



UNIVERSITY OF LEEDS

This is a repository copy of *Energy Efficient Virtual Machines Placement in IP over WDM Networks*.

White Rose Research Online URL for this paper:  
<http://eprints.whiterose.ac.uk/116824/>

Version: Accepted Version

---

**Proceedings Paper:**

Alharbi, HA, El-Gorashi, TEH, Lawey, AQ et al. (1 more author) (2017) Energy Efficient Virtual Machines Placement in IP over WDM Networks. In: 19th International Conference on Transparent Optical Networks (ICTON 2017). 19th International Conference on Transparent Optical Networks (ICTON 2017), 02-06 Jul 2017, Girona, Spain. IEEE . ISBN 978-1-5386-0859-3

<https://doi.org/10.1109/ICTON.2017.8024957>

---

©2017 IEEE. This is an author produced version of a paper published in proceedings of 19th International Conference on Transparent Optical Networks (ICTON 2017). Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works. Uploaded in accordance with the publisher's self-archiving policy.

**Reuse**

Unless indicated otherwise, fulltext items are protected by copyright with all rights reserved. The copyright exception in section 29 of the Copyright, Designs and Patents Act 1988 allows the making of a single copy solely for the purpose of non-commercial research or private study within the limits of fair dealing. The publisher or other rights-holder may allow further reproduction and re-use of this version - refer to the White Rose Research Online record for this item. Where records identify the publisher as the copyright holder, users can verify any specific terms of use on the publisher's website.

**Takedown**

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing [eprints@whiterose.ac.uk](mailto:eprints@whiterose.ac.uk) including the URL of the record and the reason for the withdrawal request.



[eprints@whiterose.ac.uk](mailto:eprints@whiterose.ac.uk)  
<https://eprints.whiterose.ac.uk/>

# Energy Efficient Virtual Machines Placement in IP over WDM Networks

Hatem A. Alharbi, Taisir E. H. El-Gorashi, Ahmed Q. Lawey and Jaafar M. H. Elmirghani

School of Electronic and Electrical Engineering, University of Leeds, Leeds, LS2 9JT, UK

E-mail: {ml13ha@leeds.ac.uk, T.E.H.Elgorashi@leeds.ac.uk, A.Q.Lawey@leeds.ac.uk,

J.M.H.Elmirghani@leeds.ac.uk}

## ABSTRACT

Virtual machines (VMs) offer an economic and scalable solution to efficiently utilize the physical resources. In this paper, we investigate the optimization of VM placement in IP over WDM core networks considering a VM workload that varies with the number of users served by the VM. Our results show that the optimal VM placement in distributed clouds yields up to 23% total power saving compared to a single cloud.

**Keywords:** Cloud computing, virtual machine, IP over WDM network, power consumption, MILP.

## 1. INTRODUCTION

Cloud Computing has dominated the information and communication technology (ICT) industry by providing efficient resource sharing solutions where an Internet-based pool of network, storage and computational resources is made available to serve a large number of users in different locations at the same time. The growing demand for cloud resources has led to a significant increase in the size and energy consumption of cloud data centres. Improving the energy efficiency of data centres has drawn a significant amount of industrial and academic attention [1]. Cloud computing is also creating an increasing burden on the network infrastructure as it relies on the network to provide connectivity between the cloud and the clients [2]. This increasing burden on the network calls for new measures to jointly optimise the network and clouds resources to reduce the total power consumption.

In [3], the authors designed a framework to evaluate the power consumption of cloud services in IP over WDM networks. They developed a Mixed Integer Linear Programming (MILP) model to consider three clouds services; content delivery, storage as a service (StaaS) and virtual machines (VMs) for processing applications. They evaluated network related factors including the centralization versus distribution of clouds and the impact of demand, content popularity and access frequency on the clouds placement, and cloud capability factors including the number of servers, switches and routers and the amount of storage required in each cloud.

To investigate VMs, the authors in [3] assumed a constant VM workload under different number of clients. In this paper, we extend the VM placement scheme in [3] by investigating the VM placement under a linear relationship between the VM workload and the number of VM users. We study the effect of users' variation on the VMs' optimal placement in core networks that support distributed clouds. The rest of this paper is organized as follows: In Section 2, we discuss the VMs placement optimisation approach. In Section 3, we investigate the power saving obtained by optimising the VMs placement in IP over WDM network clouds. Finally, the paper is concluded in Section 4.

