Strategies for Virtual Optical Network Allocation

Albert Pagès, Jordi Perelló, Salvatore Spadaro, *Member, IEEE*, and Gabriel Junyent

***Abstract*—This paper presents Integer Linear Programming (ILP) formulations to optimally allocate Virtual Optical Networks (VONs) over a transparent optical network substrate. These formulations serve the purpose of building either completely transparent VONs or opaque ones, where electrical termination capabilities are assumed at each virtual network node. In addition, a lightweight Greedy Randomized Adaptive Search (GRASP) heuristic is provided for the transparent case. The obtained results validate the accuracy of the proposed heuristic and reveal the benefits of the presented solutions against simpler shortest-path-based VON allocation strategies.**

***Index Terms*—Optical networks, virtualization, allocation.**

I. Introduction

Network virtualization will allow transport network owners to not only offer data transport services over their physical infrastructures, but also portions of such infrastructures as a service for exploitation by external service providers [1]. To this goal, virtualization techniques have been well applied to Layer-2/3 networks. However, their applicability to optical networks is still under research [2], [3].

In a virtualized optical network environment, completely isolated VONs belonging to different service providers can coexist over the same optical network substrate. VONs are composed of a set of virtual optical nodes connected together by a set of virtual links. Each virtual optical node is mapped on a particular physical device, allowing the management of the resources assigned to the VON on that device. In turn, virtual links are mapped over physical paths, allocating a portion of network resources along them.

An efficient virtual network allocation arises of paramount importance to maximize the resource utilization of the underlying physical infrastructure, which impacts directly on the resulting revenues [1]. Looking at the literature, the problem of mapping virtual network demands over a physical network substrate has been referred as the virtual network embedding problem (e.g., see [4], [5]). Solutions to this problem seek the optimal allocation of virtual network demands over a physical substrate with scarce network resources, mostly realized through a Layer-2/3 network, such as an IP network, for which virtualization techniques are quite mature. To the best of our knowledge, however, no work in the literature has addressed this problem in the context of VONs. In light of this, the present paper addresses the planning problem of optimally allocating a set of VON demands over an optical network substrate, while accounting for the peculiarities of optical networks. We denote this problem as the Virtual Optical Network Allocation (VONA) problem.

As will be illustrated, the VONA problem includes a Routing and Wavelength Assignment (RWA) that can differ depending on the service provider's needs. Two variants of the VONA problem have been considered in this paper, namely, transparent and opaque VONA. In transparent VONA, optically transparent end-to-end services are provisioned over the VON. This requires the allocation of exactly the same set of wavelengths for every virtual link. Alternatively, in opaque VONA, we assume that electronic termination capabilities are physically present at each VON node and opaque transport services are provided from the VON viewpoint. In such a case, there is no need to allocate the same set of wavelengths for each virtual link, but can differ thanks to the Optical-Electrical-Optical (OEO) conversion stages.

In both scenarios we assume an all-optical network substrate without wavelength conversion capabilities (i.e., virtual links must ensure the *wavelength continuity constraint*), where Physical Layer Impairment (PLI) degradations do not compromise the feasibility of the optical channels. Very large network scenarios may require the introduction of PLI information in both transparent and opaque VONA problems in order to ensure the feasibility of the provisioned VONs. However, this is not considered in this work and left for further study. Also, note that, in its current form, the presented VONA formulation is still valid for a wide range of transparent network scenarios. As recently published in [6], where a quite restrictive scenario was considered (i.e., impact of both linear and non-linear PLIs as well as the co-existence of different bit rates and modulation formats), feasible transparent reaches of 3000 and 1600 km were obtained for 10 and 40 Gb/s, respectively.

II. VONA Problem Formulation

Let the optical network substrate be characterized by 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. Let W denote the set of available wavelengths per physical link. Consider D as the set of VON demands to be allocated over the optical network. Each demand $d \in D$, is characterized by a graph $G'_d = (N'_d, E'_d)$, $N'_d \subseteq N$, $E'_d = \{(i, j), (j, i) : i, j \in N'_d, i \neq j\}$. We denote by $U_{\{d\}}$ the number of wavelengths per virtual link desired by demand $d \in D$.

The VONA problem consists in accommodating all or the maximum number of VONs from the demand set given the limited capacity of the underlying optical network. VONs are treated as entities instead of a composition of lightpaths, which makes VONA differ from the RWA problem with the objective

Manuscript received September 9, 2011. The associate editor coordinating the review of this letter and approving it for publication was J. Wang.

