

Energy Efficient Virtual Network Embedding for Cloud Networks

Leonard Nonde, Taisir E. H. El-Gorashi, and Jaafar M. H. Elmirghani

Abstract—Network virtualization is widely considered to be one of the main paradigms for the future Internet architecture as it provides a number of advantages including scalability, on demand allocation of network resources, and the promise of efficient use of network resources. In this paper, we propose an energy efficient virtual network embedding (EEVNE) approach for cloud computing networks, where power savings are introduced by consolidating resources in the network and data centers. We model our approach in an IP over WDM network using mixed integer linear programming (MILP). The performance of the EEVNE approach is compared with two approaches from the literature: the bandwidth cost approach (CostVNE) and the energy aware approach (VNE-EA). The CostVNE approach optimizes the use of available bandwidth, while the VNE-EA approach minimizes the power consumption by reducing the number of activated nodes and links without taking into account the granular power consumption of the data centers and the different network devices. The results show that the EEVNE model achieves a maximum power saving of 60% (average 20%) compared to the CostVNE model under an energy inefficient data center power profile. We develop a heuristic, real-time energy optimized VNE (REOVINE), with power savings approaching those of the EEVNE model. We also compare the different approaches adopting an energy efficient data center power profile. Furthermore, we study the impact of delay and node location constraints on the energy efficiency of virtual network embedding. We also show how VNE can impact the design of optimally located data centers for minimal power consumption in cloud networks. Finally, we examine the power savings and spectral efficiency benefits that VNE offers in optical orthogonal division multiplexing networks.

Index Terms—Cloud networks, energy efficient networks, IP over WDM networks, MILP, network virtualization, optical OFDM, virtual network embedding.

I. INTRODUCTION

THE ever growing uptake of cloud computing as a widely accepted computing paradigm calls for novel architectures to support QoS and energy efficiency in networks and data centers [1]. Estimates indicate that in the long term, if current trends continue, the annual energy bill paid by data center operators will exceed the cost of equipment [2]. Given the ecological and economic impact, both academia and industry are focusing efforts on developing energy efficient paradigms for cloud computing [3].

Manuscript received July 23, 2014; revised October 25, 2014 and December 6, 2014; accepted December 6, 2014. Date of publication December 17, 2014; date of current version March 13, 2015. This work was supported by the Engineering and Physical Sciences Research Council (EPSRC), INTERNET (EP/H040536/1), and STAR (EP/K016873/1).

The authors are with the School of Electronic and Electrical Engineering, University of Leeds, Leeds LS2 9JT, U.K.

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/JLT.2014.2380777

In [4], the authors stated that the success of future cloud networks where clients are expected to be able to specify the data rate and processing requirements for hosted applications and services will greatly depend on network virtualization. The form of cloud computing service offering under study here is Infrastructure as a Service (IaaS). IaaS is the delivery of virtualized and dynamically scalable computing power, storage and networking on demand to clients on a pay as you go basis. Network virtualization allows multiple heterogeneous virtual network architectures (comprising virtual nodes and links) to coexist on a shared physical platform, known as the substrate network which is owned and operated by an infrastructure provider (InP) or cloud service provider whose aim is to earn a profit from leasing network resources to its customers (Service Providers (SPs)) [5]. It provides scalability, customised and on demand allocation of resources and the promise of efficient use of network resources [6]. Network virtualization is therefore a strong proponent for the realization of an efficient IaaS framework in cloud networks. InPs should have a resource allocation framework that reserves and allocates physical resources to elements such as virtual nodes and virtual links. Resource allocation is done using a class of algorithms commonly known as “virtual network embedding (VNE)” algorithms. The dynamic mapping of virtual resources onto the physical hardware maximizes the benefits gained from existing hardware [7]. The VNE problem can be either *Offline* or *Online*. In offline problems [8] all the virtual network requests (VNRs) are known and scheduled in advance while for the online problem, VNRs arrive dynamically and can stay in the network for an arbitrary duration [9], [10]. Both online and offline problems are known to be NP-hard. With constraints on virtual nodes and links, the offline VNE problem can be reduced to the NP-hard multiway separator problem [11], as a result, most of the work done in this area has focused on the design of heuristic algorithms and the use of networks with minimal complexity when solving mixed integer linear programming (MILP) models.

Network virtualization has been proposed as an enabler of energy savings by means of resource consolidation [8], [12]–[15]. In all these proposals, the VNE models and/or algorithms do not address the link embedding problem as a multi-layer problem spanning from the virtualization layer through the IP layer and all the way to the optical layer. Except for the authors in [14], the others do not consider the power consumption of network ports/links as being related to the actual traffic passing through them. On the contrary, we take a very generic, detailed and accurate approach towards energy efficient VNE (EEVNE) where we allow the model to decide the optimum approach to minimize the total network and data centers server power consumption.

We consider the granular power consumption of various network elements that form the network engine in backbone networks as well as the power consumption in data centers. We develop a MILP model and a real-time heuristic to represent the EEVNE approach for clouds in IP over WDM networks with data centers. We study the energy efficiency considering two different power consumption profiles for servers in data centers;

An energy inefficient power profile and an energy efficient power profile. Our work also investigates the impact of location and delay constraints in a practical enterprise solution of VNE in clouds. Furthermore we show how VNE can impact the design problem of optimally locating data centers for minimal power consumption in cloud networks.

The remainder of the paper is organized as follows: We examine related work in Section II. The EEVNE MILP model and heuristic are introduced in Section III and Section IV, respectively. The EEVNE approach is evaluated and compared to the CostVNE and the VNE-EA approaches in Section V. In Section VI we re-evaluate the energy efficiency of the EEVNE approach considering an energy efficient data center power profile and in Section VII, we investigate the impact of location and delay constraints on the energy efficiency of VNE. In Section VIII we solve the VNE design problem of optimally locating data centers for minimal power consumption in cloud networks. We introduce VNE in O-OFDM cloud networks in Section IX and we finally conclude the paper in Section X.

II. RELATED WORK

The VNE problem has been extensively investigated in the literature from different perspectives such as load balancing in the substrate network [14], [16], [17] and efficient use of resources [9]. We have investigated the VNE problem for profit maximization by InPs in cloud networks [18]. We developed a MILP model and showed that higher acceptance rates do not necessarily lead to higher profit due to the high cost associated with accepting some of the VNRs. We also showed that despite the profit maximized model showing slightly higher power consumption than the power minimized model, it still performed better because the profit achieved by the profit maximized model is higher than the increase in power consumption when the latter is expressed in monetary terms. In [19] the authors developed a MILP model which attempts to minimize the bandwidth cost of embedding a VNR. In [12] the virtual network embedding energy aware (VNE-EA) model minimized the energy consumption by imposing the notion that the power consumption is minimized by switching off substrate links and nodes. The authors also assume that the power saved in switching off a substrate link is the same as the power saved by switching off a substrate node. In [13] the authors assumed that the power consumption in the network is insensitive to the number of ports used. They also seek to minimize the number of active working nodes and links. Botero and Hesselbach [14] have proposed a model for energy efficiency using load balancing and have also developed a dynamic heuristic that reconfigures the embedding for energy efficiency once it is performed. They have implemented and evaluated their MILP models and heuristic algorithms using the

ALEVIN Framework [20]. The ALEVIN Framework is a good tool for developing, comparing and analyzing VNE algorithms. We are however unable to compare our model and heuristic to the implemented algorithms on the platform for the following reasons:

1. Our input parameters are not compatible to the existing models and algorithms on the platform. Extensive extensions to the algorithms and models would be needed for them to include the optical layer. Our parameters include among others; the distance in km between links for us to determine the number of EDFA's or Regenerators needed on a link, the wavelength rate, the number of wavelengths in a fiber, the power consumption of EDFAs, transponders, regenerators, router ports, optical cross connects, multiplexers, de-multiplexers, etc.
2. The assumptions made in the calculation of power in our model and the models on the platform are different. We define the power consumption to its fine granularity to include power consumed due to traffic on each element that forms the network engine. One of our main contributions in this work is the inclusion of the optical layer in link embedding which is currently not supported by any of the algorithms on the ALEVIN platform.

Wang *et al.* [8] developed a generalized power consumption model of embedding a VNR and formulated it as a MILP model; however, they also assumed that the power consumption of the network ports is independent of traffic. In [15] the authors propose a trade-off between maximizing the number of VNRs that can be accommodated by the InP and minimizing the energy cost of the whole system. They propose embedding requests in regions with the lowest electricity cost.

III. ENERGY EFFICIENT VIRTUAL NETWORK EMBEDDING

A. Virtual Networks in IP Over WDM Networks

The VNE problem defines how virtualized resources should be realized onto the substrate network. As described in [7], VNRs are annotated with node and link demands and in the same way, the substrate network is annotated with node and link resources. Demands and resources then have to be matched in order to achieve complete embeddings. In Fig. 1, the VNRs 1, 2 and 3 with node and link demands are to be embedded onto the substrate network which is an IP over WDM network with data centers. The IP over WDM network is composed of two layers, the IP layer and the optical layer [21]. Successful embedding of VNRs' link demands will therefore need resources both in the IP layer and the optical layer. IP routers aggregate traffic from VNRs and in each substrate node IP routers are connected to optical switches which are connected by optical fiber links. The optical layer provides the large bandwidth required for communication between IP routers. On each fiber, a pair of multiplexers/demultiplexers is used to multiplex/demultiplex wavelengths. The transponders provide OEO processing for full wavelength conversion at each switching node. In addition, for long-distance transmission, erbium-doped fiber amplifiers (EDFAs) are used to amplify the optical signal in each fiber. Each substrate node is considered to host a data center in addition

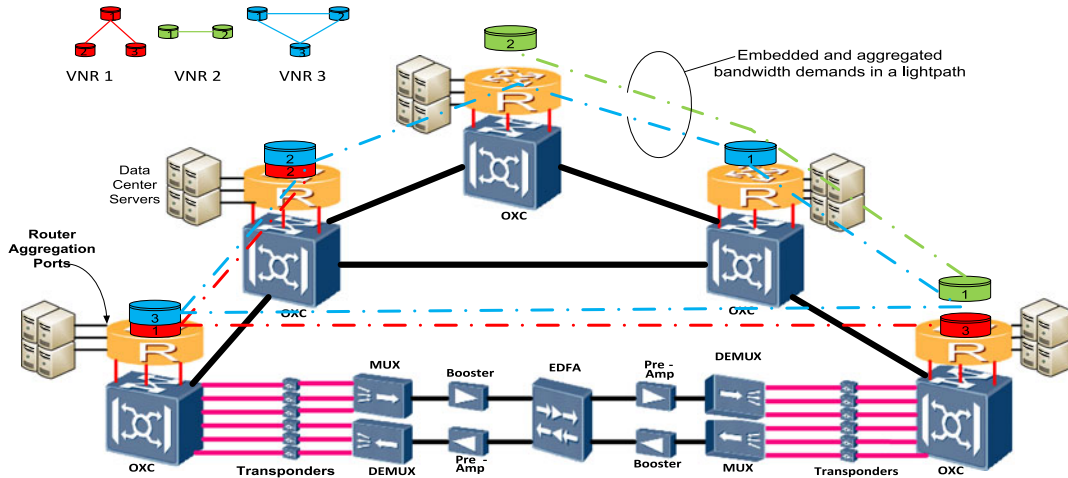


Fig. 1. Virtual network embedding in IP over WDM architecture.

to the IP and optical equipment. The node demands of the VNRs are embedded in the data centers. When a virtual node is embedded in the substrate network, its CPU demands instantiate virtual servers in the data center and its bandwidth demands instantiate a virtual router in the core router such that the requester of the service is granted control of both the virtual servers and virtual router and has the ability to configure any protocols and run any applications.

B. MILP Model for EEVNE

In this section we develop a MILP model to minimize the total power consumption of the IP over WDM architecture with data centers shown in Fig. 1 by optimizing the embedding of VNRs.

The substrate network is modeled as a weighted undirected graph $G = (N, L)$ where N is the set of substrate nodes and L is the set of substrate links. Each node or link in the substrate network is associated with its own resource attributes. The VNR v is represented by the graph $G^v = (R^v, L^v)$ where R^v is the set of virtual nodes and L^v is the set of virtual links. In Fig. 2, we illustrate how demands in a VNR are mapped onto the substrate network across multiple layers. We clearly show where some of the variables and parameters used are located across the layers. Before introducing the model, we define the sets, parameters and variables used;

Sets:

- V Set of VNRs;
- R Set of nodes in a VNR;
- N Set of nodes in the substrate network;
- N_m Set of neighbor nodes of node m in the optical layer;

Parameters:

- s and d Index the source and destination nodes in the topology of a VNR;
- b and e Index source and destination substrate nodes of an end to end traffic demand aggregated from all VNRs;
- i and j Index the nodes in the virtual lightpath topology (IP layer) of the substrate network;




Fig. 2. Illustration of VNE across multiple layers in IP/WDM.

- m and n Index the nodes in the physical topology (optical layer) of the substrate network;
- CPU_b Available CPU resources at substrate node (data center) b ;
- N_b Number of servers at substrate node b ;
- $N_{CPU}^{v,s}$ Number of the servers required by VNR v in node s ;
- $CPU_b^{v,s}$ $CPU_b^{v,s} = \frac{N_{CPU}^{v,s}}{N_b}$, is the percentage CPU utilization of virtual node s of VNR v when embedded in node b ;

$H^{v,s,d}$	Bandwidth requested by VNR v on virtual link (s,d) ;
B	Wavelength rate;
W	Number of wavelengths per fiber;
$D_{m,n}$	Length of the physical link (m,n) ;
DO_b	$DO_b = 1$ if the data center at node b has been activated by previous embeddings;
$F_{m,n}$	Number of fibers in physical link (m,n) ;
$EA_{m,n}$	Number of EDFAs in physical link (m,n) . Typically $EA_{m,n} = \left\lceil \left(\frac{D_{m,n}}{S} \right) - 1 \right\rceil + 2$, where S is the distance between two neighboring EDFAs;
PR	Power consumption of a router Port;
PT	Power consumption of a transponder;
PE	Power consumption of an EDFA;
PO_m	Power consumption of an optical switch at node m ;
PMD	Power consumption of a multi/demultiplexer;
DM_m	Number of multi/demultiplexers at node m ;
P_{idle}	Idle server power consumption of a data center;
μ	Data center server power consumption per 1% CPU load.

