Planning of Optical and IT Resources for EfficientVirtual Infrastructure Embedding

Albert Pagès, Jordi Perelló, Salvatore Spadaro
Advanced Broadband Communications Center (CCABA)
Universitat Politècnica de Catalunya (UPC)
Barcelona (Spain)
{albertpages, spadaro}@tsc.upc.eu, perello@ac.upc.edu

Abstract—**This work addresses the offline problem of efficiently planning network substrates joining geographically distributed IT resource locations over high-capacity optical resources. Given an already known set of virtual infrastructure demands to be embedded over the network substrate, we present an iterative optimization heuristic that appropriately embeds all virtual infrastructure demands in order to minimize both the physical substrate IT and optical resources needed to allocate them. For this purpose, the heuristic arranges the demands in descending order, based on the amount of resources they require. Then, it processes all demands, one after another, optimally embedding each one of them on the network substrate using Integer Linear Programming (ILP) techniques. The performance of the proposed heuristic is highlighted through illustrative results.**

Keywords: virtualization, virtual infrastructure, planning.

I. INTRODUCTION

Huge research efforts have recently been focused on extending virtualization technologies, so as to not only encompass intra-data center environments, but also network segments inter-connecting remote data center locations. A key rationale behind this is the provisioning of novel virtual infrastructure services integrating IT and network resources [1]. Customizability and manageability will become important characteristics of these Virtual Infrastructures (VIs), offering a "real physical device experience" to service providers operating them, but without having to invest large amounts of money on real physical equipment.

In a network virtualization context [2], infrastructure providers will seek to maximize profits from their physical IT and network resources by serving as many VI requests as possible. To address this need, the present paper investigates the scenario where a new infrastructure provider, willing to deliver both IT and network resources as a service, is interested in planning the amount of IT and optical resources that must be deployed to serve an already known set of VI demands. In other words, the VI embeddings (virtual-physical node and link mappings) that minimize the amount of IT and optical resources in the substrate are targeted. In general, the freedom existent in the virtual-physical node mapping (a virtual node can, in principle, be mapped to any physical node in the substrate network) makes the problem of embedding a VI on a physical network substrate to be NP-hard [3]. Moreover, the complexity of the problem is further increased when the number of VIs demands to be embedded grows up.

Based on the observations above, the remainder of this paper derives an iterative optimization heuristic to address the network substrate planning problem in reasonable execution times. The proposed heuristic and the considered scenario are described in section II, while section III obtains and discusses illustrative results. Section IV concludes the paper.

The work in this paper has been supported by the Government of Catalonia and the European Social Fund through a FI-AGAUR research scholarship grant and the Spanish Science Ministry through the ELASTIC project (TEC2011-27310).

II. PROPOSED HEURISTIC

As said before, the target of the proposed heuristic is to determine the amount of physical resources, both optical and IT, to serve a set of VI demands. As considered in this work, a VI demand is a set of virtual nodes with IT capacities that are inter-connected over a high-capacity optical network. We assume that the physical substrate (optical network + IT sites) has sufficient capacity to serve all offered demands, i.e., there is no blocking. Specifically, we model the optical network as a graph $G = (N, E)$, where N denotes the set of nodes and $E = \{(i, j), (j, i) : i, j \in N, i \neq j\}$ the set of physical links. W is the set of wavelengths available per physical link. We assume that IT sites are connected to the core transport network through a network with sufficient capacity. This is equivalent to having some nodes in the optical network substrate graph with IT resources associated to them. Thus, we define $N_{IT} \subseteq N$ as the set of nodes with IT resources. The IT resources considered per node are CPU frequency (in GHz), storage capacity (in GB), memory capacity (in GB) and number of CPU cores.

Let D denote the set of demands to be accommodated on the physical substrate. Each demand $d \in D$ is characterized by a graph $G'_d = (N'_d, E'_d)$, $N'_d \subseteq N_{IT}$, $E'_d = \{(i, j), : i, j \in N'_d, i \neq j\}$, where each of the virtual nodes $n' \in N'_d$ requests for a certain amount of IT resources. We denote as $F^d_{n'}$ (CPU frequency), $S^d_{n'}$ (storage), $M^d_{n'}$ (memory) and $C^d_{n'}$ (cores) the particular IT resources requested by virtual node $n' \in N'_d$, and F^d_{max} , S^d_{max} , M^d_{max} and C^d_{max} the maximum amount of IT resources requested by some virtual node in demand $d \in D$. Additionally, W_d denotes the number of wavelengths per virtual link requested by demand $d \in D$. The pseudo-code of the proposed heuristic, called Deterministic Ordering Virtual Infrastructure Planner (DOVIP), is displayed in Fig. 1.