## 2. VM PLACEMENT OPTIMIZATION IN IP OVER WDM NETWORKS:

Virtualisation [4] technology is widely deployed in cloud data centres to provide dynamic cloud management. It offers a flexible environment to run several VMs on a single physical machine to serve a set of users. Cloud providers can flexibly add, remove, or migrate several VMs over physical servers. VMs consolidation provides a promising paradigm in the cloud by improving the resource utilization and reducing the energy consumption. In the literature, several papers have discussed the VM energy efficiency in dynamic infrastructures. The authors in [5] proposed a topology-aware VM algorithm to choose sets of communicating groups of VMs to be migrated to other data centres, to minimise the energy consumption. The selection of virtual machines considers the data centre network topology, aiming at migrating groups of VMs that will allow the switching off of physical servers and network resources. In [6], the authors studied the balance between server energy consumption and network energy consumption to present an energy-aware VM placement inside the data centre.

In [3], the authors considered three VM placement schemes; migration, replication and slicing. In migration, only one copy of each VM is allowed in the network. In replication, more than one copy of each VM can be created and located at different locations. In slicing, a single VM can be sliced to smaller VMs to serve a smaller number of users over multiple clouds. Among the three schemes, slicing is found to be the most energy efficient scheme with savings up to 25% of the total power consumption compared to a single cloud scenario. In these schemes, the VM workload was not related to the number of users served by the VM. In the migration scheme, the VM workload is assumed to be constant at different times of the day. In the replication scheme, the different replicas of the VM

are considered to have the same workload regardless of the number of users served by the replica. In the slicing scheme, all slices are assumed to be of the same workload regardless of the number of users served by the slice. In this paper, we take a more realistic approach by considering the VM workload to vary with the number of users served by the VM and study the effect of this variation on the optimal placement of VMs.

In [7], the authors performed a CPU benchmarking for different applications over multiple VM containers. The relationship between the VM workload and the number of users served by the VM is shown to be linear with a wide range of slopes. In applications such as database VM, the slope is very small making the relationship between the VM workload and number of users almost constant while in web VM the slope is large resulting in a significant increase in the VM workload as the number of users increases. We extended the VM placement MILP model in [3] to consider the relationship between the VM workload and number of served users shown in Figure 1 where each VM requires a minimum CPU utilization to run an application.

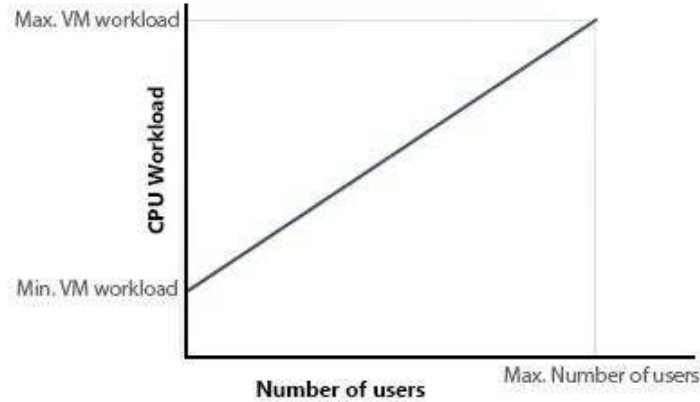


Figure 1: CPU workload of each VM versus number of users.

### 3. RESULT

The NSFNET network, depicted in Figure 2, is considered to evaluate the VMs placement in clouds supported by IP over WDM networks. The NSFNET network consists of 14 nodes with 21 bidirectional links. Cloud users are considered to be uniformly distributed over different NSFNET network nodes. We consider 200k users at 6 AM, and 1200k users at 10 PM which represent the off-peak and peak users demand, respectively.

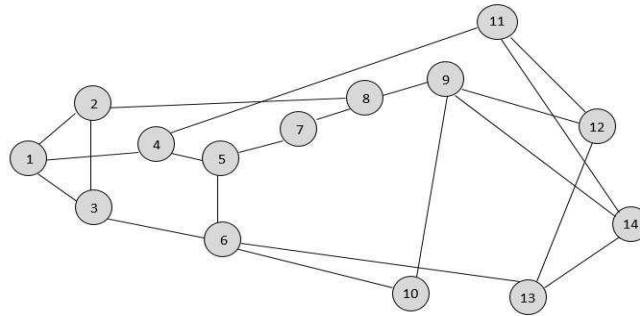


Figure 2: The architecture of NSFNET network.