Variables:

$\delta_b^{v,s}$	$\delta_b^{v,s} = 1$, if node s of VNR v is embedded in a substrate node b , otherwise $\delta_b^{v,s} = 0$;
Ψ^v	$\Psi^v = 1$, if all the nodes of a VNR v are fully embedded in the substrate network, otherwise $\Psi^v = 0$;
$\rho_{b,e}^{v,s,d}$	$\rho_{b,e}^{v,s,d} = 1$, if the embedding of virtual nodes s and d of virtual request v in substrate nodes b and e , respectively is successful and a link b,e is established if a virtual link s,d of VNR v exists.
$\omega_{b,e}^{v,s,d}$	$\omega_{b,e}^{v,s,d}$ is the XOR of $\delta_b^{v,s}$ and $\delta_e^{v,d}$, i.e. $\omega_{b,e}^{v,s,d} = \delta_b^{v,s} \oplus \delta_e^{v,d}$;
$L_{b,e}$	Traffic demand between the substrate node pair (b,e) aggregated from all embedded VNRs;
Φ^v	$\Phi^v = 1$, if all the links of VNR v are fully embedded in the substrate network, otherwise $\Phi^v = 0$;
$L_{i,j}^{b,e}$	Amount of traffic demand between node pair (b,e) passing through the lightpath (i,j) in the substrate network;
$C_{i,j}$	Number of wavelengths in lightpath (i,j) in the substrate network;
Q_b	Number of aggregation router ports at node b ;
$W_{m,n}^{i,j}$	The number of wavelengths of lightpath (i,j) passing through a physical link (m,n) ;
$W_{m,n}$	Number of wavelengths in physical link (m,n) ;
$\lambda_{m,n}$	$\lambda_{m,n} = 1$, if the physical link (m,n) is activated otherwise $\lambda_{m,n} = 0$;
K_b	$K_b = 1$ if the data center in node b is activated, otherwise $K_b = 0$.

The total power consumption is composed of the power consumption in data centers and the power consumption in the network.

As discussed in [22] the CPU utilization is the main contributor to the power consumption variations in a server. Therefore the CPU utilization of servers in a data center is also the main contributor to the variation in the power consumption of a data center. We only consider the servers power consumption when calculating the power consumption of data centers as the largest

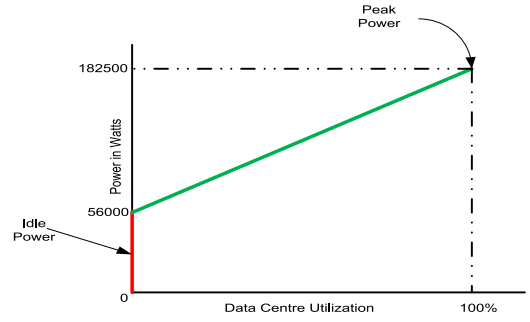


Fig. 3. Data center servers power profile with 500 Dell R720 servers.

proportion of data center power is drawn by Servers. According to [23], 59% of the power delivered in the data center goes to IT equipment. It also states that decreasing the power draw of servers clearly has the largest impact on the power cost in data centers. While we appreciate the efforts done by other researchers and institutions like Google and Microsoft in improving the energy consumption due to cooling and the use of efficient power supplies, our efforts in the paper are restricted to the power consumed in servers and high capacity network equipment in data centers. It is worth noting here that large data centers have now achieved power usage effectiveness (PUE) approaching 1, for example [24] and [25] report PUE values of 1.12 and 1.125 respectively, where it is clear that the data center servers (and communication equipment) power consumption now dominates in large data centers compared to cooling, lighting and power supply units losses. The power consumption due to the local area network (LAN) inside the data center also contributes to the power consumption in the data center. However, as we will show later, the LAN power consumption is very low compared to the power consumption due to servers. We therefore consider the energy inefficient data center power profile shown in Fig. 3 which considers only the power consumption due to servers. The power profile in Fig. 3 is constructed from a data center containing 500 Dell Power Edge R720 Servers [26] each with idle power rated at 112W and 365W at full load. The power consumption in a data center due to the embedding of node s of VNR v is given as:

$$PD_b^{v,s} = \begin{cases} P_{idle} + \mu \cdot CPU_b^{v,s}, & \text{if the data centre at } b \text{ is ON} \\ 0, & \text{otherwise.} \end{cases}$$

The power consumption of data centers due to servers is given as:

$$PD = \sum_{b \in N} \sum_{v \in V} \sum_{s \in R} CPU_b^{v,s} \cdot \delta_b^{v,s} \cdot \mu + K_b \cdot P_{idle}.$$

The network power consumption under non-bypass where lightpaths passing through an intermediate node are terminated and forwarded to the IP router [21] is composed of:

Power consumption of router ports:

$$\sum_{m \in N} PR \cdot \left(Q_m + \sum_{n \in N_m} W_{m,n} \right).$$

Power consumption of transponders:

$$\sum_{m \in N} \sum_{n \in N_m} PT \cdot W_{m,n}.$$

Power consumption of EDFAs:

$$\sum_{m \in N} \sum_{n \in N_m} PE \cdot EA_{m,n} \cdot F_{m,n} \cdot \lambda_{m,n}.$$

Power consumption of optical switches:

$$\sum_{m \in N} PO_m.$$

Power Consumption of multi/demultiplexers:

$$\sum_{m \in N} PMD \cdot DM_m$$

The model is defined as follows:

Objective:

Minimize the overall power consumption in the network and data centers given as:

$$\begin{aligned} & \sum_{m \in N} PR \cdot \left(Q_m + \sum_{n \in N_m} W_{m,n} \right) + \sum_{m \in N} \sum_{n \in N_m} PT \cdot W_{m,n} \\ & + \sum_{m \in N} \sum_{n \in N_m} PE \cdot EA_{m,n} \cdot F_{m,n} \cdot \lambda_{m,n} + \sum_{m \in N} PO_m \\ & + \sum_{m \in N} PMD \cdot DM_m \\ & + \sum_{b \in N: DO_b=0} \sum_{v \in V} \sum_{s \in R} CPU_b^{v,s} \cdot \delta_b^{v,s} \cdot \mu + K_b P_{idle}. \end{aligned} \quad (1)$$

Subject to:

VNRs constraints:

$$\sum_{v \in V} \sum_{s \in R} CPU_b^{v,s} \cdot \delta_b^{v,s} \leq CPU_b \quad b \in N. \quad (2)$$

Constraint (2) ensures that the CPU demands allocated to a substrate node do not exceed the capacity of its data center

$$\sum_{b \in N} \delta_b^{v,s} = 1 \quad \forall v \in V, \quad \forall s \in R. \quad (3)$$

Constraint (3) ensures that each node in a VNR is embedded only once in the substrate network

$$\begin{aligned} \delta_b^{v,s} + \delta_e^{v,d} &= \omega_{e,b}^{v,d,s} + 2 \cdot \rho_{b,e}^{v,s,d} \\ \forall v \in V, \quad \forall b, e \in N : b \neq e, \quad \forall s, d \in R : s \neq d. \end{aligned} \quad (4)$$

Constraint (4) ensures that virtual nodes connected in the VNR are also connected in the substrate network. We achieve this by introducing a binary variable $\omega_{e,b}^{v,d,s}$ which is only equal to 1 if $\delta_b^{v,s}$ and $\delta_e^{v,d}$ are exclusively equal to 1 otherwise it is zero

$$\rho_{b,e}^{v,s,d} = \rho_{b,e}^{v,d,s}$$

$$\forall v \in V, \quad \forall b, e \in N : b \neq e \quad \forall s, d \in R : s \neq d. \quad (5)$$

Constraint (12) ensures that the bidirectional traffic flows are maintained after embedding the virtual links

$$\sum_{v \in V} \sum_{s \in N} \sum_{d \in N: s \neq d} H^{v,s,d} \cdot \rho_{b,e}^{v,s,d} = L_{b,e} \quad \forall b, e \in N : b \neq e. \quad (6)$$

Constraint (6) generates the traffic demand matrix resulting from embedding the VNRs in the substrate network and ensures that no connected nodes from the same VNR are embedded in the same substrate node

$$\sum_{b \in N} \sum_{s \in R} CPU_b^{v,s} \cdot \delta_b^{v,s} = \Psi^v \sum_{b \in N} \sum_{s \in R} CPU_b^{v,s} \quad \forall v \in V. \quad (7)$$

Constraint (7) ensures that nodes of a VNR are completely embedded meeting their CPU demands

$$\begin{aligned} & \sum_{b \in N} \sum_{e \in N: b \neq e} \sum_{s \in R} \sum_{d \in R: s \neq d} H^{v,s,d} \cdot \rho_{b,e}^{v,s,d} \\ & = \Phi^v \sum_{s \in R} \sum_{d \in R: s \neq d} H^{v,s,d} \quad \forall v \in V. \end{aligned} \quad (8)$$

Constraint (8) ensures the bandwidth demands of a request are completely embedded

$$\Phi^v = \Psi^v \quad \forall v \in V. \quad (9)$$

Constraint (9) ensures that both nodes and links of a VNR are embedded. Constraints (7)–(9) collectively ensure that a request is not partially embedded.

Flow conservation in the IP Layer

$$\sum_{j \in N: i \neq j} L_{i,j}^{b,e} - \sum_{j \in N: i \neq j} L_{j,i}^{b,e} = \begin{cases} L_{b,e}, & \text{if } i = b \\ -L_{b,e}, & \text{if } i = e \\ 0, & \text{otherwise.} \end{cases} \quad (10)$$

$$\forall b, e, i \in N : b \neq e$$

Constraint (10) represents the flow conservation constraint for the traffic flows in the IP Layer.

Lightpath capacity constraint

$$\sum_{b \in N} \sum_{e \in N: b \neq e} L_{i,j}^{b,e} \leq C_{i,j} \cdot B \quad \forall i, j \in N : i \neq j. \quad (11)$$

Constraint (11) ensures that the sum of all traffic flows through a virtual link does not exceed its capacity.

Flow conservation in the optical layer

$$\sum_{n \in N_m} W_{m,n}^{i,j} - \sum_{n \in N_m} W_{n,m}^{i,j} = \begin{cases} C_{i,j}, & \text{if } m = i \\ -C_{i,j}, & \text{if } m = j \\ 0, & \text{otherwise.} \end{cases} \quad (12)$$

Constraint (12) ensures the conservation of flows in the optical layer.

Physical link capacity constraints

$$\sum_{i \in N} \sum_{j \in N: i \neq j} W_{m,n}^{i,j} \leq W \cdot F_{m,n} \quad \forall m \in N, n \in N_m \quad (13)$$




Fig. 4. Grouping of nodes of a VNR 1 (a) A VNR requiring a minimum of two substrate nodes (b) A VNR requiring a minimum of three substrate nodes.

$$\sum_{i \in N} \sum_{j \in N: i \neq j} W_{m,n}^{i,j} = W_{m,n} \quad \forall m \in N, n \in N_m \quad (14)$$

Constraints (13) and (14) represent the physical link capacity constraints. Constraint (13) ensures that the number of wavelengths in a physical link does not exceed the capacity of fibers in the physical links. Constraint (14) gives the total number of wavelength channels used in a physical link.

Aggregation ports constraint

$$Q_b = \sum_{e \in N: e \neq b} L_{b,e} / B \quad \forall b \in N. \quad (15)$$

Constraint (15) gives the number of aggregation router ports at each node in the substrate network.

IV. THE REAL-TIME ENERGY OPTIMIZED VIRTUAL NETWORK EMBEDDING (REOVINE) HEURISTIC

The (real-time energy optimized virtual network embedding) (REOVINE) heuristic provides real-time implementation of the EEVNE approach. The flow chart of the heuristic is shown in Fig. 5. The heuristic obtains the graph $G^v = (R^v, L^v)$ of VNR v . The nodes of v are grouped such that those that are not connected to each other in the VNR topology graph G^v are put into one group. Separate groups are created for all the other nodes that are connected to each other to avoid embedding nodes that are connected in the VNR into the same substrate node. The heuristic determines the number of candidate substrate nodes required to embed a full VNR based on the number of node groups created. Fig. 4 shows examples of nodes grouping of a VNR with four nodes and three links. In Fig. 4(a), Nodes 1 and

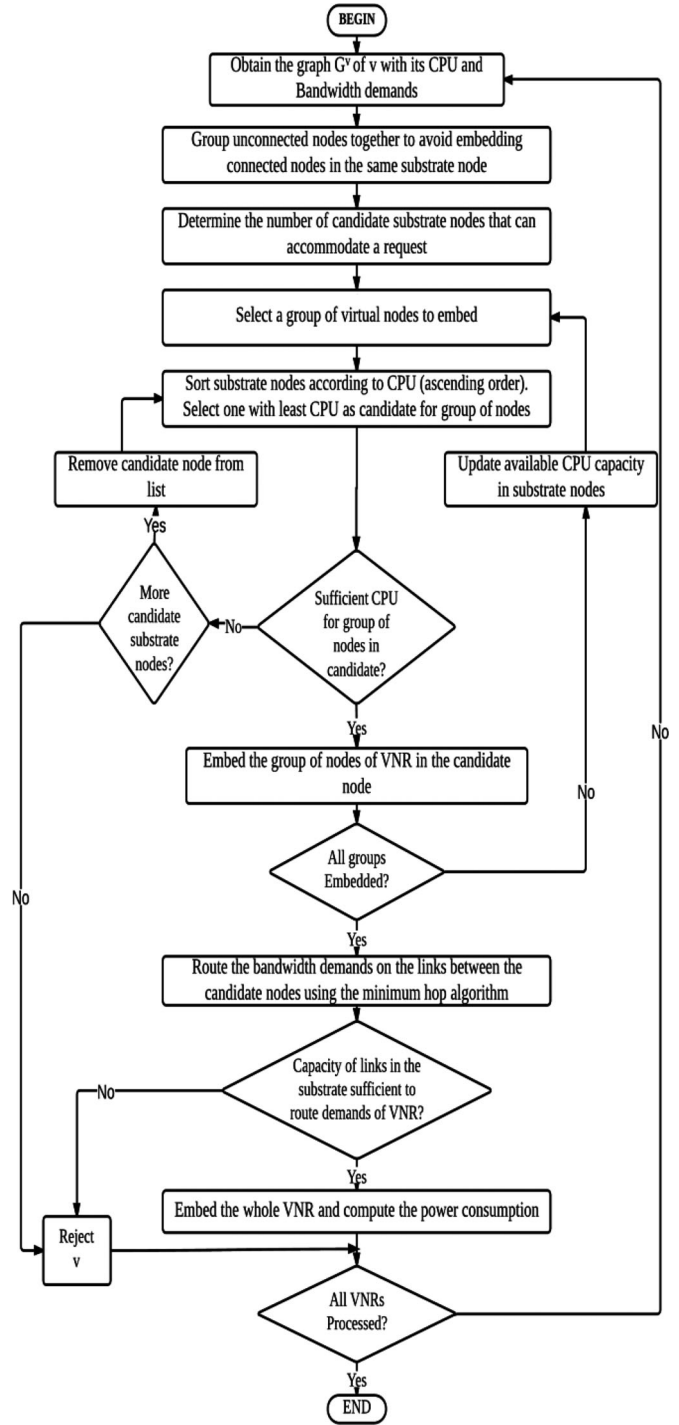


Fig. 5. REOVINE heuristic flow chart.

