University of Leeds
Faculty of Engineering
School of Electronic and Electrical Engineering

Energy Efficient Cloud-Fog Networks Architecture

Hatem A. Alharbi

Submitted in accordance with the requirements for the degree of

Doctor of Philosophy

September 2019

The candidate confirms that the work submitted is his own, except where work which has formed part of jointly-authored publications has been included. The contribution of the candidate and the other authors to this work has been explicitly indicated below. The candidate confirms that appropriate credit has been given within the thesis where reference has been made to the work of others.

Chapter 3 is based partly on the work from:

H. A. Alharbi, T. E. H. El-Gorashi, A. Q. Lawey and J. M. H. Elmirghani, "Energy efficient virtual machines placement in IP over WDM networks," 2017 19th International Conference on Transparent Optical Networks (ICTON), Girona, 2017, pp. 1-4.

Prof Elmirghani, the supervisor, suggested the study of Energy efficient virtual machines placement in IP over WDM networks. The co-supervisor Dr El-Gorashi, worked with the student on the MILP model development and revised the paper. Dr. Lawey, helped in the initial model preparation. The PhD student developed the model, obtained and analysed the results, and wrote paper.

Chapter 4 is based partly on the work from:

H. A. Alharbi, T. E. H. El-Gorashi, A. Q. Lawey and J. M. H. Elmirghani, "The Impact of Inter-Virtual Machine Traffic on Energy Efficient Virtual Machines Placement," The 2019 IEEE Sustainability through ICT Summit (StICT), Montreal, 2019.

Prof Elmirghani, the supervisor, suggested the study of The Impact of Inter-Virtual Machine Traffic on Energy Efficient Virtual Machines Placement. The co-supervisor Dr El-Gorashi, worked with the student on the MILP model development and revised the paper. Dr. Lawey, helped in the initial model preparation. The PhD student developed the model, obtained and analysed the results, and wrote paper.

Chapter 5 is based partly on the work from:

H. A. Alharbi, T. E. H. El-Gorashi and J. M. H. Elmirghani, "Energy Efficient Virtual Machine Services Placement in Cloud-Fog Architecture," 2019 21th International Conference on Transparent Optical Networks (ICTON), Angers, 2019.

Prof Elmirghani, the supervisor, suggested the study of Energy Efficient Virtual Machine Services Placement in Cloud-Fog Architecture. The co-supervisor Dr El-Gorashi, worked with the student on the MILP model development and revised the paper. The PhD student developed the model, obtained and analysed the results, and wrote paper.

Acknowledgements

I would like to express my sincere gratitude to my supervisor Prof. Jaafar Elmirghani for the support during my PhD, for his constant support, patience, comments and sharing of his immense knowledge. His guidance and advices helped me during the time of research and writing of this thesis.

I would like to express my sincere gratitude to Dr Taisir Elgorashi, my second supervisor for the support, encouragement, comments, and useful discussions I had with her during my PhD. Her insights and advices paved the road toward accomplishing the goals of this thesis.

I would like to acknowledge the Government of Saudi Arabia for funding my PhD study and for their support.

I would like to express my love and gratitude to my beloved wife Abeer who has been a constant source of support, patient and encouragement during the challenges of my PhD and life. I am truly thankful for having you in my life.

I would like to express my love to my daughter Miral who fills me with pleasure and joy.

Finally, I would like to dedicate this thesis to my parents; to my father, Ayed and my mother, Azizah. I have been extremely fortunate in my life to have parents who have shown me unconditional love and support. They have played an important role in the development of my identity and shaping who I am today. I am very grateful for their sacrifices, their dedication, for inspiring me with the desire for success, and for supporting me spiritually throughout my life. Many thanks also to my beloved brothers and sisters; Gehad, Ahmed, Nada, Suha and Raghad for their support and for making my life enjoyable and full of fun.

Hatem

Abstract

The advancements of cloud computing came as a radical transformation in the way Information and Communication Technology (ICT) services are deployed, maintained and paid for. Cloud computing provides ubiquitous on-demand access to an Internet-based pool of processing, storage, and communication resources offered to a large set of geographically distributed users. As the cloud computing infrastructure grows and demand increases, the need for a new breed of on-demand computing that can efficiently and profitably maintain Quality of Service (QoS) requirements has increased. Fog computing was proposed to address the limitations of cloud computing in terms of delay and high bandwidth requirements by extending the on-demand resources of clouds to the edge of the network bringing them closer to the users. Cloud computing and fog computing employ Virtual Machines (VMs) for efficient resource utilization. To achieve the most of the efficient environment, VMs can be optimally placed over geo-distributed physical machines to follow variations in the demand, workload of the cloud/fog resources or network condition.

The massive growth and wide use of cloud-fog services have created serious power consumption concerns. In this thesis, we study the energy efficiency and profitability of cloud-fog architectures taking into consideration different services and applications. Mixed integer linear programming (MILP) optimisation models and heuristic algorithms are developed to optimise the cloud-fog architecture. The thesis starts by analysing the energy efficiency of VMs placement over distributed clouds. The analysis addresses the impact of different factors including VM popularity, the traffic between the VM and its users, the VM workload, the workload versus number of users profile and power usage effectiveness (PUEs). The total power savings achieved are up to 51% and 38%, compared to the power consumption of the traditional cloud locations in AT&T and BT core network topologies, respectively. Furthermore, the impact of inter-VM traffic (synchronisation and cooperation between VMs) on the energy efficiency is investigated in the context of optimal VM placement. The existence of inter-VM traffic has resulted on reducing the energy efficiency of replicating VMs. Our results showed that ignoring inter-VM traffic when placing VMs can increase the total power consumption by a factor of 39 for VMs with

an inter-VM traffic data rate of 5 Gbps. Then, the thesis shows how introducing fog layers at the edge of the network in addition to the cloud layer in the core network can significantly decrease the total power consumption of VMs of certain popularity, download rates and workloads. In addition, a pricing scheme is proposed for Internet service providers (ISPs) to maximise their profit under net neutrality repeal scenarios where ISPs are allowed to treat different traffic strands differently. The results shows that the VMs of a linear workload profile with high data rate and minimum CPU utilisation of 1% allows offloading VMs with 16% popularity to the access fogs. Other VMs are optimally replicated to metro fog nodes. Significant power savings of 48% compared to optimised placement in distributed clouds and 64% compared to a placement considering traditional cloud locations, have resulted from this offloading. A techno-economic MILP model is used to find the optimum prices of different levels of service and the resulting power consumption. The proposed pricing scheme proves that, in addition to increasing ISPs profit, the repeal of net neutrality will positively influence (limit) the end-users' consumption of data-intensive content which consequently decreases the power consumption of the communication networks. The results show that the discriminatory pricing scheme can increase the ISPs profit by a factor of 8. The results also show that by influencing the way end-users consume data-intensive content, the core network traffic and consequently power consumption are reduced by up to 49% and 55%, respectively, compared to the net neutrality scenario.

Contents

List of Figures

List of Tables

List of Abbreviations

AI	Artificial intelligence
Amazon EC2	Amazon Elastic Compute Cloud
AMPL	A Mathematical Programming Language
AT&T	American Telephone & Telegraph
AWS	Amazon Web Service
B&B	Branch-and-bound
B&C	Branch-and-cut
Bps	Bit Per Seconds
BT	British Telecom
CDN	Content Delivery Network
CO	Central Office
CP	Content Providers
CPU	Central Processing Unit
CRS	Carrier Routing System
CT	Communication Technology
DiffServ	Differentiated Service
DPI	Deep Packet Inspection
EDFA	Erbium Doped Fibre Amplifier
FCC	Federal Communications Commission
Gbps	Gigabit Per Seconds
GHG	Green House Gases
GIS	Geographic Information Systems
GPON	Gigabit Passive Optical Network
HD	High Definition
IaaS	Infrastructure as a service
ICT	Information and Communication Technologies
ILP	Integer Linear Programming
IoT	Internet of Things

IP	Internet Protocol
ISP	Internet Service Provider
IT	Information Technology
KM	Kilometres
LAN	Local Area Network
LP	Linear Programming
Mbps	Mega Bits Per Seconds
Metro	Metropolitan
MILP	Mixed Integer Linear Program
MPLS	MultiProtocol Label Switching
Net Neutrality	Network Neutrality
NIST	American National Institute of Standards and Technology
NP	Nondeterministic Polynomial
NSFNET	National Science Foundation Network
OEO	Optical Electrical Optical
OLT	Optical Line Terminal
ONU	Optical Network Unit
OXC	Optical Switch
PaaS	Platform as a service
PED	Price Elasticity of Demand
PON	Passive Optical Network
PUE	Power Usage Effectiveness
QoE	Quality of Experience
QoS	Quality Of Service
RAM	Random Access Memory
SaaS	Software as a service
SD	Standard Definition
SLA	Service Level Agreement
Tbps	Tera Bits Per Seconds
UHD	Ultra-High Definition
US	The United States
VM	Virtual Machine

VMM	Virtual Machine monitor
VNI	Visual Network Index
WAN	Wide area Network
WDM	Wavelength Division Multiplexing
XGPON	10 Gigabit Passive Optical Network

List of Symbols

N	Set of IP over WDM network nodes. Set of nodes represents nodes distributed across different geographic regions in the selected core network topology.
VM	Set of VMs.
$SW^{(CB)}$	Cloud switch bit rate.
$SW^{(CP)}$	Cloud switch power consumption.
$SW^{(R)}$	Cloud switch redundancy.
$R^{(CB)}$	Cloud router port bit rate.
$R^{(CP)}$	Cloud router port power consumption.
$S^{(P)}$	Power consumption of a server.
$S^{(maxW)}$	Maximum workload of a server.
c	Cloud power usage effectiveness.
C_s	$C_s = 1$ if a cloud is hosted in node s, otherwise $C_s = 0$.
$\delta_{v,s}^{(C)}$	$\delta_{v,s}^{(C)} = 1$ if the cloud hosted in node s hosts a copy of VM v, otherwise $\delta_{v,s}^{(C)} = 0$.
$R_s^{(C)}$	Number of router aggregation ports in the cloud hosted in node s.
$SW_s^{(C)}$	Number of switches in the cloud hosted in node s.
$S_s^{(C)}$	Number of processing servers in the cloud hosted in node s.
$s\ and\ d$	Indices of source and destination nodes of a traffic flow in the distributed cloud architecture.
V	Number of VMs.
S_v	Number of VM v users.
r_v	User download rate of VM v.
L	Large enough number.
x	Maximum number of users served by a single VM replica.
W_v	Maximum workload of VM v (workload can be specified in GHz or as a ratio of the CPU capacity).
M	Workload baseline of VM (the minimum CPU utilisation needed in the absence of load). To maintain the service level agreement (SLA), each VM needs a minimum workload to run an application regardless of the number of users served by the VM.
T_v	Traffic resulting from VM replica v serving the maximum number of users. $$T_v = x\, r_v$$
$W_v^{(R)}$	Workload per traffic unit,

$$W_v^{(R)} = \frac{W_v - M}{T_v}$$

evaluated for VM replica v.

$W_{v,s}^{(CR)}$ Workload of VM replica v hosted in cloud in node s.

$W_s^{(C)}$ Total workload of cloud hosted in node s.

$D_{v,s,d}^{(C)}$ Traffic flow from VM replica v hosted in cloud of node s to users in node d.

$L_{s,d}$ Traffic from cloud node s to users in node d.

P Set of PON networks.

A_p Average broadband data rate in PON p.

Φ_v ratio of traffic due to VM v to the total PON traffic.

$OLT_{p,d}^{(B)}$ Capacity of OLT serving PON p connected to node d.

$U_{v,p,d}$ Number of users in PON p connected to core node d requesting VM v.

$$U_{v,p,d} = \left(\frac{OLT_{p,d}^{(B)}}{A_p}\right) \Phi_v$$

if typical national/regional values of A_p, Φ_v and $OLT_{p,d}^{(B)}$ are used, then $U_{v,p,d}$ determines the number of users and their VM popularity.

$OLT_{p,d}^{(N)}$ Number of OLTs in PON network p connected to node d.

$OLT^{(P)}$ OLT power consumption.

$D_{v,p,d}$ Traffic flow from VM v to users in PON network p connected to core node d given as:

$$D_{v,p,d} = U_{v,p,d}\, r_v$$

$ONU_{p,d}^{(N)}$ Number of ONUs in PON network p connected to node d.

$ONU^{(P)}$ Power consumption of an ONU.

n Network power usage effectiveness (PUE). PUE is a metric used to determine the energy efficiency of a network by dividing the total amount of energy consumed of network facility (including lighting and cooling) by the energy consumed by networking equipment.

$R^{(MB)}$ Metro router bit rate.

$R^{(MP)}$ Metro router power consumption.

$R^{(MR)}$ Metro router redundancy.

$SW^{(MB)}$ Metro Ethernet switch bit rate.

$SW^{(MP)}$ Metro Ethernet power consumption.

$R_s^{(M)}$ Number of router ports in metro network connected to node s.

$SW_s^{(M)}$ Number of Ethernet switches in metro network connected to node s.

$m\ and\ n$ Indices of the end nodes of a physical link.

$i\ and\ j$ Indices of the end nodes of a virtual link.

Nm_m	Set of neighbouring nodes of node m.
$R^{(P)}$	Core router port power consumption.
$t^{(P)}$	Transponder power consumption.
$e^{(P)}$	EDFA power consumption.
$SW_s^{(P)}$	Optical switch power consumption in node s.
$G^{(P)}$	Regenerator power consumption.
$\mathcal{W}$	Number of wavelengths per fibre.
$\mathcal{W}^{(B)}$	Wavelength data rate.
S	Maximum span distance between two EDFAs in kilometres.
$D_{m,n}$	Distance in kilometres between node pair (m,n).
$A_{m,n}$	Number of EDFAs between node pair (m,n). $A_{m,n}$= $\left\lfloor \frac{D_{m,n}}{S} - 1 \right\rfloor$ where S is the reach of the EDFA.
$G_{m,n}$	Number of regenerators between node pair (m,n). Typically $G_{m,n}$= $\left\lfloor \frac{D_{m,n}}{R} - 1 \right\rfloor$, where R is the reach of the regenerator.
$C_{i,j}$	Number of wavelengths in virtual link (i,j).
$\mathcal{W}_{m,n}$	Number of wavelengths in physical link (m,n).
$R_s^{(AC)}$	Number of router ports in node s that aggregate the traffic from/to clouds.
$R_d^{(AE)}$	Number of router ports in node d that aggregate the traffic from/to metro routers.
$F_{m,n}$	Number of fibres on physical link (m,n).
$L_{i,j}^{s,d}$	Amount of traffic flow between node pair (s,d) traversing virtual link (i,j).
$\mathcal{W}_{m,n}^{i,j}$	Number of wavelengths of virtual link (i,j) traversing physical link (m,n).
$v\ and\ y$	Indices of VMs.
$CV_{v,y}$	$CV_{v,y}$ =1, if there is cooperating traffic between VM v and VM y, otherwise $CV_{v,y}$ =0.
$CVM_{v,y}$	Cooperation traffic between VM v and VM y.
SVM_v	Synchronisation traffic between VM v replicas.
Ψ_d	Download traffic of cloud hosted in node d.
$\Phi_{v,y,s,d}$	$\Phi_{v,y,s,d} = 1$, if VM v located in node s is a candidate to cooperate with VM y located in node d, otherwise $\Phi_{v,y,s,d} = 0$.
$б_{v,y,d}$	$б_{v,y,d} = 1$, if cooperation traffic exists from VM v located at any node to VM y located in node d, otherwise $б_{v,y,d} = 0$.
$\chi_{v,y,s,d}$	$\chi_{v,y,s,d} = 1$, if cooperation traffic exists from VM x located in node s to VM y located in node d, otherwise $\chi_{v,y,s,d} = 0$.
$\beta_{v,y,s,d}$ $\boldsymbol{\alpha_{v,y,s,d}}$	Binary variables set to 1 only if one or two of the following conditions are satisfied; there is a cooperating traffic from VM $\boldsymbol{v}$ to VM $\boldsymbol{y}$, VM $\boldsymbol{v}$ is located in node $\boldsymbol{s}$ or VM $\boldsymbol{y}$ is located in node $\boldsymbol{d}$, otherwise $\boldsymbol{\beta_{v,y,s,d}}$ and $\boldsymbol{\alpha_{v,y,s,d}}$ are set to 0.

$IC_{v,y,s,d}$	Cooperating traffic from VM v to VM y located in nodes s and d.
$\Theta_{v,s,d}$	$\Theta_{v,s,d} = 1$, if VM v replicas are located in nodes s and d, respectively, otherwise $\Theta_{v,s,d} = 0$.
$\varphi_{v,s,d}$	$\varphi_{v,s,d} = 1$, if only one VM v replica is located in either node s or node d, otherwise $\varphi_{v,s,d} = 0$.
$IS_{v,s,d}$	Synchronisation traffic between VM v replicas located in nodes s and d, respectively.
$R_d^{(IV)}$	Number of router ports in node d that aggregate the traffic to clouds.
$SW^{(MFB)}$	Metro fog switch bit rate.
$SW^{(MFP)}$	Metro fog switch power consumption.
$SW^{(AFB)}$	Access fog switch bit rate.
$SW^{(AFP)}$	Access fog switch power consumption.
$R^{(MFB)}$	Metro fog router port bit rate.
$R^{(MFP)}$	Metro fog router port power consumption.
$R^{(AFB)}$	Access fog router port bit rate.
$R^{(AFP)}$	Access fog router port power consumption.
m	Metro fog power usage effectiveness.
a	Access fog power usage effectiveness.
$F_s^{(MF)}$	$F_s^{(MF)} = 1$ if a fog processing node is hosted in the metro network connected to core node s, otherwise $F_s^{(MF)} = 0$.
$\delta_{v,s}^{(MF)}$	$\delta_{v,s}^{(MF)} = 1$ if the fog processing node hosted in the metro network connected to node s hosts a replica of VM v, otherwise $\delta_{v,s}^{(MF)} = 0$.
$R_s^{(MF)}$	Number of router ports used in the fog processing node hosted in the metro network connected to node s.
$SW_s^{(MF)}$	Number of switches used in the fog processing node hosted in the metro network connected to node s.
$S_s^{(MF)}$	Number of processing servers in the fog processing node hosted in the metro network connected to node s.
$F_{p,s}^{(AF)}$	$F_{p,s}^{(AF)} = 1$ if a fog processing node is built in access network p connected to core node s, otherwise $F_{p,s}^{(AF)} = 0$.
$\delta_{v,p,s}^{(AF)}$	$\delta_{v,p,s}^{(AF)} = 1$ if the fog processing node in access network p connected to core node s, hosts a replica of VM v, otherwise $\delta_{v,p,s}^{(AF)} = 0$.
$R_{p,s}^{(AF)}$	Number of router ports used in the fog processing node located in the access network p connected to core node s.
$SW_{p,s}^{(AF)}$	Number of switches used in the fog processing node located in access network p connected to core node s.
$S_{p,s}^{(AF)}$	Number of processing servers in the fog processing node located in the access network p connected to core node s.
$W_{v,s}^{(MFR)}$	Workload of the VM replica v hosted in the fog processing node located in the metro network connected to node s.

$W_s^{(MF)}$ Total workload of the metro fog processing node located in core node s.

$D_{v,s}^{(MF)}$ Traffic from the VM replica v hosted in the fog processing node of the metro network connected to core node s.

$W_{v,p,s}^{(AFR)}$ Workload of the VM replica v hosted in the fog processing node located in the access network p connected to core node s.

$W_{p,s}^{(AF)}$ Total workload of the fog processing node located in the access network p connected to core node s.

$D_{v,p,s}^{(AF)}$ Traffic flow from the VM replica v hosted in the fog processing node located in the access network p connected to core node s.

α Set of service classes.

CN Number of clouds hosted in core network.

u Total number of users in net neutrality scenario (i.e. before net neutrality is repealed).

LB Minimum percentage of users served by CP to be maintained by the pricing scheme.

d_i Download rate of class i.

℄ The cost in US$ of provisioning a Gbps of IP over WDM network bandwidth per month.

Ɔ The cost in US$ of provisioning a Gbps of metro and access network bandwidth per month.

PS The net neutrality selling price in US$ of a Gbps of network bandwidth per month.

E_i Price elasticity of demand of class i.

$N_{d,i}$ Number of users of class i located in node d under net neutrality scenario.

δ_s $\delta_s = 1$, if a cloud datacentre is hosted in node s, otherwise $\delta_s = 0$.

F_d $F_d = 1$, if there is no fog datacentre hosted in node d, otherwise $F_d = 0$.

Ƨ Set of all possible solutions.

$\rho_{s,i}$ The price of class i under solution s and class i.

$yn_{s,d,i}$ The number of users in solution s subscribing to class i in node d as a result of its PED, where

$$\frac{PS}{\rho_{s,i} - PS} E_i = \sum_{d \in N} \left(\frac{yn_{s,d,i} - N_{d,i}}{N_{d,i}} \right)$$

$$\forall\ i \in \alpha,\ s \in Ƨ.$$

r_i ISP's revenue achieved by delivering traffic of class i to CP users.

R Total ISP's revenue in US$ of delivering networking services to CPs content.

C Total ISP cost in US$ of provisioning core network.

P_i The price in US$ per Gbps of network bandwidth per month charged to the class i.

$\boldsymbol{U_{d,i}}$	Number of users who subscribe to class $\boldsymbol{i}$ located in node $\boldsymbol{d}$.
$\boldsymbol{CD_{i,d}}$	Cloud flow from users in node $\boldsymbol{d}$ subscribed to class $\boldsymbol{i}$.
$\boldsymbol{Z_{s,i}}$	$\boldsymbol{Z_{s,i} = 1}$, if solution $\boldsymbol{s}$ is selected for class $\boldsymbol{i}$, otherwise $\boldsymbol{Z_{s,i} = 0}$.
$\boldsymbol{ys_{s,d,i}}$	The number of users in solution $\boldsymbol{s}$ subscribing to class $\boldsymbol{i}$ in node $\boldsymbol{d}$, $\boldsymbol{ys_{s,d,i} > 0}$ if solution $\boldsymbol{s}$ is selected for class $\boldsymbol{i}$, otherwise $\boldsymbol{ys_{s,d,i} = 0}$.

Chapter 1: Introduction

The significant impact of Information and Communications Technologies (ICT) services in people daily lives led to an increasing perception that cloud computing is the 5th utility after water, electricity, gas, and telephony. Cloud Computing has dominated the ICT industry by providing efficient resource sharing solutions where an Internet-based pool of network, storage and computational resources is made available to simultaneously serve a large number of geographically distributed users. Cloud computing provides the bandwidth, memory, and processing power needed to serve Big Data, Internet of things (IoT), Artificial Intelligence (AI) applications such as image recognition, video analytics, augmented and virtual reality with real-time processing of streaming data.

According to Cisco [1], in 2017, the global cloud computing traffic was 56% of the Internet traffic. Further growth is projected within the approaching years as global cloud computing traffic is expected to hit 72% of the Internet traffic in 2022. This proliferation in data volume and processing requirements increases the need for a new breed of on-demand computing placement and administration. Fog computing is proposed by academia and industry to bring cloud services closer to users. Fog computing complements the clouds by extending processing, networking and storage resources to the edge of the network. Offloading tasks to fog nodes is proposed to provide low operating cost, low latency, preserve network bandwidth, provide a real-time analytics and interactions, improve security and improve Quality-of-Service (QoS) and Quality-of-Experience (QoE) for different computing services [2]-[3] .

Cloud and fog processing employs Virtual Machines (VMs) for efficient resource utilisation. Virtualisation abstracts the server resources including the Central Processing Unit (CPU), Random Access Memory (RAM), hard disk and Input/output (I/O) network to create an isolated virtual entity that can run its operating system and applications. The existence of such a virtual environment allows the scaling up and down of server resources in a dynamic manner based on the variation in user demands [4]. Further dynamism can be achieved by migrating or replicating VMs over geo-distributed servers to achieve different features such as load balancing [5] and energy

efficiency [6]. The problem of migration and replication of VMs is referred to as VMs placement. VMs placement needs to be optimised to follow variations in the VMs demands, workload of the cloud/fog resources or network status [7].

As a result of the significant growth in demand for ICT services, energy efficiency has been recognised in the last few years as one of the core requirements needed to develop a sustainable ICT infrastructure. By 2025, the ICT industry is projected to consume 20% of the global electricity demand [8]. Virtualisation has been proposed in [6], [9]–[16] as an enabler for energy efficient cloud services through the consolidation of cloud resources. In this thesis, we present a comprehensive study of VMs placement over an end-to-end cloud-fog architecture considering the three network layers; access network and metro network layer (each equipped with a fog computing layer) and the core network layer (equipped with a cloud computing layer).

In addition to providing energy efficient cloud-fog services, delivering these services needs to be profitable. Therefore, proposing a pricing scheme to maximise the profit of Internet service providers (ISPs) is imperative. However, the maximisation of profit is challenging for ISPs as it involves the concurrent determination of equilibrium prices and the profit achieved. In this thesis, in addition to energy efficiency improvement, a profit-driven framework is introduced to increase the ISPs profit. Also, the impact of the profit-driven framework on the network power consumption and network neutrality is investigated.

The thesis delves into the energy consumption and the profitability of delivering cloud-fog services by raising headline questions related to; how significant the problem itself is, how different conditions/scenarios affect the energy consumption and the profitability of the architecture, and how to orchestrate the use of the architecture in a profitable manner and in an energy-efficient manner.