In this paper, we consider 500 VMs in the cloud. The maximum VM workload, ranges between 10% and 100% of the server's CPU utilization while the minimum VM workload is set to 5%. In our investigation, we compare the power consumption of the optimized VM placement scenario to a single cloud scenario where single copies of all VMs are placed in node 6, the node with minimum hop distance to all other nodes in NSFNET. Table 1 shows the input data for the VM placement model.

Table 1. Model Input Parameters:

Router port power consumption	825 W [8]
Transponder power consumption	167 W [8]
Regenerator power consumption	334 W, reach 2500 KM [8]
EDFA power consumption	55 W [8]
Optical switch power consumption	85 W [8]
Server maximum power consumption	339 W [9]
Number of VM in the cloud	500 VMs
Set of VM maximum CPU workload	Random {10, 20, 30, 40, 50, 60, 70, 80, 90, 100} %
Baseline CPU workload per VM	5%
VM clients	200k users at 6 AM, and 1200k users at 10PM
Cloud switch power consumption	3800 W [10]
Cloud switch capacity	320 Gbps [10]
Cloud router power consumption	5100 W [10]
Cloud router capacity	660 Gbps [10]

Figure 3 shows the optimal VM placement at the different network nodes at 6 AM and 10 PM. At 6 AM, nodes 1 and 9 are selected as the optimum cloud locations to host a copy of each VM to serve its closest nodes' users as they yield the minimum average hop count to cloud users. Node 1 serves the users' demand from the west nodes (node 1 to node 6) and node 9 serves the users' demand in the east nodes (node 7 to node 14). On the other hand, at 10 PM, due to the high traffic demand, multiple copies of each VM are distributed through the different nodes to serve users either locally or within a single hop. Each cloud hosts between 495–500 VMs.

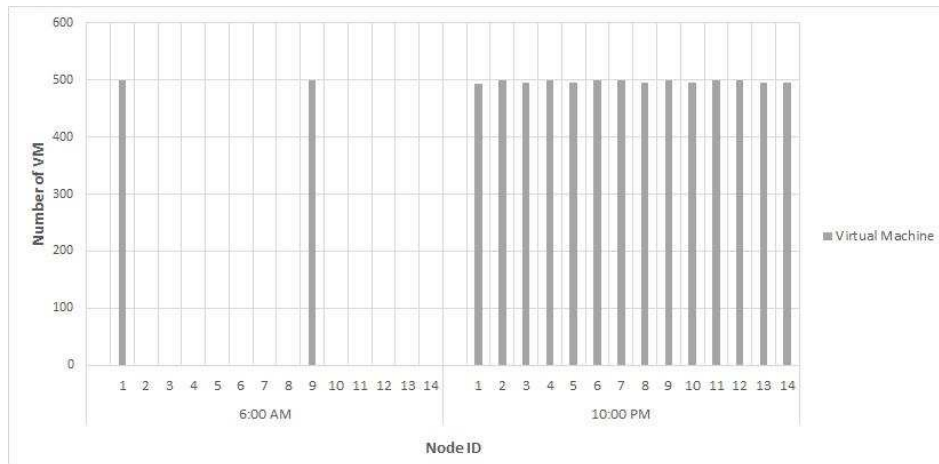


Figure 3. VM placement over core network clouds at 6 AM and 10 PM.

Figure 4 shows the cloud and IP over WDM network power consumption of a single cloud and an optimised VMs placement scenarios at low and high users demands. At low users' demand, although 22% power saving is achieved in the core network power consumption, creating two copies of each VM increases the cloud power consumption and limits the total power saving to 1% compared to a single cloud scenario. Whereas, at high users' demand, optimizing the cloud locations achieves 23% saving in the total power consumption compared to a single cloud scenario. Creating multiple copies of the VMs has increased the clouds power consumption to 39% compared to a single cloud scenario. But, we save 99% of core network power consumption since the majority of VM demands are served locally.