Algorithm 1 DOVIP
1: Calculate $Cost_d$ for every demand $d \in D$
2: Order them in descending order
3: n = 1
4: Execute model for demand d_n
5: Actualize substrate state according to model output
6: n = n+1; if n = $\|D\|$, **terminate**, otherwise **go to** step 2

Figure 1. DOVIP heuristic pseudo-code.

Firstly, demands are ordered in descending order according to the expression depicted in (1), where we denote as $m(\cdot)$ the maximum function (among all $d \in D$). In this way, the heuristic initially tries to fit the biggest demands, which may be more complicated to allocate, and then the smaller demands, more easily allocable. Factors α, β_F, β_S, β_M, and β_C are used to put more or less weight to the terms of the expression. Once the demands are ordered, the heuristic iteratively executes for each demand an ILP model whose goal is to jointly minimize the number of wavelengths per physical link and, at the same time, the amount of IT resources in the more congested node $n \in N_{IT}$ once the current demand is served. After each model execution, the state of the physical substrate is updated before starting with the next iteration.

$$Cost_d = \alpha \frac{|E'_d| W_d}{m(|E'_d|) m(W_d)} + (1-\alpha)\left(\beta_F \frac{F^d_{max}}{m(F^d_{max})} + \beta_S \frac{S^d_{max}}{m(S^d_{max})} + \beta_M \frac{M^d_{max}}{m(M^d_{max})} + \beta_C \frac{C^d_{max}}{m(C^d_{max})}\right) \frac{|N'_d|}{m(|N'_d|)}, \forall\, d \in D \quad (1)$$

The details of the ILP model executed in step 4 of the heuristic are discussed below. We define F^n_u, S^n_u, M^n_u and C^n_u as the amount of IT resources that are already in use in node $n \in N_{IT}$; F^u_{max}, S^u_{max}, M^u_{max} and C^u_{max} as the maximum amount of IT resources that are already in use for some node $n \in N_{IT}$; and W^e_u as the set of wavelengths that are already in use in physical link e. Additionally, we denote as $\delta^+(n)$ and $\delta^-(n)$ the set of outgoing and incoming edges from/to $n \in N$ respectively. Finally, $a(e')$ and $b(e')$ denote the source and destination of virtual edge $e' \in E'_d$.

The decision variables of the model are:

$x^n_{n'}$ = {1 if virtual node n' is mapped over physical node n, 0 otherwise}, $\forall\, n' \in N'_d, n \in N$.

$y^{e,w}_{e'}$ = {1 if virtual link e' is mapped over physical link e and wavelength w, 0 otherwise}, $\forall\, e' \in E'_d, e \in E, w \in W$.

t_w = {1 if wavelength w is used in any physical link, 0 otherwise}, $\forall\, w \in W$.

F_n, S_n, M_n, C_n = IT capacities in physical node n, $\forall\, n \in N_{IT}$.

$S_{max}, M_{max}, C_{max}$ = maximum values among S_n, M_n and C_n.

The ILP formulation is:

$$min\ \alpha \frac{1}{|W|} \sum_{w \in W} t_w + (1-\alpha)\left(\frac{1}{F^u_{max} |N_{IT}|} \beta_F \sum_{n \in N_{IT}} F_n + \beta_S \frac{S_{max}}{S^u_{max} + S^d_{max}} + \beta_M \frac{M_{max}}{M^u_{max} + M^d_{max}} + \beta_C \frac{C_{max}}{C^u_{max} + C^d_{max}}\right),\ s.t. \quad (2)$$

$$F^d_{n'} x^n_{n'} \le F_n,\ \forall\, n' \in N'_d,\ n \in N_{IT} \quad (3)$$

$$F^n_u \le F_n,\ \forall\, n \in N_{IT} \quad (4)$$

$$S_n = S^n_u + \sum_{n' \in N'_d} S^d_{n'} x^n_{n'},\ \forall\, n' \in N'_d,\ n \in N_{IT} \quad (5)$$

$$S_n \le S_{max},\ \forall\, n \in N_{IT} \quad (6)$$

$$M_n = M^n_u + \sum_{n' \in N'_d} M^d_{n'} x^n_{n'},\ \forall\, n' \in N'_d,\ n \in N_{IT} \quad (7)$$

$$M_n \le M_{max},\ \forall\, n \in N_{IT} \quad (8)$$

$$C_n = C^n_u + \sum_{n' \in N'_d} C^d_{n'} x^n_{n'},\ \forall\, n' \in N'_d,\ n \in N_{IT} \quad (9)$$