3 are not connected and therefore, they are grouped together to form a group. Similarly nodes 2 and 4 are grouped together to form another group. With this grouping, the VNR requires a minimum of two substrate nodes for the node embedding to be successful. Fig. 4(b) shows another grouping option where the unconnected nodes 1 and 4 are placed in one group. In this case, node 3 has to be in a group of its own as it cannot be in the same group as it is connected to node 2 and cannot be in the same group as nodes 1 and 4, as it is connected to node 4.

Similarly node 2 has to be in a group of its own. The grouping in Fig. 4(b) results in 3 substrate nodes, whereas Fig. 4(a) grouping needs only 2 substrate nodes to embed the VNR. The grouping is done to achieve maximum compaction, i.e. minimum number of substrate nodes. Therefore the grouping in Fig. 4(a) will be selected as the optimal grouping. From the example above we can summarize the rules of grouping as follows:

- i. Nodes that are not connected in the virtual request can be grouped in one group and embedded in the same substrate node.
- ii. Nodes that are connected to each other must belong to different groups.
- iii. The grouping process is only finished when all nodes in the virtual request have found a group.
- iv. There may be more than one grouping option. The grouping option that leads to the minimum number of substrate nodes is selected.

The substrate nodes are sorted according to the available CPU capacity in ascending order. Because of the energy inefficient power profile of the data centers (large idle power), the heuristic tries to consolidate the embeddings in data centers by filling the ones with least residual capacity before switching on others. The substrate node with the least available CPU capacity is selected as the candidate to embed the first group of nodes. The order in which the groups are selected is not important because it does not affect the success or failure of the embedding of nodes. However, we select the grouping option that requires the minimum number of substrate nodes. If the candidate substrate node has sufficient capacity, the virtual nodes will be embedded in that substrate node; otherwise the heuristic will try to embed the virtual nodes in another substrate node. If, however, we have exhausted all the candidate substrate nodes, the request will be rejected and the heuristic will proceed to embed another request. After embedding all the nodes of the VNR, the heuristic tries to route the bandwidth demands on the links between the candidate substrate nodes using the minimum hop algorithm to minimize the network power consumption. The request will be blocked if any of the links does not have sufficient capacity to accommodate the traffic. If the traffic is successfully routed, the request is accepted. The heuristic calculates the power consumption of the accepted request and continues to embed other VNRs.

V. PERFORMANCE EVALUATION

To evaluate the performance of the proposed model and heuristic, the NSFNET network is used as the substrate network. NSFNET comprises 14 nodes and 21 links as shown in Fig. 6. We consider a scenario in which each node in NSFNET hosts a small data center of 500 servers to offer cloud services. Table I shows the parameters used. The power consumption of the network devices we have used are consistent with our previous work in [27] which are derived from [28]–[32]. The IP router ports are the most energy consuming devices in the network.

We have adopted the Dell Power Edge R720 [26] server power specifications. We adapted the CostVNE model [20] and the VNE-EA model [12] for the IP over WDM network architecture




Fig. 6. The NSFNET network.

TABLE I
EVALUATION SCENARIO PARAMETERS

Distance between two neighboring EDFAs	80 (km)
Number of wavelengths in a fiber (W)	32
Number of Fibers per link ($F_{m,n}$)	1
Power consumption of a transponder (PT)	73 (W)
Power consumption of a single router port (PR)	1000(W)
Power consumption of an EDFA (PE)	8 (W)
Power consumption of an optical switch (PO)	85(W)
Power consumption of a multi/demultiplexer (PMD)	16(W)
Dell Server full load power consumption [26]	365(W)
Dell Server idle power consumption [26]	112(W)
Data Center idle power consumption (500 servers)	56 000(W)

and compared their performance to our EEVNE model and the REOVNE heuristic in terms of power consumption and number of accepted requests. The objective functions of the two models are as follows:

CostVNE Model Objective

The CostVNE model minimizes the substrate resources allocated to a VNR. Since the amount of CPU resources allocated to a request cannot be reduced through consolidation, the model only optimizes the use of bandwidth on the links by consolidating wavelengths. The objective of the CostVNE model is given as:

$$\text{minimize } \sum_{m \in N} \sum_{n \in N_m} W_{m,n}$$

VNE-EA Objective

The VNE-EA model [12] minimizes the energy consumption by minimizing the number of substrate links and nodes that are activated when embedding VNRs. The objective function of this model is given as:

$$\text{minimize } \sum_{m \in N: NO_m=0} \sigma_m + \sum_{m \in N} \sum_{n \in N_m: LO_{m,n}=0} \beta_{m,n}$$

where σ_m and $\beta_{m,n}$ are binary variables that indicate the active nodes and links in the substrate network, respectively. Note that NO_m and $LO_{m,n}$ are binary parameters indicating the already activated substrate nodes and links, respectively before the embedding of requests.

A. Embedding of VNRs of Uniform Load Distribution

The complexity of the MILP model grows exponentially with the number of nodes in the VNR. Therefore we consider nodes in a VNR of a limited number uniformly distributed between

2 and 6. The CPU demand of nodes in the VNR is uniformly distributed between 2% and 10% of the total CPU resources of a data center. The bandwidth on links of the VNR is also uniformly distributed between 10 and 130 Gb/s. We have considered embedding 50 VNRs onto the 14 node NSFNET substrate network. The AMPL software with the CPLEX 12.5 solver is used as the platform for solving the MILP models and REOVINE was implemented in Matlab. All the Models and REOVINE were executed on a PC with an Intel Core i5–2500 CPU, running at 3.30 GHz, with 8 GB RAM.

The results in Fig. 7 show the network power consumption, the data centers power consumption and the overall power consumption versus the number of VNRs. The requests arrive two at an instance and the existing requests are not reconfigured when embedding arriving requests.

The CostVNE model has resulted in the minimum network power consumption as it optimizes the use of bandwidth of the substrate network by consolidating wavelengths regardless of the number of data centers activated (see Fig. 7(a)). Compared to the EEVNE model, the CostVNE model has saved a maximum of 5% (average 3%) of the network power consumption. The EEVNE model, where the energy consumption is minimized by jointly optimizing the use of network resources and consolidating resources in data centers, has resulted in better power savings compared to the VNE-EA model where the power consumption is minimized by switching off substrate links and nodes. This is because the network power consumption is mainly a function of the number of wavelengths rather than the number of active links as the number of wavelengths used determines the power consumption of router ports and transponders, the most power consuming devices in the network (see Table I). The REOVINE heuristic approaches the EEVNE model in terms of the network power consumption.

Fig. 7(b) shows the power consumption of data centers under the different models and heuristic. As mentioned above, the CostVNE model does not take into account the number of activated data centers, therefore it performs very poorly as far as the power consumption in data centers is concerned. However, as the network gets fully loaded and all the data centers are activated, the EEVNE model loses its merit over the CostVNE model. For a limited number of requests, the VNE-EA model performs just as good as the EEVNE model. However as the number of requests increases, the VNE-EA model tends to route the virtual links through multiple hops to minimize the number of activated links and data centers and therefore consumes more power. The REOVINE heuristic also approaches the EEVNE performance in terms of the data centers power consumption.

Since the power consumption of the data centers is the main contributor to the total power consumption (given the power consumption of servers and network devices (see Table I)), the total power consumption, shown in Fig. 7(c), follows similar trends to that of the data centers. In the best case over a span of 50 VNRs, the EEVNE model saves 60% of the overall power consumption compared to the CostVNE model with an average saving of 21%. Compared to the VNE-EA model, the EEVNE model achieves maximum combined data center and network power savings of 9% (3% on average).

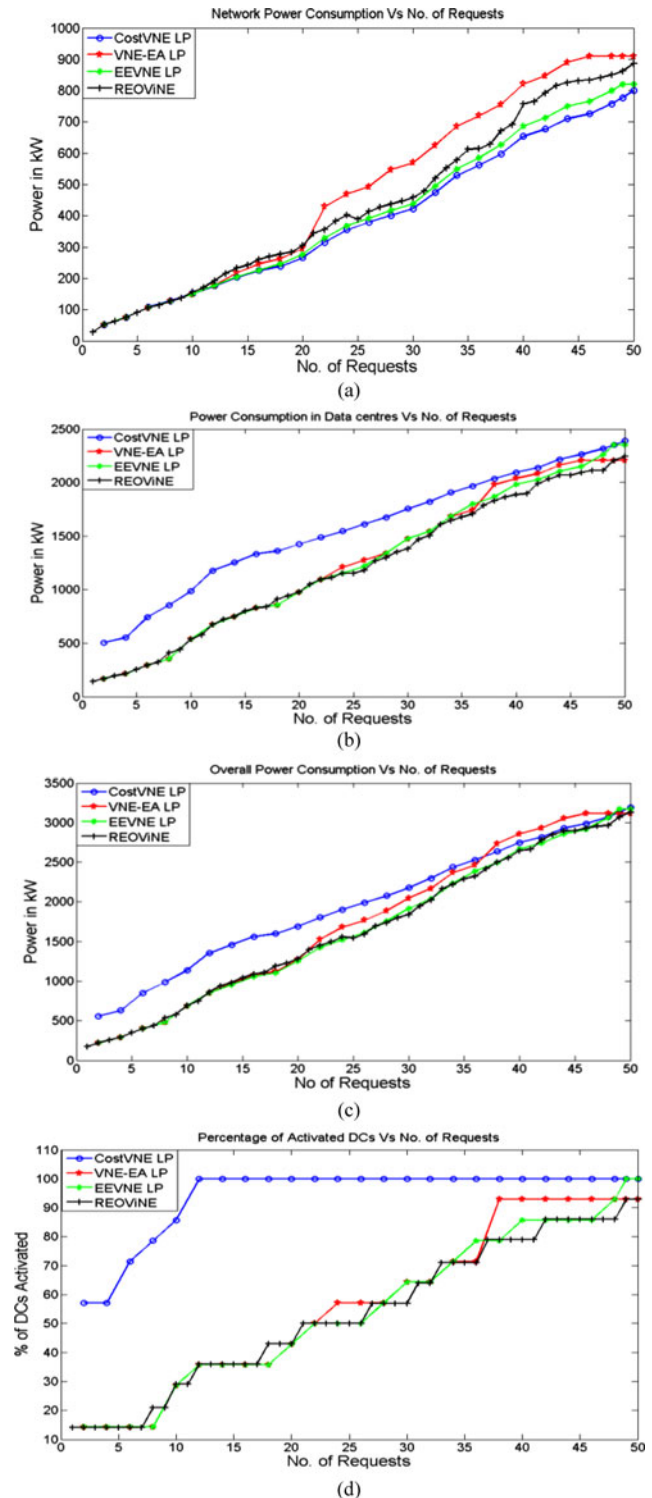


Fig. 7. (a) Network power consumption of the different approaches, (b) DCs power consumption of the different approaches, (c) Overall power consumption of the different approaches, (d) Data center activation of the different approaches.

Fig. 7(d) shows the activation rate of data centers as the mapping of VNRs increases in the network. The CostVNE model readily activates all the data centers while the other models gradually activate them. This ability to consolidate resources in the already active data centers is what leads to low power consumption in both the EEVNE and VNE-EA models.

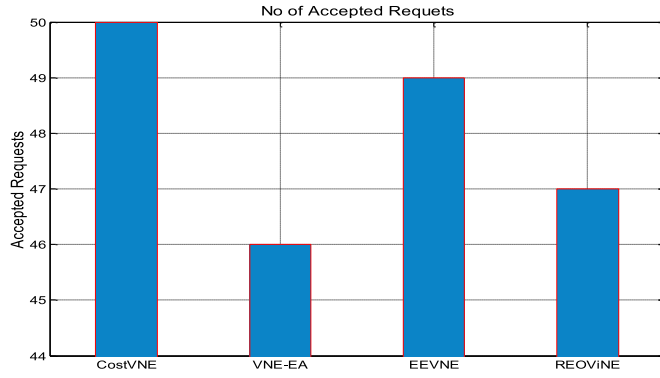


Fig. 8. Number of accepted requests.

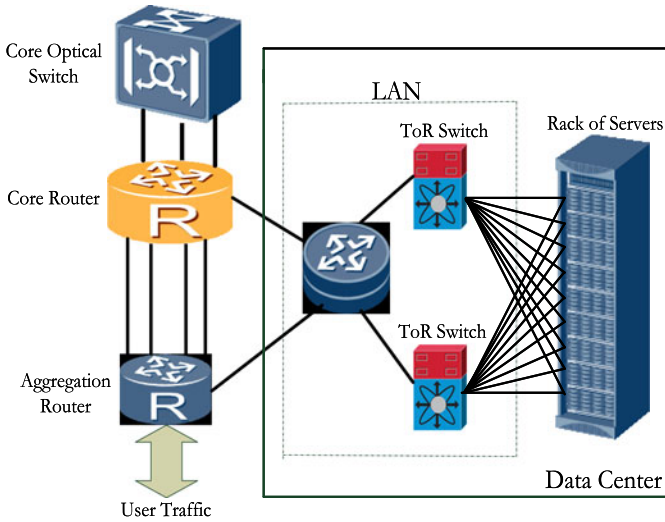


Fig. 9. Cloud data center architecture.

The revenue of a virtual network service provider can be greatly affected by its inability to accept new requests. Fig. 8 shows the number of accepted requests for the different models and the REOVINE heuristic. In an evaluation scenario with sufficient CPU resources to embed all virtual nodes, the CostVNE model has accepted all the 50 requests as it efficiently utilizes the network bandwidth. The improved energy efficiency achieved by the EEVNE model has a limited effect on its ability to accept VNRs. The EEVNE model has however only rejected one request. The REOVINE heuristic rejected three requests. The VNE-EA model is the worst performer in terms of accepted requests with four rejected requests. This arises as the VNE-EA model tries to fully use activated links before turning on more links therefore depleting the capacity on links to certain destinations. It therefore fails to map VNRs despite the substrate network having sufficient CPU resources.

As mentioned earlier in Section III-B, the power consumption of the LAN inside the data centre contributes to the overall power consumption of the data centre. Fig. 9 below shows a typical architecture of the cloud data centre and how it is connected to the core network. Typically, data centres are built very close to the core network nodes to benefit from the large band-

width capacities available from such nodes. If the data centre is serving users located at this node, the traffic flows from the servers onto the top of the rack (ToR) switches, then to the LAN router located inside the data centre and eventually to the aggregation router. The traffic serving users located at other core nodes flows to the core router and is then switched by the high capacity optical switch to the desired destination. The server power consumption is proportional to the amount of processing required by the embedded VNRs while the network power consumption inside the data center is proportional to the traffic generated by the requests. For example a low bandwidth and processing intensive VNR such as gaming will present low traffic and thus low power consumption in the network inside a data center but very high server power consumption. On the other hand, the network inside the data center will consume high power when embedding VNRs demanding high bandwidth and low processing like video streaming.