The authors are with the Advanced Broadband Communications Center (CCABA), Universitat Politècnica de Catalunya (UPC), Barcelona, Spain (e-mail: {albertpages, spadaro, junyent}@tsc.upc.edu, perello@ac.upc.edu).

Digital Object Identifier 10.1109/LCOMM.2011.120211.111900

to maximize the number of lightpaths established. Indeed, a specific demand $d \in D$ is accommodated if and only if all its virtual links in E'_d can be mapped over available resources. Moreover, in transparent VONA, the same set of wavelengths must be allocated to each of its virtual links, allowing transparent network services across multiple virtual links. The rest of this section presents optimal ILP formulations for transparent and opaque VONA and a heuristic for the transparent case.

A. Transparent VONA

This subsection presents an ILP model of the transparent VONA problem called TVONA_ILP. For this, we define P as the set of paths in the physical network, $S_{\{e,e'\}}$ as the set of $p \in P$ associated with virtual link e' that traverse edge $e \in E$, and $S_{\{e',d\}}$ as the set of $p \in P$ associated with virtual link e' in demand d. The decision variables of TVONA_ILP are:

$x(d,e',p,w) = \{1$ if for demand d the virtual link e' is supported through path p and wavelength w, 0 otherwise$\}$

and the auxiliary variables are:

$y(d,w) = \{$an integer number equal to the minimum number of times wavelength w is used to serve demand $d\}$

$z(d) = \{1$ if demand d can be satisfied, 0 otherwise$\}$

Objective function (1) aims at maximizing the number of VONs to be allocated in the underlying optical network. Constraints (2) are the wavelength clashing constraints, which avoid that two virtual links are supported over the same wavelength in the same physical link. Constraints (3) ensure that at most $U_{\{d\}}$ different (p,w) will be assigned to every virtual link belonging to demand d. Constraints (4) discriminate if wavelength $w \in W$ is being used by all virtual links in demand d. Constraints (5) discriminate whether demand d is supported over the requested number of wavelengths $U_{\{d\}}$.

$$max \sum_{d \in D} z(d), s.t. \tag{1}$$

$$\sum_{d \in D} \sum_{e' \in E'_d} \sum_{p \in S_{\{e,e'\}}} x(d,e',p,w) \leq 1, \forall e \in E, w \in W \tag{2}$$

$$\sum_{p \in S_{\{e',d\}}} \sum_{w \in W} x(d,e',p,w) \leq U_{\{d\}}, \forall d \in D, e' \in E'_d \tag{3}$$

$$y(d,w) \leq \sum_{p \in S_{\{e',d\}}} x(d,e',p,w), \forall d \in D, e' \in E'_d, w \in W \tag{4}$$

$$z(d) \leq \sum_{w \in W} y(d,w)/U_{\{d\}}, \forall d \in D \tag{5}$$

As will be shown in the following section, the execution time of the TVONA_ILP model grows substantially up as the number of VONs in D increases. In view of this, a heuristic for the transparent VONA problem based on the GRASP metaheuristic [7] is presented in this paper, which ensures practical execution times even when the number of VONs in D is large. The pseudo-code of this heuristic, called TVONA_GRASP, is depicted in Fig. 1.

Looking at the figure, TVONA_GRASP builds a feasible solution in phase 2, considering the same constraints as the model. The purpose of the *multiplicative_factor* is to assign higher costs to variables associated with demands not under construction, thus favoring those demands with elements already in the solution (i.e., the probability to build full demands increases). In this work we have fixed this factor to 4, which provided us with the best accuracy in the scenarios under study. Nonetheless, its configuration is left to the transport network owner's discretion. Phase 3 tries to improve this solution by local search in the solution space around the solution from phase 2. After each iteration, if a solution with a better objective value than the overall best solution found so far is found, this solution becomes the new best solution. At the end of the process, the best solution is returned.

```
Phase 1: Preprocessing
Eliminate d ∈ D with nodes with degree greater than physical ones
BestSol = ∅
for i = 0 to i = MaxIterations and Obj(Sol) ≠ |D| do
    Sol = ∅; Candidate list C = ∅
    Phase 2: Solution construction
    C ← ∪x(d, e', p, w) ∀d ∈ D, e' ∈ E'_d, p ∈ S_{e',d}, w ∈ W
    Assign cost equal to physical hops ∀c ∈ C
    W_d = ∅ wavelengths already allocated to d, ∀d ∈ D
    while C ≠ ∅ do
        cost_min ← min cost from C
        RCL ← {c ∈ C | cost(c) = cost_min}
        Select an element c from the RCL at random
        Sol ← Sol ∪ {c}
        if |W_d| ≠ U_{d} then
            W_d ← W_d ∪ {wavelength associated to c}
        Erase from C c and candidates causing wavelength clash
        if num_wavelengths of e' in Sol = U_{d} then
            Erase from C all candidates of e'
        if |W_d| = U_{d} then
            Erase from C candidates of d with w not in W_d
        for all c ∈ C do
            if some elements of d are in Sol then
                cost(c) = num_hops
            else
                cost(c) = num_hops × multiplicative_factor
    if Obj(Sol) ≠ |D| then
        Phase 3: Solution improvement
        Temporally extract partially satisfied demands d_p
        for less constructed d_p to more constructed do
            Find combination of (p, w) that satisfies d_p
            if found then
                Sol ← Sol ∪ {d_p}
    if Obj(Sol) > Obj(BestSol) then
        BestSol ← Sol
```