$$C_n \le C_{max},\ \forall\, n \in N_{IT} \quad (10)$$

$$y^{e,w}_{e'} \le t_w,\ \forall\, e' \in E'_d,\ e \in E,\ w \in W \quad (11)$$

$$1 \le t_w,\ \forall\, e \in E,\ w \in W^e_u \quad (12)$$

$$y^{e,w}_{e'} = 0,\ \forall\, e' \in E'_d,\ e \in E,\ w \in W^e_u \quad (13)$$

$$\sum_{n \in N_{IT}} x^n_{n'} = 1,\ \forall\, n' \in N'_d \quad (14)$$

$$\sum_{n' \in N'_d} x^n_{n'} \le 1,\ \forall\, n \in N_{IT} \quad (15)$$

$$x^n_{n'} = 0,\ \forall\, n' \in N'_d,\ n \notin N_{IT} \quad (16)$$

$$\sum_{e' \in E'_d} (y^{e_{i,j},w}_{e'} + y^{e_{j,i},w}_{e'}) \le 1,\ \forall\, e \in E,\ w \in W,\ i \ne j \quad (17)$$

$$\sum_{e \in \delta^+(n)} \sum_{w \in W} y^{e,w}_{e'} - \sum_{e \in \delta^-(n)} \sum_{w \in W} y^{e,w}_{e'} = W_d (x^n_{a(e')} - x^n_{b(e')}),\ \forall e' \in E'_d, n \in N \quad (18)$$

$$-x^n_{b(e')} \le \sum_{e \in \delta^+(n)} y^{e,w}_{e'} - \sum_{e \in \delta^-(n)} y^{e,w}_{e'} \le x^n_{a(e')},\ \forall\, e' \in E'_d,\ n \in N,\ w \in W \quad (19)$$

$$y^{e,w}_{e'} \le 1 - x^n_{b(e')},\ \forall\, e' \in E'_d,\ n \in N,\ e \in \delta^+(n),\ w \in W \quad (20)$$

$$y^{e,w}_{e'} \le 1 - x^n_{a(e')},\ \forall\, e' \in E'_d,\ n \in N,\ e \in \delta^-(n),\ w \in W \quad (21)$$

$$W_d y^{e,w}_{e'} \le \sum_{w \in W} y^{e,w}_{e'},\ \forall\, e' \in E'_d,\ e \in E \quad (22)$$

The ILP objective function is depicted in (2). Constraints (3), (4), (5), (7) and (9) serve to give value to the IT resource variables, taking into account the current state of the physical resources and those resources requested by the demand. Constraints (6), (8) and (10) are the constraints that enable the minimization of the IT resources in the most congested IT node. Constraints (11) account for the wavelengths that are being used to serve the demand, while constraints (12) account for the wavelengths that are already in use due to previous allocations. Constraints (13) avoid the use of pairs (e, w) that are already in use when allocating the demand. Constraints (14) ensure that all virtual nodes $n' \in N'_d$ are being mapped over a unique physical node, while constraints (15) ensure that a physical node $n \in N_{IT}$ does not support more than one virtual node per demand. Constraints (16) avoid using physical nodes $n \notin N_{IT}$ for mapping virtual nodes. Constraints (17) are the wavelength clashing constraints, implicitly accounting for the true bidirectional nature of the demands through considering both directions of physical arc (i, j), (j, i), although unidirectional demands where defined; in this way, the complexity of the model is reduced and, hence, its execution time. Constraints (18) ensure that all virtual links $e' \in E'_d$ are provided with the required number of wavelengths. Constraints (19) are the flow conservation and wavelength continuity constraints. Constraints (20) and (21) serve to avoid loops at the remote endpoints of a virtual link; in the source remote endpoint, traffic can only flow outside the node, while in the destination remote endpoint, traffic can only enter at the node. Finally, constraints (22) ensure that for a virtual link, all requested wavelengths follow the same route; this is done to avoid delays.

While the proposed formulation only takes a single VI demand as input, placing it optimally according to the current substrate state and the characteristics of the VI itself, we alternatively evaluated the feasibility of an offline ILP planning model that takes the entire demand set D, and tries to place all demands over the underlying physical substrate in an optimal way. Nevertheless, we experienced extremely long execution times, even for small demand set sizes. For example, for $|D| = 5$, we registered execution times in the order of 10^6 seconds. This became a key motivation behind DOVIP heuristic as a way to solve the addressed problem in reasonable time.