1.1 Research Objectives:

This thesis tackles challenges related to providing energy-efficient and profitable cloud-fog architectures. It has the following objectives:

1. To propose an energy efficient VMs placement approach over distributed clouds in IP over WDM networks with datacentres to minimise the total power consumption in the core network and in datacentres.
2. To investigate the impact of inter-VM traffic, in addition to users download traffic, on the energy efficient VM placement considering different VM workloads.
3. To propose an energy efficient VMs placement approach over cloud-fog architectures to minimise the total power consumption considering the impact of the proximity of fog nodes to users.
4. To investigate how the repeal of network (net) neutrality will affect the ICT market (pricing, core network power consumption and core network traffic) considering different Content Providers (CP) delivery scenarios in cloud-fog architectures.

1.2 Original Contributions:

The key contributions of this thesis can be outlined as follows:

- A framework for energy efficient VMs placement over distributed clouds in IP over WDM networks has been developed. The framework considers numerous core factors including the traffic between the VM and its users, the VM workload, the workload versus number of users profile and the cloud power usage effectiveness (PUE). With the objective of minimising the total power consumption of providing the VM services, a MILP model and a heuristic are developed and are shown to minimise the total power consumption of providing the VM services. The optimal VMs placement and total power consumptions savings are investigated considering American Telephone & Telegraph (AT&T) and British Telecom (BT) core network topologies.
- A MILP model and a heuristic are developed to study the impact of inter-VMs traffic on the energy efficient placement of VMs in geo-distributed datacentres. The model and heuristic take into consideration cooperation traffic between different VMs and synchronisation traffic between replicas of the same VM in addition to the download traffic from VM to users. The National Science

Foundation Network (NSFNET) is used as a core network topology example to investigate the energy efficient placement of VMs here.

- A comprehensive framework analysis based on MILP mathematical modelling and heuristics is developed to study the optimal VMs placement over cloud-fog architectures taking into consideration the minimisation of the total power consumption of the architecture. The placement of VMs is optimised over an end-to-end cloud-fog architecture that traverses the core network, metro network and access network. The optimal VMs placement and total power consumptions savings are investigated considering AT&T and BT network topologies.
- A techno-economic MILP model is developed to optimise the pricing scheme in the ISP-CP relationship under the repeal of net neutrality to maximise the ISP profit considering could-based and fog-based delivery schemes. The impact of the optimised pricing scheme on reducing the network traffic and total power consumption of the AT&T core network is studied.

1.3 Related Publications:

The following journal and conference papers have been published / are to be submitted for publications:

1. H. A. Alharbi, T. E. H. El-Gorashi, A. Q. Lawey and J. M. H. Elmirghani, "Energy Efficient Virtual Machines Placement in IP over WDM networks," 2017 19th IEEE International Conference on Transparent Optical Networks (ICTON), Girona, Spain 2017, pp. 1-4.
2. H. A. Alharbi, T. E. H. El-Gorashi, A. Q. Lawey and J. M. H. Elmirghani, "The Impact of Inter-Virtual Machine Traffic on Energy Efficient Virtual Machines Placement," The 2019 IEEE Sustainability through ICT Summit (StICT), Montreal, Canada, 2019.
3. H. A. Alharbi, T. E. H. El-Gorashi and J. M. H. Elmirghani, "Energy Efficient Virtual Machine Services Placement in Cloud-Fog Architecture," 2019 21st IEEE International Conference on Transparent Optical Networks (ICTON), Angers, France, 2019.

4. Hatem A. Alharbi, T. E. H. El-Gorashi, and Jaafar M. H. Elmirghani "Energy Efficient Cloud-Fog Architectures", submitted to IEEE Communication Magazine, 2020.
5. Hatem A. Alharbi, T. E. H. El-Gorashi, and Jaafar M. H. Elmirghani " Energy Efficient Virtual Machines Placement over Cloud-Fog Network Architecture", submitted to IEEE Access, 2020.
6. H. A. Alharbi, T. E. H. Elgorashi and J. M. H. Elmirghani, "Impact of the Net Neutrality Repeal on Communication Networks," in IEEE Access.
7. Hatem A. Alharbi, M. Musa, T. E. H. El-Gorashi and Jaafar M. H. Elmirghani IEEE P1928.1 - Standard for a Mechanism for Energy Efficient Virtual Machine Placement to be submitted to IEEE Standard Association; IEEE P1928.1 Project Authorisation Request (PAR) accepted and published.

1.4 Organisation of the Thesis:

The remainder of this thesis is organised as follows:

Chapter 2 is a background chapter. It presents a review of the core network, metro network, access network and their components. It also presents a review of cloud and fog computing as well as the ongoing efforts to optimise the placement of VMs to improve the energy efficiency. A brief review of the net neutrality concept and work related to it in the literature are presented.

Chapter 3 presents the MILP optimisation model developed for energy efficient VMs placement in distributed clouds, discusses its results, and proposes a real-time heuristic.

Chapter 4 introduces the MILP optimisation model developed to study the impact of inter-VM traffic on energy efficient VMs placement in distributed clouds and describes the associated real-time heuristic.

Chapter 5 studies the energy efficiency of VMs placement over a cloud-fog architecture using MILP modelling and a real-time heuristic.

Chapter 6 studies the impact of the net neutrality repeal on communication networks considering different CP delivery scenarios over cloud-fog architectures using a MILP model.

Chapter 7 summarises the thesis main contributions and gives a set of recommendation for future directions of work.

Chapter 2: Background and Related Work

2.1 Introduction

In this chapter, an overview of the topics related to the work in this thesis is presented. We briefly review the layers of the communication networks considered and their architecture. We follow that by an overview of cloud and fog computing architectures and the role of VMs in providing dynamicity to these architectures. Related work from the literature on optimising VMs placement in cloud-fog architectures is also reviewed. The concept of neutrality is introduced and the impact of repealing net neutrality on the ICT market is discussed. Furthermore, MILP optimisation is briefly introduced.

2.2 Communication Networks

The traditional Internet protocol (IP) network structure [17], [18] as used by ISPs, can logically be split into three main layers; the core network, the metro networks and the access network. Fig. 2-1 illustrates the three layers architecture. The core network represents the backbone infrastructure of any telecom network as it interconnects major regions/cities. IP over wavelength division multiplexing (WDM) technology is widely deployed in the core network due to its ability to provide high capacity, scalability and transfer speed [19]. Based on communication networks hierarchy, each core node is connected to a metro network, which covers a metropolitan area. Metro Ethernet is the dominant technology used in enterprise metro network. It provides direct connectivity between residential users in access networks and core network node. The access network represents the last mile of the telecom network by connecting telecom offices and end-users. Passive optical network (PON) is the main technology deployed broadly in the access network.

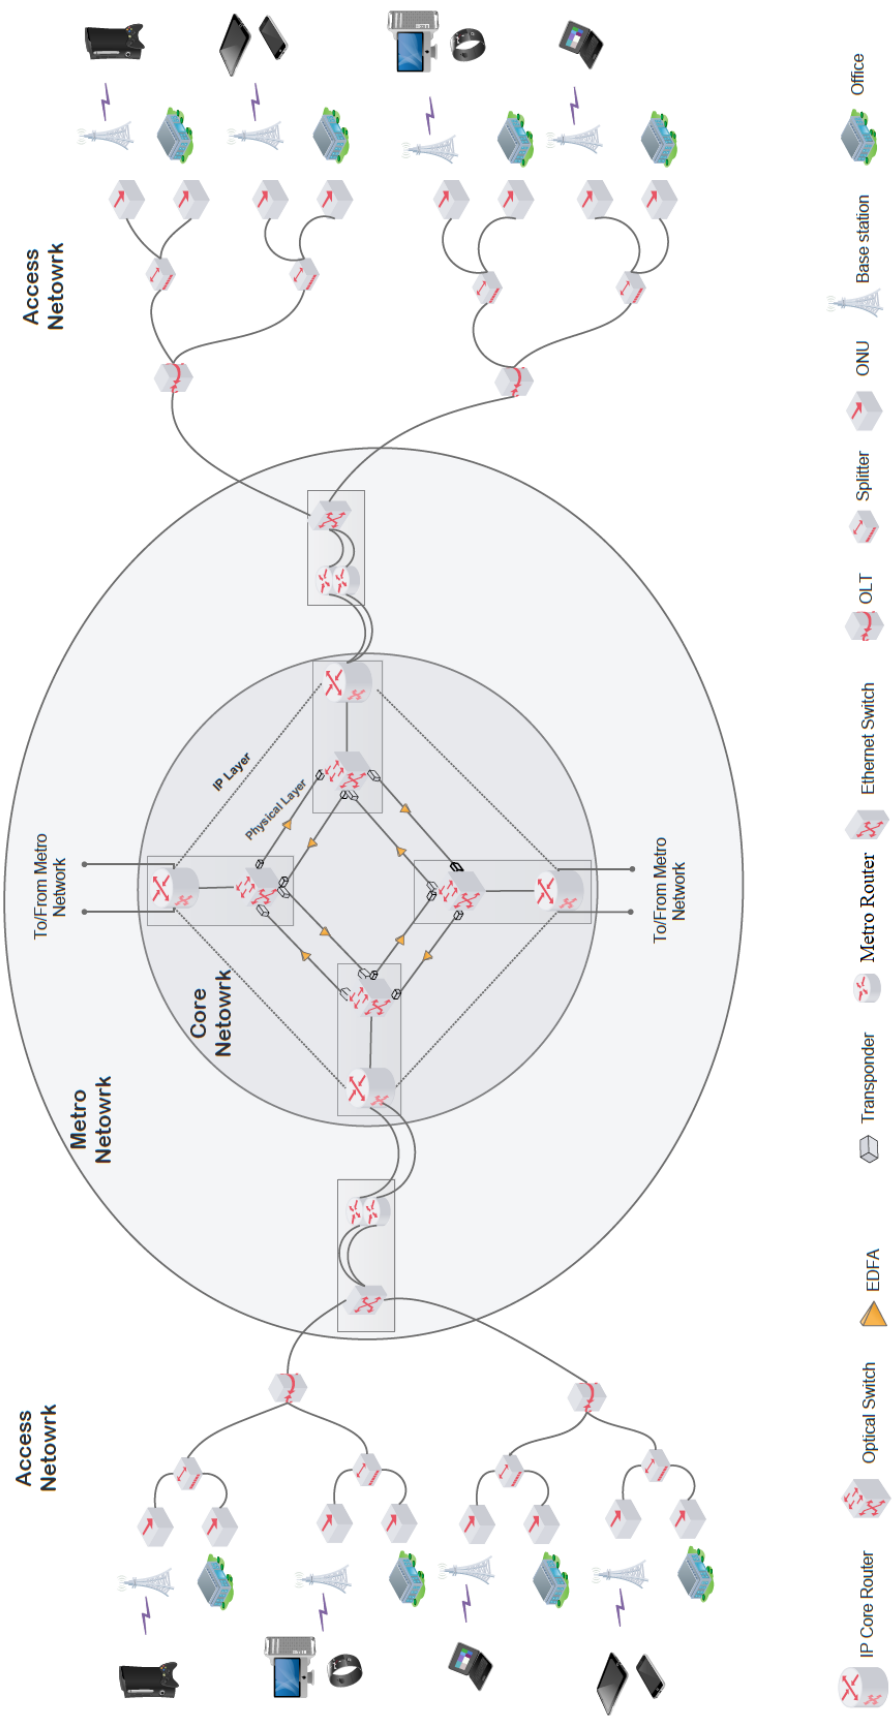

Figure 2-1: The telecom network architecture three layers: the core network, the metropolitan network and the access network.

2.2.1 IP over WDM

The IP over WDM [19] core network is composed of two layers: the IP layer and the optical layer. In the IP layer, the main function, which is done by the IP router, is to process the IP packet on its route from the access end to its destination nodes. Meanwhile, on the optical layer, an optical switch / cross connect (OXC) connects the IP and the optical layers. Moreover, there are the transponders, which provide optical-electronic-optical (OEO) processing and full wavelength conversion in every node. Erbium-doped fibre amplifiers (EDFAs) are used in WDM to amplify optical signals in the fibre links. Sometimes, regenerators may be used if the link's length exceeds a certain distance dictated by the data rate and modulation format used, for example 2500 km for PSK at 10Gb/s [19], [20]. Multiprotocol Lambda Switching protocol (MPLS) over WDM is the routing protocol used in IP over WDM optical core networks that route the data packets from source to destination nodes through a connection-oriented service [21]. Using MPLS, the IP over WDM routed traffic is implemented by either bypass or non-bypass light paths. In the case of bypass, the intermediate nodes IP routers do not process the packets. The packets takes a cut through path to the destination node. Meanwhile, on a non-bypass path, the packets are processed by every intermediate node in the journey from the source to the destination. Implementing a bypass lightpath can significantly reduce the power consumption of the network as it eliminates the need for intermediate IP router ports (the major contributor to IP over WDM networks power consumption) between source and destination nodes. However, a non-bypass light path allows for scanning, inspecting, and monitoring of packets to check for any security threats and allows traffic engineering to be implemented at intermediate nodes [19], [22] and [23].

2.2.2 Metro Ethernet Network

A metro network functions as a gateway for the access networks into the core network. Metro Ethernet is the dominant technology used in enterprise metro networks. The basic components of metro Ethernet are Ethernet switch and edge routers as shown in Fig. 2-1. The Ethernet switch interconnects several access networks. It also connects the access networks to edge routers. The best practice in ISP metro networks is to use two edge routers to provide reliability and redundancy in the network [17] [18].

2.2.3 PON Network

The PON [24] has proven its performance in the access network in terms of high bandwidth and reliability compared to Ethernet access networks. PON technology is used to provide fibre to the end-users. A unique feature of a PON is that it implements a point-to-multipoint architecture, where passive (unpowered) optical fibre splitters are used to allow multiple end-points to be served by a single link [25]. In PON architectures, there are two main active components deployed; the optical network unit (ONU) and the optical line terminal (OLT). The ONU is the end-user interface to the PON network and the OLT serves as a central office (CO) node to connect multiple ONUs. At present, gigabit PON (GPON) is the architecture and technology widely deployed among service providers [26]. GPON is capable of delivering data rates up to 2.5 Gbps over a single link. As the demand for high Internet speed continues to rise, a faster PON technology is spawned from the existing technology. 10G-PON (also known as XGPON) [27] provides extra link capability to telecom providers, which can deliver up to 10 Gbps per single fibre link.

2.3 Cloud Computing

Cloud computing provides ubiquitous on-demand access to an Internet-based pool of computing, storage, and communication resources. These resources can be provided to a large set of geographically distributed users. Typically, the cloud interface resides within a single window in the users' Internet browser. According to American National Institute of Standards and Technology (NIST) [28], the cloud computing model is essentially composed of five features; on-demand self-service, online availability through any platform (e.g. smartphones, laptops), a pool of computing resources available to end-users through a multitenancy architecture, and elastic and measurable service.

Cloud computing service models can be classified into three categories [28]; software as a service (SaaS), platform as a service (PaaS) and infrastructure as a service (IaaS). SaaS is the process of delivering an application to the end-users through cloud datacentre, which can be accessed through users’ web interfaces (e.g. websites, Health care systems, Geographic Information Systems (GIS), Microsoft

office). PaaS offers end-users the environment needed to create an application on cloud infrastructure using a set of online tools (e.g. programming languages) to provide facilities for designing, developing, testing and deployment of applications. Examples of this model include Google App Engine [29] and Microsoft Azure [30]. Further abstracted computing resources are available to end-users through the IaaS model. In addition to the deployed application, the end-user has control over the operating system, middleware (for enabling communication between two applications [31]), runtime and data, in addition, to directly using networking, storage and compute resources, which are usually made available on a subscription basis. Amazon Elastic Compute Cloud (Amazon EC2) [32] is a widely known solution based on the IaaS service model. Fig. 2-2 illustrates the architecture of the three service models compared with a traditional on-premises model, where the user has full control over datacentre. The virtualisation concept is employed in the three cloud computing models to satisfy the rapid growth/shrinkage in the usage of datacentre physical resources by cloud users. Virtualisation [4] simulates hardware functionality by creating a virtual version of cloud hardware resources such as CPU, RAM and hard disks to increase flexibility and scalability. This allows Cloud organisations to run multiple virtual entities on a single server. In Section 2.5, the virtualisation concept is reviewed intensively.

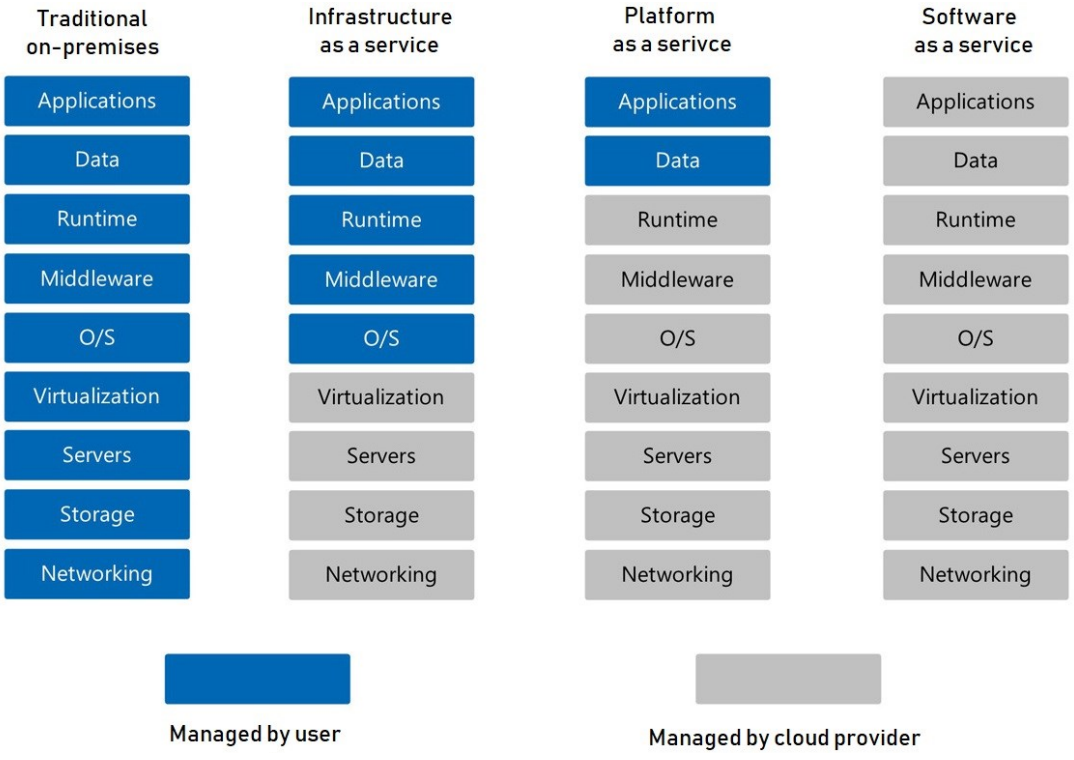

Figure 2-2: Cloud computing service models

Cloud computing can be mainly deployed as; public cloud, community cloud, private cloud, or hybrid cloud. In the public cloud, the cloud sources, which reside in the cloud provider premises, are rented over the Internet for public use. In the community cloud, the cloud infrastructure is provisioned for the use of organisations or end-users that have shared resources and goals. In the private cloud, the cloud infrastructure is operated exclusively for a single organisation, whether managed by the organisation itself or by a third party. The combination of multiple clouds deployment models (e.g. public, community, private) is called a hybrid cloud.

Cloud datacentres are typically consist of wide range of servers arranged in multiple racks and a Local Area Network (LAN) made of two switch layers used to connect racks to each other (in order to enable inter rack communication), and an aggregation router to connect the datacentre to the outside world, i.e. users and other datacentres (inter-datacentres communication). The architecture of cloud datacentre is illustrated in Fig. 2-3.

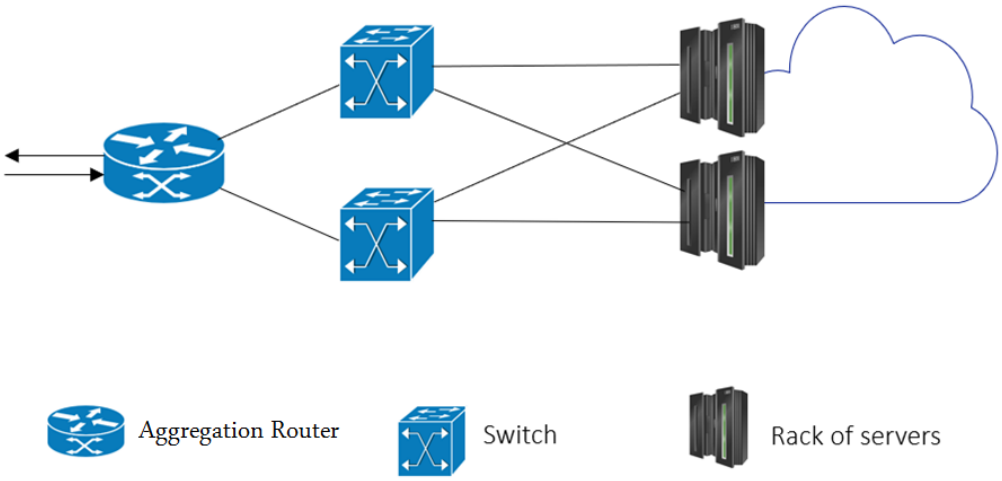

Figure 2-3: Typical cloud datacentre architecture.

2.4 Migrating to Fog

The concept of fog computing was introduced by Cisco in 2014 [33] to bring cloud services closer to the users. According to the OpenFog Consortium [34], a group of industry and academic organisations advocating the development and standardisation of fog computing in various facets, fog computing is an architecture that complements the clouds by extending processing, networking and storage resources to the edge of

the network. The advent of fog computing together with cloud computing has introduced a promising paradigm shift. As illustrated in Fig. 2-4, the fog computing general framework consists of three layers organised in a hierarchical architecture [35], [36]. In this architecture, the fog computing layer resides in the middle between end-users and cloud computing layers.

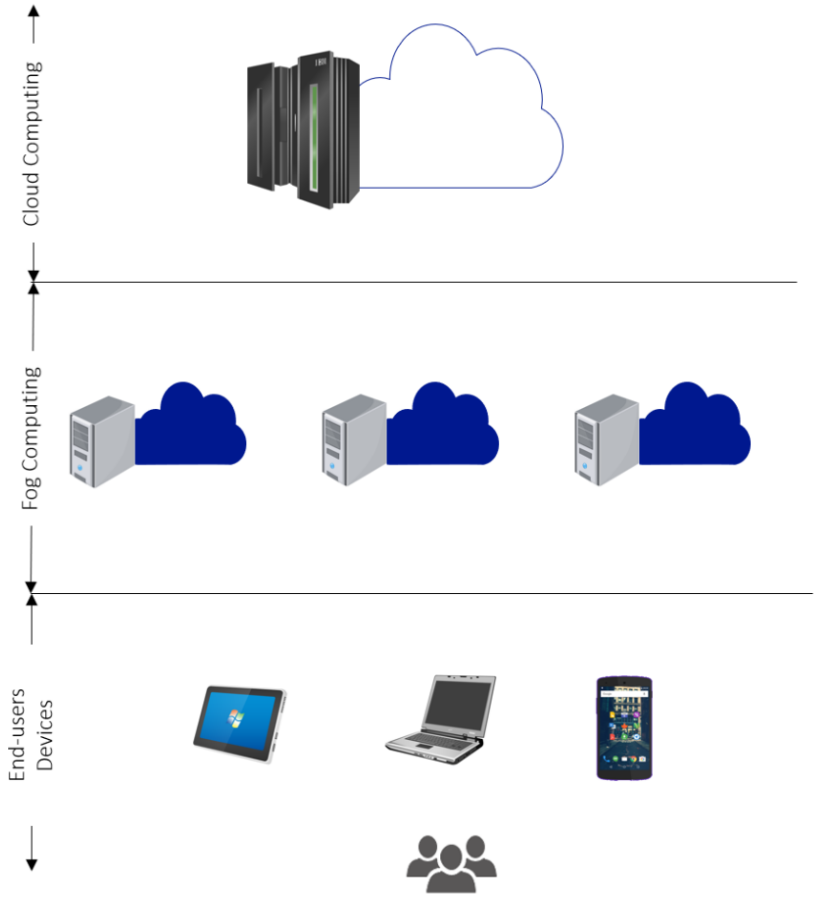

Figure 2-4: Fog computing hierarchical architecture.

The research efforts studying fog computing have mainly focused on illustrating its potential advantages over cloud computing. Offloading tasks to fog nodes is proposed [2], [3], [42]–[49], [33], [35]–[41] to provide:

- Realtime analytics: moving processing capabilities closer to users' premises to provide a real time data analytics.
- Low operating expense: reduce core network bandwidth by processing data locally in fog nodes instead of processing in the cloud.

- Low latency: processing users requests closer to their devices can significantly reduce the latency compared to when the traffic travels longer distances between the clouds and users.
- Privacy: processing sensitive data locally instead of processing at distant clouds which may be in a different organisation or country.
- Improve Qos and QoE for different computing services. QoS is the total compatibility of telecommunications service characteristics to the user needs, whereas, QoE focuses on the entire users' service experience of an application or service.