We have extended the EEVNE model to take into account the power consumption due to the LAN inside the data centre. We define the following additional variables and parameters in the model:

- ToR_PC LAN ToR switch power consumption;
- ToR_C LAN ToR switch switching capacity;
- ToR_PPB LAN ToR switch energy per bit, $ToR_PPB = ToR_PC/ToR_C$;
- R_PC LAN Router power consumption;
- R_C LAN Router routing capacity;
- R_PPB LAN Router energy per bit, $R_PPB = R_PC/R_C$;

The power consumption of both the LAN ToR switches and routers inside the data center (LANPC_DC) is proportional to the traffic inside the data center [33] and can be given as

$$LANPC_DC = \sum_{b \in N} \sum_{e \in N: e \neq b} L_{b,e} \cdot (ToR_PPB + R_PPB).$$

Therefore, the data center power consumption which includes both servers and the LAN (PDC) becomes:

$$PDC = \sum_{b \in N} \sum_{v \in V} \sum_{s \in R} CPU_b^{v,s} \cdot \delta_b^{v,s} \cdot \mu + K_b \cdot P_{idle} + \sum_{b \in N} \sum_{e \in N: e \neq b} L_{b,e} \cdot (ToR_PPB + R_PPB).$$

The Cisco Catalyst 4900M switch [34] with a capacity of 320 Gb/s and power rating of 1 kW was considered as the ToR switch. The Cisco 7613 Router [35] rated at 720 Gb/s and 4 kW was considered as the data center LAN router.

Fig. 10 shows the power consumption in data centers due to servers and due to the LAN. It can be seen that the power consumption due to servers dominates the power consumption due to the LAN. The LAN power consumption only accounts for 4% of the total power inside the data center. The inclusion of the LAN power consumption has also had no effect on the way the nodes and links are embedded across the network.

In Fig. 11 we see that the power consumption in the core network is much higher than the LAN power consumption inside the data centers. The LAN power consumption is only 15% of

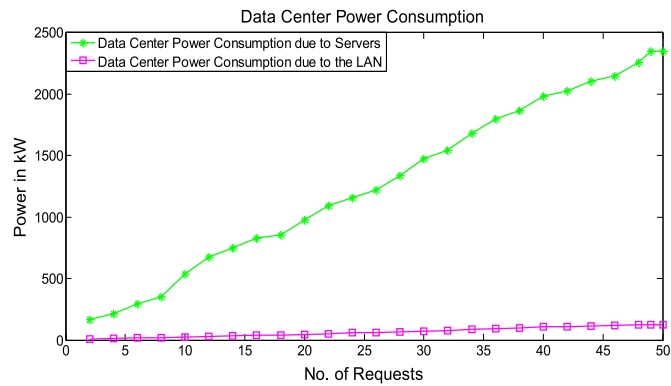


Fig. 10. Power consumption in dc due to servers and LAN.

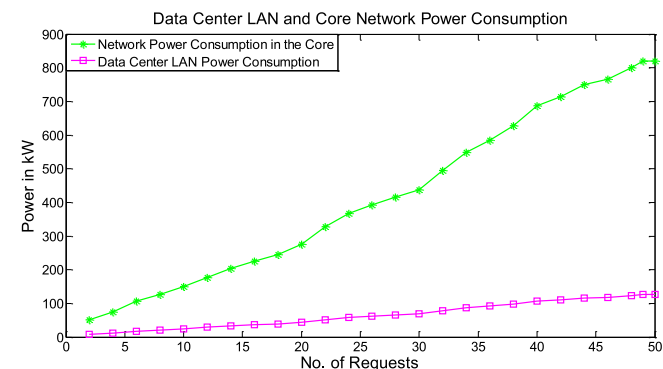


Fig. 11. Network power consumption in the Core and the dc LAN.

TABLE II
PERFORMANCE OF REOVINE AND EEVNE IN DIFFERENT NETWORKS

Substrate Network	Method	Request Acceptance	Running Time
NSFNET	EEVNE	49/50, 98%	85 000 s
Random Topology 10 Nodes 23 Links	EEVNE	34/50, 68%	64 800 s
Random Topology 8 Nodes 14 Links	EEVNE	28/50, 56%	39 600 s
NSFNET	REOVINE	47/50, 94%	7 s
USNET	REOVINE	49/50, 98%	40 s
ITALIA	REOVINE	49/50, 98%	28 s

the power consumption in the core when handling the same amount of traffic because the core network comprises more equipment and routes traffic over much longer distances.

In order to investigate the scalability of our energy efficient VNE approach, we evaluated our model (EEVNE) in small randomly generated topologies and our heuristic (REOVINE) in large network topologies since testing complex MILP models like ours on large networks often leads to huge complexities and unaffordable execution times. The topology matrices and associated distances of the randomly generated networks were produced in MATLAB with link distances randomly distributed between 100 and 1500 km. For the large networks we considered the 24-node 43-link USNET [21] substrate network and the 21-node 36-link Italian network [36]. Table II shows the acceptance ratio and the running times for the different topologies under the same load.




Fig. 12. REOVINE and EEVNE model overall power consumption in different networks.

TABLE III
LOAD DISTRIBUTIONS

Load	Distribution
High bandwidth and high CPU demands (Load 1)	CPU (9% to 10% of data center capacity) Bandwidth demand (100 to 130 Gb/s)
High bandwidth and low CPU demands (Load 2)	CPU (2% to 3% of data center capacity) Bandwidth demand (100 to 130 Gb/s)
Low bandwidth and low CPU demands (Load 3)	CPU (2% to 3% of data center capacity) Bandwidth demand (10 to 15 Gb/s)
Low bandwidth high CPU demands (Load 4)	CPU (9% to 10% of data center capacity) Bandwidth demand (10 to 15 Gb/s)

Fig. 12 shows the overall power consumption of the EEVNE model in the small random topologies and the REOVINE heuristic in the large network topologies. The power consumption in the USNET and Italian networks is slightly higher than the power consumption in the NSFNET network. This is due to the fact that both the USNET and the Italian networks have accepted more requests than the NSFNET network as seen in Table II because they have more network and data center resources. The distance between links in a particular network also has a bearing on the power consumption since longer distances call for the deployment of more EDFAs. The power curves for the small randomly generated topologies saturate after reaching the maximum number of requests that they can accommodate given their resources. This is because their resources get depleted much more quickly under the same load compared to the larger networks. For all the substrate networks, however, the general power consumption performance of both the EEVNE model and REOVINE heuristic is very consistent regardless of the substrate network used.

B. Embedding of VNRs of Non Uniform Load Distribution

In addition to the uniform bandwidth and CPU load distribution considered above, we compare the performance of the different models under different load conditions. We study the performance under different combinations of high and low bandwidth and CPU demands. Table III shows the characteristics of the different loads studied. Similar to Section V (A), the number of nodes in a VNR is uniformly distributed between 2 and 6.

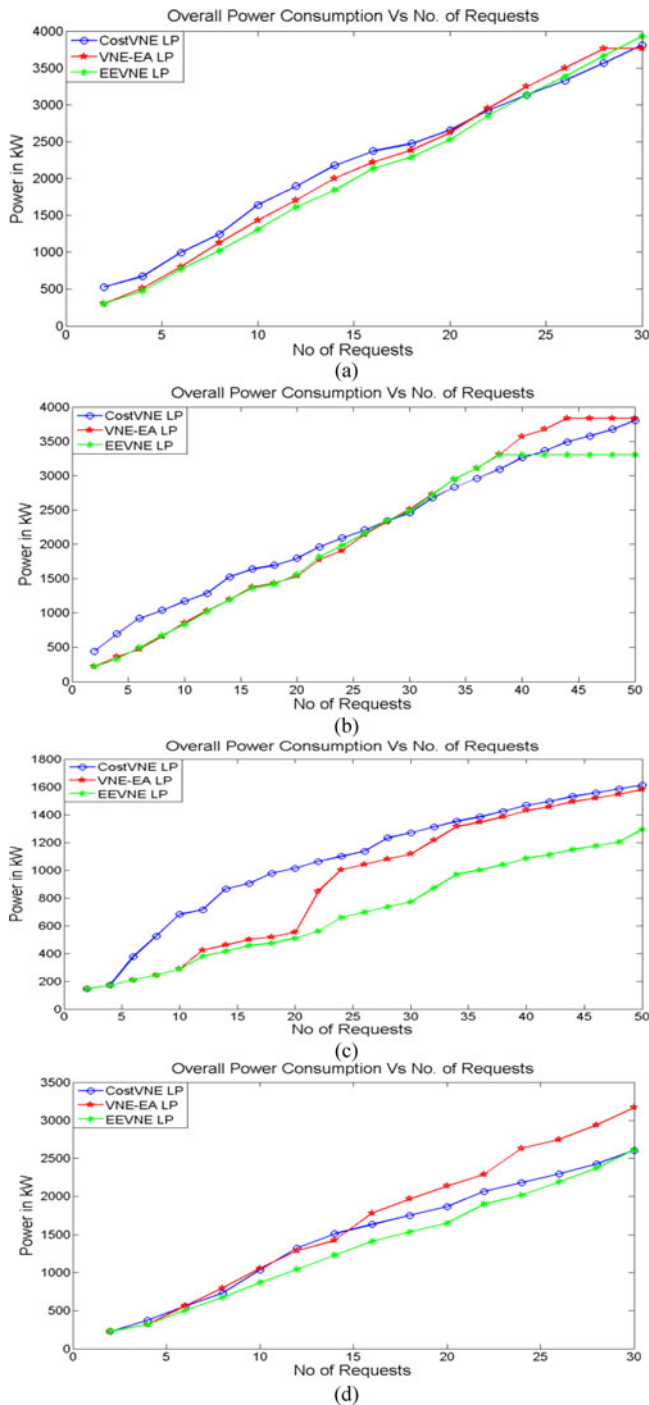


Fig. 13. (a) Power consumption under a load of high bandwidth and high CPU demands (Load 1), (b) power consumption under a load of high bandwidth and low CPU demands (Load 2), (c) power consumption under a load of low bandwidth and low CPU demands (Load 3), (d) power consumption under a load of low bandwidth high CPU demands (load 4).

Fig. 13 shows the overall power consumption of the different models under the different load conditions of Table II. Fig. 13(a) shows the performance of the different models under Load 1 (high bandwidth and high CPU demands). At this load, the high CPU demands overwhelm the CPU resources of the data centers allowing the substrate network to only accommodate a maximum 30 VNRs out of 50 as will be seen in Fig. 14. The




Fig. 14. Number of accepted requests at various loads. The loads refer to those in Table IV.

EEVNE model outperformed the other models up to a load of 22 VNRs. However, as the load on data centers approaches their full capacity, the CostVNE model starts to perform better as no further savings can be acquired by consolidating resources in data centers and the model with higher savings in the network, in this case the CostVNE model, will have the minimum power consumption.

Fig. 13(b) shows the performance under high bandwidth and low CPU demands. As the links are readily used up to support the high bandwidth demands, the EEVNE model loses its advantage over the VNE-EA model as the bandwidth demands will be routed through routes of more hops. The CostVNE model eventually starts to perform better than both the VNE-EA and EEVNE models. The higher propensity of the VNE-EA and EEVNE models to consolidate resources in data centers for minimal power consumption depletes network resources resulting in future requests being embedded far away and thus increasing network power consumption. Notice that this happens as requests are served sequentially and the new requests are accommodated in the spare capacity. If the network is re-optimized following each new request, the EEVNE model will always be better (or equal to) the CostVNE model. Note that the EEVNE model starts rejecting VNRs earlier than the VNE-EA as it has higher propensity to consolidate resources in data centers.

For request with low bandwidth and low CPU demands (see Fig. 13(c)), the substrate network can successfully accommodate all the requests as will be seen in Fig. 14. The superior performance of the EEVNE model over both the CostVNE and the VNE-EA model is very clear with average power savings of 35% and 23%, respectively as the EEVNE model is able to consolidate resources in data centers and minimize the network power consumption without being restricted by the depletion of the data centers or network resources. Only 30 of the VNRs of low bandwidth and high CPU demands were successfully mapped as seen in Fig. 14. We however still see a superior performance in power savings from the EEVNE model over the other two models as shown in Fig. 13(d).

Fig. 14 shows the number of accepted requests for each model at different loads. At the load representing high CPU and high bandwidth demands (Load 1), both the CostVNE and EEVNE models accept a total of 30 VNRs. Here 20 VNRs have been rejected due to the depletion of CPU resources in data centers.

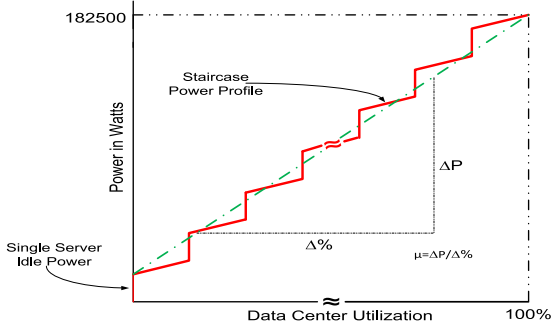


Fig. 15. Energy efficient data center servers power profile.

The VNE-EA model has however only successfully mapped 28 VNRs. This is as a result of the model depleting bandwidth resources in the network before even depleting CPU resources in data centers. For Load 2, the load of high bandwidth and low CPU demands, the CostVNE model has successfully mapped all the VNRs but both the VNE-EA and the EEVNE models have fallen short as they have depleted the network resources in order to consolidate resources in data centers. For the load representing low bandwidth and low CPU demands (Load 3), all the three models have successfully mapped all the VNRs because the substrate network has sufficient resources to accommodate them. Finally for Load 4, the load of low bandwidth and high CPU demands, all the three models only successfully map 30 VNRs as they have run out of CPU resources in the data centers.

VI. ENERGY EFFICIENT VIRTUAL NETWORK EMBEDDING WITH ENERGY EFFICIENT DATA CENTERS

In the previous section we evaluated the power consumption of VNE in IP over WDM networks considering energy inefficient data centers where unused servers are set to idle state. In this section, we adopt data centers with an energy efficient (EE) power profile where only servers needed to serve a given workload are activated, i.e., unused servers are switched off. The EE power profile is given as a staircase curve as shown in Fig. 15. Considering a large number of data centers as shown by the authors in [37] reduces this relationship between the power consumption and the workload to a linear profile given as:

$$PD_b^{v,s} = \begin{cases} \mu \cdot CPU_b^{v,s}, & \text{if the data centre is ON} \\ 0, & \text{otherwise.} \end{cases} \quad (16)$$

The server power consumption in the data centers is given as

$$PD = \sum_{b \in N} \sum_{v \in V} \sum_{s \in R} CPU_b^{v,s} \cdot \delta_b^{v,s} \cdot \mu. \quad (17)$$

Under the EE data center power profile, the EEVNE model no longer seeks the consolidation of data centers as the power consumption of data centers is a function of the workload only not the number of activated data centers. Replacing the data center power consumption in the objective function of the model in Section III by (17) and considering the evaluation scenario in Section V, we evaluate the power consumption of the three VNE models under the energy efficient data center power profile.

