

Energy efficiency considerations in integrated IT and optical network resilient infrastructures

A. Tzanakaki¹, M. Anastasopoulos¹, K. Georgakilas¹, J. Buysse², M. De Leenheer², C. Develder², S. Peng³, R. Nejabati³, E. Escalona³, D. Simeonidou³, N. Ciulli⁴, G. Landi⁴, M. Brogle⁵, A. Manfredi⁵, E. Lopez⁶, J. Ferrer Riera⁶, J. A. García-Espín⁶, P. Donaldio⁷, G. Parladori⁷, J. Jimenez⁸, A. Tovar De Duenyas⁸, P. Vicat-Blanc⁹, J. van der Ham¹⁰, C. de Laat¹⁰, M. Ghijsen¹⁰, B. Belter¹¹, A. Binczewski¹¹, M. Antoniak-Lewandowska¹²

¹Athens Information Technology (AIT), Greece, Email: atza@ait.edu.gr, ²IBBT, Belgium, ³University of Essex, UK, ⁴Nextworks, Italy, ⁵SAP, Switzerland, ⁶i2cat, Spain, ⁷Alcatel-Lucent, Italy, ⁸Telefonica, Spain, ⁹Lyatis, France, ¹⁰University of Amsterdam, Netherlands, ¹¹Poznań Supercomputing and Networking Center, Poland, ¹²Telekomunikacja Polska S.A, Poland

ABSTRACT: The European Integrated Project GEYSERS - Generalised Architecture for Dynamic Infrastructure Services - is concentrating on infrastructures incorporating integrated optical network and IT resources in support of the Future Internet with special emphasis on cloud computing. More specifically GEYSERS proposes the concept of Virtual Infrastructures over one or more interconnected Physical Infrastructures comprising both network and IT resources. Taking into consideration the energy consumption levels associated with the ICT today and the expansion of the Internet in size and complexity, that incurring increased energy consumption of both IT and network resources, energy efficient infrastructure design becomes critical. To address this need, in the framework of GEYSERS, we propose energy efficient design of infrastructures incorporating integrated optical network and IT resources, supporting resilient end-to-end services. Our modeling results quantify significant energy savings of the proposed solution by jointly optimizing the allocation of both network and IT resources.

1. INTRODUCTION

As the scale of information processing is increasing, from Petabytes of Internet data to the projected Exabytes in networked storage at the end of this decade [1], novel network solutions are required to support the Future Internet and its new emerging applications such as UHD IPTV, 3D gaming, virtual worlds etc. These high-performance applications cannot be intrinsically delivered by the current Best Effort Internet. On the other hand, optical networking is offering a very high capacity transport with increased dynamicity and flexibility. In this context, the European project GEYSERS is developing a novel optical network solution, capable of provisioning “Optical Network and IT resources” for end-to-end service delivery. GEYSERS is focusing on the concept of Virtual Infrastructures (VIs) over one or more Physical Infrastructures (PIs) comprising both network and IT resources. The introduction of VIs facilitates sharing of physical resources among various virtual operators, introducing a new business model suitable for the Future Internet and enables new exploitation opportunities for the underlying physical infrastructures.

At the same time ICT accounts for 4% of the primary energy consumption worldwide [2]. The expansion of the Internet incurs increased energy consumption, thus attracting a lot of attention on energy efficient networking [3]. Designing and operating infrastructures comprising IT resources and optical networks in a power-aware manner becomes critical, as operation of IT resources requires very high levels of power and their conventional operating window is commonly not optimized for energy efficiency. Hence, allocating IT processing jobs in an energy-aware manner through a relatively low energy-consuming optical network infrastructure is expected to offer significant energy savings.

In addition, resilience of the optical network and IT resources to any kind of failures is an aspect of major importance in these infrastructures as they commonly support a huge amount of data that may also have very tight availability requirements. It is true to say that although supporting resilience can have a significant impact in the resource requirements and hence the corresponding power consumption, no attention has been given to the impact of resilience in the overall power consumption of these infrastructures including both network and IT resources to date.