In term of QoS, the authors in [42] studied IoT service placement in fog node resources, taking into consideration its QoS requirements. They found that such a fog optimisation model maintains QoS and can achieve 35% less execution cost compared to cloud computing placement. Also, a mathematical model of fog computing was proposed in [47] to investigate the possibility of reducing the latency of IoT applications. The results showed that the overall latency associated with processing applications in fog computing is 50% lower than the latency associated with processing in cloud datacentres. In [37], the authors considered improving websites performance by connecting users to the Internet via fog servers. They found that the performance of fog-based websites is enhanced beyond what has been achieved by cloud servers.

The energy consumption of the fog computing paradigm has been given limited attention in the literature. In [35], the authors built a theoretical model of fog computing architecture and compared it with the conventional cloud computing model. In addition to the low latency, they found that offloading applications to the fogs can significantly reduce power consumption by 41%. However, their investigation did not consider a detailed model of the telecom network architecture. The work in [46] found that the number of hops between the user and the content has little impact on the total energy consumption compared to the type of application running on servers and factors such as the number of downloads and the number of updates. The authors in [50] studied the interplay and cooperation between the fog and the cloud to achieve a trade-off between power consumption and delay in a cloud-fog computing system by developing a mathematical model to formulate the workload allocation problem. They found that by allocating fog computing for processing a workload, the total delay reduces while the power consumption rises.

2.5 Machine Virtualisation

Cloud and fog scalability is highly dependent on the efficient provisioning of the datacentre physical resources. Virtualisation was introduced to enable efficient use of resources in the cloud-fog architecture. The first implementation of the idea of virtualisation was done by IBM in the mid-1960s, where numerous users were able to create, execute, modify and terminate their tasks on a shared system [51]. In the 1970s and 1980s, virtualisation was almost inexistent as the computing hardware purchasing price became cheaper. Then, in the 1990s, due to the broad variability in computing capabilities, the virtualisation concept was resuscitated [52]. At present, virtualisation techniques are swelling due to the cloud computing revolution.

The principal architectural characteristic of virtualisation is the solid resource allocation management approach where the datacentre physical resources are abstracted into several logical entities called VMs [4], [53]. Each VM is allocated its resources of CPU, memory, network bandwidth and storage to run a logically isolated application from other VMs. Using VMs, multiple applications can be consolidated into a fewer number of servers by allowing multiple heterogeneous virtual entities, each serving a different client, to coexist on a shared physical resource (servers) owned and operated by an infrastructure/service provider. Each physical server can host up to hundreds of VMs and each VM-hosted application allows multiple tenants or users to share a single application as it runs on a dedicated environment [54]. These virtual entities are created and tore down on demand to cater for the cloud clients' needs allowing for scalable growth and efficient use of resources through consolidating virtual entities in fewer physical resources [55], [16]. Fig. 2-5 illustrates the VM consolidation technique.

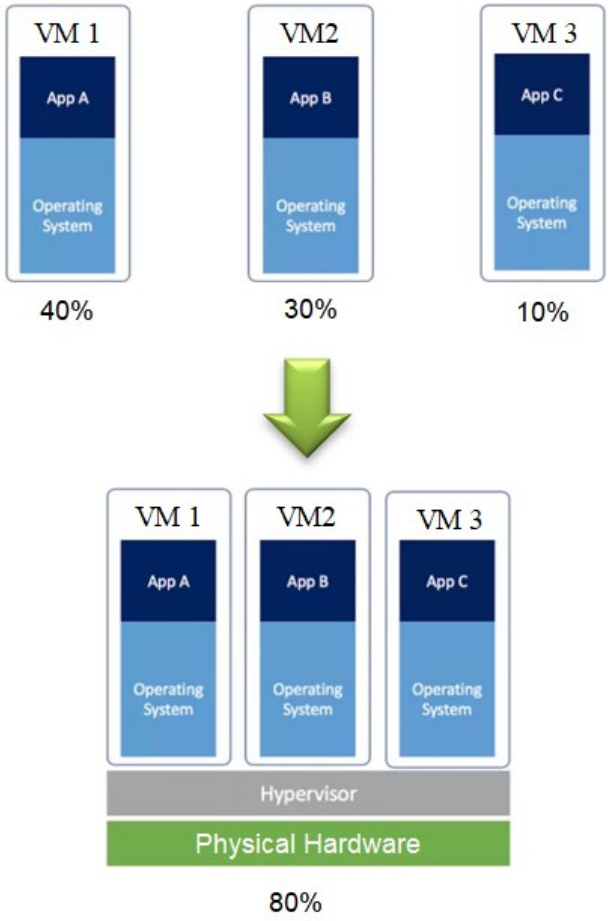

Figure 2-5: VMs consolidation.

Consider an example, where a CP wants to run an application on a cloud architecture. At time $\boldsymbol{t = 1}$, the application is unpopular and a single VM is enough to serve all users. At time $\boldsymbol{t = 2}$, the application becomes popular, so, a single machine is no longer enough to serve all users. Thus, an estimation can be made to find how many VM replicas are required to run the application by taking into consideration the server resources and the number of concurrent users accessing the application at that time. An elastic cloud and fog architecture should instantly react to the increased load and serve all application users responsively.

Further dynamism in resource management can be achieved by placing or relocating VMs within or across distributed datacentres through either replication or migration. VMs placement is a central operation in cloud and fog computing infrastructure where the most suitable server is found to host the VM based on workload balancing, datacentre maintenance, failover recovery or energy efficiency. In Section 2.5.2, the VMs placement problem is discussed.

In a virtualisation environment, VMs are orchestrated by the VM monitor (VMM) or a hypervisor. The hypervisor is the virtualisation software in a system that aims to create and run one or multiple VMs (guest machine) over physical hardware (host machine), which is typically a server. The hypervisor has the authorisation to control all the VM's operating system and its hosted application as well as server resources.

2.5.1 VMs Categories

Based on computation and network bandwidth requirement, cloud applications can be classified into three categories [56], [57]; 1) CPU-intensive applications; only require computation resources and produce a low data rate over the communication network. e.g. high-performance computing applications, 2) data-intensive applications; require fewer computation resources and produce high traffic between the communicating nodes. e.g. video applications. 3) balanced applications; both computation and communication network resources are required. e.g. GISs applications.

From a CPU perspective, studies in the literature have shown that the workload of VM versus the number of users served by VM mostly follows one of two profiles; constant or linear profiles as seen in Fig. 2-6. In [58], the authors presented a CPU performance benchmark study for web application VMs serving a varying number of users with constant CPU workload as illustrated in Fig. 2-6 (a). Also, various benchmarking studies in the literature have demonstrated linear workload profiles for database applications [59], web-based video conferencing systems [60] and multiplayer games [61] with different slope coefficients. To maintain the service level agreement (SLA), each VM needs a minimum workload to run an application regardless of the number of users served by the VM, resulting in the workload profile shown in Fig. 2-6 (b). The minimum workload required to serve a user in a VM varies from as low as 1% to 60% [59] - [61].

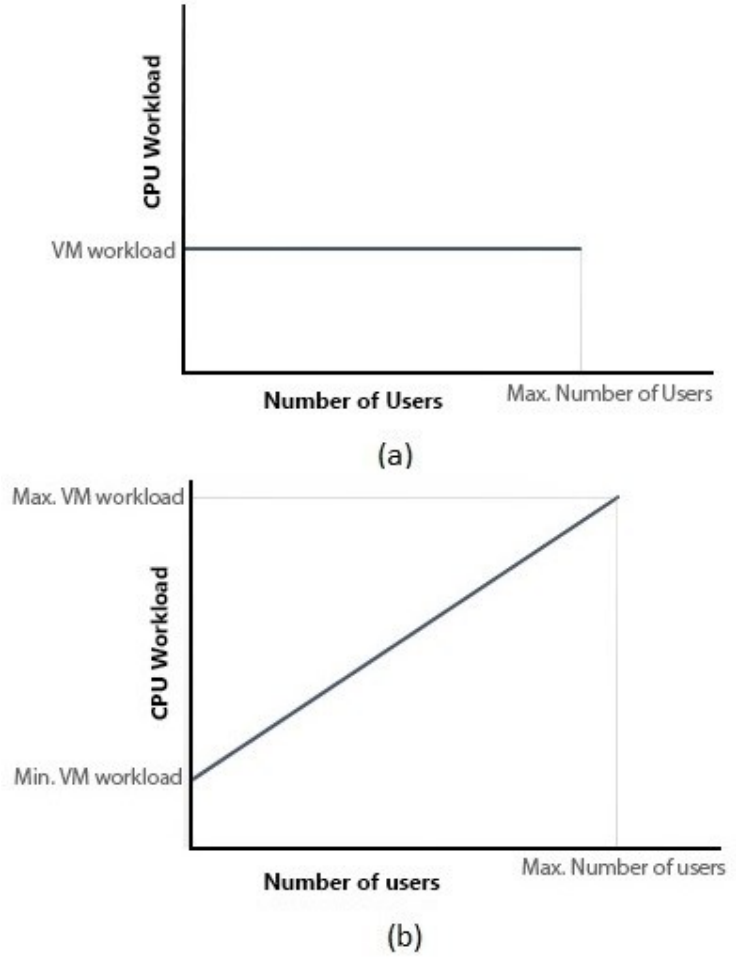

Figure 2-6: Relationship between VM workload and the number of users; (a) constant (b) linear relationship between VM workload and number of users.

From network bandwidth perspective, cloud service applications based on user network bandwidth requirements are classified into three main categories: (i) Mbps for light web browsing [62] (emails, Google docs [63] and websites with lower definition video content [64]), (ii) 10 Mbps for applications processing high-definition video quality [65] and online multiplayer games [66], and (iii) 25 Mbps for applications processing ultra-high video quality [67].

This inter-VM traffic is a major contributor to the east-west traffic (server to server traffic) which is expected to be responsible for 85% of the global cloud traffic by 2021 as opposed to north-south (traffic between server and client), which accounts for the remaining traffic [62]. In a cloud environment, different VMs may need to communicate to complete their processing jobs as seen in Fig. 2-7(a) [53]. As well, in case of replication, replicas of a VM need to communicate to ensure synchronisation (see Fig. 2-7(b)) [68]. Co-located VMs in the same datacentre, can communicate with

each other through LAN, whereas, if they are located in geo-distributed datacentres, the communication traffic will pass through the core network backbone [69].

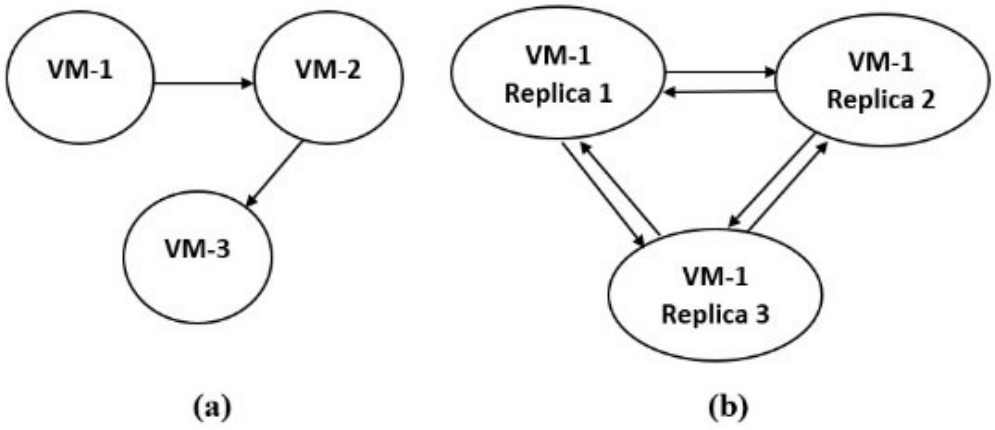

Figure 2-7: Illustrative example of inter-VM traffic, (a) VM-VM cooperation traffic (b) VM replicas synchronisation.

Inter-VM traffic has been intensively investigated in the literature. The authors in [70] studied the traffic of communicating VMs hosted by a group of servers. The trace analysis shows that inter-VM traffic varies significantly between different VMs pairs. In [71], the authors developed a system that measures the throughputs between data-intensive VM pairs inside Amazon EC2 and Rackspace clouds. They found that the throughput varies from as low as 100 Mbps to almost 4.5 Gbps. Also, they developed an integer linear programming (ILP) model and an algorithm to formulate the problem of intra-datacentre network-aware VM placement.

Although virtualisation comes with plentiful advantages, resource sharing when running multiple virtual entities over a single physical machine may reduce QoS and cause unstable performance. To reduce the impact of resource sharing and achieve scalable cloud architecture, some techniques have been discussed in the literature and are employed by cloud providers. For example, in order to ensure QoS, cloud providers either implement traffic shaping techniques [34], to limit the traffic bandwidth which is allocated to each VM or enforce a minimum configuration of each resource [35] (e.g. bandwidth capacity or CPU usage) in order to preserve SLAs established with the end-users by providing a satisfying performance.

2.5.2 VMs Placement

The VMs placement is a vital process where the most suitable server is selected to host a VM. Selecting an appropriate VM container is critical to improving QoS, energy efficiency and optimal usage of physical resources. However, VM placement in cloud computing was demonstrated to be a very complex task. The simultaneous requests arrival patterns of the VM are usually unpredictable. Also, the datacentre size is typically enormous and for a certain load, optimising the VMs placement is considered to be a nondeterministic polynomial (NP)-hard problem. For example, if $\boldsymbol{v}$ is the number of VMs and $\boldsymbol{s}$ is the number of servers, then the number of possible VM placement in different servers is $\boldsymbol{vs}$. In case of replicating VMs into multiple datacentres ($\boldsymbol{N}$), exhaustive search of datacentre distributed locations require the evaluation of $\left(\sum_{\boldsymbol{i=1}}^{\boldsymbol{N}} \frac{\boldsymbol{N!}}{\boldsymbol{(N-i)!}}\right)$ placement combinations in order to find the optimal number and locations of VM copies needed.

The VM placement schemes are classified as static and dynamic schemes. In the static scheme, the placement of the VM is fixed and is not adjusted for a long period of time whereas in the dynamic scheme, the VM placement can change periodically based on change in the status i.e. server load, change in the network or cloud components. Dynamic schemes can be categorised into reactive and proactive VM placement. In the reactive scheme, the placement of the VM is only changed if the system reaches undesired conditions, so, the placement of VM will only be changed for example in case of emergency or when applying maintenance routines, restoring QoS or to reduce power consumption. On the other hand, in the proactive scheme, VM placement is changed before the system reaches undesired condition [7], [15], [53], [72]–[74].

In the literature, several papers discussed the VMs placement considering various factors. To reduce the server load, improve the QoS and meet the SLAs, the VMs can be migrated or replicated to another server/servers within the same datacentre [6], [7] or in geographically distributed datacentres [76]. Virtualised cloud architectures are also able to provide efficient disaster resilience in case of physical machine failure by migrating VMs into different host machines [77] or by replicating VMs content to distributed datacentres. From an energy efficiency perspective, under-utilised servers can significantly increase the energy consumption, and consequently increase the carbon emissions and operating costs of cloud datacentres. VMs consolidation by bin

packing them into a fewer number of servers can significantly improve the energy efficiency. In general, VM placement algorithms should consider the following:

- Cloud and fog computing resources (servers, datacentre networks) [41], [52], [84]–[86], [55], [72], [78]–[83] .
- VM-users traffic and VM-VM traffic [11], [14], [92]–[97], [55], [57], [71], [87]–[91].
- Communication network layers (core, metro and access network) [6], [9], [98]–[101], [11], [14], [39], [46], [49], [55], [94], [97].
- QoS and SLAs [42], [79], [102]–[106].
- Cloud and fog computing resources (servers, datacentre networks) [49], [52], [55], [68] .
- VM-users traffic and VM-VM traffic [11], [14], [55], [57], [67].
- Communication network layers (core, metro and access network) [6], [9], [11], [14], [41], [44], [47], [55].
- QoS and SLAs [37], [78].

2.6 Energy Efficient VMs Placement

The power consumption of a server involves two parts; the idle and active power consumptions. The idle power consumption of a server is consumed when it is switched on but not processing any workload whereas the active power consumption is the additional power consumed when the server is running a certain workload. The total power consumption of the server is equal to the sum of idle and active power consumptions. In this thesis, as the difference between idle power consumption and full load power consumption is very small in a server [107], we consider an on-off power profile for servers, i.e. if a server is activated, it will operate at maximum power consumption.

The power consumption of VM is based on the hosting server. The authors in [108] found that the CPU utilisation and power consumption of a server are highly correlated. Another work in [80] studied the relationship between the power consumption of a server and the CPU utilisation and found that the power

consumption and the CPU utilisation are related linearly. Thus, the work introduced in this thesis follows the same approach and takes in consideration the CPU utilisation only in modelling the power consumption of VMs placement.

The problem of energy-efficient VMs placement has been investigated thoroughly in the literature considering different parameters and pursuing various objectives by using several optimisation methods. In [6], the authors compared the energy efficiency of three schemes that place VMs over distributed datacentres in the core network; 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 copies to serve a smaller number of users over multiple clouds. The slicing scheme was found to be the most energy-efficient scheme with savings up to 25% of the total power consumption. However, the size of the slice in the work proposed in [6] was not proportional to the number of users served. The authors in [100] proposed a datacentre network topology-aware VM migration algorithm aiming at migrating groups of VMs to allow the switching off of physical servers and network resources. In [109], the authors designed algorithms to study the impact of datacentre architectures, constraints of servers and application dependencies on energy-aware VM placement. From an inter-VM communication perspective, the authors in [109]–[113] studied the energy-efficient placement of VMs inside a datacentre taking into consideration intra-VM traffic.

The majority of studies of VMs placement in the fog have been limited to evaluating the reduction of overall network overhead [96], optimising the placement of physical resources in the edge network [99] and the scheduling of VMs to share the limited fog resources to minimise SLA violations [43].

Despite the diverse factors affecting the power consumption of cloud-fog architectures, the majority of VMs placement research in the literature is limited to considering few parameters. Table 2-1 summarises the parameters that have been considered in the literature when studying energy efficient VM placement decisions. The table classifies the factors into cloud-fog computing layers, cloud/fog datacentres networks, the three network architecture layers, users traffic, inter-VM traffic, CPU requirement and PUE value. It also identifies the type of model/algorithm used to solve the problem.

The problem of providing energy-efficient VMs placement over end-to-end cloud-fog architecture considering the above-mentioned architecture has not received any attention. Thus, the objective of this thesis is to develop a novel framework that covers all these parameters in optimising the energy efficiency of VMs placement.

Table 2-1: Classification of Energy Efficient VMs placement research in the literature

Paper	Computing Layer		Traffic		Computing		Communication Networks			PUE	Based on
	Cloud Computing	Fog Computing	VM-Users	VM-VM	CPU	Datacentre Networks	Core Network	Metro Network	Access Network		
Lawey et al. [6]	■	×	■	×	■	■	■	×	×	■	MILP and Heuristic
Silva and Fonseca [100]	■	×	■	×	■	■	■	×	×	■	Heuristic
Huang et al. [109]	■	×	■	■	■	■	×	×	×	×	Heuristic
Yu et al. [114]	■	×	×	■	■	■	×	×	×	×	Heuristic
Zheng et al. [82]	■	×	×	×	■	×	×	×	×	×	Heuristic
Tsai et al. [115]	■	×	×	×	■	×	×	×	×	×	Heuristic
Zhang et al. [89]	■	×	×	■	×	×	×	×	×	×	Heuristic
Tziritas et al. [113]	■	×	×	■	■	×	×	×	×	×	Heuristic
Wang et al. [81]	■	×	×	×	■	×	×	×	×	×	ILP and Heuristic
Buyya et al. [88]	■	×	×	×	■	×	×	×	×	×	Heuristic and Simulation
Gao et al. [103]	■	×	×	×	■	×	×	×	×	×	LP and Heuristic

Dong and Herbert [105]	■	×	×	×	■	×	×	×	×	×	Heuristic
Fang et al. [110]	■	×	×	■	×	■	×	×	×	×	Heuristic
Li et al. [83]	■	×	×	×	■	×	×	×	×	×	Heuristic and Simulation
Arroba et al. [116]	■	×	×	×	■	×	×	×	×	×	Heuristic and Optimisation
Katsaros et al. [102]	■	×	×	×	■	×	×	×	×	■	Heuristic and Simulation
Farahnakian et al. [13]	■	×	×	×	■	×	×	×	×	×	Heuristic and Simulation
Beloglazov and Buyya [117]	■	×	×	×	■	×	×	×	×	×	Heuristic
Horri et al. [79]	■	×	×	×	■	×	×	×	×	×	Heuristic and Simulation
Wang et al. [90]	■	×	×	×	■	×	×	×	×	×	Optimisation and Simulation
Dabbagh et al. [84]	■	×	×	×	■	×	×	×	×	×	Optimisation and Heuristic
Shen et al. [85]	■	×	×	×	■	×	×	×	×	×	Heuristic
Farahnakian et al. [72]	■	×	×	×	■	×	×	×	×	×	Heuristic and Simulation
Feller et al. [118]	■	×	×	×	■	×	×	×	×	×	Heuristic and Simulation
Van et al. [119]	■	×	×	×	■	×	×	×	×	×	Simulation
Mishra et al. [44]	×	■	×	×	■	×	×	×	×	×	Heuristic and Simulation

Mishra et al. [78]	■	■	×	×	×	×	×	×	×	×	LP and Heuristic
This Thesis	■	■	■	■	■	■	■	■	■	■	MILP and Heuristic

2.7 Net neutrality in ICT

Network (net) neutrality regulations prohibit ISPs from applying different treatment to IP packets based on their content e.g. prioritising, blocking or throttling certain Internet content or allowing quality differentiation. Net neutrality, which was scrapped by the US Federal Communications Commission (FCC) in December 2017, has been the subject of remarkable debate in recent years ISPs and CPs with each side trying to exploit their assets and expand their profit and influence. The debate is fuelled by the rapidly escalating demand for CPs services as a result of the interconnection between Internet and broadcasting markets. Cisco forecasts [120] that by 2021, annual global Internet traffic will hit 2.2 Zettabytes per month and CPs datacentres will be the source of 71% of this traffic. Online video services are the primary cause of this accelerated growth in Internet traffic. Video streaming is poised to consume 78% of the total CPs bandwidth with 75% of Internet video traffic originating from higher video services quality (HD and UHD).

Proponents of preferential treatment of Internet traffic complain that the increasing demand for data-intensive content creates a significant burden on the communication network. They argue that removing net neutrality will give ISPs further control of their infrastructure, which is crucial in order to improve QoS and reduce security threats. Another argument is that a significant fraction of the profit of this tremendously growing market is seized by CPs whereas ISPs act as a transit or transport medium into CPs customers. In the US, the quarterly profit margin of AT&T (an ISP) has been almost stable over the last six years whereas Netflix (a CP) profit margin has risen up in rapid pace from 0.7% to 9.8% within the same period [121], [122]. In contrast, advocates warn that removing net neutrality will slow down the innovation in the Internet and its content and will limit the content competition by disadvantaging small businesses, and subsequently, diminish online services.

Deploying traffic discrimination in video delivery services has many challenges, e.g. detecting video packets and enforcing a policy on a certain video quality. Traffic discrimination in IP communication networks has been surveyed intensively in the literature. Several traffic management practices have been surveyed in [123]. The

authors highlighted that traffic discrimination taxonomy has four features: (i) characteristics or condition of the traffic (e.g. based on content, protocol or source/destination); (ii) traffic classification (e.g. based on flow rate, header information or routing); (iii) mechanism of discrimination (e.g. modify, delay, drop or block); and (iv) perceived discrimination by end-users. Video traffic can be analysed using two mechanism; deep packet inspection (DPI) [124], [125] or traffic profiling [126], [127]. DPI examines the data packets that is sent over the network and traffic profiling detects abnormal network traffic by comparing new traffic against previous traffic profile. For example, an alarm can be triggered if the data rate transmitted over the network (measured in bps) spikes above the desired data rate, which could indicate an increase in data rate. QoS for video services delivery can be applied either by reserving network bandwidth for video packets (e.g. using IntServ) or labelling video content as high priority e.g. by applying Differentiated Services (e.g. using DiffServ) [128].

The Internet ecosystem is complex with many stakeholders. As illustrated in Fig. 2.8, the main stakeholders in the Internet ecosystem are; ISPs, CPs, content delivery networks (CDNs) and end-users. Users pay ISPs a subscription fee to get Internet access and subscribe to CPs (if required) to access their content. CPs subscribe to a CDN to access storage and processing capacity and to deliver their content to customers. CDNs are responsible for sending CPs content at large scale over ISPs network infrastructure, e.g. the CP Netflix collaborates with the CDN Amazon Web Services (AWS) to reach their customers [129]. ISPs play as the key intermediary in the delivery process as they provide the required connectivity between users and content. Most ISPs such as AT&T [130] and Comcast [131] are now providing CDN services in additional to networking services. To simplify our analysis, we consider a direct relationship between ISP and CPs.

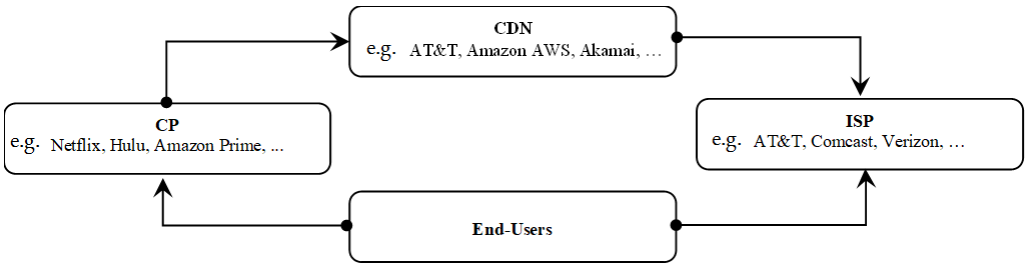

Figure 2-8: Main stakeholders in Internet ecosystem. Arrows represent customer-provider relationship.