Fig. 1. TVONA_GRASP heuristic pseudo-code

B. Opaque VONA

This subsection presents an ILP formulation of the opaque VONA problem called OVONA_ILP. In OVONA_ILP, variables $y(d,w)$, unnecessary here, are suppressed along with constraints (4). The rest of the formulation is almost identical to the TVONA_ILP formulation. Objective function (6) aims at maximizing the number of VONs to be allocated in the underlying optical network. Constraints (7) are the wavelength clashing constraints, which avoid that two virtual links are supported over the same wavelength in the same physical link. Constraints (8) ensure that at most $U_{\{d\}}$ different (p,w) will be assigned to every virtual link belonging to demand d.

Constraints (5) are slightly modified to fit the characteristics of such an opaque scenario, now transformed into constraints (9), whose purpose is to discriminate whether demand d is fully served or not.

$$max \sum_{d \in D} z(d), s.t. \tag{6}$$

$$\sum_{d \in D} \sum_{e' \in E'_d} \sum_{p \in S_{\{e,e'\}}} x(d, e', p, w) \leq 1, \forall e \in E, w \in W \tag{7}$$

$$\sum_{p \in S_{\{e',d\}}} \sum_{w \in W} x(d, e', p, w) \leq U_{\{d\}}, \forall d \in D, e' \in E'_d \tag{8}$$

$$z(d) \leq \sum_{p \in S_{\{e',d\}}} \sum_{w \in W} x(d, e', p, w)/U_{\{d\}}, \forall d \in D, e' \in E'_d \tag{9}$$

The removal of variables $y(d, w)$ and constraints (4), ensuring that each virtual link $e' \in E'_d$ uses the exact same subset of wavelengths in TVONA_ILP, lightens OVONA_ILP when compared to its transparent counterpart. This makes OVONA_ILP solvable to optimality within a reasonable time span, as will be illustrated in the following section. Hence, no heuristic approach has been considered as necessary to solve the opaque VONA problem.

III. Results and Discussion

As mentioned in section II, the TVONA_ILP model becomes impracticable for large D sizes, which motivated the proposal of the TVONA_GRASP heuristic, as a way to obtain results close to optimality but in a much shorter time. In order to highlight the accuracy of TVONA_GRASP, as well as its running times compared to the exact TVONA_ILP model, we have executed both of them on the 16-Node EON core network topology [8] with 8 wavelengths per physical link. In particular, $|D|$ equal to 10, 20 and 30 has been considered, assuming that each demand requests 1 wavelength, that is, $U_{\{d\}} = 1, \forall d \in D$.

The generation of the demand sets for all experiments throughout this section follows a 2-step process. Firstly, 3 or 4 physical network nodes (with equiprobability) are randomly selected as virtual nodes for each demand. In this way, we obtain reasonable medium-sized virtual networks compared to the underlying physical network size. Next, the selected virtual nodes are then randomly connected using the Erdős-Rényi algorithm [9], here slightly modified to prevent the generation of non-connected graphs. The parameter p has been set to 0.5 in the algorithm, which leads to the generation of any connectivity matrix with equiprobability.