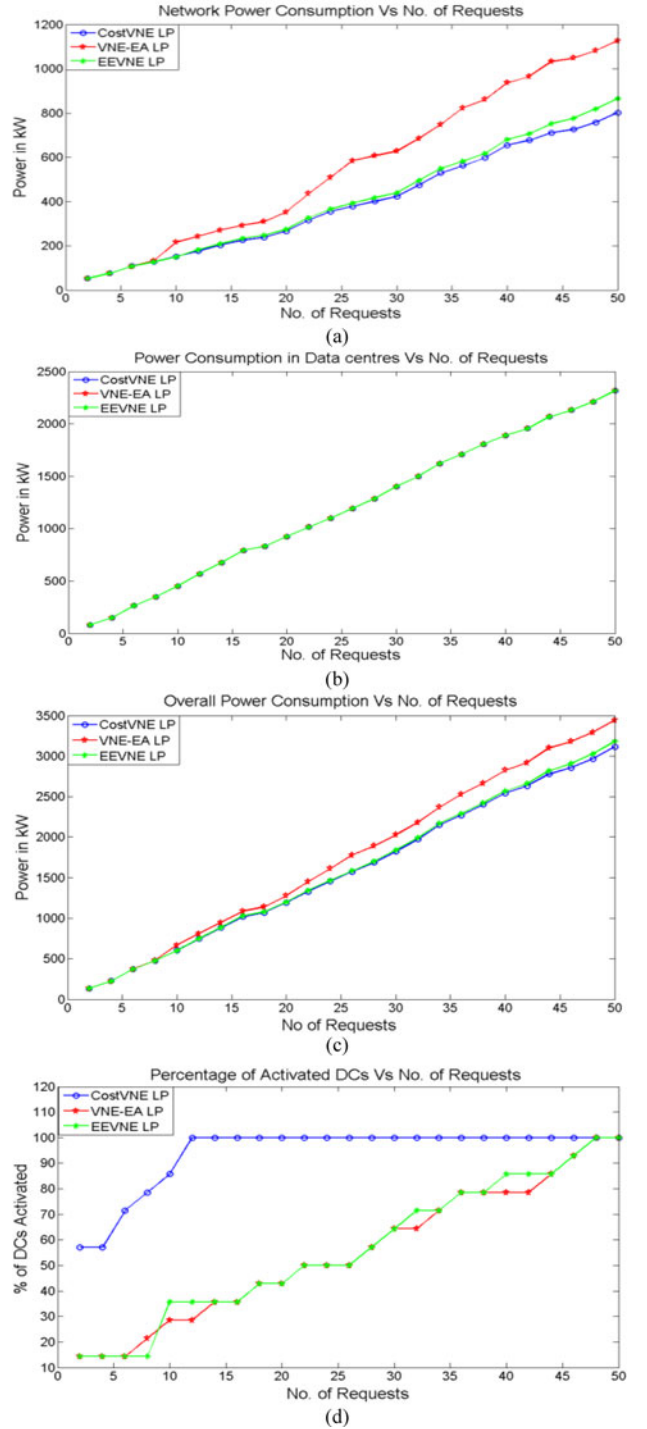


Fig. 16. (a) Network power consumption under the EE data center power profile, (b) Data center servers power consumption under EE data center servers power profile, (c) overall power consumption under EE data center servers power profile, (d) Data center activation under EE data center power profile.

Fig. 16(a) shows the network power consumption of the different approaches under the EE data center power profile. The performance of the CostVNE model will not gain any benefit in terms of network power consumption from adopting the energy efficient data center power profile. This is because the CostVNE model does not take data center consolidation into

account, i.e. the network resource utilization is not restricted by consolidation of resources in data centers. The EEVNE model marginally outperforms the CostVNE model at low loads (not very clear in the figure). This is because in addition to minimizing the hop count between nodes chosen for embedding, the EEVNE model tries to minimize the distance between nodes to reduce number of EDFAs. This results in the EEVNE model clustering embedded nodes close to each other and running out of capacity. As a result, future requests are embedded far away thus increasing network power consumption. As such, ignoring distance and concentrating on hop count (number of transponders and routers but not EDFAs) has given the CostVNE model an advantage. As mentioned above, this happens as requests are served sequentially and the new requests are accommodated in the spare capacity. The CostVNE model has maximum power saving of 2% and an average saving of 1% over the EEVNE model, similar to the energy inefficient data center power profile results in Fig. 7(a). The EEVNE model however performs better than the VNE-EA model with a maximum saving of 30% and average of 22%.

Fig. 16(b) shows that all the three models have similar power consumption trends in the data centers because the power consumption of energy efficient data centers is a function of the workload and not the number of activated data centers as there is no idle power associated with the activation of data centers. Their total power consumption therefore follows the trends of the network power consumption as seen in Fig. 16(c). This leads us to the conclusion that with the energy efficient data center power profile, the optimal VNE approach with the minimum power consumption is the one that only minimizes the number of hops, in this case, the CostVNE model. As mentioned earlier, this is only useful if requests are served sequentially and the new requests are accommodated in the spare capacity.

Fig. 16(d) shows the rate at which all the three models activate data centers with increasing number of VNRs. The CostVNE model has activated all the data centers after embedding the 12th VNR. The EEVNE and VNE-EA models however have a much more gradual activation of data centers which only reaches its peak after embedding the 46th VNR.

VII. LOCATION AND DELAY CONSTRAINTS

In the previous sections, node embedding was not subject to any location constraints, i.e. nodes could be embedded anywhere in the cloud as long as connected nodes of a VNR are not embedded in the same substrate node. In this section we evaluate scenarios where the node embedding is subject to location and delay constraints. This scenario represents, for example, an enterprise solution for a service that requires running applications on a virtual machine(s) in a given fixed location (e.g., company headquarter or branch) with one or more redundant virtual machines for protection or load balancing. The virtual machine(s) with a fixed location is referred to as the master node and the redundant virtual machines are referred to as the slave nodes. Such VNRs typically form star topologies with the master node at the center. The slave nodes can be located anywhere in the cloud subject to delay constraints between them and the

master node in addition to the bandwidth capacity constraints. Depending on the service level agreement (SLA), the redundant virtual machine can be located in the same data center but in different racks or full geographical redundancy can be supported by locating each virtual machine in a different substrate node. Locating all the nodes of a VNR in one substrate node eliminates the need for network connections and therefore reduces the network power consumption and allows the service provider to charge the clients less. On the other hand the cost increases with the provision of full geographical redundancy. In the following results we investigate the impact of different redundancy formats on the power consumption by defining the node consolidation factor (α) as a measure of how many nodes of a VNR can be embedded in the same substrate node. For $\alpha = 1$, the model only allows one virtual node from the same request to be embedded in the same substrate node while for $\alpha = 5$, the model allows 5 virtual nodes from the same request to be embedded in the same substrate node.

A. MILP Model With Location and Delay Constraints

We extend the model in Section III, to introduce location and delay constraints. In addition to the sets, parameters and variables defined in Section III-B, we define the following parameters and variables:

Parameters:

- LOC_b^v $LOC_b^v = 1$ if the master node of VNR v must be located at substrate node b , otherwise $LOC_b^v = 0$;
- DEL^v The threshold propagation delay on the links of request v ;
- α Node consolidation factor;
- $\nabla^{m,n}$ Propagation delay on the physical link m, n , given as: $\nabla^{m,n} = \frac{3D_{m,n}}{2C}$, where C is the speed of light.

Variables:

- $Z_{m,n}^{v,b,e}$ $Z_{m,n}^{v,b,e} = 1$ if the IP over WDM virtual link (b, e) of VNR v traverses the physical link (m, n) , otherwise $Z_{m,n}^{v,b,e} = 0$.

In addition to constraints (2)–(14) presented in the model in Section III-B, the following constraints are introduced:

$$\delta_b^{v,1} = LOC_b^v \quad \forall v \in V, b \in N \quad (18)$$

$$\sum_{s \in R} \delta_b^{v,s} \leq \alpha \forall v \in V, b \in N \quad (19)$$

$$\sum_{n \in N_m} Z_{m,n}^{v,b,e} \leq 1 \quad \forall v \in V, b \in N, e \in N, m \in N : b \neq e \quad (20)$$

$$\sum_{m \in N} \sum_{n \in N_m} Z_{m,n}^{v,b,e} \cdot \nabla_{m,n} \leq DEL^v \quad \forall v \in V, b \in N, e \in N : b \neq e. \quad (21)$$

Constraint (18) fixes the location of the master node of request v at node b in the substrate network. Constraint (19) defines the format of node consolidation. Constraint (20) ensures that virtual network links are not bifurcated when mapped in the IP layer. Note that we do not allow traffic bifurcation to facilitate the calculation of propagation delay. Constraint (21) ensures

that the propagation delay of the virtual links does not exceed the delay threshold.

B. MILP Model Results and Analysis

The NSFNET network is also considered as the substrate network to evaluate the energy efficient VNE with delay and location constraints. We consider a scenario where 15 VNRs for the service discussed above are embedded in the substrate network. The number of nodes in the request range between 2 and 5. The propagation delay threshold of the 15 requests is uniformly distributed between 7.5 ms and 75 ms. This range adequately covers the tolerable propagation delay in most applications [38].

The five most populated cities (nodes) in the NSFNET network are selected to be the location of master nodes of the VNRs. The concentration of master nodes on a substrate node is proportional to its population. The master nodes in a set of 15 requests are distributed as follows: Houston Texas (Node 6) can host a maximum of 5 master nodes of the set of 15 requests, San Diego (Node 3) can host 4, and Pittsburg (Node 9), Seattle (Node 2) and Atlanta (Node 10) can each host 2 master nodes.

The power consumption of embedding the 15 VNRs is evaluated versus an increasing CPU and bandwidth load. The range of loads considered is shown in Table III represented by the numbers $\{1, 2, 3, 4, 5, 6, 7 \text{ and } 8\}$. The CPU load is related to the bandwidth requirement according to the assumption that a server working at its maximum CPU utilizes 1 Gb/s of its network resources. Therefore a data center of 500 servers with a 2% workload will generate 10 Gb/s of traffic. This scenario depicts applications whose networking requirements vary proportionally to CPU processing requirements. Examples include online video gaming and augmented reality applications for live video streaming for sports. In the following results we compare the energy efficiency of the CostVNE and EEVNE models, the best performing approaches in the previous section, considering constraints on the virtual nodes location and link delay. We refer to the two models with location and delay constraints as CostVNE-LD and EEVNE-LD. In Fig. 17 we evaluate the network power consumption, datacenter power consumption and the overall power consumption, versus the increasing loads in Table IV at $1 \leq \alpha \leq MAX$, where $MAX = 5$.

Fig. 17(a) shows the network power consumption. At $\alpha = 1$ each virtual machine of a VNR is embedded in a different substrate node which increases network traffic thus allowing the CostVNE-LD model to achieve higher power savings in the network compared to the EEVNE-LD model by efficiently using the network wavelengths when embedding VNRs. With an increasing α , however, the EEVNE-LD model starts to match the CostVNE-LD model as the network traffic is reduced by embedding more virtual machines of a VNR in the same substrate node. The ability of the CostVNE-LD and EEVNE-LD to embed all the virtual machines of a VNR in one substrate node at $\alpha = 5$ allows them to minimize the network power consumption to zero at load 1. As the load increases the CPU resources become insufficient to embed all the virtual machines of each VNR in one substrate node thus virtual links have to be established between data centers thus increasing the network power consumption.

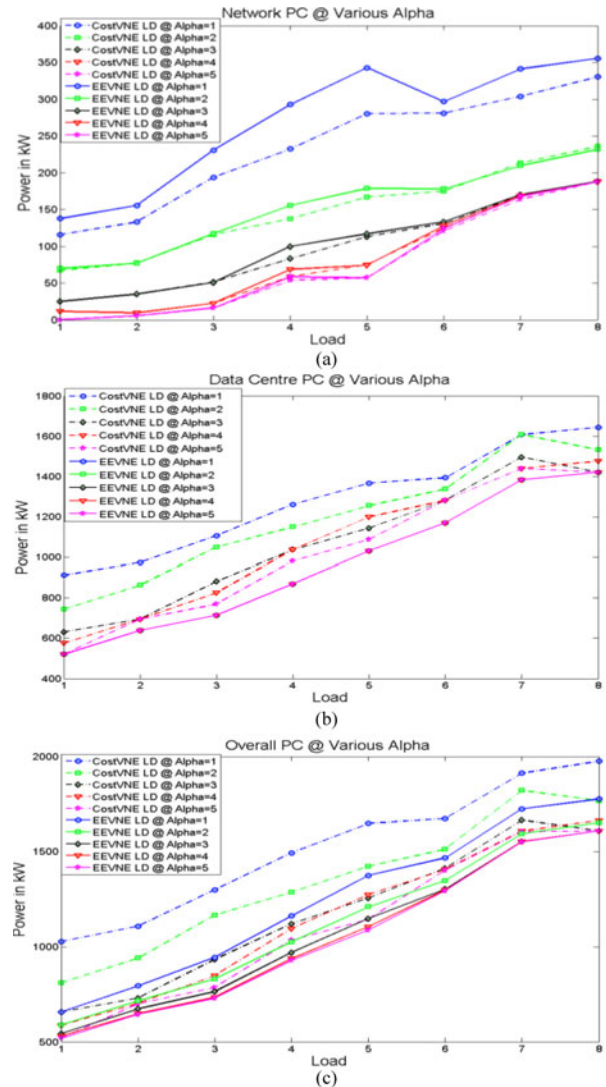


Fig. 17. (a) Network power consumption at various values of α , (b) Data center power consumption at various values of α . (c) Overall power consumption at various values of α .

In Fig. 17(b), we extend the analysis to power savings in the data centers. The virtual machine location constraints do not stop the EEVNE-LD model from efficiently consolidating CPU resources and maintaining the same power consumption in data centers as α increases. This is because the EEVNE-LD model manages to efficiently consolidate virtual machines of different requests at various values of α .

Introducing further restrictions on the virtual machine locations has, however, increased the power consumption in data centers for the CostVNE-LD model as more data centers will be activated to embed CPU resources while maintaining the minimum network power consumption. The EEVNE-LD model has saved 18% of the data center power consumption at $\alpha = 1$ and 5% at $\alpha = 5$ over the CostVNE-LD model.