In this paper, a modeling approach suitable for the planning of resilient VIs, formed over an integrated IT and optical network infrastructure, is proposed and implemented, extending our previous work on energy efficient VI planning [4]. The resilience scheme considered is that of 1:1 protection for both optical network and IT resources. More specifically, in case of a failure of the primary IT server, demands are forwarded to a secondary (protection) IT server, while in case of an optical link failure demands are routed to their destination through a secondary (protection) path. To identify the least energy consuming resilient VI, the detailed power consumption models and figures, including both optical network and IT resources, are considered [4], [5]. Mapping the virtual to physical resources and defining the energy consumption parameters of the VIs themselves is also part of the VI planning.

2. ENERGY AWARE RESILIENT VIRTUAL INFRASTRUCTURE PLANNING

The problem is formulated using an Integer Linear Programming (ILP) model. The network comprises one resource layer that contains the physical infrastructure and will produce as an output the virtual infrastructure layer illustrated in figure 1. The physical infrastructure is described through an eleven-node topology corresponding to the Pan-European optical network in which randomly selected nodes generate demands d ($d = 1, 2, \dots, D$) to be served by a set of IT servers s ($s = 1, 2, \dots, S$). The granularity of demands is the wavelength. The IT locations (demand destinations) at which the services will be handled, are not specified and are of no importance to the services themselves. However, uninterrupted service provisioning is of crucial importance in the deployment of transport optical networks. To this end, countermeasures against failures of the IT servers and optical links of the PI that may lead to service disruption should be taken into account during the VI planning process. In the proposed planning algorithm, a possible failure of the primary IT server, $s_i, s_i \in S$, is treated by forwarding demands to a secondary IT server, $s_j, s_j \in S$ with $s_i \neq s_j$. In case of an optical link failure, demands are routed through a secondary path.

In order to formulate the problem of service protection from possible failure of an IT server, the binary variable $a_{ds_i s_j}$ is introduced to indicate whether demand d is assigned to one of the IT servers s_i, s_j or not. This variable takes value equal to 1 if and only if demand d is processed on primary server, s_i , or in case of its failure on the secondary, s_j . Moreover, it is assumed that each demand can be assigned only to one IT server at a time

$$\sum_{s_i} \sum_{s_j} a_{ds_i s_j} = 1 \quad d = 1, 2, \dots, D \quad (1)$$

For each demand d , its demand volume h_d is realized by means of a number of lightpaths assigned to paths of the VI. Let $p_i = 1, 2, \dots, P_{ds_i s_j}$ be the candidate path list in the VI for the lightpaths required to support demand d at the primary server s_i or the secondary, s_j , and $x_{dp_{s_i s_j}}$ the non-negative number of lightpaths allocated to path $p_{s_i s_j}$. The following demand constraints should be satisfied in the VI [6]:

$$\sum_{s_i} \sum_{s_j} \sum_{p_{s_i s_j}} a_{ds_i s_j} x_{dp_{s_i s_j}} = h_d \quad d = 1, 2, \dots, D \quad (2)$$

Summing up the lightpaths through each link e ($e = 1, 2, \dots, E$) of the VI we can determine the required link capacity y_e for link e :

$$\sum_d \sum_{s_i} \sum_{s_j} \sum_{p_{s_i s_j}} \delta_{edp_{s_i s_j}} x_{dp_{s_i s_j}} \leq y_e \quad e = 1, 2, \dots, E \quad (3)$$

where $\delta_{edp_{s_i s_j}}$ is a binary variable taking value equal to 1 if link e of VI belongs to path p realizing demand d at server s_i or s_j ; 0 otherwise. Using the same rationale, the capacity of each link e in the VI is allocated by identifying the required lightpaths in the PI. The resulting PI lightpaths z determine the load of each link g ($g = 1, 2, \dots, G$) of the PI, and hence its capacity u_g . Assuming that $q = 1, 2, \dots, Q_e$ is used for denoting the PI's candidate path list realizing link e , then, the following demand constraint for link e should be satisfied:

$$\sum_q z_{eq} = y_e \quad e = 1, 2, \dots, E \quad (4)$$

where the sum is taken over all paths q on the routing list Q_e of link e . Introducing the link-path incidence coefficients γ_{geq} for the PI taking values equal to 1 if link g of PI belongs to path q realizing link e , 0 otherwise, the general formula specifying the PI capacity constraint can be stated as:

$$\sum_e \sum_q \gamma_{geq} z_{eq} \leq u_g \quad g = 1, 2, \dots, G \quad (5)$$