Due to net neutrality regulations, current pricing policy of ISP networking services applies a fixed charge which is not linked with bit rate usage. For example, in the US, AT&T uses a fixed pricing model by charging CPs $3,282 per 10 Gbps per month [132] regardless of the content type transferred to users (either UHD video content or a simple text message).

Many papers in the literature discussed and analysed various aspects of net neutrality. From a legalisation and regulation perspective, net neutrality in the Internet ecosystem has been surveyed by the authors in [133] and [134]. They emphasised that cloud computing has initiated the net neutrality battle between ISPs and CPs. In [135] the authors analysed the Internet video streaming contest, taking into account all of ISPs and CPs assets (e.g. content rights, network access, users, …etc). They stated that video distribution makes the dilemmas of net neutrality solid and perceptible. Their analysis demonstrates that net neutrality correlates highly with video service delivery at different points including competition between CPs and ISPs, competition between stand-alone CP and CP owned by ISPs in providing video delivery services, growth of video traffic …etc.

A number of papers in the literature focuses on providing mathematical models to investigate the influence of the repeal of net neutrality on the communication networks. Paid service differentiation where CPs voluntarily pay a monopoly ISP for prioritising their traffic under shared network infrastructure was investigated by the authors in [136]. The differentiation occurs where ISPs offer service classes for CPs to choose from where traffic of a higher-priority class will be processed before those of a lower-priority. They studied the optimal pricing based on either maximising the CPs' choices of service classes or minimising system delays. Consequently, they highlighted that ISPs optimal pricing strategy can result in an efficient differentiation among CPs maximising social welfare. Also, they found that applying paid

prioritisation can lead to money flows (profit) from CPs to ISPs. The authors in [128] modelled the competition of video services delivery market between an ISP's own integrated CP and stand-alone CP. They studied the impact of applying different QoS (marking video traffic as high priority) pricing strategies either by selling QoS to CPs, selling QoS to users, or choosing to not provide QoS at all. They investigated the impact of QoS pricing on the video service prices and CPs profit. The analysis showed that ISPs can sell QoS to CPs at a higher price than when QoS is sold to users, and the CPs are able to make more profit when QoS is directly sold to users than the case when QoS is sold to CPs. Also, they found that ISP is more likely to use QoS exclusively for its own video services when it provides a similar content of CPs. The work in [137] considered the impact of maximising profit of CDN providers considering users who access CPs content from either cloud or fog server. In the case of competitive CPs, the CDN always places the content of the popular CP in fog servers, even when a less popular CP pays more, as the CDN tries to reduce core network transit cost.

2.8 Mixed Integer Linear Programming (MILP)

Mathematical representation is the most concise and accurate representation of a problem and can help understand and solve the problem in hand. Linear Programming (LP) is a mathematical optimisation technique where a maximum or a minimum of a linear function is to be found subject to a set of linear constraints and bounds. A LP model consists of the following four elements;

- The objective function, which is maximised or minimised to find the optimum value of the variables.
- Parameters, which are values that remain unchanged throughout the optimisation process.
- Variables, which are values that vary are to be identified within upper or lower constraints. An important subclass of variables are binary variables, where a variable can take a value of either 0 or 1 to model a specific decision.
- Constraints are a set of mathematical inequality equations, which define the solution feasibility region.

The variables of LP model are normally continuous, making the problem solvable in polynomial time. However, if they are constrained to integer values, the problem turns into an NP hard problem where the optimal solution cannot be found in polynomial time.

In MILP, the variables are a mix of integer and non-integer values. The use of MILP was shown to be efficient in optimisation problems in numerous applications [138] including; economic and transportation. The use MILP optimization has been used also used intensively in the literature to solve network design problem such as maximizing the profit achieved by infrastructure providers [139] and minimizing the power consumption of delivering clouds services [6].

To illustrate the MILP principle, the following standard form is presented; Objective function;

Maximise or Minimise: $\boldsymbol{f = x_{1,}y_1 + x_{2,}y_2 + \cdots + x_{n,}y_n}$

Constraints;

$$\boldsymbol{\alpha_{11,}y_1 + \alpha_{12,}y_2 + \cdots + \alpha_{1n,}y_n \leq a_1}$$

$$\boldsymbol{\alpha_{21,}y_1 + \alpha_{22,}y_2 + \cdots + \alpha_{2n,}y_n \leq a_2}$$

$$\vdots$$

$$\boldsymbol{\alpha_{n1,}y_1 + \alpha_{n2,}y_2 + \cdots + \alpha_{1n,}y_n \leq a_n}$$

and non-negativity constraints:

$$\boldsymbol{\forall\ y_1 \geq 0, y_2 \geq 0, \ldots\ y_n \geq 0}$$

To solve any MILP problem, the mathematical model is represented in computer systems in several ways. A Mathematical Programming Language (AMPL) [140] is one of the most popular ways to make a connection between the MILP model and its data file by providing a high-level language for describing optimisation programmes. The AMPL formulation is solved by MILP solvers, for example CPLEX [138], and lpsolve [141]. The solver most widely used, and the one used in this thesis, is the IBM ILOG CPLEX optimisation software. CPLEX provides a high-performance, robust, reliable, and flexible mathematical solver to find an optimal solution to the problem.

In MILP, the optimal solution uses an NP-hard discrete optimisation. MILP optimisation problems are typically non-covex. Non-convex problems may have several feasible regions, and within each region, various local optimal points can be found. Therefore, a systematic approach is used to find a solution. The most common approach used to find the optimum solution is the branch-and-bound (B&B) approach. As shown in Fig. 2-9, this approach divides the problem into a rooted tree of nodes (stages). Every node has its branches of variables. The complexity of the problem equals 2^n, where $\boldsymbol{n}$ is the number of variables in the problem. B&B solves the problem by eliminating the non-worthy nodes within its branches if proven that they will not provide a feasible optimal solution.

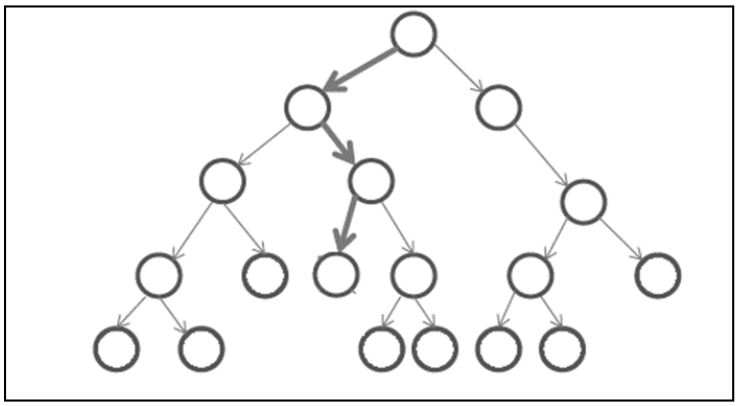

Figure 2-9: The branch and bound (B&B) approach.

The branch-and-cut (B&C) approach is also used to solve MILP problems. It works by removing the non-feasible branches, according to the information given by MILP. It provides a fast solution but with a less reliable result. A genetic approach is also used sometimes to find a solution by performing sequence iterations on randomly created solutions. The potential solutions are continuously optimised over generations by performing a fitness test where only one solution can survive [142]. However, although these approaches provide near optimal solutions, there is no guarantee that the solution is optimal.

2.9 Summary

In this article, we have provided a review of energy efficient cloud-fog architecture, especially for . It provided an overview of communication networks focusing on IP

over WDM networks, metro Ethernet networks and PON networks. These are the essential elements that represent the wired network of interest in this thesis. Attention was then given to cloud computing and its recent introduction as an integral element of communication networks. The recent migration of processing and storage services from the cloud to the fog was outlined. A key enabling technology for the efficient utilisation of cloud resources is virtualisation and the creation of virtual machines. These concepts were reviewed including VM categories, VM placement and energy efficient VM placement. Net neutrality was then introduced as it impacts what network operators and service providers can do and can impact the placement of VMs. Finally, mixed integer linear programming (MILP), the main optimisation tool used in this thesis was introduced and discussed.

Chapter 3: Energy Efficient Virtual Machines Placement over Distributed Clouds

3.1 Introduction

In this chapter, a framework for optimising the placement of VMs over distributed clouds to minimise the networking and computing resources power consumption is introduced. A MILP model is developed for this purpose. The MILP model is used to analyse the impact of different factors including VM popularity, the traffic between the VM and its users, the VM workload, the workload versus number of users profile, the impact of core network topology on VM placement and the power usage effectiveness (PUE).

The results show that high traffic of highly popular VMs and VMs with high user data rates make network power consumption more important and consequently determine the placement. The results also indicate the tendency to distribute VMs with linear workload profile compared to VMs with constant workload. Furthermore, the results show that at high PUE, processing power consumption becomes more dominant, which makes it play a bigger role in VMs placement.

Based on the MILP model, a heuristic is developed to mimic the MILP model in a real-time environment.

3.2 MILP model

In this section, a MILP model is developed to optimise the placement of VMs over distributed clouds in IP over WDM core networks so the total power consumption of providing the VMs is minimised. The architecture in Fig. 3-1 is considered. Unlike [6] where certain placement schemes (migration, replication and slicing) are used to place all types of VMs, in this work, the MILP model is developed to decide the

optimum placement scheme over the distributed cloud architecture based on the VM popularity, the VM workload and users data rate. The model aims to achieve a trade-off between network power saved by replicating VMs in multiple clouds and the power consumed by these replicas. The creation of a VM replica will result in power savings if the network power consumption exceeds the VM replication power consumption.

A typical cloud datacentre, as illustrated in Fig. 3.1, consists of servers arranged in multiple racks and a LAN network, made of routers and switches, to connect racks to each other (inter rack communication) and to users outside the datacentre.

Before introducing the MILP model, the sets, parameters and variables related to different network layers (core, metro and access) and clouds resources are defined. Note that, the metro and the access networks are not part of the optimisation problem, namely they do not host candidate VM sites. They are considered in the calculation of the power consumption in this model to compare its results to the cloud-fog model in Chapter 5.

The sets, parameters and variables representing cloud computing resources are defined in Table 3-1, Table 3-2 and Table 3-3:

Table 3-1: List of sets used in the MILP model

Set	Description
$\boldsymbol{N}$	Set of IP over WDM network nodes. Set of nodes represents nodes distributed across different geographic regions in the selected core network topology.
$\boldsymbol{VM}$	Set of VMs.

Table 3-2: List of cloud parameters used in the MILP model

Parameter	Description
$\boldsymbol{SW^{(CB)}}$	Cloud switch bit rate.
$\boldsymbol{SW^{(CP)}}$	Cloud switch power consumption.
$\boldsymbol{SW^{(R)}}$	Cloud switch redundancy.
$\boldsymbol{R^{(CB)}}$	Cloud router port bit rate.
$\boldsymbol{R^{(CP)}}$	Cloud router port power consumption.
$\boldsymbol{S^{(P)}}$	Power consumption of a server.
$\boldsymbol{S^{(maxW)}}$	Maximum workload of a server.
$\boldsymbol{c}$	Cloud power usage effectiveness.

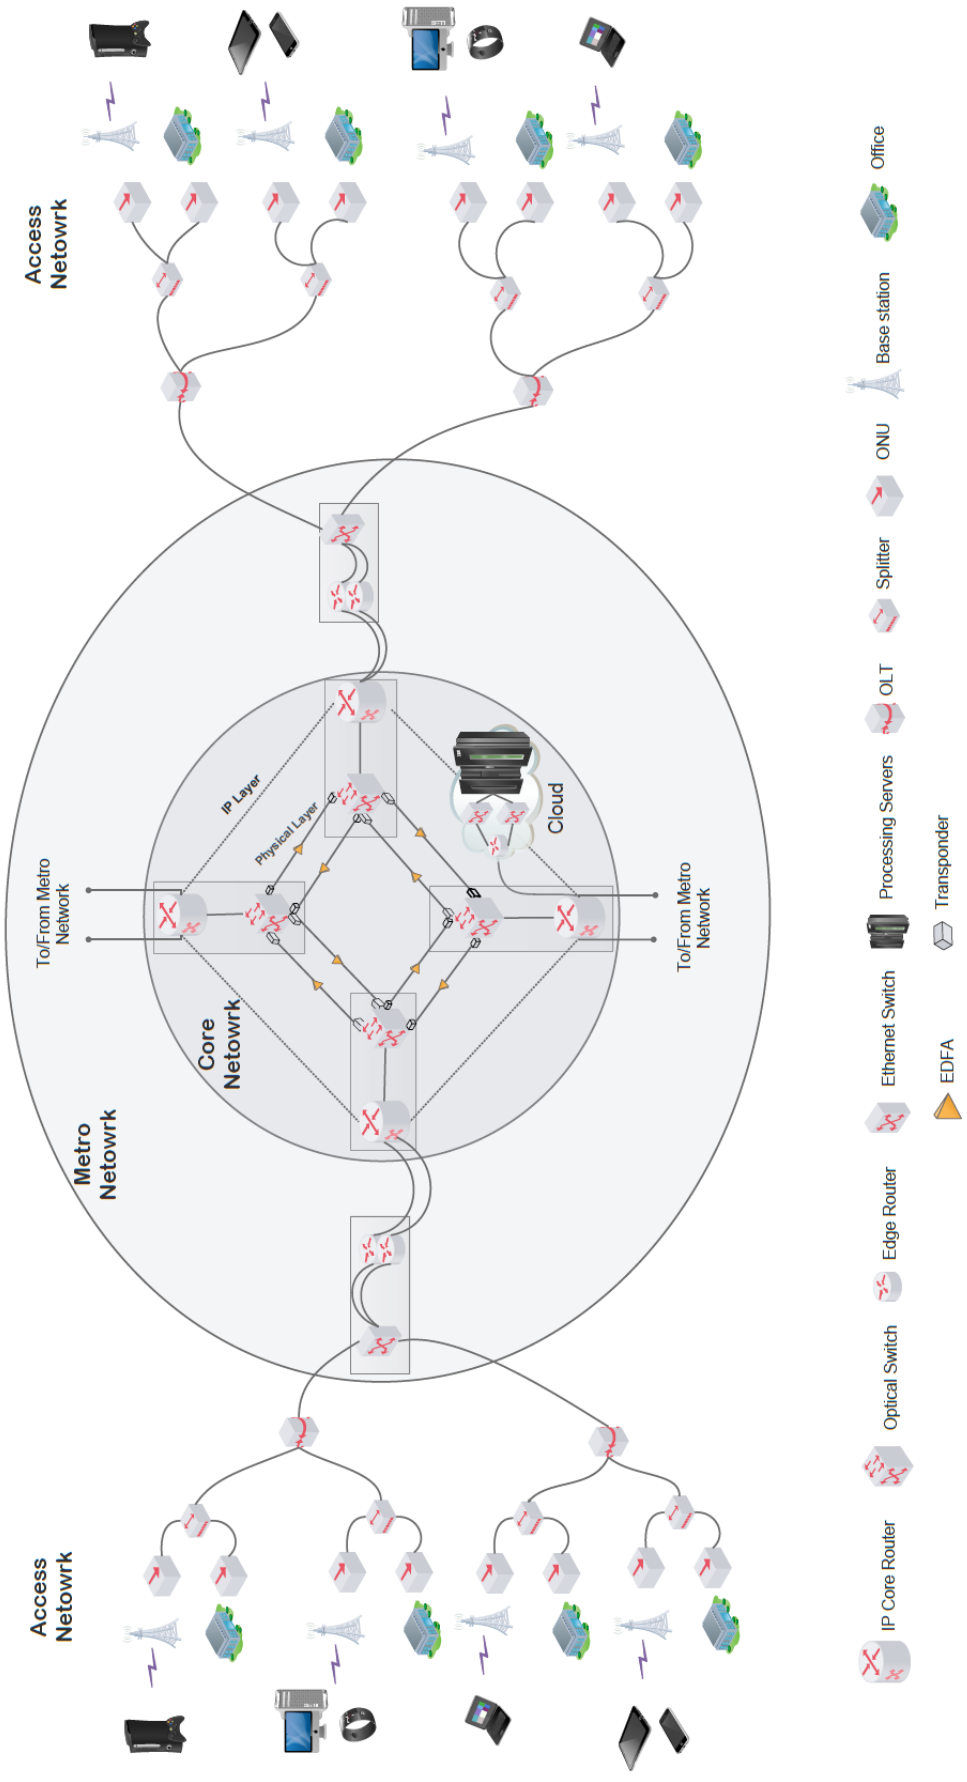

Figure 3-1: Network layers supported by clouds.

Table 3-3: List of cloud variables used in the MILP model

Variable	Description
C_s	$C_s = 1$ if a cloud is hosted in node s, otherwise $C_s = 0$.
$\delta_{v,s}^{(C)}$	$\delta_{v,s}^{(C)} = 1$ if the cloud hosted in node s hosts a copy of VM v, otherwise $\delta_{v,s}^{(C)} = 0$.
$R_s^{(C)}$	Number of router aggregation ports in the cloud hosted in node s.
$SW_s^{(C)}$	Number of switches in the cloud hosted in node s.
$S_s^{(C)}$	Number of processing servers in the cloud hosted in node s.

The VMs to be hosted in the cloud and the traffic resulting from them are defined by the parameters and variables in Table 3-4 and Table 3-5:

Table 3-4: List of VM parameters used in the MILP model

Parameter	Description
s and d	Indices of source and destination nodes of a traffic flow in the distributed cloud architecture.
V	Number of VMs.
S_v	Number of VM v users.
r_v	User download rate of VM v.
L	Large enough number.
x	Maximum number of users served by a single VM replica.
W_v	Maximum workload of VM v (workload can be specified in GHz or as a ratio of the CPU capacity).
M	Workload baseline of VM (the minimum CPU utilisation needed in the absence of load). To maintain the service level agreement (SLA), each VM needs a minimum workload to run an application regardless of the number of users served by the VM.
T_v	Traffic resulting from VM replica v serving the maximum number of users. $$T_v = x\, r_v$$
$W_v^{(R)}$	Workload per traffic unit, $$W_v^{(R)} = \frac{W_v - M}{T_v}$$ evaluated for VM replica v.

Table 3-5: List of VM variables used in the MILP model

Variable	Description
$W_{v,s}^{(CR)}$	Workload of VM replica v hosted in cloud in node s.
$W_s^{(C)}$	Total workload of cloud hosted in node s.
$D_{v,s,d}^{(C)}$	Traffic flow from VM replica v hosted in cloud of node s to users in node d.
$L_{s,d}$	Traffic from cloud node s to users in node d.

The clouds power consumption is composed of:

- The power consumption of cloud servers:

$$c \sum_{s \in N} S_s^{(C)} S^{(P)} \qquad (3.1)$$

Note that, as the difference between the server idle power and full load is very small (as shown in Fig. 3-2) , we consider an on-off power profile for servers, i.e. if a server is activated, it operates at maximum power consumption, e.g. idle power consumption of S814-8286-41A IBM server is 300 Watt and power consumption at the full server utilization is 333W [47] Moreover, the parameter c represents the PUE of the cloud datacentre. PUE is a metric used to determine the energy efficiency of a cloud/fog data centre by dividing the total amount of energy consumed of cloud/fog data centre facility (including lighting and cooling) by the energy consumed by cloud/fog computing and networking equipment. For example, the average power efficiency or PUE for Google data centres is 1.11 [143].

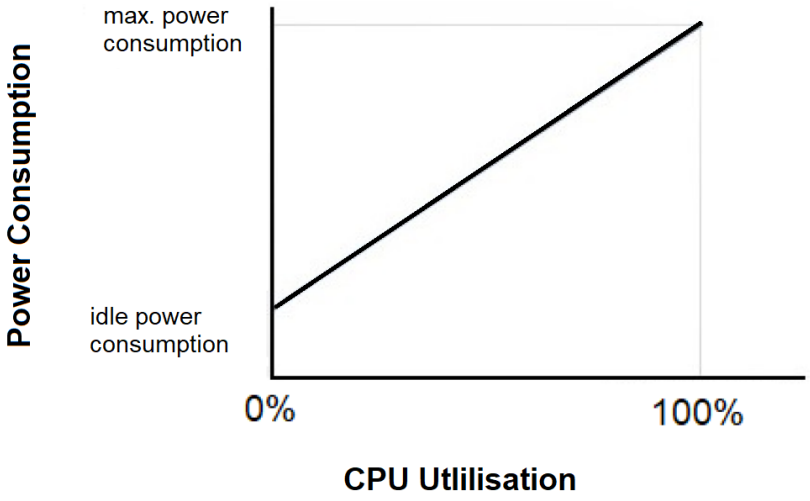

Figure 3-2: Power Consumption of a server versus CPU utlilisation.

- The power consumption of cloud routers and switches:

$$c\left(\sum_{s \in N}\left(\left(SW_s^{(C)} SW^{(R)} SW^{(CP)}\right) + R_s^{(C)} R^{(CP)}\right)\right) \qquad (3.2)$$

Note that, as the difference between the server idle power and full load is very small [107], we consider an on-off power profile for servers, i.e. if a server is activated, it will operate at maximum power consumption.

The parameters defined in Table 3-6 represent access networks:

Table 3-6: List of access network input parameters used in the MILP model

Parameter	Description
$\boldsymbol{P}$	Set of PON networks.
$\boldsymbol{A_p}$	Average broadband data rate in PON $\boldsymbol{p}$.
$\boldsymbol{\Phi_v}$	ratio of traffic due to VM $\boldsymbol{v}$ to the total PON traffic.
$\boldsymbol{OLT^{(B)}_{p,d}}$	Capacity of OLT serving PON $\boldsymbol{p}$ connected to node $\boldsymbol{d}$.
$\boldsymbol{U_{v,p,d}}$	Number of users in PON $\boldsymbol{p}$ connected to core node $\boldsymbol{d}$ requesting VM $\boldsymbol{v}$. $$\boldsymbol{U_{v,p,d}} = \left(\frac{\boldsymbol{OLT^{(B)}_{p,d}}}{\boldsymbol{A_p}}\right) \boldsymbol{\Phi_v}$$ if typical national/regional values of $\boldsymbol{A_p}$, $\boldsymbol{\Phi_v}$ and $\boldsymbol{OLT^{(B)}_{p,d}}$ are used, then $\boldsymbol{U_{v,p,d}}$ determines the number of users and their VM popularity.
$\boldsymbol{OLT^{(N)}_{p,d}}$	Number of OLTs in PON network $\boldsymbol{p}$ connected to node $\boldsymbol{d}$.
$\boldsymbol{OLT^{(P)}}$	OLT power consumption.
$\boldsymbol{D_{v,p,d}}$	Traffic flow from VM $\boldsymbol{v}$ to users in PON network $\boldsymbol{p}$ connected to core node $\boldsymbol{d}$ given as: $$\boldsymbol{D_{v,p,d}} = \boldsymbol{U_{v,p,d}}\, \boldsymbol{r_v}$$
$\boldsymbol{ONU^{(N)}_{p,d}}$	Number of ONUs in PON network $\boldsymbol{p}$ connected to node $\boldsymbol{d}$.
$\boldsymbol{ONU^{(P)}}$	Power consumption of an ONU.
$\boldsymbol{n}$	Network power usage effectiveness (PUE). PUE is a metric used to determine the energy efficiency of a network by dividing the total amount of energy consumed of network facility (including lighting and cooling) by the energy consumed by networking equipment.

The PON networks power consumption is composed of:

- Total power consumption of OLTs:

$$\boldsymbol{n}\left(\sum_{\boldsymbol{p}\in\boldsymbol{P}}\sum_{\boldsymbol{d}\in\boldsymbol{N}}\left(\boldsymbol{OLT^{(P)}}\,\boldsymbol{OLT^{(N)}_{p,d}}\right)\right) \qquad (\mathbf{3.3})$$

- Total power consumption of ONUs:

$$n\left(\sum_{p\in P}\sum_{d\in N}\left(ONU^{(P)}\ ONU_{p,d}^{(N)}\right)\right) \tag{3.4}$$

The parameters and variables defined in Table 3-7 and Table 3-8 represent the metro network.

Table 3-7: List of metro network parameters used in the MILP model

Parameters	Description
$R^{(MB)}$	Metro router bit rate.
$R^{(MP)}$	Metro router power consumption.
$R^{(MR)}$	Metro router redundancy.
$SW^{(MB)}$	Metro Ethernet switch bit rate.
$SW^{(MP)}$	Metro Ethernet power consumption.

Table 3-8: List of metro network variables used in the MILP model

Parameters	Description
$R_s^{(M)}$	Number of router ports in metro network connected to node s.
$SW_s^{(M)}$	Number of Ethernet switches in metro network connected to node s.