For the TVONA_GRASP heuristic, we impose some limitations to restrict its execution time while still producing good results: 1) the $MaxIterations$ field is set to 125; 2) the number of paths associated to the virtual links of the demands is limited to 30 paths per virtual link; 3) the number of combinations to check during phase 3 is limited to 10^6. Table I compares the performances of TVONA_ILP and TVONA_GRASP in terms of execution time and number of successfully allocated demands. The presented results have

TABLE I
TVONA_ILP vs TVONA_GRASP

	TVONA_ILP		TVONA_GRASP		
	Time (s.)	Result	Time (s.)	Result	% Error
$\lvert D\rvert = 10$	1533.52	8.68	366.49	8.64	0.46
$\lvert D\rvert = 20$	1.64×10^6	14.48	601.67	13.96	3.59
$\lvert D\rvert = 30$	2.38×10^6	17.76	612.5	16.52	6.98

been averaged over 25 executions, randomly generating a new set of demands at the beginning of each execution. The experiments have been launched on Intel Core Duo at 3 GHz PCs with 4 GB RAM memory.

As seen, the execution times of TVONA_GRASP stay largely below the ones of TVONA_ILP. This reduction is especially important as the number of offered demands increases (e.g., a four orders of magnitude reduction is achieved for $|D| = 30$). Moreover, TVONA_GRASP still provides accurate results, showing relative errors between 0.46% and 6.98% in the worst case. Note, however, that such errors represent one more demand blocked in average barely. As stated before, the execution times of OVONA_ILP are much lower than those of TVONA_ILP. Indeed, additional experiments to the ones presented in Table I showed us average execution times of OVONA_ILP around 14.2 s, 116.1 s and 131.2 s for $|D| = 10$, $|D| = 20$ and $|D| = 30$, respectively.

Aiming to quantify the benefits of the proposed contributions for efficient VON allocation, we benchmark them against more simpler allocation approaches. To this end, we consider a Shortest-Path (SP) strategy which serves the demands in D on a one-by-one basis, mapping their virtual links to the shortest physical path that connects both endpoints. On these physical paths, a first-fit wavelength selection is performed. Note, however, that in the transparent case, the wavelength selection in the virtual link firstly allocated for a demand constrains those selections in the remainder ones.

Fig. 2 shows the number of allocated demands as a function of the demand set size in both transparent and opaque VON scenarios. Due to the impracticality of TVONA_ILP for large D sizes, TVONA_GRASP has been used in the transparent case. The presented results have been averaged over 100 executions with newly generated demand sets per execution, which provides us with statistically relevant results. Again, $U_{\{d\}} = 1, \forall d \in D$. As observed, both OVONA_ILP and TVONA_GRASP outperform their shortest-path-based counterparts, showing more pronounced improvements as the number of offered demands increases. In the transparent case, for example, while a 12.8% improvement is achieved for $|D| = 10$, this one raises up to around 30% for $|D| = 30$. Even more pronounced improvements are observed in the opaque case, ranging from 9.3% for $|D| = 10$ to 42.8% for $|D| = 30$. This is due to the fact that more flexibility is given to the VON allocation in the opaque case, as different sets of wavelengths can be used to allocate each VON virtual link.

So far, we concluded that TVONA_ILP is strongly affected by D size. To complete our study, we have analyzed how TVONA_ILP and OVONA_ILP execution times scale as $U_{\{d\}}$ increases. To this goal, we have fixed $|D| = 10$ and 100 executions of both models have been launched, with newly gen-

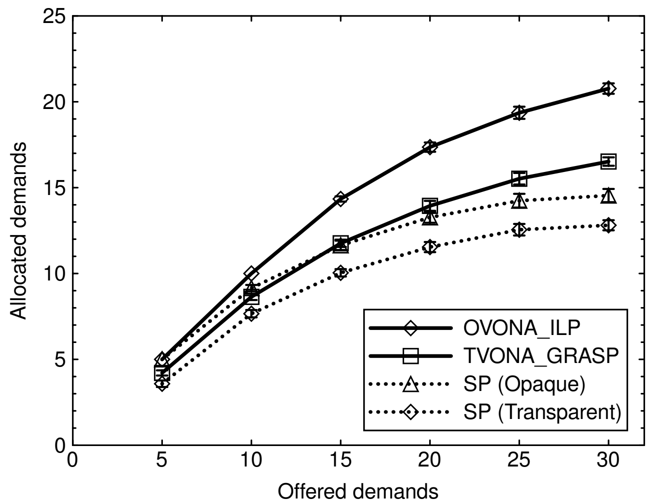

Fig. 2. Average number of allocated demands as a function of the offered demands set size.