where G is the total number of links in the PI and the summation for each link g is taken over all lightpaths in the PI layer. Finally, in order to protect the planned network from a possible failure of link g , alternative protection (secondary) paths are allocated. Assuming that $r = 1, 2, \dots, R_g$ are the candidate restoration paths for link g , then the following protection capacity constraints should be satisfied:

$$\sum_r z'_{gr} = u_g \quad g = 1, 2, \dots, G \quad (6)$$

Finally, the required protection capacity for the remaining operating links l of the PI is given by

$$\sum_r \beta_{lgr} z'_{gr} = u'_l \quad l, g = 1, 2, \dots, G, \quad l \neq g \quad (7)$$

where β_{lgr} is a binary variable taking value equal to 1 if link l belongs to path r restoring link g ; 0 otherwise. Apart from link capacity constraints (3), (5) for the VI and PI, respectively, the total demands that are assigned to each server should not exceed its capacity p_s . The latter capacity corresponds to the underlying physical resources, such as CPU, memory, disk storage etc. The inequality specifying servers' capacity constraints is given by

$$\sum_d \sum_{p_{s_j}} a_{ds,s_j} c_{ds}(x_{dp_{s_j}}) \leq p_s, \quad s_i, s_j = 1, 2, \dots, S \quad (8)$$

where the summation is taken over all demands that arrive at server s and $c_{ds}(x_{ds})$ is a parameter specifying the computational requirements for demand d on server s . In practice, this parameter is determined by the set of relevant benchmarks for computer systems provided by the Standard Performance Evaluation Corporation (SPEC).

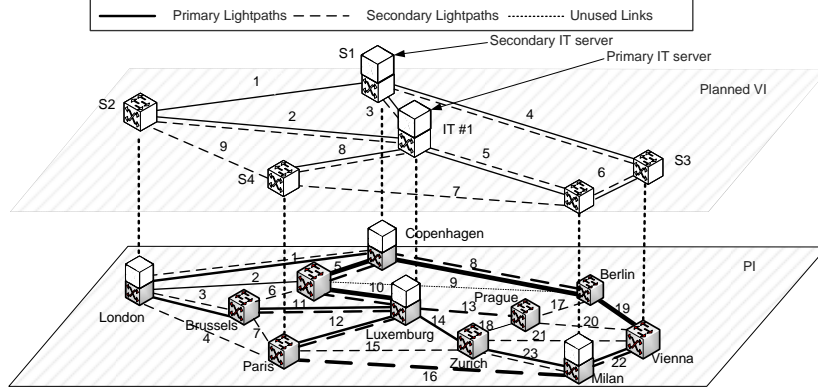


Figure 1 Resilient Virtual Infrastructure architecture over a converged optical network and IT servers

The objective of the current problem formulation is to minimize the total cost of the resulting network configuration as this cost consists of the following components: a) k_g the cost of the capacity of PI link g and b) E_s the energy consumption for processing u_s wavelengths in the IT server s . k_g corresponds to the energy consumed by each lightpath due to transmission and reception of the optical signal, optical amplification at each fiber span and switching. The IT power consumption model adopted in this paper mainly concentrates on the power consumption associated with the CPU load of IT resources and is described via the following linear equation

$$E_s(u_s) = P_s^{idle} + P_s^{busy} u_s \quad (9)$$

where P_s^{idle} , P_s^{busy} are parameters describing the energy consumption of the IT server s at idle state and per wavelength, respectively [7]. In addition to the power consumption due to data processing and network services, a 100% power overhead due to cooling has been incorporated. In this context, minimum energy consuming VIs are obtained by minimizing the following cost function

$$\text{Minimize } F = \sum_g k_g (u_g + u'_g) + \sum_s E_s \left[\sum_d \sum_p a_{ds} c_{ds}(x_{dps}) \right] \quad (10)$$

3. NUMERICAL RESULTS AND CONCLUSIONS

To investigate the energy efficiency of the proposed VI design scheme, the GEYSERS architecture depicted in Figure 1 is considered. For the PI the COST239 Pan-European reference topology has been used in which four randomly selected nodes generate demands to be served by two IT servers. Further, we assume a single fiber per link, 40 wavelengths per fiber, and wavelength channels of 10Gb/s each. It is also assumed that each IT server can process up to 2Tb/s and its power consumption ranges from 6.6 to 13.2KW, under idle and full load, respectively.