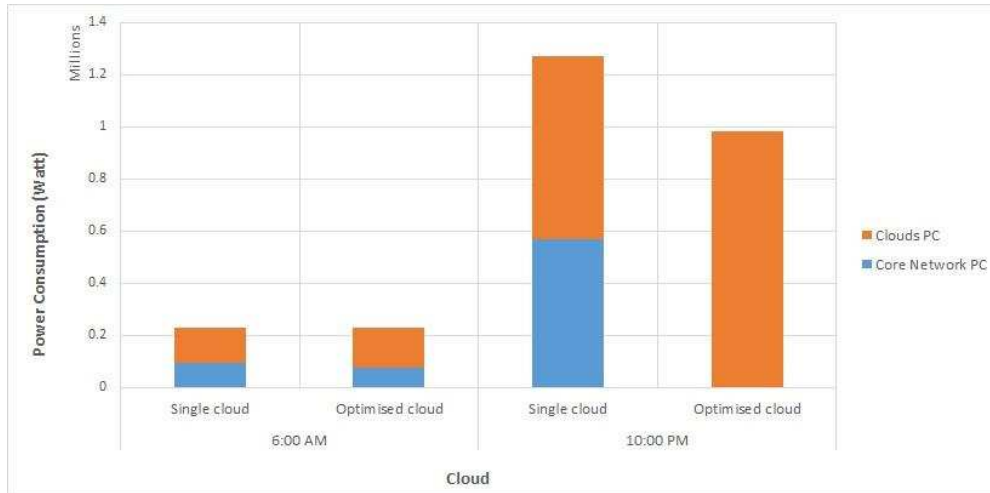


Figure 4. Cloud and IP over WDM core network power consumption (PC) at 6 AM and 10 PM.

#### 4. CONCLUSIONS

This paper has optimized the placement of VMs over cloud data centres hosted in IP over WDM networks. In our optimization, we considered a VM workload that varies with the number of users served by the VM. Specifically, we considered a linear relationship between the VM workload and the number of VM users. Our results show that at low users' demand the total power savings obtained by optimizing the VM placement is very limited due to the low network traffic. At high users' demand, a significant total power saving of 23% is achieved by optimally distributing multiple copies of the VM in IP over WDM networks.

#### ACKNOWLEDGEMENTS

The authors would like to acknowledge funding from the Engineering and Physical Sciences Research Council (EPSRC), INTERNET (EP/H040536/1) and STAR (EP/K016873/1) projects. The first author would like to thank the Government of Saudi Arabia and Taibah University for funding his PhD scholarship. All data are provided in full in the results section of this paper.

#### REFERENCES

- [1] A. Beloglazov, R. Buyya, Y. C. Lee, and A. Zomaya, *A Taxonomy and Survey of Energy-Efficient Data Centers and Cloud Computing Systems*, vol. 82. 2011.
- [2] N. Fonseca and R. Boutaba, *Cloud Services, Networking, and Management*. JohnWiley & Sons, Inc., 2015.
- [3] A. Q. Lawey, T. E. H. El-Gorashi, and J. M. H. Elmirghani, "Distributed energy efficient clouds over core networks," *J. Light. Technol.*, vol. 32, no. 7, pp. 1261–1281, 2014.
- [4] VMware, "Virtualization Essentials," 2014. [Online]. Available: <http://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/ebook/gated-vmw-ebook-virtualization-essentials.pdf>. [Accessed: 22-Mar-2017].
- [5] R. A. C. Silva and N. L. S. Fonseca, "Energy-aware migration of groups of virtual machines in distributed data centers," in *IEEE Global Communications Conference (GLOBECOM)*, 2016.
- [6] D. Huang, D. Yang, H. Zhang, and L. Wu, "Energy-aware virtual machine placement in data centers," *GLOBECOM - IEEE Glob. Telecommun. Conf.*, pp. 3243–3249, 2012.
- [7] P. Padala, X. Zhu, Z. Wang, S. Singhal, and K. G. Shin, "Performance Evaluation of Virtualization Technologies for Server Consolidation," 2007.
- [8] "GreenTouch Green Meter Research Study : Reducing the Net Energy Consumption in Communications Networks by up to 90 % by 2020," White Pap., pp. 1–25, 2015.
- [9] "IBM System Energy Estimator," S814-8286-41A. [Online]. Available: <http://see.ausyd.mybluemix.net/see/EnergyEstimator>. [Accessed: 15-Mar-2017].
- [10] J. Baliga, R. W. A. Ayre, K. Hinton, and R. S. Tucker, "Green Cloud Computing: Balancing Energy in Processing, Storage, and Transport," *Proc. IEEE*, vol. 99, no. 1, pp. 149–167, 2011.