. Hjelm, R. Inkret, G. H. A. Jäger, M. Joindot, A. Kuchar, E. Le Coquil, P. Leuthold, G. de Marchis, F. Matera, B. Mikac, H.-P. Nolting, J. Späth, F. Tillerot, B. Van Caenegem, N. Wauters, and C. Weinert, "Study on the implementation of optical transparent transport networks in the European environment—Results of the research project COST 239," *Photonic Network Commun.*, vol. 2, no. 1, pp. 15–32, 2000.

Konstantinos N. Georgakilas earned the diploma in Telecommunications Science and Technology from the Telecommunications Science and Technology Department, University of Peloponnese, Tripolis, Greece, in 2007 and the M.Sc. in Information Networking from Carnegie Mellon University, Pittsburgh, PA, and Athens Information Technology in 2009.

He is currently a Research Assistant in the Network Design and Services (NDS) group at Athens Information Technology, and he is pursuing his Ph.D. degree from Aalborg University. His main research interests include optical networking, network optimization and traffic engineering.

Anna Tzanakaki is an Associate Professor at Athens Information Technology (AIT), where she leads the Network Design and Services (NDS) research group. She is also an adjunct faculty member of the Information Networking Institute of Carnegie Mellon University, Pittsburgh, PA.

She has obtained a B.Sc. degree in Physics from the University of Crete, Greece, and an M.Sc. and a Ph.D. both from the University of Essex, UK. She is a co-author of over 120 publications in international journals and conferences. She is a co-inventor of one granted and eleven published patents. She is a technical referee for various journals and conferences. Her research interests include resource and energy-efficient optical networks, cross-layer resilient network design and traffic provisioning.

Dr. Tzanakaki is a senior member of the IEEE and has served as a member of several Technical Program Committees.

Markos P. Anastasopoulos received the Diploma degree in Electrical and Computer Engineering, the M.Sc. in Techno-economics and the Dr.Eng. degree from the National Technical University of Athens, Athens.

Since October 2010 he has been a research scientist at the Network Design and Services group of AIT, working for EU FP7 funded projects. His main research interests lie in the areas of optical and wireless communication

networks, mobile and distributed computing, network design and management. He is a co-author of over 30 papers in international journals and conference proceedings.

Dr. Anastasopoulos has been awarded scholarships for his academic achievements by the Kyprianides, Eugenides and Propondis founda-

tions. He is a TPC member and co-chair for several international conferences and also an Associate Editor for Wiley's *International Journal of Communication Systems*.

Jens Myrup Pedersen is an Associate Professor and Head of the Networking and Security Section, Department of Electronic Systems, Aalborg University. His current research interests include network planning, traffic monitoring, and network security. He obtained his M.Sc. in Mathematics and Computer Science from Aalborg University in 2002 and his Ph.D. in Electrical Engineering also from Aalborg University in 2005. He is author/co-author of

more than 70 publications in international conferences and journals and has participated in Danish, Nordic, and European funded research projects. He is also a board member for a number of companies within technology and innovation.