An example of the optimal VI topology design for a scenario, in which four source nodes generate demands equal to 15 wavelengths each, is depicted in figure 1. In this scenario, the generated VI topology consists of 9 virtual links and 6 virtual nodes, while all demands are routed to a single IT server. In case of failure of the primary IT server demands are routed to the secondary IT server. Furthermore, in case of a failure of a working path, additional capacity has been reserved to a link-disjoint backup path. The working and protecting capacity of each virtual link along with its mapping to the PI is given in Table I. For example it is observed that the virtual link Y5 connecting Milan and Luxemburg is realized via the working physical layer path u14-u23 with capacity 15 and the protecting paths u12-u16 and u13-u20-u22.

TABLE I. SAMPLE VIRTUAL TO PHYSICAL MAPPING

Virtual link	Capacity (wavelengths)		Physical Layer Paths realizing virtual links	PI paths Capacity(wavelengths)
	Primary	Secondary		
Y1	30	0	(P) u1	30
Y2	15	10	(P) u3-u11	15
			(S) u2-u10	10
Y3	30	20	(P) u5-u10	30
			(S) u1-u3-u11	10
			(S) u8-u17-u13	10
Y4	30	10	(P) u8-u19	30
			(S) u5-u6-u7-u15-u21	10
Y5	15	30	(P) u14-u23	15
			(S) u12-u16	15
			(S) u13-u20-u22	15
Y6	15	15	(P) u22	15
			(S) u21-u23	15
Y7	0	35	(S) u16	25
			(S) u15-u23	10
Y8	15	15	(P) u12	15
			(S) u7-u11	15
Y9	0	5	(S) u4	5

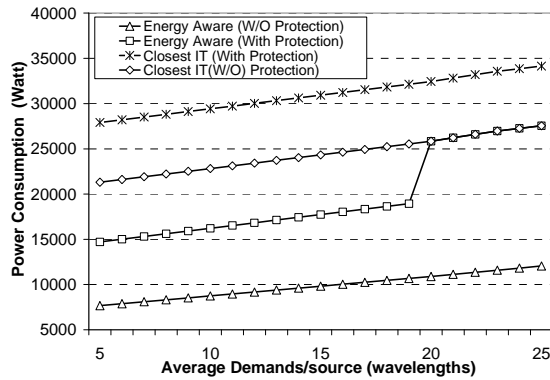


Figure 2 Comparison of the energy aware scheme with the closest IT server demand allocation scheme

In figure 2, the performance of the proposed energy aware VI design with and without (w/o) resilience considerations (1:1 protection mechanisms) is compared to the demand allocation scheme where demands from each source node are assigned to its closest IT server. Note that “closest” refers to the shortest distance between a source node and a data center. Comparing these two schemes, it is observed that the energy aware VI design without resilience consumes significantly lower energy for serving the same amount of demands compared to the closest IT scheme: in the former approach only one IT server is activated to serve the same amount of demands. Given that the power consumption required for the operation of the IT servers is dominant in this type of networks, avoiding using the unnecessary IT resources achieves significant reduction of energy consumption. Furthermore, the energy aware scheme including resilience considerations achieves significantly lower power consumption or for high traffic demands the same, with the closest IT scheme without protection.

Acknowledgements

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.

References

[1] B. Swanson et al., “Estimating the Exaflood”, <http://www.discovery.org/a/4428>, accessed October 2008
 [2] M. Pickavet et al., “Worldwide energy needs for ICT: The rise of power-aware networking,” in Proc. IEEE ANTS, Dec. 2008, pp. 1–3.
 [3] J. Baliga et al., “Energy Consumption in Optical IP Networks”, Journal of Lightwave Technology, Vol 27, No. 13, July 1, 2009.
 [4] A.Tzanakaki et al., “Energy Efficiency in integrated IT and Optical Network Infrastructures: The GEYSERS approach”, in proc. of IEEE INFOCOM 2011, Workshop on Green Communications and Networking (2011)
 [5] A.Tzanakaki et al, “Dimensioning the future Pan-European optical network with energy efficiency considerations,” JOCN 3, 272-280 (2011)
 [6] E. Kubilinskas, P. Nilsson, M. Piore, “Design Models for Robust Multi-Layer Next Generation Internet Core Networks Carrying Elastic Traffic”, in proc. of DRCN 2003, 61-68 (2003).
 [7] X. Fan et al., “Power provisioning for a warehouse-sized computer”, SIGARCH Comput. Archit. News 35, 2, 13-23 (2007)