The metro network power consumption is composed of:

- Total power consumption of metro router ports:

$$n\ \left(\sum_{s\in N} R_s^{(M)}\ R^{(MR)}R^{(MP)}\right) \tag{3.5}$$

- Total power consumption of metro Ethernet switches:

$$n\left(\sum_{s\in N} SW_s^{(M)}\ SW^{(MP)}\right) \tag{3.6}$$

The parameters and variables defined in Tables 3-9 and 3-10 represents the IP over WDM core network:

Where number of metro routers ports and ethernet switches in metro network:

$$R_s^{(M)} \geq \frac{\sum_{v\in VM}\sum_{s\in N} D_{v,s,d}^{(C)}}{R^{(MB)}}$$

$$\forall\ s \in N \tag{3.7}$$

$$SW_s^{(M)} \geq \frac{\sum_{v \in VM} \sum_{s \in N} D_{v,s,d}^{(C)}}{SW^{(MB)}}$$

$$\forall\, s \in N \qquad (3.8)$$

Equations (3.7) and (3.8) calculate the number of routers ports and switches, respectively, in each metro network.

Table 3-9: List of IP over WDM parameters used in the MILP model

Variable	Description
m and n	Indices of the end nodes of a physical link.
i and j	Indices of the end nodes of a virtual link.
Nm_m	Set of neighbouring nodes of node m.
$R^{(P)}$	Core router port power consumption.
$t^{(P)}$	Transponder power consumption.
$e^{(P)}$	EDFA power consumption.
$SW_s^{(P)}$	Optical switch power consumption in node s.
$G^{(P)}$	Regenerator power consumption.
W	Number of wavelengths per fibre.
$W^{(B)}$	Wavelength data rate.
S	Maximum span distance between two EDFAs in kilometres.
$D_{m,n}$	Distance in kilometres between node pair (m,n).
$A_{m,n}$	Number of EDFAs between node pair (m,n). $A_{m,n}= \left\lfloor \frac{D_{m,n}}{S} - 1 \right\rfloor$ where S is the reach of the EDFA.
$G_{m,n}$	Number of regenerators between node pair (m,n). Typically $G_{m,n}= \left\lfloor \frac{D_{m,n}}{R} - 1 \right\rfloor$, where R is the reach of the regenerator.

Table 3-10: List of IP over WDM variables used in the MILP model

Variable	Description
$C_{i,j}$	Number of wavelengths in virtual link (i,j).
$W_{m,n}$	Number of wavelengths in physical link (m,n).
$R_s^{(AC)}$	Number of router ports in node s that aggregate the traffic from/to clouds.
$R_d^{(AE)}$	Number of router ports in node d that aggregate the traffic from/to metro routers.
$F_{m,n}$	Number of fibres on physical link (m,n).
$L_{i,j}^{s,d}$	Amount of traffic flow between node pair (s,d) traversing virtual link (i,j).
$W_{m,n}^{i,j}$	Number of wavelengths of virtual link (i,j) traversing physical link (m,n).

Under the non-bypass approach, the IP over WDM network power consumption is composed of [19]:

- The power consumption of router ports:

$$\boldsymbol{n}\left(\sum_{\boldsymbol{s}\in\boldsymbol{N}}\boldsymbol{R}^{(\boldsymbol{P})}\,\boldsymbol{R}_{\boldsymbol{s}}^{(\boldsymbol{AC})}+\sum_{\boldsymbol{d}\in\boldsymbol{N}}\boldsymbol{R}^{(\boldsymbol{P})}\boldsymbol{R}_{\boldsymbol{d}}^{(\boldsymbol{AE})}+\sum_{\boldsymbol{m}\in\boldsymbol{N}}\sum_{\boldsymbol{n}\in\boldsymbol{N}\boldsymbol{m}_{\boldsymbol{m}}:\boldsymbol{n}\neq\boldsymbol{m}}\boldsymbol{R}^{(\boldsymbol{P})}\,\boldsymbol{\mathcal{W}}_{\boldsymbol{mn}}\right) \qquad (\mathbf{3.9})$$

Equation (3.9) calculates the power consumption of router ports including the power consumption of router ports in each core node that aggregate the traffic from/to the clouds ($\sum_{s\in N}\boldsymbol{R}^{(P)}\,\boldsymbol{R}_{s}^{(AC)}$), the power consumption of router ports in each core node that aggregate the traffic from/to metro routers ($\sum_{d\in N}\boldsymbol{R}^{(P)}\boldsymbol{R}_{d}^{(AE)}$) and the power consumption of intermediate router ports that transmit the traffic between the source and destination nodes ($\sum_{m\in N}\sum_{n\in N m_m:n\neq m}\boldsymbol{R}^{(P)}\,\boldsymbol{\mathcal{W}}_{mn}$).

Where total number of aggregation ports in a core node:

$$\boldsymbol{R}_{\boldsymbol{s}}^{(\boldsymbol{AC})}=\frac{\mathbf{1}}{\boldsymbol{\mathcal{W}}^{(\boldsymbol{B})}}\sum_{\boldsymbol{d}\in\boldsymbol{N}}\boldsymbol{L}_{\boldsymbol{s,d}}$$

$$\forall\,\boldsymbol{s}\in\boldsymbol{N} \qquad (\mathbf{3.10})$$

$$\boldsymbol{R}_{\boldsymbol{d}}^{(\boldsymbol{AE})}=\boldsymbol{R}^{(\boldsymbol{MR})}\left(\frac{\mathbf{1}}{\boldsymbol{\mathcal{W}}^{(\boldsymbol{B})}}\sum_{\boldsymbol{s}\in\boldsymbol{N}}\boldsymbol{L}_{\boldsymbol{s,d}}\right)$$

$$\forall\,\boldsymbol{d}\in\boldsymbol{N} \qquad (\mathbf{3.11})$$

Equation (3.10) calculates the total number of router ports in each core node that aggregate the traffic from/to the clouds. Equation (3.11) calculates the total number of router ports in each core node that aggregate the traffic from/to metro routers.

The power consumption of transponders:

$$\boldsymbol{n}\left(\sum_{\boldsymbol{m}\in\boldsymbol{N}}\sum_{\boldsymbol{n}\in\boldsymbol{N}\boldsymbol{m}_{\boldsymbol{m}}:\boldsymbol{n}\neq\boldsymbol{m}}\boldsymbol{t}^{(\boldsymbol{P})}\,\boldsymbol{\mathcal{W}}_{\boldsymbol{mn}}\right) \qquad (\mathbf{3.12})$$

- The power consumption of EDFAs:

$$n\left(\sum_{m\in N}\sum_{n\in N m_m : n\neq m} e^{(P)}\, F_{mn}\, A_{mn}\right) \tag{3.13}$$

- The power consumption of optical switches:

$$n\left(\sum_{s\in N} SW_s^{(P)}\right) \tag{3.14}$$

Equation (3.14) calculates the power consumption of optical switches located in each node s. Each node IP routers are connected to optical switches which are connected by optical fibre links. The optical layer (including optical switches and fibre links) provides the large bandwidth required for communication between IP routers.

- The power consumption of regenerator:

$$n\left(\sum_{m\in N}\sum_{n\in N m_m : n\neq m} G^{(P)}\, G_{m,n}\, \mathcal{W}_{m,n}\right) \tag{3.15}$$

The model is defined as follows:

Objective: Minimise the total power consumption given as:

$$c\sum_{s\in N} S_s^{(C)}\, S^{(P)} + c\left(\sum_{s\in N}\left(\left(SW_s^{(C)}\, SW^{(R)}\, SW^{(CP)}\right) + R_s^{(C)}\, R^{(CP)}\right)\right)$$

$$+ n\left(\sum_{p\in P}\sum_{d\in N}\left(OLT^{(P)}\, OLT_{p,d}^{(N)}\right)\right) + n\left(\sum_{p\in P}\sum_{d\in N}\left(ONU^{(P)}\, ONU_{p,d}^{(N)}\right)\right)$$

$$+ n\left(\sum_{s\in N} R_s^{(M)}\, R^{(MR)} R^{(MP)}\right) + n\left(\sum_{s\in N} SW_s^{(M)}\, SW^{(MP)}\right)$$

$$+ n\left(\sum_{s\in N} R^{(P)}\, R_s^{(AC)} + \sum_{d\in N} R^{(P)} R_d^{(AE)} + \sum_{m\in N}\sum_{n\in N m_m : n\neq m} R^{(P)}\, \mathcal{W}_{mn}\right)$$

$$+ n\left(\sum_{m\in N}\sum_{n\in N m_m : n\neq m} t^{(P)}\, \mathcal{W}_{mn}\right) + n\left(\sum_{m\in N}\sum_{n\in N m_m : n\neq m} e^{(P)}\, F_{mn}\, A_{mn}\right)$$

$$+ \boldsymbol{n}\left(\sum_{\boldsymbol{s}\in\boldsymbol{N}} \boldsymbol{SW}_{\boldsymbol{s}}^{(\boldsymbol{P})}\right) + \boldsymbol{n}\left(\sum_{\boldsymbol{m}\in\boldsymbol{N}} \sum_{\boldsymbol{n}\in\boldsymbol{N}\boldsymbol{m}_{\boldsymbol{m}}:\boldsymbol{n}\neq\boldsymbol{m}} \boldsymbol{G}^{(\boldsymbol{P})}\, \boldsymbol{G}_{\boldsymbol{m},\boldsymbol{n}}\, \boldsymbol{\mathcal{W}}_{\boldsymbol{m},\boldsymbol{n}}\right) \tag{3.16}$$

Equation (3.16) gives the total power consumption as the sum of the power consumption of the clouds, the access network, the metro network and IP over WDM core network. This Equation is the sum of equations (3.1) to (3.11) which described above.

Subject to:

Serving VMs demand:

$$\sum_{\boldsymbol{p}\in\boldsymbol{P}} \sum_{\boldsymbol{d}\in\boldsymbol{N}} \boldsymbol{D}_{\boldsymbol{v},\boldsymbol{p},\boldsymbol{d}} = \sum_{\boldsymbol{s}\in\boldsymbol{N}} \sum_{\boldsymbol{d}\in\boldsymbol{N}} \boldsymbol{D}_{\boldsymbol{v},\boldsymbol{s},\boldsymbol{d}}^{(\boldsymbol{C})}$$

$$\forall\ \boldsymbol{v} \in \boldsymbol{VM} \tag{3.17}$$

Constraint (3.17) ensures that the users demand for a VM in all nodes are served by VM placed in the distributing clouds connected to IP over WDM networks.

Placing VMs in clouds constraint:

$$\boldsymbol{L} \sum_{\boldsymbol{d}\in\boldsymbol{N}} \boldsymbol{D}_{\boldsymbol{v},\boldsymbol{s},\boldsymbol{d}}^{(\boldsymbol{C})} \geq \boldsymbol{\delta}_{\boldsymbol{v},\boldsymbol{s}}^{(\boldsymbol{C})}$$

$$\forall\, \boldsymbol{s} \in \boldsymbol{N}\,, \boldsymbol{v} \in \boldsymbol{VM} \tag{3.18}$$

$$\sum_{\boldsymbol{d}\in\boldsymbol{N}} \boldsymbol{D}_{\boldsymbol{v},\boldsymbol{s},\boldsymbol{d}}^{(\boldsymbol{C})} \leq \boldsymbol{L}\, \boldsymbol{\delta}_{\boldsymbol{v},\boldsymbol{s}}^{(\boldsymbol{C})}$$

$$\forall\, \boldsymbol{s} \in \boldsymbol{N}\,, \boldsymbol{v} \in \boldsymbol{VM} \tag{3.19}$$

Constraints (3.18) and (3.19) relate the binary variable that indicates whether a VM is hosted in a cloud or not, ($\boldsymbol{\delta}_{\boldsymbol{v},\boldsymbol{s}}^{(\boldsymbol{C})}$), to the traffic between users of this VM and the cloud ($\sum_{\boldsymbol{d}\in\boldsymbol{N}} \boldsymbol{D}_{\boldsymbol{v},\boldsymbol{s},\boldsymbol{d}}^{(\boldsymbol{C})}$) by setting $\boldsymbol{\delta}_{\boldsymbol{v},\boldsymbol{s}}^{(\boldsymbol{C})} = \boldsymbol{1}$ if $\sum_{\boldsymbol{d}\in\boldsymbol{N}} \boldsymbol{D}_{\boldsymbol{v},\boldsymbol{s},\boldsymbol{d}}^{(\boldsymbol{C})} > \boldsymbol{0}$ and $\boldsymbol{\delta}_{\boldsymbol{v},\boldsymbol{s}}^{(\boldsymbol{C})} = \boldsymbol{0}$ otherwise. The constraints replicate VM $\boldsymbol{v}$ to cloud $\boldsymbol{s}$ if cloud $\boldsymbol{s}$ is selected to serve requests for VM $\boldsymbol{v}$.

Clouds workload:

$$W_{v,s}^{(CR)} = \delta_{v,s}^{(C)}\ W_v \quad (\textbf{Constant Workload Profile})$$

$$\forall\ v \in VM, s \in N \tag{3.20}$$

$$W_{v,s}^{(CR)} = \left(\frac{\sum_{d \in N} D_{v,s,d}^{(C)}}{r_v\ x}\ M\right) + \left(W_v^{(R)} \sum_{d \in N} D_{v,s,d}^{(C)}\right) \quad (\textbf{Linear Workload Profile})$$

$$\forall\ v \in VM, \text{s} \in N \tag{3.21}$$

$$W_s^{(C)} = \sum_{v \epsilon VM} W_{v,s}^{(CR)}$$

$$\forall\ s \in N \tag{3.22}$$

Constraint (3.20) calculates the workload of a VM replica under a constant workload profile. Constraint (3.21) calculates the workload of a VM replica under a linear workload profile. It gives the workload of a VM as a linear function of the traffic resulting from serving users of the replica $\left(W_v^{(R)} \sum_{d \in N} D_{v,s,d}^{(C)}\right)$ plus the workload baseline $\left(\frac{\sum_{d \in N} D_{v,s,d}^{(C)}}{r_v\ x}\ M\right)$. Constraint (3.22) calculates the total workload of a cloud by summing the workload of VMs hosted in it.

Number of servers in clouds:

$$S_s^{(C)} \geq \frac{W_s^{(C)}}{S^{(maxW)}}$$

$$\forall\ s \in N \tag{3.23}$$

Constraint (3.23) calculates the number of servers in each cloud based on the CPU utilisation.

Number of router ports and switches in cloud constraints:

$$R_s^{(C)} \geq \frac{\sum_{v \in VM} \sum_{d \in N} D_{v,s,d}^{(C)}}{R^{(CB)}}$$

$$\forall\ s \in N \tag{3.24}$$

$$SW_s^{(C)} \geq \frac{\sum_{v \in VM} \sum_{d \in N} D_{v,s,d}^{(C)}}{SW^{(CB)}}$$

$$\forall\, s \in N \qquad (3.25)$$

Constraints (3.24) - (3.25) calculate the number of routers ports and switches in each cloud, respectively.

Traffic flow on IP over WDM core network constraint:

$$L_{s,d} = \sum_{v \in VM} D_{v,s,d}^{(C)}$$

$$\forall\, s, d \in N \qquad (3.26)$$

Constraint (3.26) calculates the demand between the IP over WDM nodes by summing the demand due to VMs placed in the clouds.

Flow conservation constraint in the IP layer:

$$\sum_{j \in N: i \neq j} L_{i,j}^{s,d} - \sum_{j \in N: i \neq j} L_{i,j}^{s,d} = \begin{cases} L_{s,d} & i = s \\ -L_{s,d} & i = d \\ 0 & otherwise \end{cases}$$

$$\forall\, s, d, i \in N : s \neq d \qquad (3.27)$$

Constraint (3.27) represents the flow conservation for IP layer on the IP over WDM network. It ensures that the total incoming traffic is equal to the total outgoing traffic in all nodes; excluding the source and destination nodes.

Virtual link capacity constraint:

$$\sum_{s \in N} \sum_{d \in N: s \neq d} L_{i,j}^{s,d} \leq C_{i,j}\, W^{(B)}$$

$$\forall\, i, j \in N : s \neq d \qquad (3.28)$$

Constraint (3.28) ensures that the traffic transmitted through a virtual link does not exceed its maximum capacity.

Flow conservation constraint in the optical layer:

$$\sum_{n \in N m_m} w_{m,n}^{i,j} - \sum_{n \in N m_m} w_{m,n}^{i,j} = \begin{cases} C_{i,j} & m = i \\ -C_{i,j} & m = j \\ 0 & otherwise \end{cases}$$

$$\forall\, i, j, m \in N : i \neq j \qquad (3.29)$$

Constraint (3.29) represents the flow conservation for the optical layer. It ensures that the total number of incoming wavelengths in a virtual link is equal to the total number of outgoing wavelengths in all nodes excluding the source and destination nodes of the virtual link.

Physical link capacity:

$$\sum_{i \in N} \sum_{j \in N : i \neq j} w_{m,n}^{i,j} \leq W F_{m,n}$$

$$\forall\, m, n \in N \qquad (3.30)$$

$$w_{mn} = \sum_{i \in N} \sum_{j \in N : i \neq j} w_{m,n}^{i,j}$$

$$\forall\, m, n \in N \qquad (3.31)$$

Constraints (3.30) and (3.31) represent the physical link capacity limit. Constraint (3.30) ensures that the number of wavelengths in virtual links traversing a physical link does not exceed the maximum capacity of fibres in the physical link. Constraint (3.31) calculates the number of wavelengths in a physical link as the sum of wavelength channels in virtual links traversing the physical link.

3.3 Results and Discussions:

In this section, the optimal VMs placement over distributed clouds in IP over WDMs is investigated considering the AT&T network topology of the US and the BT network topology of the UK as use cases.

3.3.1 AT&T Network Use Case:

In this section, the optimal VMs placement over AT&T distributed cloud architecture is investigated. The AT&T core networks topology is illustrated in Fig.

3-3 [144]. The AT&T core network consists of 25 nodes and 54 bidirectional links [144]. We consider an architecture where each core node is connected to two PON networks through a metro network consisting of a single ethernet switch and two metro routers (illustrated in Fig 3.1). The PON access network is considered to connect 512 different locations. The total capacity of each OLT is 1280 Gbps [145].

We start by considering optimising a single VM as the simplest representative problem. Then we consider optimisation in a realistic scenario with multiple VMs.

3.3.1.1 Simple Representative Scenario:

We investigate how the energy efficient placement of a single VM over cloud-fog architecture varies based on three factors; the CPU requirements, download traffic and PUE values.

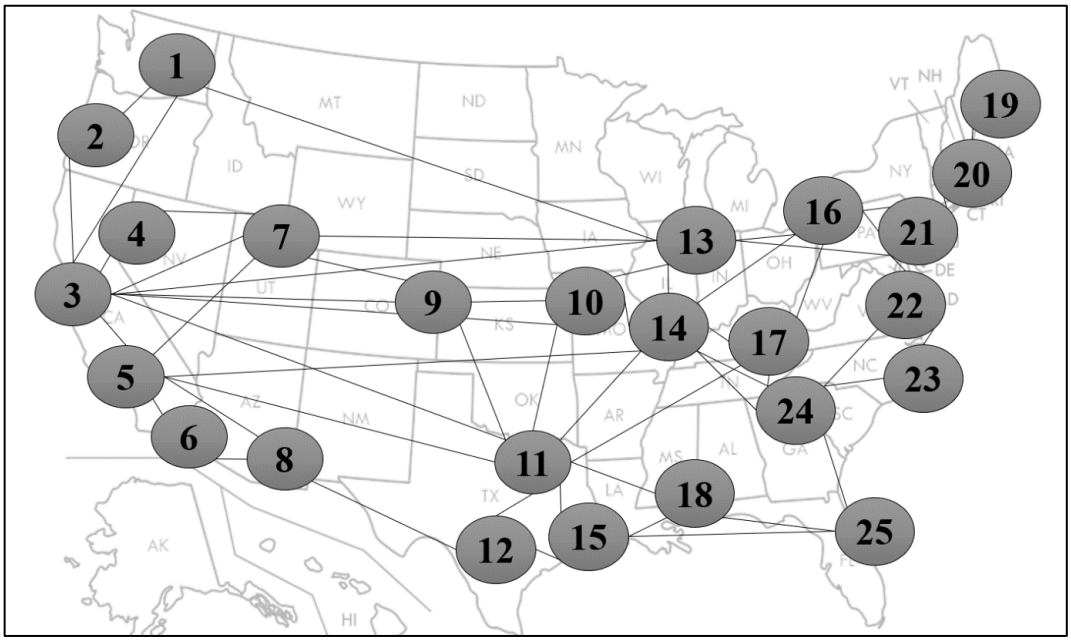

Figure 3-3: AT&T core network topology.

The impact of the VM workload profile on the VM placement is examined by considering constant and linear workload profiles. For the linear workload profile, a simple linear profile with no baseline is considered. The workload of VM of a constant workload profile and the workload of serving the maximum number of users of VM of a linear workload are considered to take one of three workloads: 10%, 50% and 100% of the server CPU capacity. The users are considered to access VM with one of following download rates; 0.1 Mbps, 1 Mbps, 10 Mbps, 20 Mbps, 50 Mbps, 100 Mbps or 200 Mbps. Each VM is considered to have 800 users. The PUE is a metric used to

determine the total energy consumption required by clouds, fogs or networks (including IT cooling, lighting, etc.) in relation to IT infrastructure components. Based on a recent study [146], PUE of cloud datacentres are 1.3 and 1.7 for best practice and 2014 values, respectively.

The MILP model is solved using the CPLEX solver, as discussed in Section 2.8, over the University of Leeds high-performance computer (Polaris) using 16 nodes (256 cores) with 16 GByte of RAM per core. Each node comprises two eight cores of the Intel 2.6 GHz Sandy Bridge E5-2670 processors [147].

Table 3-11 to Table 3-14 show the IP over WDM core network, metro network, access network and cloud input parameters.

Table 3-11: IP Over WDM core network input parameters used in the model

40 Gbps router port power consumption ($\boldsymbol{R^{(P)}}$)	638W [148]
40 Gbps transponder power consumption ($\boldsymbol{t^{(P)}}$)	129W [149]
40 Gbps regenerator power consumption ($\boldsymbol{G^{(P)}}$)	114W, reach 2000 km [150]
EDFA power consumption ($\boldsymbol{e^{(P)}}$)	11W [151]
Optical switch power consumption ($\boldsymbol{SW^{(P)}}$)	85W [152]
Number of wavelengths in a fibre ($\boldsymbol{\mathcal{W}}$)	32 [20]
Bit rate of each wavelength ($\boldsymbol{\mathcal{W}^{(B)}}$)	40 Gbps [20]
Span distance between two EDFAs ($\boldsymbol{S}$)	80 km [151]
Network power usage effectiveness ($\boldsymbol{n}$)	1.5 [6]

Table 3-12: Metro network input parameters used in the model

Metro aggregation router redundancy ($\boldsymbol{R^{(MR)}}$)	2
Metro aggregation router port bit rate ($\boldsymbol{R^{(MB)}}$)	40 Gbps
Metro aggregation router port power consumption ($\boldsymbol{R^{(MP)}}$)	30W [153]
Metro ethernet switch bit rate ($\boldsymbol{SW^{(MB)}}$)	600 Gbps [154]
Metro ethernet switch power consumption ($\boldsymbol{SW^{(MP)}}$)	470W [154]

Table 3-13: Access network input parameters used in the model

Number of PON networks in a node ($\boldsymbol{P}$)	2
Maximum number of single VM users ($\boldsymbol{x}$)	800 concurrent users
Number of ONU devices in a PON network ($\boldsymbol{ONU}_{\boldsymbol{p,d}}^{\boldsymbol{(N)}}$)	512
Power consumption of ONU device ($\boldsymbol{ONU}^{\boldsymbol{(P)}}$)	5W [155]
Number of OLTs in a PON network ($\boldsymbol{OLT}_{\boldsymbol{p,d}}^{\boldsymbol{(N)}}$)	1
OLT Capacity ($\boldsymbol{OLT}^{\boldsymbol{(B)}}$)	1280 Gbps [145]

Table 3-14: Clouds input parameters used in the model

Number of VMs ($\boldsymbol{V}$)	1
User download rate ($\boldsymbol{r_v}$)	{0.1, 1, 10, 20, 50, 100 or 200 Mbps}
Maximum workload of VM ($\boldsymbol{W_v}$)	10%, 50% and 100%
Server power consumption ($\boldsymbol{S}^{\boldsymbol{(P)}}$)	333W [156]
Maximum server workload ($\boldsymbol{S}^{\boldsymbol{(maxW)}}$)	100%
Cloud switch bit rate $\left(\boldsymbol{SW}^{\boldsymbol{(CB)}}\right)$	600 Gbps [154]
Cloud switch power consumption ($\boldsymbol{SW}^{\boldsymbol{(CP)}}$)	470W [154]
Cloud switch redundancy ($\boldsymbol{SW}^{\boldsymbol{(R)}}$)	2
Cloud router port bit rate ($\boldsymbol{R}^{\boldsymbol{(CB)}}$)	40 Gbps
Cloud router port power consumption ($\boldsymbol{R}^{\boldsymbol{(CP)}}$)	30 W [153]
Cloud power usage effectiveness ($\boldsymbol{c}$)	1.3 or 1.7 [146]

Fig. 3-4 (a), (b) and (c) represent the optimal placement of the VM in distributed clouds under 10%, 50% and 100% CPU requirements, respectively, considering the best practice PUE value. In each figure, the x-axis represents the VM workload profile, the y-axis represents the data rates which range from 0.1 Mbps to 200 Mbps. These download rates represents current (0.1 Mbps, 1 Mbps, 10 Mbps and 20 Mbps)

and futuristic (50 Mbps, 100 Mbps and 200 Mbps) scenarios of VM-hosted applications. The z-axis represents the number of cloud locations in the distributed cloud architecture. The placement of VM of linear workload profile is not affected by the VM workload as serving users will consume the same power whether centralised in a single VM or distributed among multiple replicas of smaller workloads. For data rates ≥ 10 Mbps, VMs are fully replicated considering different workloads. For constant workload profile, replicas are less energy efficient, therefore, the number of replicas decreases as the VM workload increases. While 10% workload VM under 200 Mbps is fully replicated, the number of replicas is limited to 6 and 5 under 50% and 100% workloads, respectively. The results also show that higher data rates justify the creation of more VM replicas. For example, under the linear workload profile, 1 replica, 4 replicas and 25 replicas are created under 0.1 Mbps, 1 Mbps and ≥ 10 Mbps users data rates, respectively.