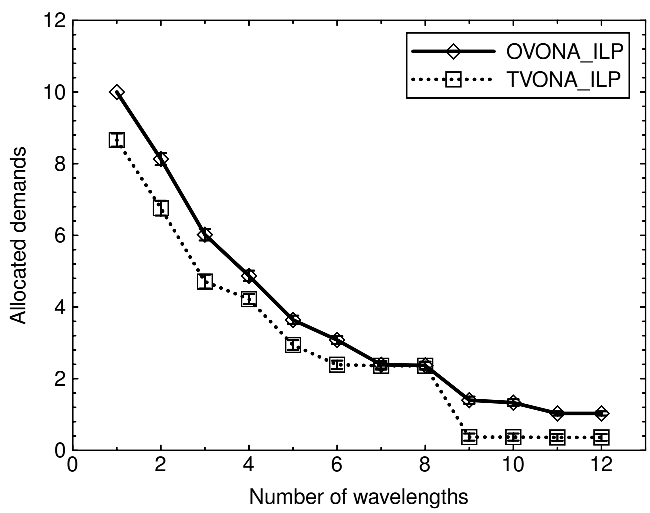

Fig. 3. Average number of allocated demands as a function of the number of wavelengths/demand.

erated demand sets per execution. Interestingly, the obtained results reveal us that neither TVONA_ILP nor OVONA_ILP show noticeable scalability issues in this regard. For example, while the execution times of TVONA_ILP remain around 1530 s for $U_{\{d\}} = 1$, they even decrease down to 52.6 s and 43 s for $U_{\{d\}}$ equal to 6 and 12, respectively. In the case of OVONA_ILP, execution times remain in the same order of magnitude for all $U_{\{d\}}$ values (e.g., 14.2 s, 42.3 s and 31.7 s for $U_{\{d\}}$ equal to 1, 6 and 12, respectively). Fig. 3 shows the number of VONs finally allocated in all experimented scenarios. As expected, both models present similar behavior in this sense, due to the resource scarcity in the optical network substrate as each VON requests more and more resources. A noteworthy point in the graph appears when $U_{\{d\}}$ changes from 8 to 9. Indeed, having 8 wavelengths per physical link, virtual links must unavoidably be mapped over multiple physical paths between the virtual link endpoints, which limits to a large extent the number of VONs that can be allocated over the optical network substrate.

IV. Concluding Remarks

This paper proposed exact ILP formulations to solve the VONA problem over a transparent optical network substrate. Such models target at different variants of the problem, depending on whether transparent or opaque services have to be offered over the VONs. For the transparent case, a light-weight heuristic was additionally proposed. The obtained results validated the benefits of these contributions against simpler VON allocation strategies.

Acknowledgment

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 project Enhanced Optical Networks Featuring Adaptable and Highly Scalable Multi-Granular Transport (ELASTIC) (TEC2011-27310).

References

[1] N. M. Mosharaf Kabir Chowdhury and R. Boutaba, “Network virtualization: state of the art and research challenges,” *IEEE Commun. Mag.*, 2009.

[2] P. V.-B. Primet, *et al.*, “Virtualizing and scheduling optical network infrastructure for emerging IT services,” *J. Optical Commun. and Networking*, 2009.

[3] S. Peng *et al.*, “An impairment-aware virtual optical network composition mechanism for future Internet,” *2011 European Conference and Exhibition on Optical Communication*.

[4] N. M. Mosharaf Kabir Chowdhury *et al.*, “Virtual network embedding with coordinated node and link mapping,” *2009 IEEE International Conference on Computer Communications*.

[5] Z. Cai *et al.*, “Virtual network embedding For evolving networks,” *2010 IEEE Global Telecommunications Conference*.

[6] O. Rival and A. Morea, “Cost-efficiency of mixed 10-40-100Gb/s networks and elastic optical networks,” *2011 Optical Fiber Communication Conference and Exposition*.

[7] M. G. C. Resende and C. C. Ribeiro, *GRASP: Greedy Randomized Adaptive Search Procedures*, 2nd edition. Springer, 2011.

[8] S. De Maesschalck, *et al.*, “Pan-European optical transport networks: an availability-based comparison,” *Photonic Network Commun.*, 2003.

[9] D. Cordeiro, *et al.*, “Random graph generation for scheduling simulations,” *2010 International ICST Conference on Simulation Tools and Techniques*.