The impact of introducing restrictions on the virtual machine locations on the overall power consumption of the EEVNE-LD model (see Fig. 17(c)) follows the same trends as the network power consumption, since the power consumption in data

TABLE IV
CPU AND BANDWIDTH LOAD DISTRIBUTION

Load	CPU Percentage Workload Distribution	Link Bandwidth Distribution
1	1–5%	10–40 Gb/s
2	3–7%	20–50 Gb/s
3	5–9%	30–60 Gb/s
4	7–11%	40–70 Gb/s
5	9–13%	50–80 Gb/s
6	11–15%	60–90 Gb/s
7	13–17%	70–100 Gb/s
8	14–19%	80–110 Gb/s

centers does not change as discussed above. The EEVNE-LD model outperforms the CostVNE-LD model in power consumption over all the values of α . Allowing more virtual machines to be embedded in the same substrate node has resulted in limited power savings. As far as the EEVNE-LD model is concerned, the highest savings in overall power consumption of 10% takes place during the transition from $\alpha = 1$ to $\alpha = 2$. The transition from $\alpha = 2$ to $\alpha = 3$ only gives a saving of 4% and the subsequent increases in α account only for very insignificant power savings. From these results we can therefore conclude that for an infrastructure service provider offering a service such as the one presented here, the option of co-location of virtual machines belonging to the same enterprise customer will save power and subsequently cost less, but will not introduce any additional savings by increasing the number of embedded virtual machines of a VNR to more than 2 in the same data center. The power savings acquired from relaxing the virtual nodes location constraints can guide service providers in cost reductions offered to enterprise customers.

VIII. OPTIMIZATION OF DATA CENTER LOCATIONS TAKING INTO ACCOUNT VIRTUAL NETWORK EMBEDDING

In the previous sections, we have considered situations where each node in the network hosts a small data center that can be activated on demand. However, practical considerations such as infrastructure cost and security may limit the number of data centers in a network. In our previous work in [27], we have studied the optimization of data center locations in IP over WDM networks. Here, we show how VNE can impact the design problem of optimally locating data centers for minimal power consumption in cloud networks. We investigate scenarios with a single data center and multiple data centers.

A. MILP Model for Optimal Data Center Locations in Cloud Networks

We have adapted the model in Section III to solve the design problem of optimally locating data centers in cloud networks. In addition to the sets, parameters and variables defined in Sections III-B and VII-A, we define the following:

Parameters:

NDC The total number of data centers in the network;

$C^{v,s}$ The number of CPU cores requested at virtual node s of VNR v ;

γ The power consumption per CPU core embedded in a data center;

Variables:

DP_b $DP_b = 1$ if substrate node b is a data center, otherwise $DP_b = 0$;

$\Delta_b^{v,s}$ $\Delta_b^{v,s} = 1$ if virtual node s of VNR v has been embedded at data center node b otherwise $\Delta_b^{v,s} = 0$;

C_b The total number of CPU cores at data center b ;

$\sigma_b^{v,s}$ $\sigma_b^{v,s}$ is the XOR of DP_b and $\delta_b^{v,s}$, i.e., $\sigma_b^{v,s} = DP_b \oplus \delta_b^{v,s}$, $\sigma_b^{v,s} = 1$ if either DP_b or $\delta_b^{v,s}$ is equal to 1, otherwise $\sigma_b^{v,s} = 0$.

The network power consumption is calculated in the same manner as Section III-B. Note that we do not consider minimizing the power consumption of data centers as only the location of a data center is optimized to minimize the network power consumption. This does not affect the data center size (which is governed by the processing requests made) and consequently its power consumption. The extended model is defined as follows:

Objective:

Minimize the network power consumption given as:

$$\begin{aligned} & \sum_{m \in N} \sum_{n \in N_m} PR \cdot W_{m,n} + \sum_{m \in N} \sum_{n \in N_m} PT \cdot W_{m,n} \\ & + \sum_{m \in N} \sum_{n \in N_m} PE \cdot EA_{m,n} \cdot F_{m,n} \\ & + \sum_{m \in N} PO_m + \sum_{m \in N} PMD \cdot DM_m. \end{aligned} \quad (22)$$

In addition to Constraints (3)–(14) defined in Section III-B and Constraint (18) in Section VII-A, the model is subject to the following constraints:

$$C_b = \sum_{v \in V} \sum_{s \in R} C^{v,s} \cdot \Delta_b^{v,s} \quad \forall b \in N \quad (23)$$

$$DP_b + \delta_b^{v,s} = 2\Delta_b^{v,s} + \sigma_b^{v,s} \quad \forall v \in V, \quad \forall b \in N, \quad \forall s \in R \quad (24)$$

$$\sum_{b \in N} \Delta_b^{v,s} = 1 \quad \forall v \in V, \quad \forall s \in R \quad (25)$$

$$\sum_{b \in N} DP_b = NDC \quad (26)$$

$$\sum_{s \in R} \Delta_b^{v,s} \leq \alpha \quad \forall v \in V, b \in N. \quad (27)$$

Constraint (23) replaces Constraint (2) in Section III-B. It calculates the capacity of each data center in terms of the number of cores. Note that since our problem is a design problem, the data centers and links are both un-capacitated. Constraint (24) ensures that virtual machines are embedded in nodes with data centers by implementing the AND operation of DP_b and $\delta_b^{v,s}$ ($DP_b + \delta_b^{v,s}$). Constraint (25) ensures that each virtual machine is only embedded in a data center once. Constraint (26) gives the number of data centers. Constraint (27) replaces Constraint (19) in Section VII-A.

To extend the model to represent the bypass approach where all the lightpaths whose destination is not the intermediate node are directly bypassed via a cut-through, the power consumption

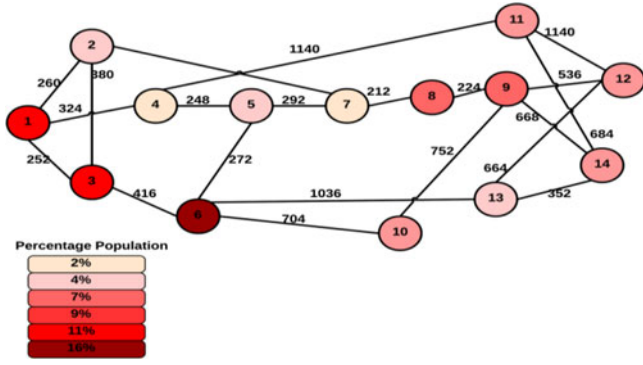


Fig. 18. NSFNET with population percentage information.

in the router ports is modelled as follows:

$$\sum_{i \in N} \sum_{j \in N: j \neq i} PR \cdot C_{i,j}. \quad (28)$$

As such, the objective function becomes:

Minimize the network power consumption given as:

$$\begin{aligned} & \sum_{i \in N} \sum_{j \in N: j \neq i} PR \cdot C_{i,j} + \sum_{m \in N} \sum_{n \in N_m} PT \cdot W_{m,n} \\ & + \sum_{m \in N} \sum_{n \in N_m} PE \cdot EA_{m,n} \cdot F_{m,n} + \sum_{m \in N} PO_m \\ & + \sum_{m \in N} PMD \cdot DM_m. \end{aligned}$$

B. MILP Model Results and Analysis

The NSFNET network is also used as the substrate network to evaluate the performance of the data center locations optimization model. We investigate a scenario where the client's entry point in the network is fixed but virtual machines could be embedded anywhere and virtual links may also be requested to connect virtual machines. The concentration of enterprise clients in a substrate node is based on the population of the states where the cities (nodes) of the NSFNET network are located (see Fig. 18). In the case of California where we have two cities in one state (nodes 1 and 3), we have evenly distributed the population of the state between the two cities. We have depicted an enterprise cloud service solution where enterprise clients request virtual machines with a specific number of CPU cores per virtual machine. The power consumption in the data centers is a function of the number of embedded CPU cores. A total of 45 enterprise clients (VNRs) have been considered where the number of virtual machines per enterprise client is uniformly distributed between 1 and 5 and the number of CPU cores required by each machine is uniformly distributed between 1 and 10. The bandwidth requirement of virtual links is uniformly distributed between 10 and 100 Gb/s.

We optimized the location of a single data center in the NSFNET network under both non-bypass and bypass approaches. The common practice in 2010 was to locate data centers to minimize cost including the data centers and network

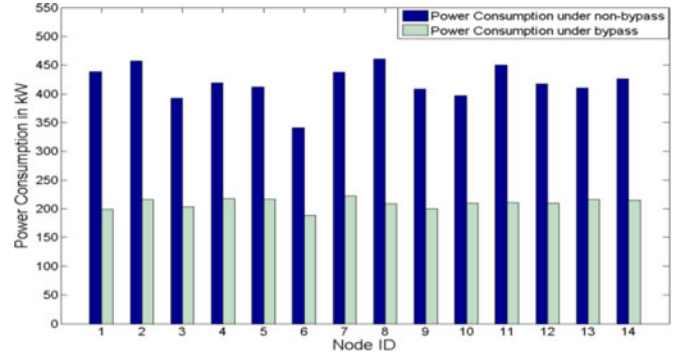


Fig. 19. The network power consumption of a single data center scenario at different locations under non-bypass and bypass approaches.

infrastructure cost and security cost. Considering the cost of the network, the optimal location of the single data center is Node 6. This however also happens to be the optimal location of the data center with the goal of minimizing power consumption in the network due to VNE. To verify the results, we ran the model while fixing the location of the data center at the different NSFNET nodes and evaluated the network power consumption as seen in Fig. 19. The model picks a data center such that the embeddings of the virtual machines in the data center will create virtual links that traverse routes with minimal hops. The hop count determines the number of router ports and transponders (the most energy consuming network devices) used under the non-bypass approach. Under the bypass approach, the IP layer is bypassed and therefore the hop count determines only the number of transponders used. Note that selecting a highly populated node and/or a node with highly populated neighbours to locate a data center can significantly reduce the average hop count as more clients are served locally as a result or through a single hop. Therefore the most populous and easily accessible (through minimum hop routes) Node 6 is selected to host the single data center. The optimal location has achieved power savings of 26% and 15% compared to the worst location for the non-bypass and bypass approaches, respectively.

The current and future criterion for designing data centers infrastructure for cloud services is to distribute the content among a number of data centers to minimize the delay experienced by the users and to avoid the scenario of having a single hot node in the network. We have therefore also considered a scenario where cloud services are supported by a total of five data centers in the network and having $\alpha = 5$, i.e., we allow co-location of virtual machines from the same VNR, i.e., allow such virtual machines to be embedded in the same data center. The optimal locations of the data centers in this case are nodes (1, 6, 8, 9 and 11) for the non-bypass approach and nodes (1, 3, 6, 8 and 9) for the bypass approach as shown in Fig. 20. When minimizing the network power consumption, the model optimizes two factors: the traffic requests served locally and the average hop count traversed by traffic demands. The optimal locations under the non-bypass approach are selected in such a way that all requests are either served locally or within a single hop. The bypass approach however can compromise on the number of hops

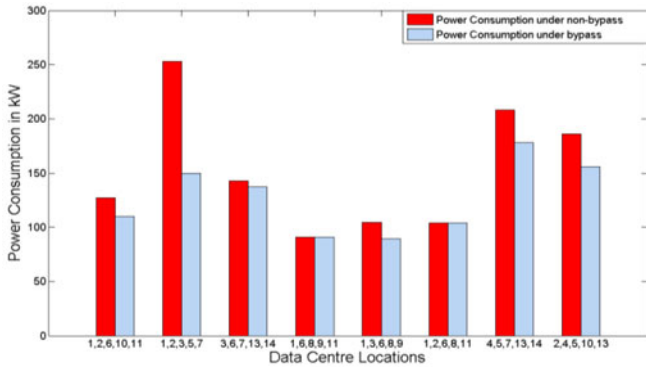


Fig. 20. The network power consumption of a 5 data center scenario under non-bypass and bypass approaches at $\alpha = 5$.




Fig. 22. The network power consumption of a data center scenario under non-bypass and bypass approaches at $\alpha = 1$.

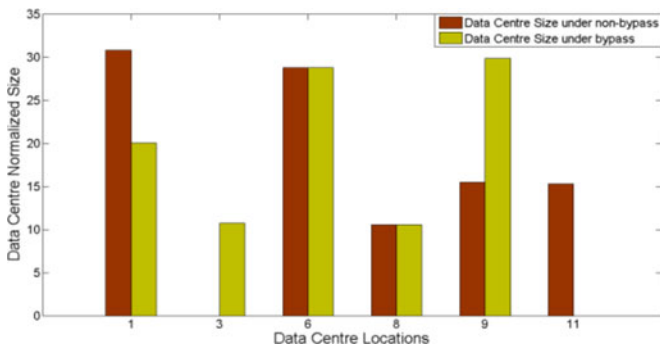


Fig. 21. The normalized size of five data centers in the optimal locations at $\alpha = 5$.

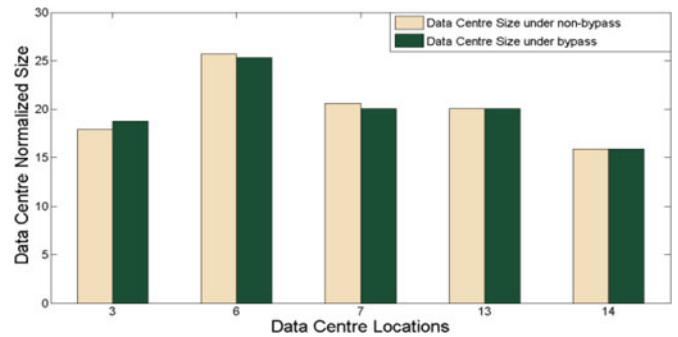


Fig. 23. The normalized size of a five data centers in the optimal locations at $\alpha = 1$.

traversed by low bandwidth demands in order to serve high demands locally. The power consumption of transponders used by the selection of routes with higher hop count is compensated by the significant power savings obtained by embedding VNRs of higher bandwidth demands locally. Therefore Node 11 in the optimal data centers locations under the non bypass approach is replaced with the highly populated Node 3 under the bypass approach. The optimal locations have achieved network power savings of 43% and 55% compared to the worst locations for the bypass and non-bypass approaches, respectively. Compared to the single data center solution, distributing the content among multiple data centers has saved 53% and 73% of the network power consumption for bypass and non-bypass, respectively.

Fig. 21 shows the normalized size of optimally located data centers at $\alpha = 5$. Under non-bypass, the data center located at Node 1 is used to embed the highest number of virtual machines. This is because Node 1 has taken up all the virtual machines of its neighboring populous Node 3 due to the distance between Node 1 and Node 3 being much shorter than the distance between Node 3 and Node 6. Under the bypass approach, a data center is created at Node 3 instead of Node 11. This has significantly reduced the number of virtual machines served by the data center at Node 1 and increased the size of the data center at Node 9 as it has taken up all the virtual machines served by the data center at Node 11 under the non-bypass routing approach.

We have also investigated optimizing the location of data centers in a scenario with geographical redundancy constraints

where virtual machines that belong to the same client cannot be embedded in the same data center ($\alpha = 1$), i.e., a client requesting five virtual machines will have one embedded in each of the five data centers. The optimal locations of the five data centers as obtained from the model are (3, 6, 7, 13, and 14) for both non-bypass and bypass. As discussed above the selection of a location to host a data center is governed by two factors: the average hop count between the data centers and clients and the client population of the candidate node and its neighbors. In this scenario where virtual machines of the same client cannot be collocated, the model will also try to minimize the hop count between data centers. Fig. 22 compares the network power consumption of embedding the VNRs under the optimal data center locations to the other random data center locations for both non-bypass and bypass approaches. The optimal locations have achieved power savings of 19% and 2% compared to the worst locations for the non-bypass and bypass approaches, respectively.