Placing the VM in a cloud architecture with higher PUE (2014 PUE), as in Fig. 3-5, increases the replicas power consumption and therefore limits their number. This is clear for the constant workload profile, e.g. VM of 10% workload and 20 Mbps data rate is fully replicated considering clouds of best practice PUE, but is limited to 2 replicas with 2014 PUE. Under the linear workload profile, the impact of increasing PUE value is very limited. The only decrease in number of VM replicas can be observed at 1 Mbps users data rate (3 replicas instead of 4).

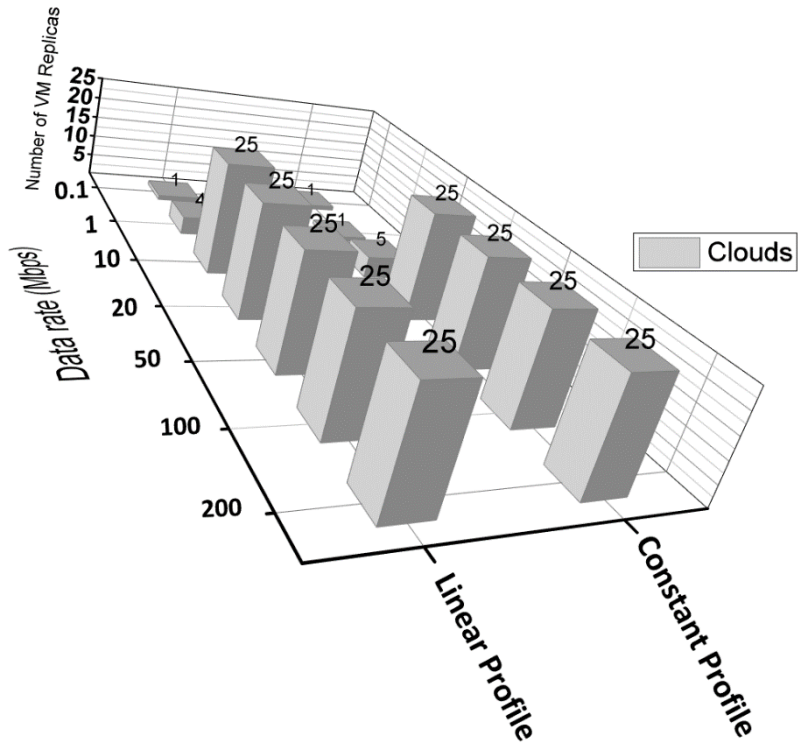
Number of VM Replicas
25
20
15
10
5
Data rate (Mbps)
0.1
1
10
20
50
100
200
Clouds
Linear Profile
Constant Profile

(a)

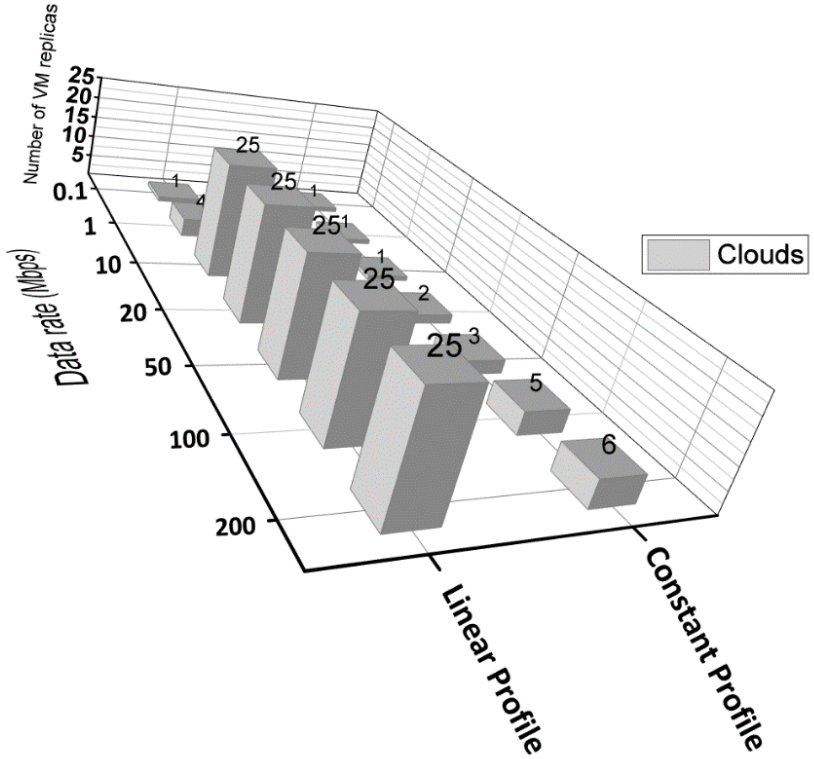
Number of VM replicas
25
20
15
10
5
Data rate (Mbps)
0.1
1
10
20
50
100
200
Clouds
Linear Profile
Constant Profile

(b)

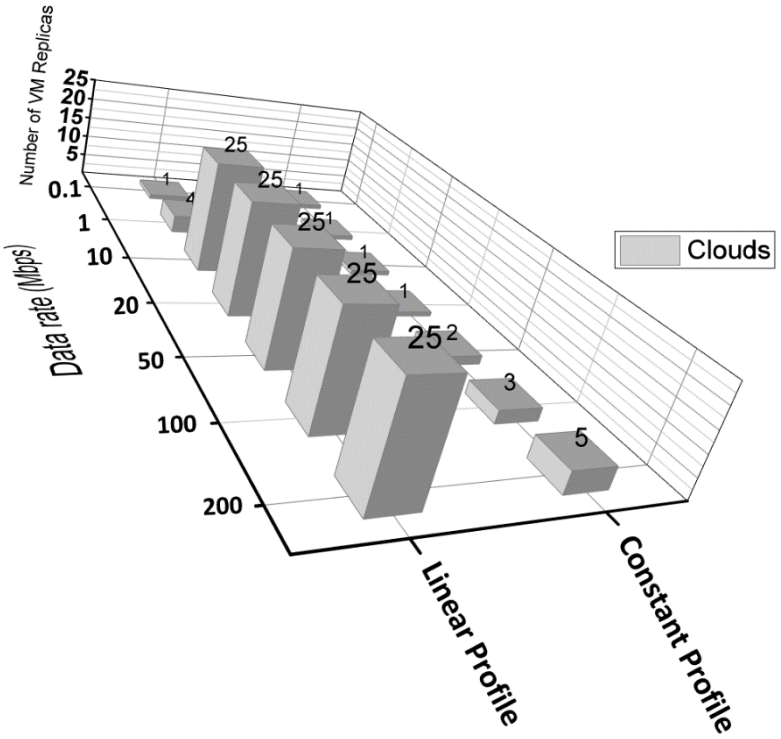

(c)

Figure 3-4: Optimal VM placement (a) constant profile at 10% of CPU and linear profile with peak utilisation of 10%, (b) 50% case, (c) 100% case at different data rates considering best practice PUE value ($\boldsymbol{c}$=1.3).

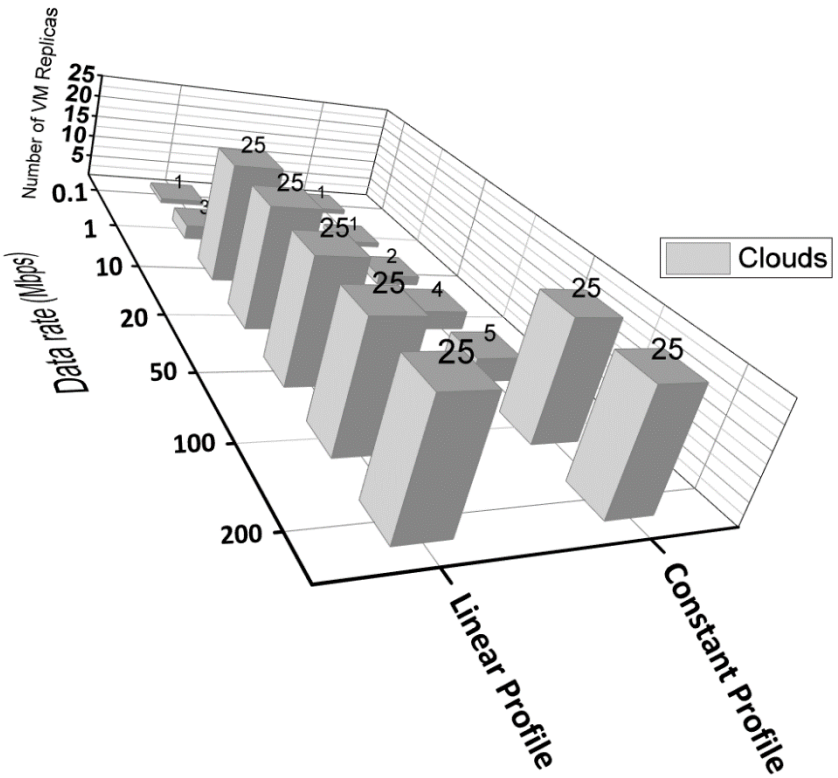
Number of VM Replicas
25
20
15
10
5
Data rate (Mbps)
0.1
1
10
20
50
100
200
Clouds
Linear Profile
Constant Profile

(a)

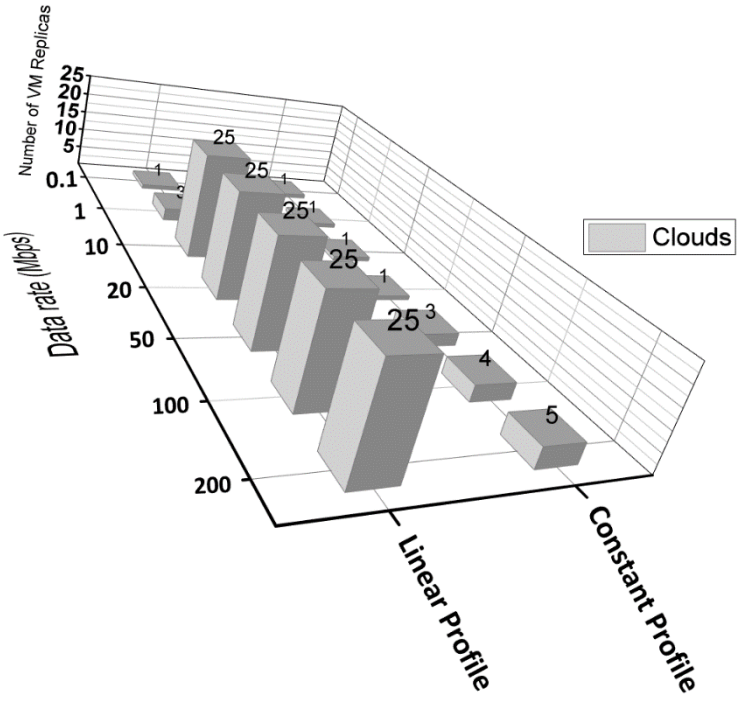
Number of VM Replicas
25
20
15
10
5
Data rate (Mbps)
0.1
1
10
20
50
100
200
Clouds
Linear Profile
Constant Profile

(b)

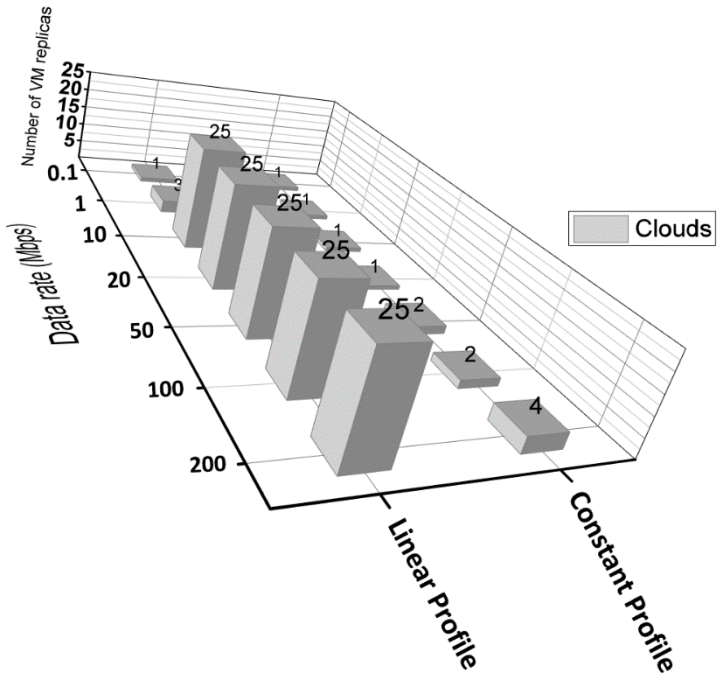

(c)

Figure 3-5: Optimal VM placement (a) constant profile at 10% of CPU and linear profile with peak utilisation of 10%, (b) 50% case, (c) 100% case at different data rates considering 2014 PUE value (***c***=1.7).

3.3.1.2 Realistic Scenario

In this scenario, analysis based on realistic number of users and VM popularity is studied. According to Cisco Visual Network Index (VNI) [157], in 2016, the average broadband data rate in US is 36 Mbps. Therefore, each OLT is assumed to be able to serve ~35k connections (or users). Cisco VNI also reports that 76% of all Internet traffic had crossed clouds in 2016. SimilarWeb [158], an online tool which provides Internet traffic statistics and analytics, shows that the top 300 applications or websites have a 50% share of all traffic. Accordingly, 13k users are considered in each PON (~50% of clouds traffic, i.e. 38% of the total traffic) to access the VMs hosting the top 300 applications or websites. The popularity of these VMs is considered to follow a Zipf distribution [159]. To simplify the analysis, VMs' popularity is divided into 6 groups as follows; 16%, 5%, 2%, 1%, 0.5% and 0.05% of the total users. The number of VMs in each popularity group are 1, 3, 5, 16, 65 and 210, respectively.

Each VM is assumed to require 50% of the CPU's server capacity in order to serve 800 users. Based on literature [59] - [61], [160], [161], a VM can serve 800 users with low error rate.

VMs of a linear workload are considered to have a workload baseline of 1%, 5% or 40% of the total server CPU capacity based on the CPU requirements for state of the art applications [59], [61], [160] (e.g. 1% workload baseline for database applications, 5% for website applications, and 40% for video games and web conference applications). The users are considered to access the VMs with one of the following data rates; 1 Mbps (low), 10 Mbps (medium) or 25 Mbps (high). Such data rates represent the recommended download speed to access the content of the state of the art applications, e.g. 1 Mbps for light web browsing [62] (emails, Google docs [63] and websites with lower definition video content [64]), 10 Mbps for applications processing high-definition video quality [65] and online multiplayer games [66], and 25 Mbps for applications processing ultra-high video quality [67].

The PUE value for the cloud is considered to be 1.5 based on the best practice PUE value [146]. For network infrastructure, a typical telecom office PUE is 1.5 [75].

The optimised VMs placement over the distributed clouds, referred to as Optimised clouds (OC) approach, is compared to the AT&T clouds (AT&T) where the VMs are placed in nodes 1, 3, 5, 6, 8, 11, 13, 17, 19, 20, 22, and 25 according to AT&T topology [144].

The MILP model is solved using the CPLEX solver over the University of Leeds high-performance computer (Polaris) using 16 nodes (256 cores) with 16 GByte of RAM per core. Each node comprises two eight cores of the Intel 2.6 GHz Sandy Bridge E5-2670 processors [147].

The Cisco Carrier Routing System 1 (CRS-1) [148] is considered as a core IP router. CRS-1 provides 160 Gbps routing capacity in 4 ports while consuming 2551W. Therefore, the power consumption of each 40 Gbps router port is 638W. In metro network and cloud aggregation routers, we considered the 1.92 Tbps Cisco NCS 5502 router [153], which is designed for global datacentres and service provider Wide area Network (WAN) aggregation networks with 1450W power consumption (30W per 40 Gbps port). Furthermore, the Cisco Nexus 93180YC-EX [154] switch is considered as metro, cloud LAN Ethernet switch with upload capacity of 600 Gbps and power rating at 470W.

In addition to the parameters in Table 3-11 to Table 3-14, Table 3-15 shows the additional/modified parameters considered for the following results.

Table 3-15: Input parameters used in the model

Number of VM users in each PON based on VMs popularity groups ($\boldsymbol{U_{v,p,d}}$)	13,000 users in each PON, six VM popularity groups; 16%, 5%, 2%, 1%, 0.5% and 0.05%.
Number of VMs ($\boldsymbol{V}$)	300
User download rate ($\boldsymbol{r_v}$)	{1, 10 or 25 Mbps}
Maximum workload of VM ($\boldsymbol{W_v}$)	50%
Cloud power usage effectiveness ($\boldsymbol{c}$)	1.3 [146]

a) Linear Workload Profile (1% Workload Baseline):

Fig. 3-6 shows the power consumption resulting from placing VMs of 1% workload baseline under 1, 10 and 25 Mbps user data rates. The savings achieved by the OC approach compared to the AT&T clouds (where clouds are fixed in the current locations found in the AT&T network) are 2%, 9% and 16%, respectively.

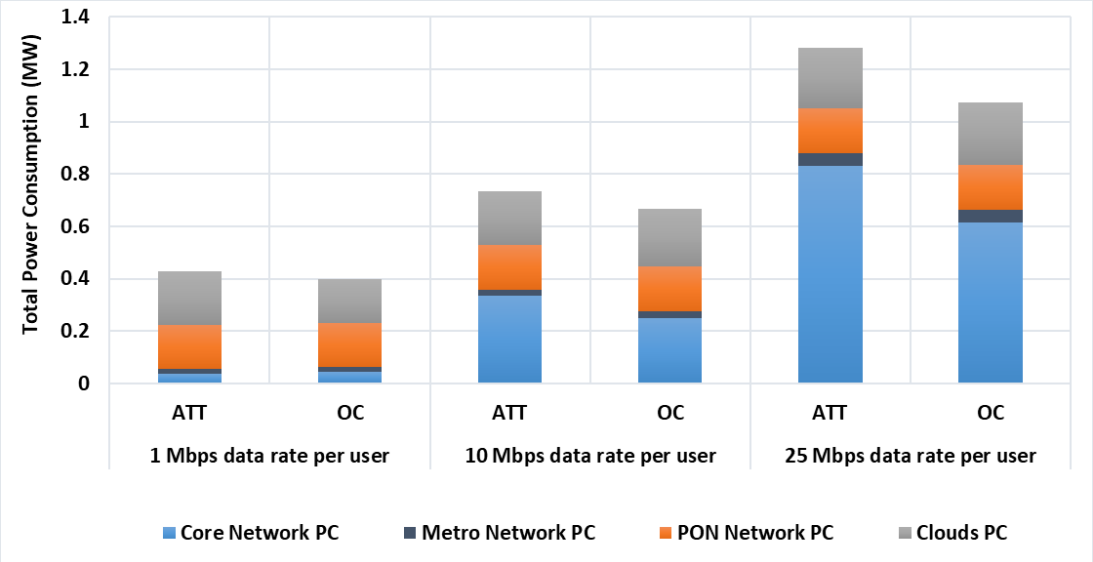

Figure 3-6: The power consumption of different VMs placement approaches considering VMs of 1% workload baseline.

Fig. 3-7 shows the optimal VMs placement under the OC approach. Note that the different colours indicate the creation of a replica of the VM in the cloud, i.e. not the number of replicas. The efficiency of VMs has allowed the creation of multiple replicas as the workload is proportional to the number of users served by the VM with a limited workload baseline. The efficient workload profile of the VMs has justified

the replication of VMs of popularity of greater than 0.5% into 10 clouds for 1 Mbps data rates and into 25 clouds (full replication) for 10 Mbps data rates. VMs of 0.05% popularity are only replicated into 2 clouds. The high traffic of VMs of 25 Mbps data rate allows full replication for the different popularity groups across all clouds.

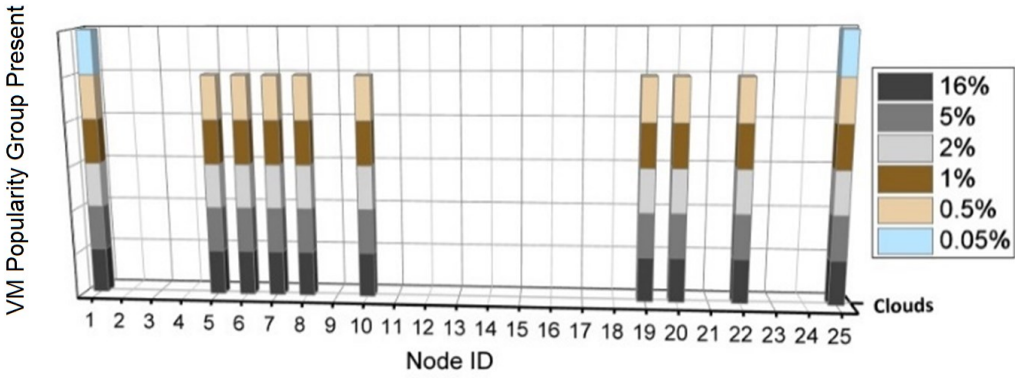

(a)

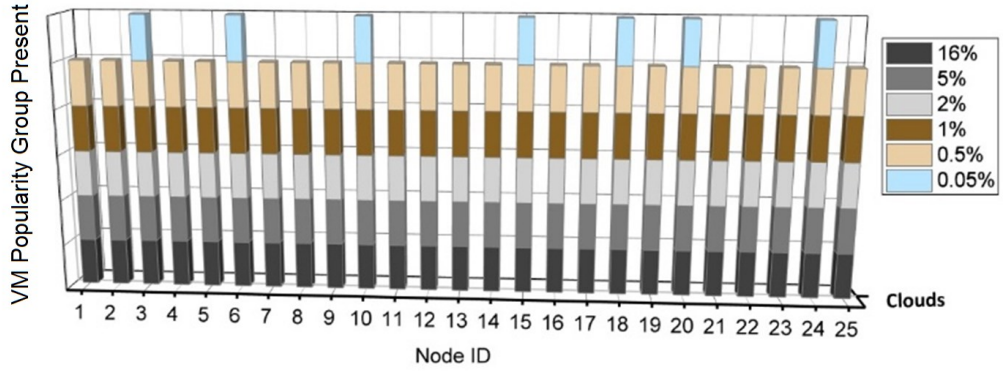

(b)

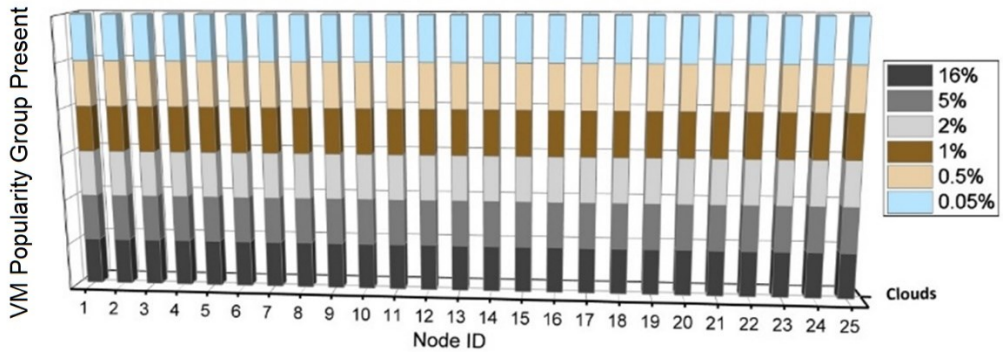

(b)

Figure 3-7: Optimal placement of different VMs popularity groups of 1% workload baseline under the OC approach with (a) 1 Mbps data rate per user, (b) 10 Mbps data rate per user and (c) 25 Mbps data rate per user.

Fig. 3-8 (a), (b) and (c) show the number of servers required to host VMs replicas considering the OC approach. The number of servers is a function of the number of VMs replicas hosted and their workload.