III. Results and Discussion

In order to illustrate the performance of the DOVIP heuristic, we have carried out a series of executions for various sizes of the demand set D and values of the factor α; factors β_F, β_S, β_M and β_C are fixed to 0.25. Results throughout this section have been averaged over 20 random generations of the demand sets, being executed in an Intel Core i5 PC with 4 cores at 3.3 GHz and 4 GB of RAM PCs, running CPLEX v.12.2.

The Deutsche Telekom network topology [4] is used as the physical substrate network, assuming that IT resources are present in Berlin, Munchen, Stuttgart, Dusseldorf, Hamburg and Hannover. Demands in the set are generated randomly, considering 3 or 4 nodes per demand and connecting them with probability 0.5, preventing the generation of non-connected graphs. We assume $W_d = 1$ for all demands. As for the IT resources requested by virtual nodes of the demands, these are chosen with equiprobability between 1.0 – 2.2 with steps of 0.2 for the CPU frequency, 10 – 30 with steps of 5 for the storage capacity, 2 – 8 with steps of 2 for the memory capacity and 1 – 4 with steps of 1 for the number of CPU cores.

Table I displays the average values for the physical substrate resources that take a role in objective function (2), along with the average execution time. We obtain these results for various values of $|D|$ and α. Note that $\alpha = 1$ and $\alpha = 0$ identify extreme situations, only focusing on the minimization of optical resources or IT resources, respectively.

TABLE I. Substrate Resources Comparison

		α = 1	**α = 0.75**	**α = 0.5**	**α = 0.25**	**α = 0**
|D| = 10	***Sum(t_w)***	5.55	5.8	5.9	6.65	14.65
	Sum(F_n)	12.43	12.12	12.1	12.06	11.95
	S_max	180	137.25	132	129.75	128.75
	M_max	43.7	30.8	30.5	30.1	29.9
	C_max	21.85	15.4	15.25	15.05	14.95
	Time (s.)	11.74	66.44	79.43	61.94	48.57
|D| = 15	***Sum(t_w)***	7.9	8.5	8.95	9.95	19.8
	Sum(F_n)	12.8	12.63	12.53	12.46	12.4
	S_max	262.5	192.5	191.5	188	185.75
	M_max	68	47.1	46.5	45.2	45
	C_max	34	23.55	23.25	22.7	22.5
	Time (s.)	25.87	150.49	150.76	138.4	86.24
|D| = 20	***Sum(t_w)***	10.2	11.2	11.2	12.1	20
	Sum(F_n)	12.83	12.61	12.56	12.55	12.29
	S_max	351.5	255	250.5	248.25	245.5
	M_max	87.4	61.3	61.1	60.1	59.7
	C_max	43.7	30.65	30.55	30.05	29.85
	Time (s.)	31.31	146.65	161.8	158.5	99.35

It can be appreciated that $\alpha = 0.5$ working point offers a fair trade-off between the amount of network and IT resources needed for any of the considered $|D|$ sizes. It uses much less optical resources than $\alpha = 0$ (around a 50% less) and less IT resources than $\alpha = 1$ (around a 2% less for the CPU frequency and a 30% less for the rest of IT resources), while it only uses around a 2% more IT resources than $\alpha = 0$ and a 10% more optical resources than $\alpha = 1$. In this sense, it can be appreciated that the joint optimization of optical and IT resources results in a good trade-off between the two parts.

IV. Conclusions and Future Work

This work investigates the problem of optical and IT resources planning for VI embedding. An iterative ILP-based heuristic has been presented as a tool to attack the problem. The presented results demonstrate that the presented heuristic is capable to solve the problem in practical execution times. Multiple working point scenarios have been tested, being $\alpha = 0.5$ the one that offers the best trade-off between optical and IT resource usage. As future work, we will devote our efforts on evaluating the goodness of the heuristic by benchmarking it against lower bounds of the optimal ILP formulation that takes all requested VI demands, jointly optimizing their placement.

References

[1] M. P Anastasopoulos and A. Tzanakaki, “Adaptive virtual infrastructure planning over interconnected IT and optical network resources using evolutionary game theory,” on proceedings of ONDM 2012, April 2012, in press.

[2] A. Khan, A. Zugenmaier, D. Jurca and W. Kellerer, “Network virtualization: a hypervisor for the Internet?,” Communications Magazine, vol.50, no.1, pp. 136-143, January 2012.

[3] N.M.M.K. Chowdhury, M.R. Rahman and R. Boutaba, “Virtual network embedding with coordinated node and link mapping,” on proceedings of INFOCOM 2009, April 2009.

[4] F. Agraz, et al., “Experimental demonstration of centralized and distributed impairment-aware control plane schemes for dynamic transparent optical networks,” on proceedings of OFC/NFOEC 2010, March 2010.