Fig. 23 shows the normalized size of data centers in the optimal locations at $\alpha = 1$. The data center located at Node 6 has the highest number of virtual machines embedded for both the non-bypass and bypass approaches. This is because Node 6, as mentioned above, is the most populous and easily accessible (through minimal hop routes) to other nodes as discussed above. The data center sizes of Fig. 23 are much more balanced compared to those of Fig. 21. This is because the restriction to only put one virtual machine belonging to the same VNR in one data

center (see Fig. 23) forces the distribution of virtual machines across the different data centers.

IX. VNE WITH OPTICAL OFDM BASED CLOUD NETWORKS

In the earlier sections, we have examined VNE in substrate networks whose transport network employs wavelength division multiplexing (WDM) technology. In this section, we consider VNE in optical orthogonal frequency division multiplexing (O-OFDM) transport networks and explore the added benefits. In our earlier work in [39] we introduced VNE in O-OFDM networks and gave the power savings obtained in the optical layer. Here we go further by presenting the spectral gains associated with VNE in O-OFDM networks. We also introduce the details of the MILP model for VNE in O-OFDM networks.

A. MILP Model for VNE in O-OFDM Cloud Networks

Before we introduce the model, we define the following sets, parameters and variables in addition to those defined in the earlier sections.

Sets:

Q Set of modulation formats of subcarriers;

Parameters:

TP_{max} Maximum power consumption of an O-OFDM transponder working at maximum line rate LR_{max} ;

NSC Maximum number of subcarriers supported by an O-OFDM transponder;

LR_{max} Maximum line rate;

M_q The capacity of a single subcarrier with modulation format q ;

W Number of wavelengths per fiber;

Variables:

$nsc_q^{b,e}$ Number of O-OFDM subcarriers using modulation format q to serve the traffic demand (b, e) ;

$SC_{m,n,q}^{b,e}$ Number of O-OFDM subcarriers using modulation format q to serve the demand (b, e) that traverses a physical link network (m, n) ;

$x_{m,n}^{b,e}$ $x_{m,n}^{b,e} = 1$, if O-OFDM subcarriers of traffic demand (b, e) traverse a physical link (m, n) otherwise $x_{m,n}^{b,e} = 0$;

$F_{m,n}$ Number of fibers on physical link (m, n) .

In modelling the power consumption of the network, we only consider the three most power consuming components in the network; IP router ports, transponders and erbium doped fibre amplifiers (EDFAs). To enable bandwidth flexible transmissions, the fixed-bandwidth components used in WDM networks need to be replaced with network components that can work at flexible rates. In [40], we studied the power consumption of these components in an optical OFDM-based network and compared it with the power consumption of similar components in IP over WDM networks. We assumed an adaptive line rate (ALR) power profile for the IP ports to calculate the power consumption of router ports supporting flexible wavelength rates. EDFAs can be used in O-OFDM-based optical networks as they can simultaneously amplify many data channels at different wavelengths within its gain region [41]. The power consumption of the O-OFDM transponder mainly depends on the electronic processing, modulation format used and the number of supported sub-

carriers. Therefore, a transponder using the highest modulation format with the maximum number of subcarriers, i.e. working at maximum rate, will have the highest power consumption. We assumed an ALR power profile to estimate the power consumption of transponders working at lower rates [42]. Therefore, the total power consumption of the O-OFDM substrate network under the bypass approach is composed of:

Power consumption of router ports:

$$\sum_{b \in N} \sum_{e \in N: b \neq e} PR \cdot L_{b,e}.$$

Power consumption of OFDM transponders:

$$\sum_{b \in N} \sum_{e \in N: b \neq e} \left(\sum_{m \in N} \sum_{n \in N_m} x_{m,n}^{b,e} \left(ALR \left(\frac{\sum_{q \in Q} SC_{m,n,q}^{b,e} \cdot M_q}{LR_{max}} \right) \cdot TP_{max} \right) \right)$$

where, $ALR()$ is the adaptive link rate power function.

Power consumption of EDFAs:

$$\sum_{m \in N} \sum_{n \in N_m} PE \cdot EA_{m,n} \cdot F_{m,n}.$$

The model is defined as follows:

Objective: Minimize

$$\begin{aligned} & \sum_{b \in N} \sum_{v \in V} \sum_{s \in R} C^{v,s} \cdot \Delta_b^{v,s} \cdot \gamma + \sum_{b \in N} \sum_{e \in N: b \neq e} PR \cdot L_{b,e} \\ & + \sum_{b \in N} \sum_{e \in N: b \neq e} \left(\sum_{m \in N} \sum_{n \in N_m} x_{m,n}^{b,e} \left(ALR \left(\frac{\sum_{q \in Q} SC_{m,n,q}^{b,e} \cdot M_q}{LR_{max}} \right) \cdot TP_{max} \right) \right) \\ & + \sum_{m \in N} \sum_{n \in N_m} PE \cdot EA_{m,n} \cdot F_{m,n}. \end{aligned} \quad (29)$$

In addition to constraints (10)–(16) defined in Section II-B, constraint (25) in Section VI-A and constraints (31)–(34) in Section VII-A, we define the following constraints:

$$L_{b,e} \leq \sum_{q \in Q} nsc_q^{b,e} \cdot M_q \quad \forall b, e \in N : b \neq e \quad (30)$$

$$\sum_{n \in N_m} x_{m,n}^{b,e} - \sum_{n \in N_m} x_{n,m}^{b,e} = \begin{cases} 1, & m = b \\ -1, & m = e \\ 0, & \text{otherwise} \end{cases} \quad (31)$$

$$\forall b, e, m \in N : b \neq e$$

$$\begin{aligned} & \sum_{n \in N_m} SC_{m,n,q}^{b,e} - \sum_{n \in N_m} SC_{n,m,q}^{b,e} \\ & = \begin{cases} nsc_q^{b,e}, & m = b \\ -nsc_q^{b,e}, & m = e \\ 0 & \text{otherwise} \end{cases} \end{aligned} \quad (32)$$

$$\forall b, e, m \in N : b \neq e, q \in Q$$

TABLE V
LOAD DISTRIBUTION

Load	CPU Cores Distribution	Link Bandwidth Distribution
1	1–5	10–40 Gb/s
2	3–7	20–50 Gb/s
3	5–9	30–60 Gb/s
4	7–11	40–70 Gb/s
5	9–13	50–80 Gb/s
6	11–15	60–90 Gb/s
7	13–17	70–100 Gb/s
8	14–19	80–110 Gb/s

TABLE VI
POWER CONSUMPTION OF NETWORK DEVICES

Power consumption of a 100 Gb/s WDM transponder	135 (W)
Power Consumption of an OFDM transponder at Maximum Line Rate TP_{max}	200 (W)
Power consumption per Gb/s of an IP router port	25 W/Gb/s
Power consumption of an EDFA	8 (W)

$$\sum_{b \in N} \sum_{e \in N: b \neq e} \left(\sum_{q \in Q} SC_{n,m,q}^{b,e} / NSC \right) \leq W \cdot F_{m,n} \quad \forall m, n \in N : m \neq n. \quad (33)$$

Constraint (30) ensures that the total capacity of the subcarriers allocated to a traffic demand is large enough to support it. Constraint (31) gives the flow conservation in the optical layer and ensures that traffic demands do not bifurcate. Constraint (32) ensures that for each traffic demand, the number of O-OFDM subcarriers using modulation format q entering node m is equal to the number of subcarriers using modulation format q leaving node m as long as node m is not the source or destination of the traffic demand. Constraint (33) ensures that the number of O-OFDM wavelengths allocated to a traffic demand does not exceed the number of wavelengths in an optical fiber.

We have also considered spectrum efficient VNE in IP over O-OFDM networks where the objective is to minimize the spectrum used to embed the VNRs. The objective in this case is defined as:

Objective: Minimize

$$\sum_{b \in N} \sum_{e \in N: b \neq e} \sum_{q \in Q} SC_{n,m,q}^{b,e} / NSC \quad \forall m \in N, n \in N_m. \quad (34)$$

B. MILP Model Results Analysis

We have depicted a similar scenario as the one in Section VII-B. Except in this case we have to determine the power consumption of the network as the load presented to the network changes due to the mapping of virtual networks. For the same reasons as in Section VIII, we only consider minimizing the network power consumption in our objective. We consider embedding 45 virtual networks at each load. The load distribution of CPU cores and bandwidth is shown in Table V. We have also capacitated all the links in the network to carry a maximum of 32 wavelength channels. The network is comprised of a max-




Fig. 24. Optical layer power consumption.

imum of seven data centers whose location is determined by the model. The data centers in this scenario are not capacitated.

Similar to our work in [43], we have considered a channel bandwidth of 50 GHz for both conventional WDM and O-OFDM networks. The maximum number of subcarriers for each O-OFDM channel is 10, each of 5 GHz where two of the channels are used as guard bands. In conventional IP over WDM networks, the available capacity is limited by the worst-case optical path. We assumed a 100 Gb/s line rate per WDM channel. The modulation format for O-OFDM subcarriers has been adapted according to the end to end physical conditions of the optical path which is measured by the distance dependent optical signal to noise ratio (OSNR) [44]. As discussed in [40], BPSK (1 bit/symbol) is used to modulate subcarriers over the largest transmission distance in NSFNET (2000 km). The modulation format increases by 1 bit/symbol as the transmission distance decreases to half. So, for the optical OFDM-based network at 1000 km, QPSK (2 bits/symbol) can be used and at 500 km 8QAM (3 bits/symbol) can be used. With 8QAM, the highest modulation format for optical OFDM, the maximum line rate for an O-OFDM transponder LR_{max} is: $5(\text{GHz}) \times 3(\text{Bits}/\text{Hz}) \times 8 = 120 \text{ Gb/s}$.

The power consumption of an O-OFDM transponder working at the maximum rate of 120 Gb/s was estimated in [40] as 204 W. Table VI shows the power consumption values used. We assumed a cubic ALR profile to estimate the power consumption of O-OFDM transponders working at lower rates. A linear power profile is considered to estimate the power consumption of router ports working at different rates given the power consumption of 25 W per Gb/s [40].

We compared power minimized and spectrum minimized VNE in IP over O-OFDM cloud networks to VNE in conventional IP over WDM networks in terms of power consumption and spectral efficiency considering the lightpath bypass approach. The different models have selected the most populous nodes (1, 3, 6, 8, 9, 11 and 12) as the optimal data center locations to serve as many demands as possible locally to minimize the use of network resources.

Fig. 24 shows the power consumption in the optical layer. The embedding over power and spectrum minimized IP over O-OFDM has saved 63% and 17%, respectively of the optical layer power consumption compared to VNE in IP over WDM.

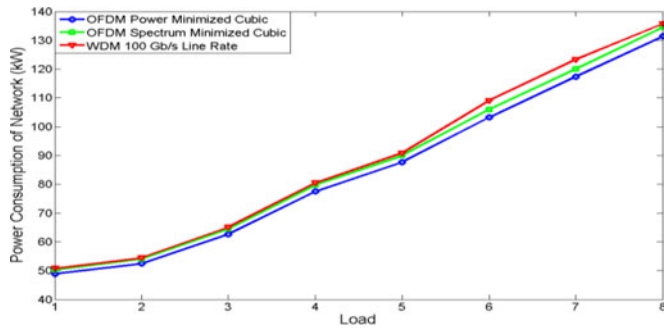


Fig. 25. Network power consumption.




Fig. 26. Network spectral efficiency.

The power minimized O-OFDM network selects routes and modulation formats that result in minimum power consumption while the spectrum minimized O-OFDM selects routes that can support the highest modulation format to minimise the used spectrum. However, since the capacity of the links is limited, the two models might consume the same amount of power to embed some of the VNRs at high loads under two scenarios. The first scenario is when the number of available subcarriers on the only available route is limited so the two models have to adopt the maximum modulation format to create enough capacity to embed the bandwidth requirement. The second scenario is when the bandwidth requirement has to be embedded over a longer route that can only support the lowest modulation format so the two models will have to adopt the lowest modulation format. Fig. 25 shows the total network power consumption including the power consumed by the IP router ports. The power savings achieved by the power and spectrum minimized IP over O-OFDM considering the total network power consumption are 4% and 2%, respectively and they are all due to the savings in the optical layer.

Fig. 26 compares the spectral efficiency of the three models. As discussed above the power minimized O-OFDM network selects the modulation format that results in the minimum power consumption to embed virtual networks. Therefore the spectrum minimized O-OFDM network is 120% more efficient in spectrum utilization than the power minimized O-OFDM network. The results in Fig. 26 also show that the WDM network and the power minimized O-OFDM network at low loads have similar spectral efficiency. This is because the O-OFDM net-

work opts for the lowest modulation format (BPSK) to minimize power which is the same modulation format employed by WDM 100 Gb/s OOK. However, as the load increases in the network, there is an improvement in spectral efficiency by the power minimized O-OFDM model due to its adaptation to higher modulation formats in order to efficiently accommodate the high load in the network. We have assumed 100 Gb/s WDM is implemented using OOK with VSB Filtering as found in [45]. The upside to this is that it makes our assumption of reach at 100 G much more realistic to cover the distances in the NSFNET network. However, the downside is that the channel spacing though improved by VSB filtering to 100 GHz is still greater than the current ITU 50 GHz channel spacing.

X. CONCLUSION

This paper has investigated the energy efficiency of virtual network embedding in IP over WDM networks. We developed a MILP model (EEVNE) and a heuristic (REOVNE) to optimize the use of wavelengths in the network in addition to consolidating the use of resources in data centers. The results show that the EEVNE model achieves a maximum power saving of 60% (average 20%) compared to the CostVNE model which minimizes the bandwidth cost of embedding a VNR. The EEVNE model has also higher power savings compared to the virtual network embedding energy aware (VNE-EA) model from the literature. We have demonstrated that when it comes to energy savings in the network, it is not sufficient to develop models that just turn off links and nodes in the network but it is important to consider all the power consuming devices in the network and then minimize their power consumption as a whole. The REOVNE heuristic's power savings and number of accepted requests approaches those of the MILP model. We have also investigated the performance of the models under non uniform load distributions showing that EEVNE model has superior power savings in most load conditions. We have gone further to show the energy efficiency of VNE considering an energy efficient data center power profile. The results show that the optimal VNE approach with the minimum power consumption is the one that only minimizes the use of network bandwidth, in this case, the CostVNE model. This however only applies when it is assumed that the network is not reconfigured when embedding new requests. We have also studied the power savings achieved by removing geographical redundancy constraints when embedding protection and load balancing virtual nodes and observed that the power savings obtained as a result can guide service providers in determining cost reductions offered to enterprise customers not requiring full geographical redundancy. We have shown how VNE can impact the optimal locations of data centers for minimal network power consumption in cloud networks. The results show that the selection of a location to host a data center is governed by two factors: the average hop count to other nodes and the client population of the candidate node and its neighbours (assuming a given average rate per user). Finally, we have developed a MILP model for VNE in O-OFDM based cloud networks and shown that they have improved power and spectral efficiency compared to conventional WDM based networks.