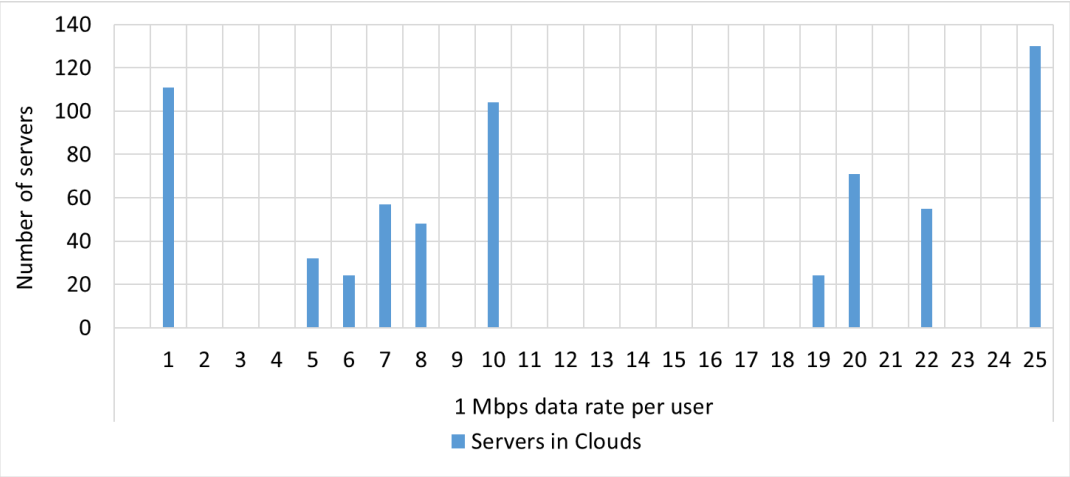

(a)

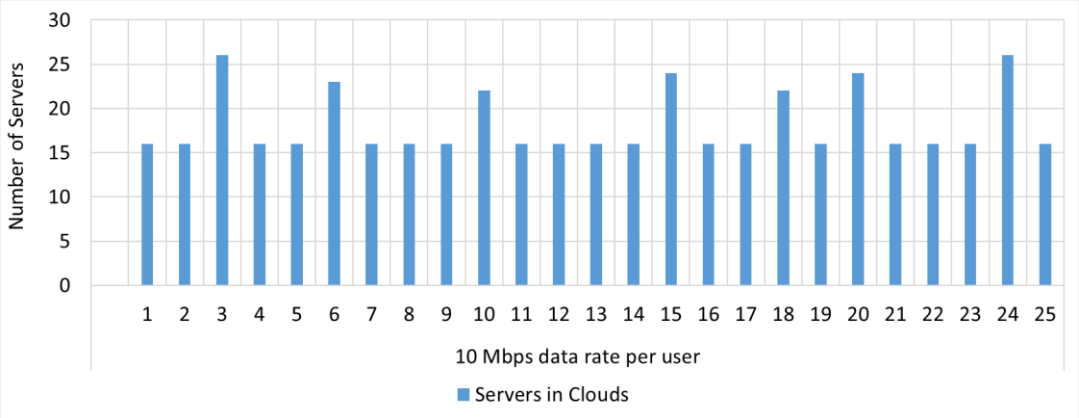

(b)

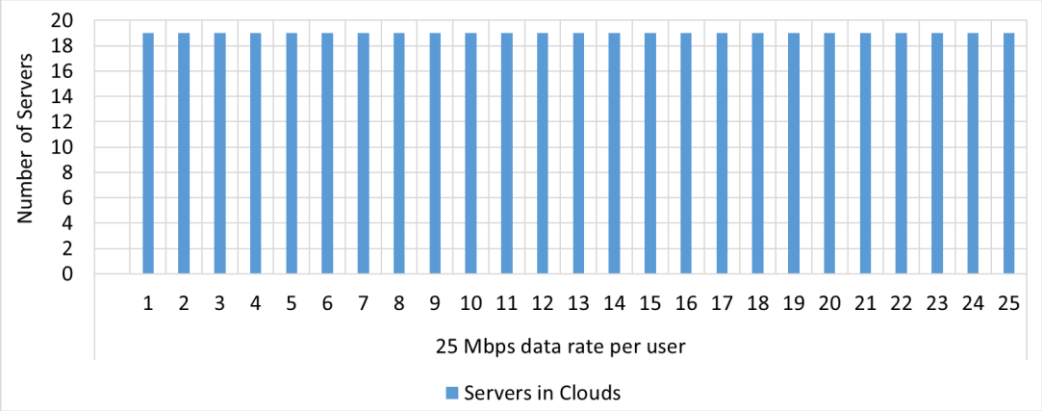

(c)

Figure 3-8: Number of servers in OC approach required to host VMs of 1% workload baseline with (a) 1 Mbps data rate per user (b) 10 Mbps data rate per user (c) 25 Mbps data rate per user.

b) Linear Workload Profile (5% Workload Baseline):

Fig. 3-9 shows that the power consumption resulting from the optimum placement of VMs of a linear workload profile with 5% workload baseline. It can be observed that increasing the workload baseline to 5% reduces the efficiency of creating more VMs replicas. The total savings achieved under the OC approach compared to the AT&T cloud are 14%, 10% and 15% under the low, medium and high data rates, respectively.

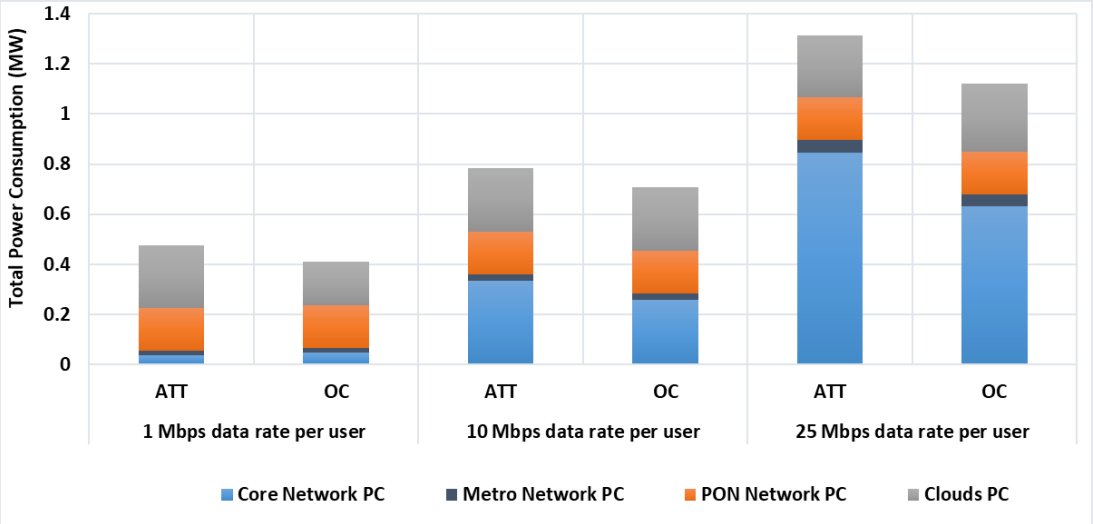

Figure 3-9: The power consumption of different VMs placement approaches considering VMs of 5% workload baseline.

Fig. 3-10(a) and (b) illustrate the placements of the VMs of 5% workload baseline considering the OC approach, under the low and high data rates, respectively. Under the low data rate, five cloud locations are created to host VM replicas of popularity $\geq$ **1**% whereas, VMs of 0.5% popularity are replicated to three locations. The potential network power consumption savings obtained by replicating VMs to multiple clouds do not compensate for creating any replica of VMs of 0.05% popularity. For VMs of 25 Mbps data rate, the high traffic allows full replication for different popularity groups across all clouds, except VMs of 0.05% popularity which are placed optimally in 4 cloud locations.

The number of activated servers in each node under the OC approach shown in Fig. 3-11 reflects the number and workload of the VMs hosted in the nodes.

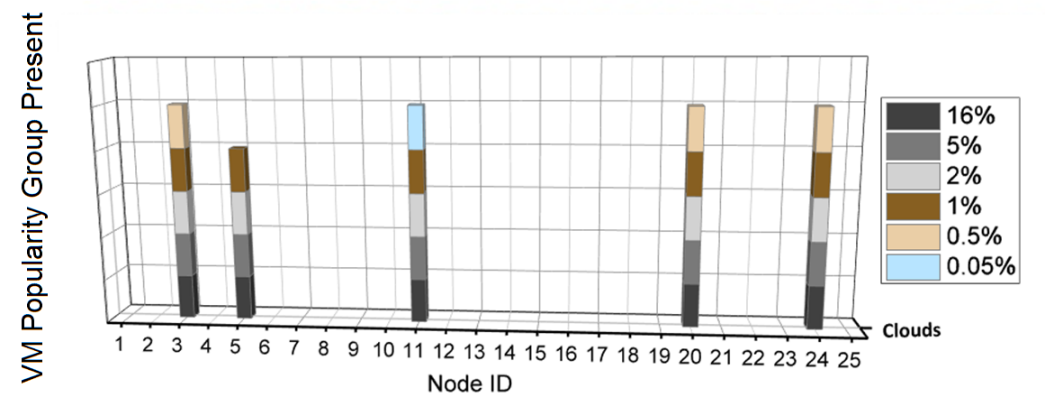

(a)

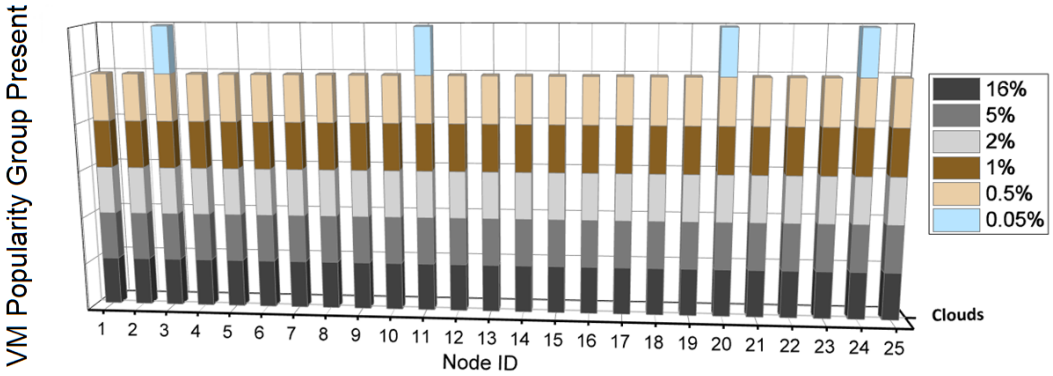

(b)

Figure 3-10: Optimal placement of different VMs popularity groups of 5% workload baseline under the OC approach with (a) 1 Mbps data rate per user (b) 25 Mbps data rate per user.

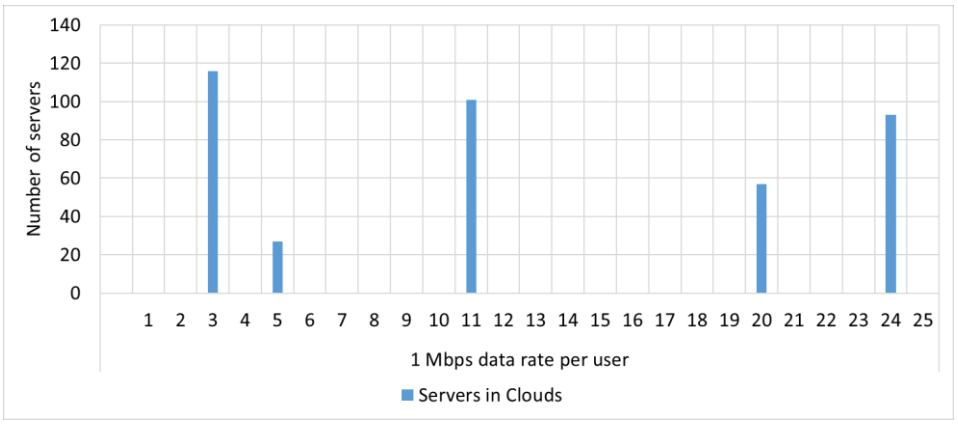

(a)

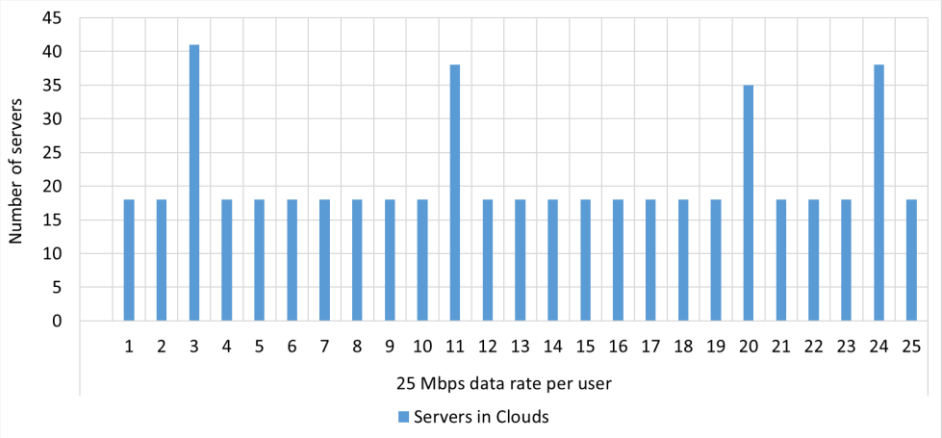

(b)

Figure 3-11: Number of servers required to host VMs of 5% workload baseline under the OC approach with with (a) 1 Mbps data rate per user (b) 25 Mbps data rate per user.

c) Linear Workload profile (40% Workload Baseline):

Fig. 3-12 shows the power consumption achieved by optimally placing VMs of 40% workload baseline. It can be observed that increasing the workload baseline to 40% results in higher power consumption compared to the previous scenarios of 1% and 5% workload baselines. The total power savings achieved under the OC approach compared to the AT&T cloud are 51%, 38% and 24% under the low, medium and high user data rates, respectively.

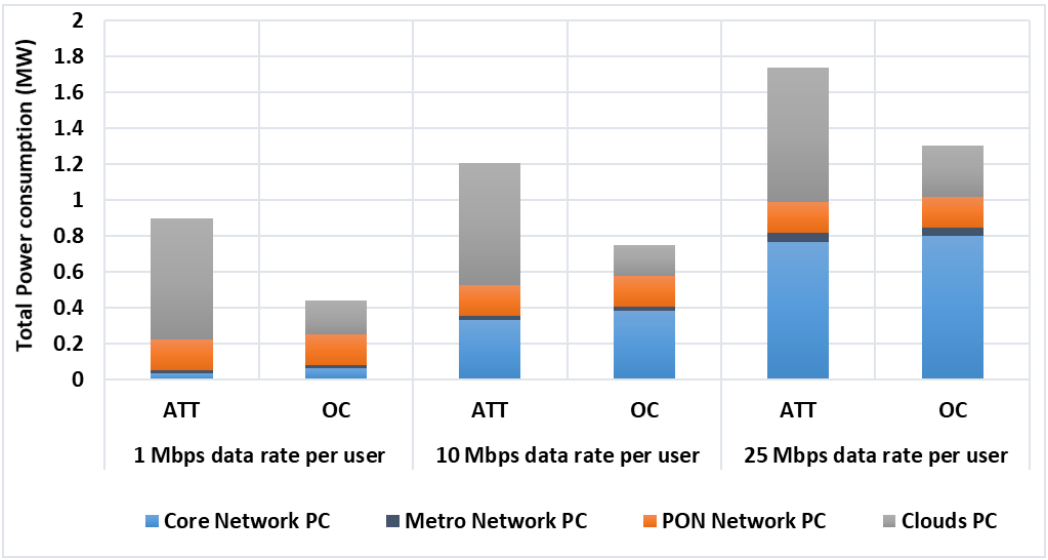

Figure 3-12: The power consumption of different VMs placement approaches considering VMs of 40% workload baseline.

Fig. 3-13(a) and (b) show the placements of the VMs of 40% workload baseline considering the OC approach under the low and high user data rates, respectively. It can be observed that increasing the workload baseline to 40% reduces the efficiency of creating more VMs replicas compared to 1% and 5% workload baseline scenarios. Under the low user data rate, VMs of $\leq$ **1**% popularity are centralised in one cloud, the node of the minimum average hop count to all the AT&T network nodes (node 11) given that VM users are uniformly distributed among 25 nodes. VMs of $\geq$ **2**% popularity are replicated into 2 nodes (node 3 to serve users of the west part of the network and node 14 to serve users in the east part). For VMs of 25 Mbps, popularity greater than **1**% justifies full replication, whereas, VMs of 0.5% and 0.05% popularity are replicated to four and one clouds, respectively.

The number of activated servers in each node under the OC approach shown in Fig. 3-14 reflects the number and workload of the VMs hosted in the nodes.

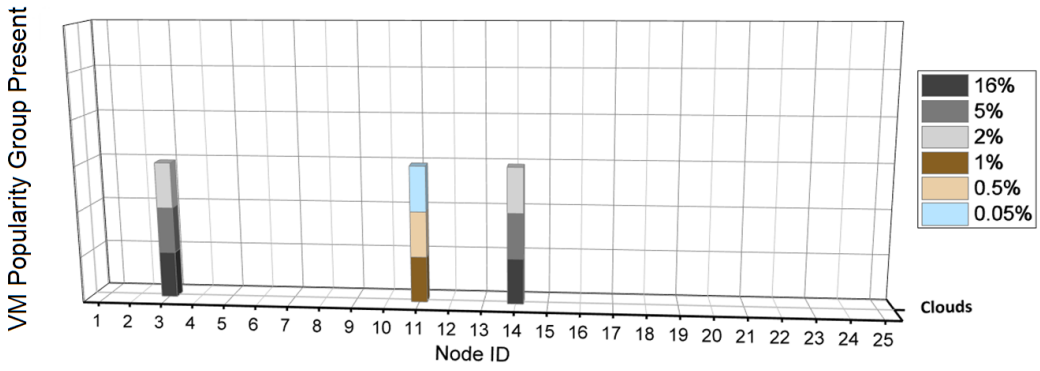

(a)

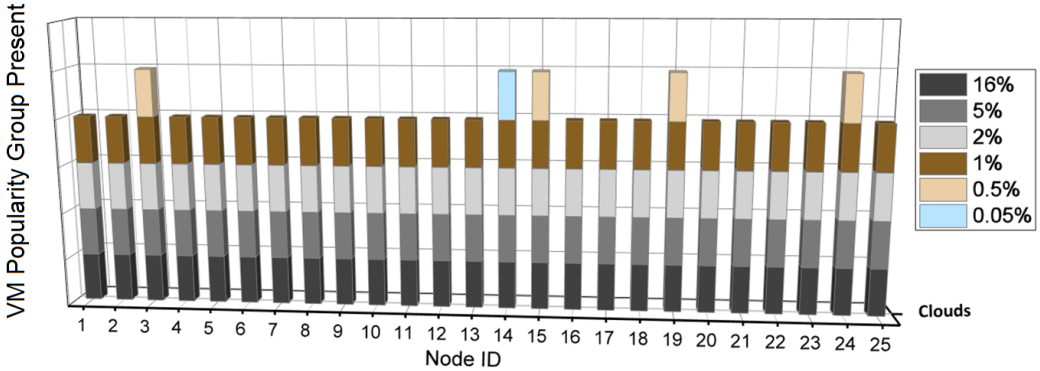

(b)

Figure 3-13: Optimal placement of different VMs popularity groups of 40% workload baseline under the OC approach with (a) 1 Mbps data rate per user (b) 25 Mbps data rate per user.

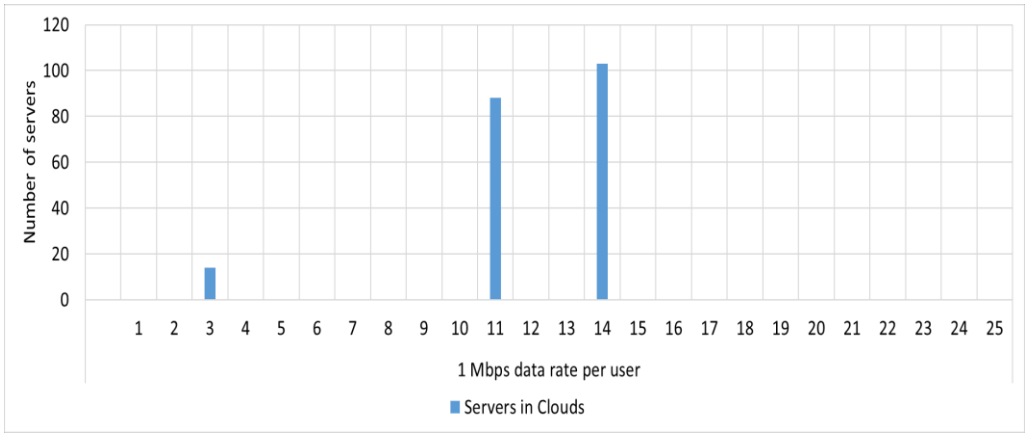

(a)

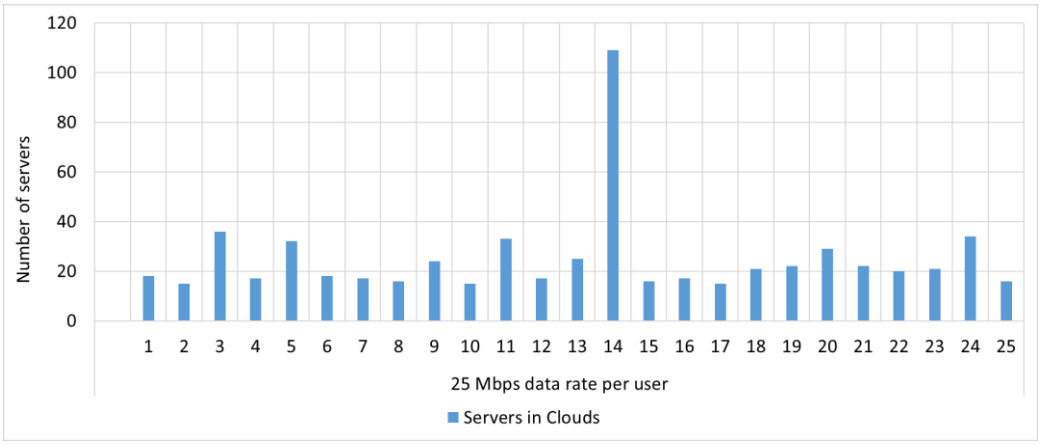

(b)

Figure 3-14: Number of servers required to host VMs of 40% workload baseline under the OC approach with (a) 1 Mbps data rate per user (b) 25 Mbps data rate per user.

3.3.2 BT Networks Use Case:

In this section, the impact of the core network topology on the VMs placement over the BT network topology [36] (illustrated in Fig. 3-15) is studied. The BT core network topology consists of 20 nodes and 68 bidirectional links. The networks and cloud architecture are considered to follow the same architecture of the previous subsection (illustrated in Fig 3.1).

According to Cisco VNI [157], in 2017, the average broadband data rate in the UK was 33.7 Mbps. Consequently, each OLT is assumed to be able to serve ~38k connections (or users). Cisco VNI also reported that in 2017, 75% of all UK Internet traffic had crossed clouds. SimilarWeb analytics tools [158] reports that the top 300 applications/websites have a share of 66% of all traffic. Accordingly, 18.8k users is considered in each PON (~66% of clouds traffic, i.e. 49% of the total traffic) to access

the VMs hosting the top 300 applications/websites. The popularity of these applications / websites VMs is considered to follow a Zipf distribution [40] with the same popularity groups discussed in the previous subsection. The same parameters of network and cloud components in previous subsection are considered.

The Optimised Clouds (OC) approach, is compared to the single clouds (SC) where all the VMs are placed in node 6 (City of London). Node 6 is selected to host the cloud in SC approach as major cloud operators base their central cloud in the UK in London (e.g. Microsoft Azure [45], Amazon AWS [46] and Google Cloud [47]). Similar to the previous section, linear VM workloads with 1%, 5% and 40% baseline are studied.

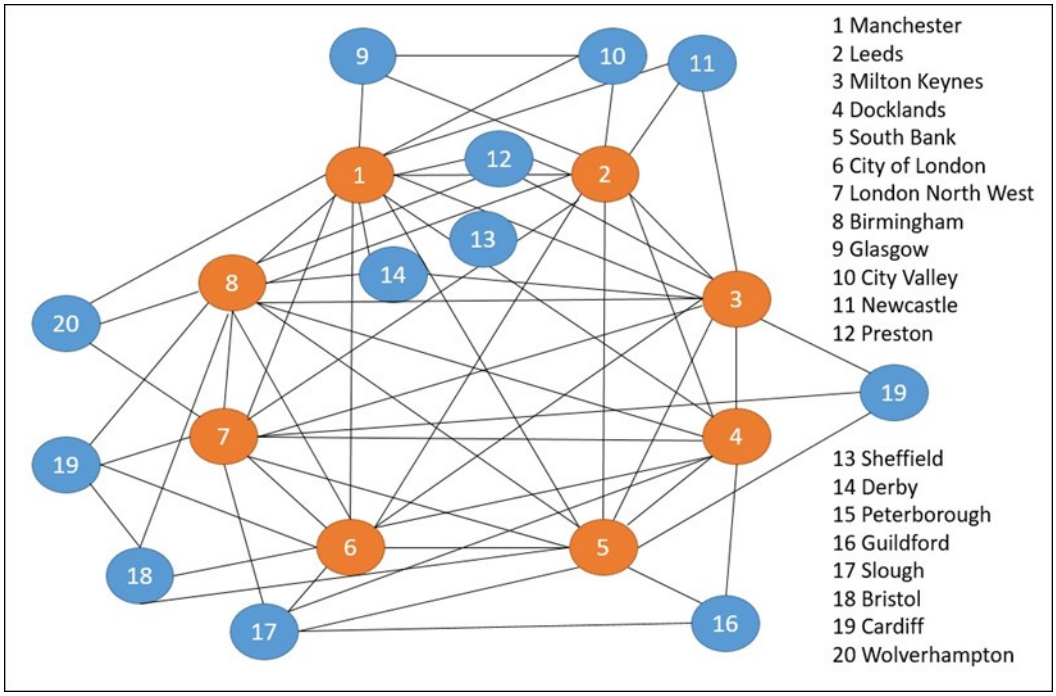

Figure 3-15: BT core network topology.

a) Linear Workload Profile (1% Workload Baseline):

Fig. 3-16 shows the power consumption resulting from placing VMs of 1% workload baseline considering the OC and SC approaches under 1, 10 and 25 Mbps users data rates. The power savings achieved by the OC approach compared to the SC approach are 2%, 26% and 37% under the low, medium and high data rates, respectively.