ACKNOWLEDGMENT

L. Nonde would like to thank the Commonwealth Scholarship Commission funded by the UK government for the award of the Ph.D. scholarship.

REFERENCES

- [1] M. A. Sharkh, M. Jammal, A. Shami, and A. Ouda, "Resource allocation in a network-based cloud computing environment: Design challenges," *IEEE Commun. Mag.*, vol. 51, no. 11, pp. 46–52, 2013.
- [2] United States Environmental Protection Agency. Report to Congress on Server and Data Center Energy Efficiency Public Law 109-431. (2007). [Online]. Available: http://www.energystar.gov/ia/partners/prod_development/downloads/EPA_Datacenter_Report_Congress_Final1.pdf.
- [3] A. Q. Lawey, T. E. H. El-Gorashi, and J. M. H. Elmirghani, "Distributed energy efficient clouds over core networks," *IEEE J. Lightw. Technol.*, vol. 32, no. 7, pp. 1261–1281, Jan. 2014.
- [4] R. Jain and S. Paul, "Network virtualization and software defined networking for cloud computing: A survey," *IEEE Commun. Mag.*, vol. 51, no. 11, pp. 24–31, Nov. 2013.
- [5] A. Belbekkouche, M. M. Hasan, and A. Karmouch, "Resource discovery and allocation in network virtualization," *IEEE Commun. Surveys Tuts.*, vol. 14, no. 4, pp. 1114–1128, Oct.–Dec. 2012.
- [6] N. M. M. K. Chowdhury and R. Boutaba, "Network virtualization: State of the art and research challenges," *IEEE Commun. Mag.*, vol. 47, no. 7, pp. 20–26, Jul. 2009.
- [7] A. Fischer, M. T. Beck, H. De Meer, and X. Hesselbach, "Virtual network embedding: A survey," *IEEE Commun. Surveys Tuts.*, vol. 15, no. 4, pp. 1888–1906, Oct.–Dec. 2013.
- [8] B. Wang, X. Chang, J. Liu, and J. K. Muppala, "Reducing power consumption in embedding virtual infrastructures," in *Proc. IEEE Globecom Workshops*, Dec. 3–7, 2012, pp. 714–718.
- [9] M. Yu, Y. Yi, J. Rexford, and M. Chiang, "Rethinking virtual network embedding: Substrate support for path splitting and migration," *SIGCOMM Comput. Commun. Rev.*, vol. 38, no. 2, pp. 17–29, 2008.
- [10] C. Yang, J. Li, T. Wo, C. Hu, and W. Liu, "Resilient virtual network service provision in network virtualization environments," in *Proc. 16th Int. Conf. Parallel Distrib. Syst.*, 2010.
- [11] D. G. Anderson, (2002). Theoretical Approaches to Node Assignment. [Online]. Available: <http://repository.cmu.edu/cgi/viewcontent.cgi?article=1079&context=compsci>.
- [12] J. F. Botero, X. Hesselbach, M. Duelli, D. Schlosser, A. Fischer, and H. De Meer, "Energy efficient virtual network embedding," *IEEE Commun.*, vol. 16, no. 5, pp. 756–759, May 2012.
- [13] S. Sen, Z. Zhongbao, C. Xiang, W. Yiwen, L. Yan, and W. Jie, "Energy-aware virtual network embedding through consolidation," in *Proc. IEEE Comput. Commun. Workshops*, 2012, pp. 127–132.
- [14] J. F. Botero and X. Hesselbach, "Greener networking in a network virtualization environment," *Comput. Netw.*, vol. 57, no. 9, pp. 2021–2039, 2013.
- [15] S. Su, Z. Zhang, A. X. Liu, X. Cheng, Y. Wang, and X. Zhao, "Energy-aware virtual network embedding," *IEEE/ACM Trans. Netw.*, vol. 22, no. 5, pp. 1607–1620, Oct. 2014.
- [16] I. Houidi, W. Louati, and D. Zeghlache, "A distributed virtual network mapping algorithm," in *Proc. IEEE Int. Conf. Commun.*, 2008, pp. 5634–5640.
- [17] M. Chowdhury, M. R. Rahman, and R. Boutaba, "ViNEYard: Virtual network embedding algorithms with coordinated node and link mapping," *IEEE/ACM Trans. Netw.*, vol. 20, no. 1, pp. 206–219, Feb. 2012.
- [18] L. Nonde, T. E. H. El-Gorashi, and J. M. H. Elmirghani, "Cloud virtual network embedding: Profit, power and acceptance," submitted to *IEEE Commun.*, 2015.
- [19] I. Houidi, W. Louati, W. Ben Ameer, and D. Zeghlache, "Virtual network provisioning across multiple substrate networks," *Comp. Netw.*, vol. 55, no. 4, pp. 1011–1023, 2011.
- [20] A. Fischer, J. F. Botero, M. Duelli, D. Schlosser, X. Hesselbach, and H. De Meer, "ALEVIN—A framework to develop, compare, and analyse virtual network embedding algorithms," *Electron. Commun. EASST*, vol. 37, pp. 1–12, 2011.
- [21] G. Shen and R. S. Tucker, "Energy-minimised design for IP over WDM networks," *IEEE/OSA J. Opt. Commun. Netw.*, vol. 1, no. 1, pp. 176–186, Jun. 2009.
- [22] X. Fan, W. D. Weber, and L. A. Barroso, "Power provisioning for a warehouse-sized computer," *ACM SIGARCH Comput. Archit. News*, vol. 35, no. 2, pp. 13–23, 2007.
- [23] A. Greenberg, J. Hamilton, D. A. Maltz, and P. Patel, "The cost of a cloud: Research problems in data center networks," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 39, no. 1, pp. 68–73, 2008.
- [24] Google. Efficiency: How do we do it. (2014). [Online]. Available: <http://www.google.co.uk/about/datacenters/efficiency/internal/>
- [25] Microsoft. Microsoft's Cloud Infrastructure—By the Numbers. (2013). [Online]. Available: http://www.microsoft.com/eu/PRESSRELEASE_Microsoft_cloud_services_growth_drives_Expansion.aspx?path=10-1
- [26] Principled Technologies. Advanced Power Management with Dell Open Manage Power Center. (2012). [Online]. Available: http://www.principledtechnologies.com/Dell/R720_power_0312.pdf.
- [27] X. Dong, T. El-Gorashi, and J. M. H. Elmirghani, "Green IP over WDM networks with data centers," *IEEE J. Lightw. Technol.*, vol. 29, no. 12, pp. 1861–1880, Jun. 2011.
- [28] Data sheet of CRS-1 16 slots chassis power systems.
- [29] Data sheet of Cisco ONS 15454 10-Gb/s multi-rate transponder card
- [30] Data sheet of Cisco ONS 15501 EDFA.
- [31] Data sheet of Cisco ONS 15454 100-GHz 4-CH Multi/Demultiplexer.
- [32] Data sheet of Glimmerglass Intelligent Optical System 500.
- [33] A. Hammadi and L. Mhamdi, "A survey on architectures and energy efficiency in data center networks," *Comput. Commun.*, vol. 40, no. 0, pp. 1–21, 2014.
- [34] Cisco. (2014). Cisco Catalyst 4900M Switch. [Online]. Available: http://www.cisco.com/c/en/us/products/collateral/switches/catalyst-4900-series-switches/Data_Sheet_Cat_4900M.html.
- [35] Cisco. (2014). Cisco 7613 Router Data Sheet. [Online]. Available: http://www.cisco.com/c/en/us/products/collateral/routers/7613-router-product_data_sheet09186a008015cfeb.html.
- [36] A. Q. Lawey, T. E. H. El-Gorashi, and J. M. H. Elmirghani, "BitTorrent content distribution in optical networks," *J. Lightw. Technol.*, vol. 32, no. 21, pp. 3607–3623, Nov. 2014.
- [37] V. Valancius, L. Massoulie, C. Diot, and P. Rodriguez, "Greening the internet with nano data centers," in *Proc. 5th Int. Conf. Emerg. Netw. Exp. Technol. ACM*, 2009.
- [38] Broadband Forum. (2014). Triple-play Services Quality of Experience (QoE) Requirements. Tech. Rep. TR-126 2006. [Online]. Available: <http://www.broadband-forum.org/technical/download/TR-126.pdf>.
- [39] L. Nonde, T. E. H. Elgorashi, and C. M. Machuca, "Green virtual network embedding in optical OFDM cloud networks," in *Proc. 16th Int. Conf. IEEE Conf. Transp. Opt. Netw.*, 2014.
- [40] T. E. H. El-Gorashi, X. Dong, and J. M. H. Elmirghani, "Green optical orthogonal frequency-division multiplexing networks," in *IET Optoelectron.*, vol. 8, no. 3, pp. 137–148, 2014.
- [41] F. D. Mahad, M. Supa'at, and A. Sahmah, "EDFA gain optimization for WDM system," *Elektrika*, vol. 11, no. 1, pp. 34–37, 2009.
- [42] J. C. Restrepo, C. G. Gruber, and C. M. Machuca, "Energy profile aware routing," in *Proc. IEEE Commun. Workshops*, 2009.
- [43] X. Dong, T. E. H. El-Gorashi, and J. M. H. Elmirghani, "Energy efficiency of optical OFDM-based networks," in *Proc. IEEE Commun.*, 2013.
- [44] A. Bocoli, M. Schuster, F. Rambach, M. Kiese, C. Bunge, and B. Spinnler, "Reach-dependent capacity in optical networks enabled by OFDM," in *Proc. Opt. Fiber Commun.*, 2009.
- [45] E. Lach and W. Idler, "Modulation formats for 100G and beyond," *Opt. Fiber Technol.*, vol. 17, no. 5, pp. 377–386, 2011.

Leonard Nonde received the B.Eng. degree in electrical and electronic engineering from the Copperbelt University, Kitwe, Zambia, in 2005, the M.Sc. degree (with Distinction) in modern digital and radio frequency wireless communications from the University of Leeds, Leeds, U.K., in 2007. He is a Commonwealth Scholar currently working toward the Ph.D. degree in the School of Electronic and Electrical Engineering, University of Leeds, Leeds, U.K.

Prior to the Ph.D. degree, he has worked for MTN Zambia, a mobile network operator in transmission design and planning of backbone and radio access networks from 2008 to 2011. He also worked for Huawei Technologies Managed Services Division in Zambia as Network Operations Manager from 2011 to 2012. His research interests include energy aware optical networks, energy efficient routing protocols in optical networks, network virtualization and energy efficient systems.

Taisir E. H. El-Gorashi received the B.S. degree (first-class Hons.) in electrical and electronic engineering from the University of Khartoum, Khartoum, Sudan, in 2004, the M.Sc. degree (with distinction) in photonic and communication systems from the University of Wales, Swansea, U.K., in 2005, and the Ph.D. degree in optical networking from the University of Leeds, Leeds, U.K., in 2010. She is currently a Lecturer of optical networks in the School of Electrical and Electronic Engineering, University of Leeds. Previously, she held a Postdoctoral Research post at the University of Leeds (2010–2014), where she focused on the energy efficiency of optical networks investigating the use of renewable energy in core networks, green IP over WDM networks with data centers, energy-efficient physical topology design, energy efficiency of content distribution networks, distributed cloud computing, network virtualization and Big Data. In 2012, she was a BT Research Fellow, where she developed energy efficient hybrid wireless-optical broadband access network and explored the dynamics of TV viewing behavior and program popularity. The energy efficiency techniques developed during her postdoctoral research contributed 3 out of the 8 carefully chosen core network energy efficiency improvement measures recommended by the GreenTouch consortium for every operator network worldwide. Her work led to several invited talks at GreenTouch, Bell Labs, Optical Network Design and Modelling conference, Optical Fibre Communications, International Conference on Computer Communications and EU Future Internet Assembly, 2013 and collaboration with Alcatel Lucent and Huawei.

Jaafar M. H. Elmirghani received the B.Sc. degree (first-class Hons.) in electrical and electronic engineering from the University of Khartoum, Khartoum, Sudan, in 1989 and the Ph.D. degree in the synchronization of optical systems and optical receiver design from the University of Huddersfield, Huddersfield, U.K., in 1994. He is currently the Director of the Institute of Integrated Information Systems within the School of Electronic and Electrical Engineering, University of Leeds, U.K. He joined Leeds in 2007 and prior to that (2000–2007) as the Chair in optical communications at the University of Wales Swansea he founded, developed, and directed the Institute of Advanced Telecommunications and the Technium Digital, a technology incubator/spin-off hub. He has provided outstanding leadership in a number of large research projects at the IAT and TD. He has coauthored *Photonic Switching Technology: Systems and Networks*, (Wiley) and has published more than 400 papers. His research interests include optical systems and networks and signal processing. He was the Chairman of IEEE Comsoc Transmission Access and Optical Systems technical committee and was the Chairman of IEEE Comsoc Signal Processing and Communications Electronics technical committee, and an Editor of IEEE Communications Magazine. He was the Founding Chair of the Advanced Signal Processing for Communication Symposium which started at IEEE GLOBECOM'99 and has continued since at every ICC and GLOBECOM. He was also the Founding Chair of the first IEEE ICC/GLOBECOM optical symposium at GLOBECOM'00, the Future Photonic Network Technologies, Architectures and Protocols Symposium. He chaired this Symposium, which continues to date under different names. He was the Founding Chair of the first Green Track at ICC/GLOBECOM at GLOBECOM 2011, and is the Chair of the IEEE Green ICT committee within the IEEE Technical Activities Board (TAB) Future Directions Committee (FDC), a pan IEEE Societies committee responsible for Green ICT activities across IEEE, 2012–2015. He is and has been on the technical program committee of 33 IEEE ICC/GLOBECOM conferences between 1995 and 2014 including 14 times as Symposium Chair. He received the IEEE Communications Society Hal Sobol Award, the IEEE Comsoc Chapter Achievement Award for excellence in chapter activities (both in 2005), the University of Wales Swansea Outstanding Research Achievement Award, 2006, the IEEE Communications Society Signal Processing and Communication Electronics outstanding service award, 2009 and a best paper award at IEEE ICC'2013. He is currently an Editor of *IET Optoelectronics*, an Editor of *Journal of Optical Communications*, the Co-Chair of the GreenTouch Wired, Core and Access Networks Working Group, an Adviser to the Commonwealth Scholarship Commission, member of the Royal Society International Joint Projects Panel and member of the Engineering and Physical Sciences Research Council (EPSRC) College. He received in excess of £22 million in grants to date from EPSRC, the EU and industry and has held prestigious fellowships funded by the Royal Society and by BT. He is an IEEE Distinguished Lecturer. He is a Fellow of the IET and a Fellow of the Institute of Physics.