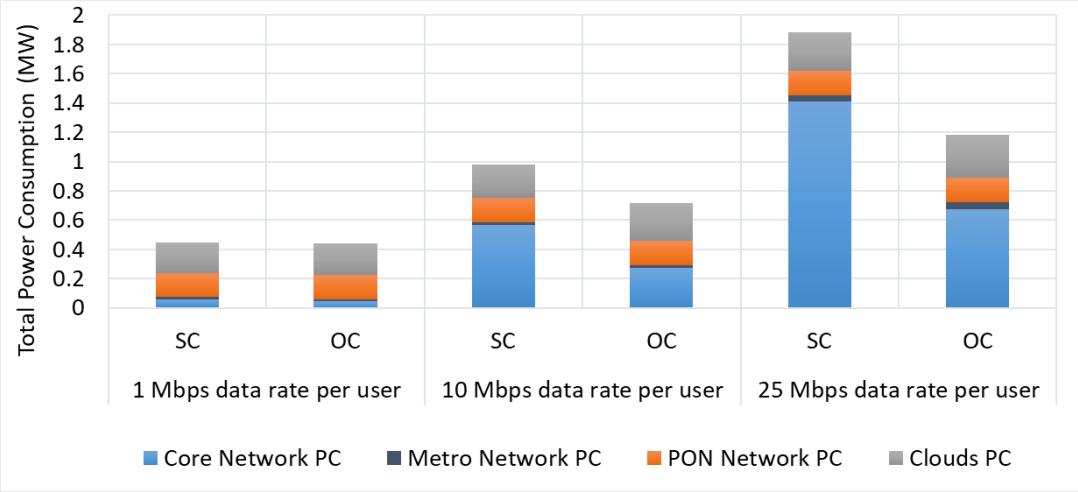

Figure 3-16: The power consumption of different VMs placement approaches considering VMs of 1% workload baseline.

Fig. 3-17 shows the optimal placement of different VMs popularity groups considering the OC approach. The efficient workload profile of VMs has justified the creation of 7 cloud locations under 1 Mbps user data rate. However, in these locations, VMs of 0.05% popularity have only justified the creation of replica copies only in a single node. The medium and high traffic of VMs with 10 and 25 Mbps data rates allow full replication for all VMs popularity groups across all cloud locations.

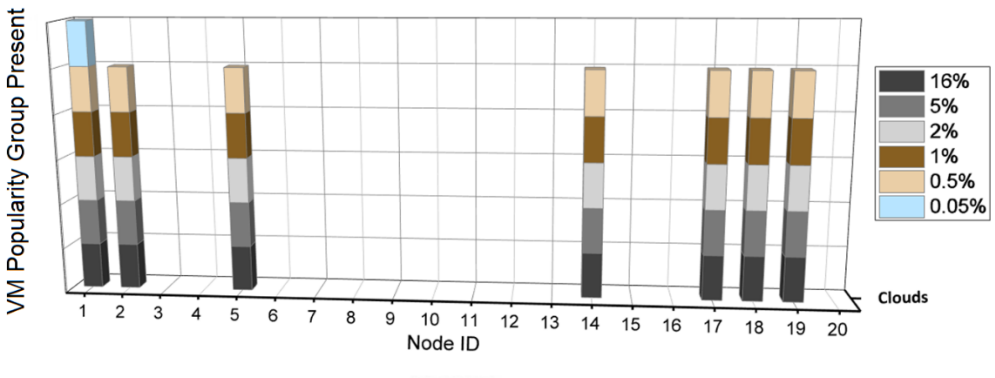

(a)

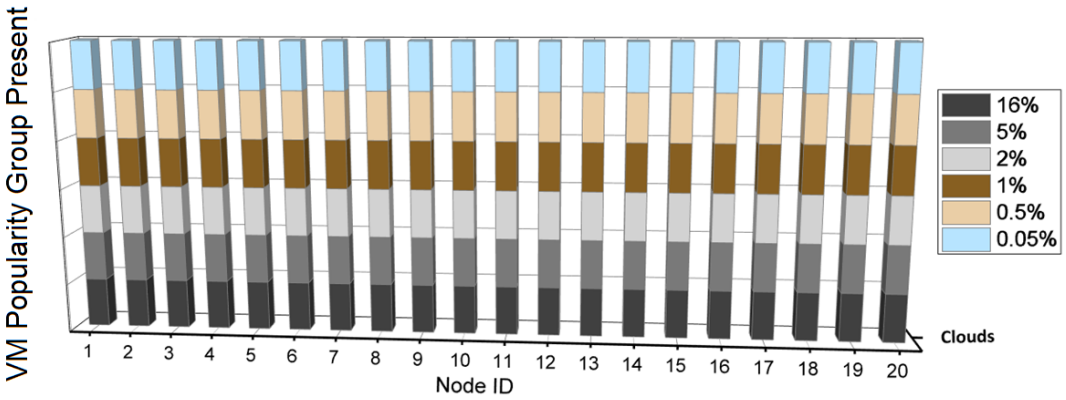

(b)

Figure 3-17: Optimal placement of different VMs popularity groups of 1% workload baseline under the OC approach with (a) 1 Mbps data rate per user (b) 10 and 25 Mbps data rates per user.

b) Linear Workload profile (5% Workload Baseline):

Fig. 3-18 shows the power consumption resulting from placing VMs of 5% workload baseline considering the two placement approaches under 1, 10 and 25 Mbps users data rates. The savings achieved by the OC approach compared to the SC approach are 5%, 25% and 34% under the low, medium and high data rates, respectively.

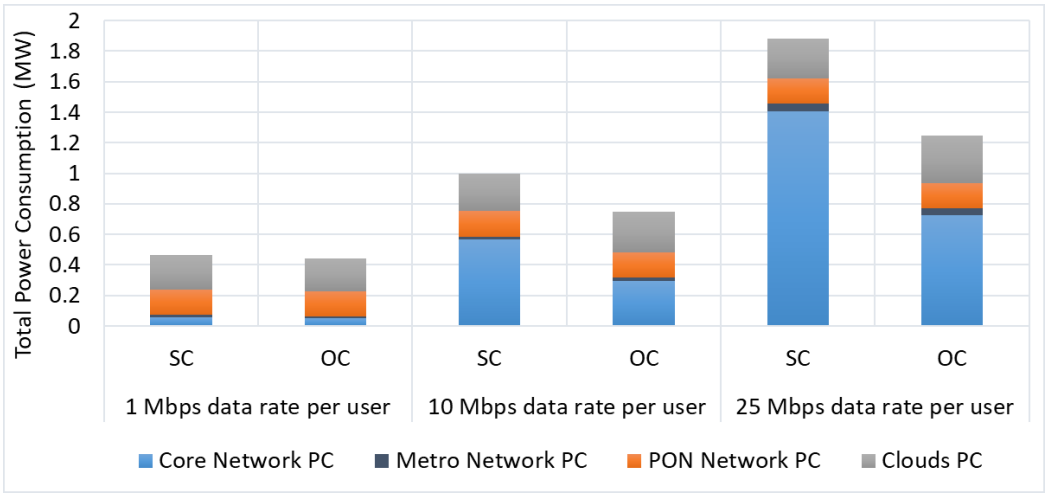

Figure 3-18: The power consumption of different VMs placement approaches considering VMs of 5% workload baseline.

Fig. 3-19 shows the optimal placement of different VM popularity groups of 5% workload baseline under the OC approach. VMs of 1 Mbps have justified the creation of 4 cloud locations under 1 Mbps user data rate. However, in these locations, VMs of 0.05% popularity have only justified the creation of a copy in a single node. The

medium and high traffic of VMs with 10 and 25 Mbps data rate allow full distribution for different popularity groups across all cloud locations except VMs with 0.05% popularity, which has been placed in 2 and 5 cloud locations, respectively.

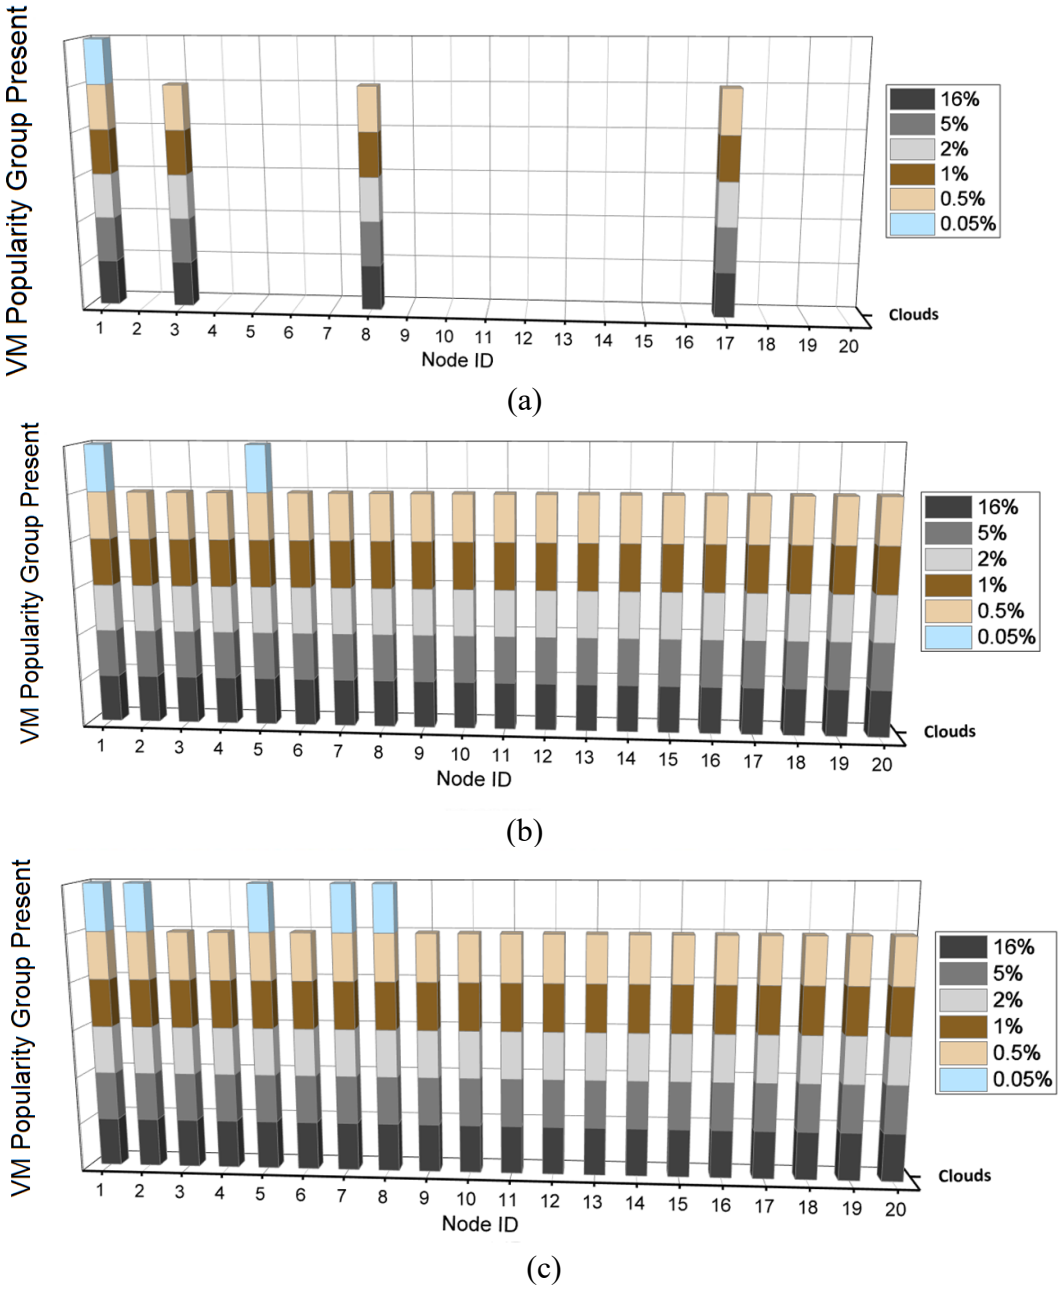

Figure 3-19: Optimal placement of different VMs popularity groups of 5% workload baseline under the OC approach with (a) 1 Mbps (b) 10 Mbps and (c) 25 Mbps data rate per user.

c) Linear Workload profile (40% workload baseline):

Fig. 3-20 shows the power consumption resulting from placing VMs of 40% workload baseline considering the two placement approaches under 1, 10 and 25 Mbps users data rates. Increasing the workload baseline to 40% has resulted in

increasing the total power consumption compared to 1% and 5% workload baseline scenarios. The savings achieved by the OC approach compared to the SC approach are 2%, 16% and 26% under the low, medium and high data rates, respectively.

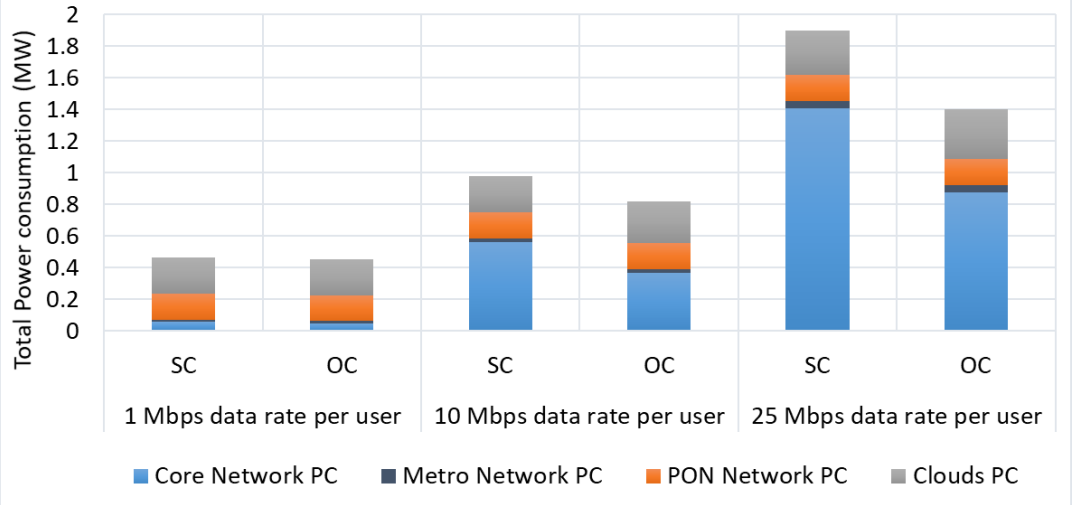

Figure 3-20: The power consumption of different VMs placement approaches considering VMs of 40% workload baseline.

Fig. 3-21 shows the optimal placement of different VM popularity groups of 40% workload baseline under the OC approach. VMs of 1 Mbps have justified the creation of 4 cloud locations under 1 Mbps user data rate. However, in these locations, VMs of 0.05% popularity have only justified the creation of replica copies in a single node. The medium and high data rate of VMs with 10 and 25 Mbps data rate allow full distribution for VMs with $\geq$ **2**% and $\geq$ **1**%, respectively, whereas, VMs of less popularity have been replicated to few locations in distributed clouds over IP over WDM networks.

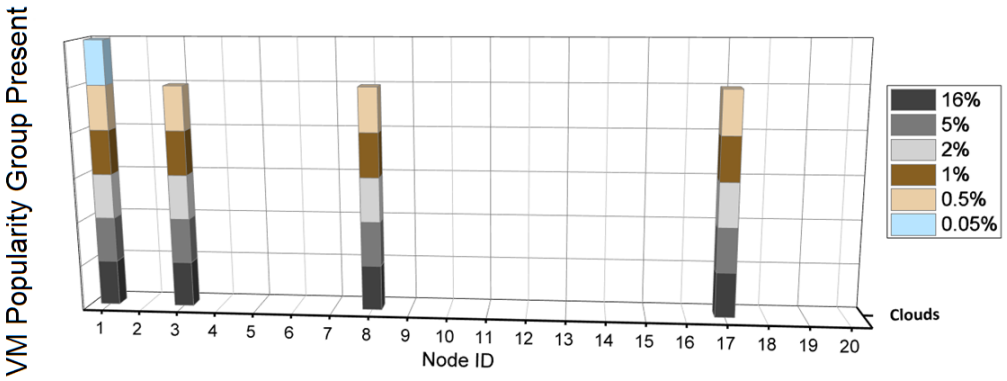

(a)

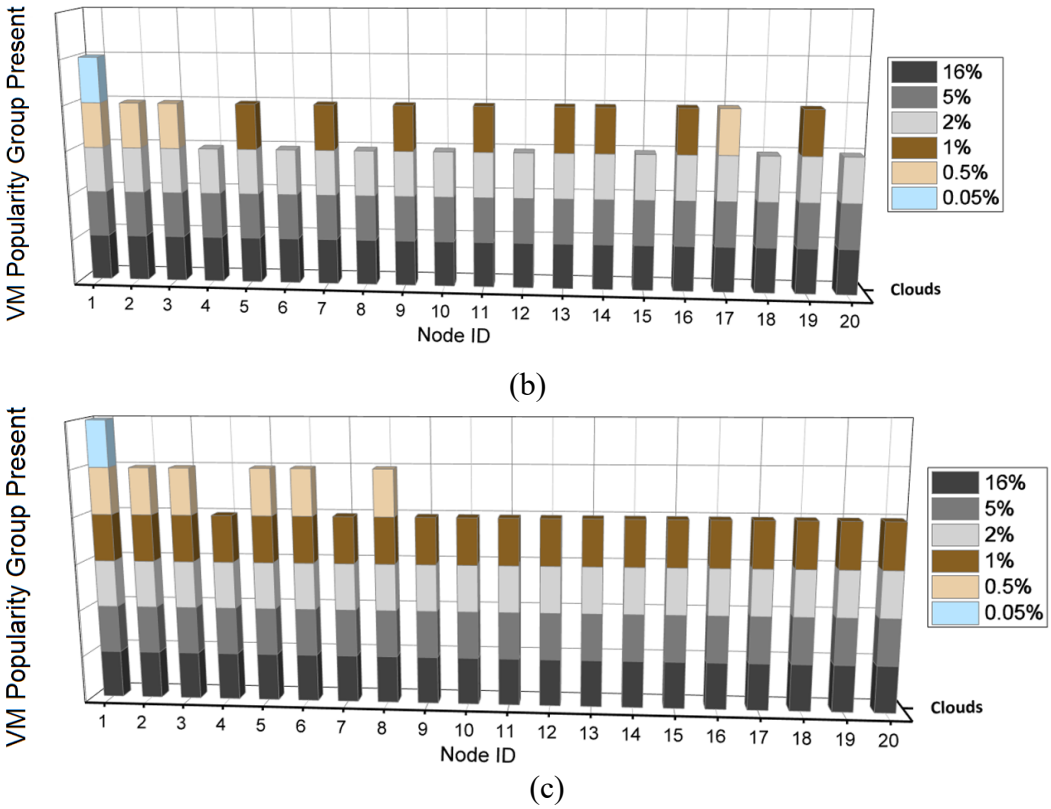

Figure 3-21: Optimal placement of different VMs popularity groups of 40% workload baseline under the OC approach with (a) 1 Mbps (b) 10 Mbps and (c) 25 Mbps data rate per user.

The VMs placement in the BT networks discussed above highlights the same trends observed in the AT&T network. The longer distance between AT&T nodes does not result in high power consumption as regenerators and EDFAs consume relatively low power.

3.4 EEVM-C Heuristic

The VMs placement over distributed clouds is NP-Hard. Therefore, the MILP model cannot be used to obtain solutions in real time. A real time heuristic referred to as energy-efficient VMs placement heuristic for the distributed clouds (EEVM-C) is developed to mimic the MILP model behaviour. The optimal solutions obtained from the MILP model can offer a benchmark to evaluate the performance of the heuristics developed. The reasons for adopting this implementation are to (i) establish a classification criterion for different VMs based on their different communication and

processing requirements and (ii) reduce the real-time implementation required to find the optimal VM placement in a real-time implementation.

A supervised learning algorithm is adopted in this heuristic. Supervised learning is a branch of machine learning where an input is matched with an output based on a sample of input-output pairs. VMs are classified into different types based on user download rates, VM workloads and VMs popularity. The optimum placement of different VM types are found in an offline phase. VMs are matched to their type in real time (online phase) and placed according to the placement obtained in the offline.

The flowchart of the offline phase is shown in Fig. 3-22 (a). The offline phase starts by classifying VMs based on their popularity, CPU usage and user data rate. The most energy efficient placement for each possible numbers of replicas, i.e. the most energy efficient placement for 1 replica, 2 replicas…N replicas (N is the number of nodes in the network), are found through an exhaustive search over all the possible placements for this number of replicas. These will create a search space, defined as $\boldsymbol{P}$, to find the optimum placement for different VMs types. Note that the search space is a function of the network topology not the VM type. For each VM type, each placement in P is examined. The traffic between users and VMs resulting from this placement are routed based on minimum hop routing and the workload of VMs of a linear workload profile is calculated based on the number of users each VM serves. The network power consumption and the cloud power consumption are calculated and the optimum placement of a VM type is the placement that results in the minimum total power consumption.

In the online phase, each VM is matched to a type and placed according to the placement obtained in the offline phase. Traffic demands resulting from this placement are routed on the core network using multi-hop routing [19] and workload of clouds where the VM replicas are placed is updated. After placing all VMs, the total power consumption of the distributed clouds is calculated.

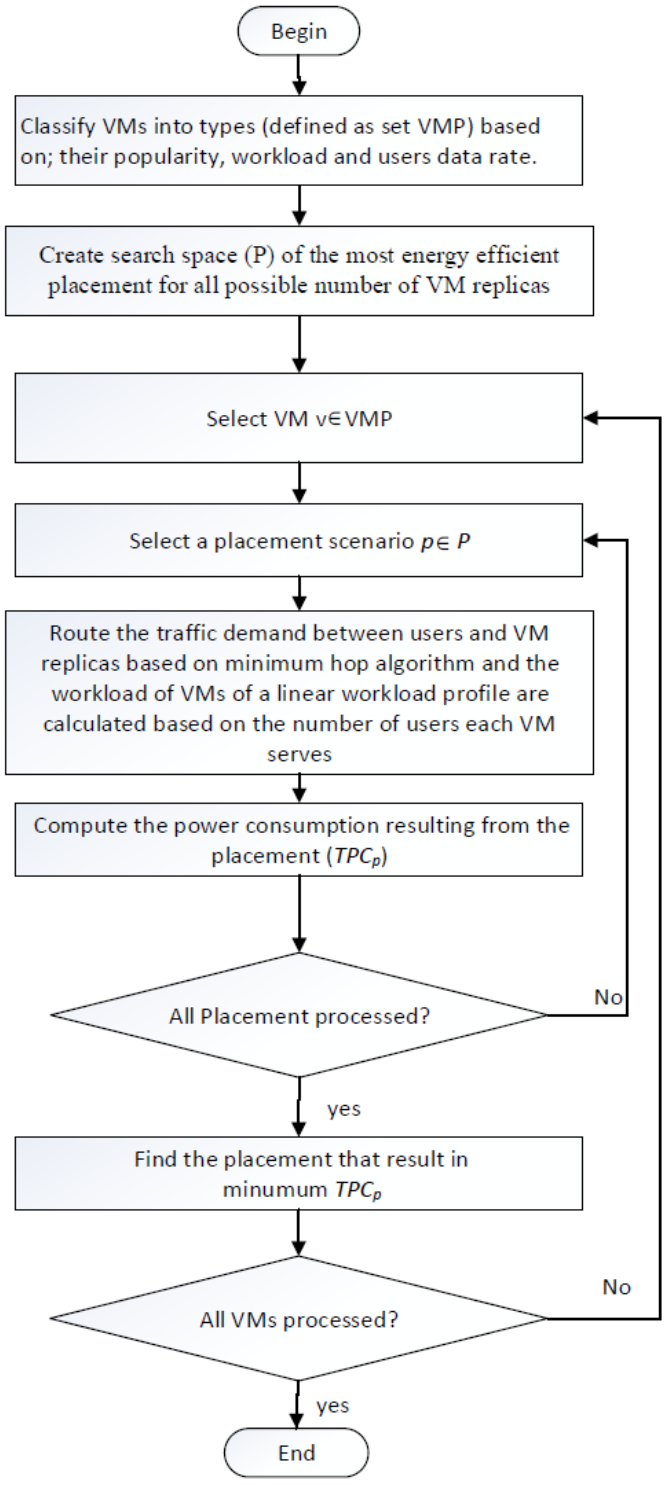
Begin
Classify VMs into types (defined as set VMP) based on; their popularity, workload and users data rate.
Create search space (P) of the most energy efficient placement for all possible number of VM replicas
Select VM v∈VMP
Select a placement scenario p∈ P
Route the traffic demand between users and VM replicas based on minimum hop algorithm and the workload of VMs of a linear workload profile are calculated based on the number of users each VM serves
Compute the power consumption resulting from the placement (TPC_p)
All Placement processed?
No
yes
Find the placement that result in minumum TPC_p
All VMs processed?
No
yes
End

(a)

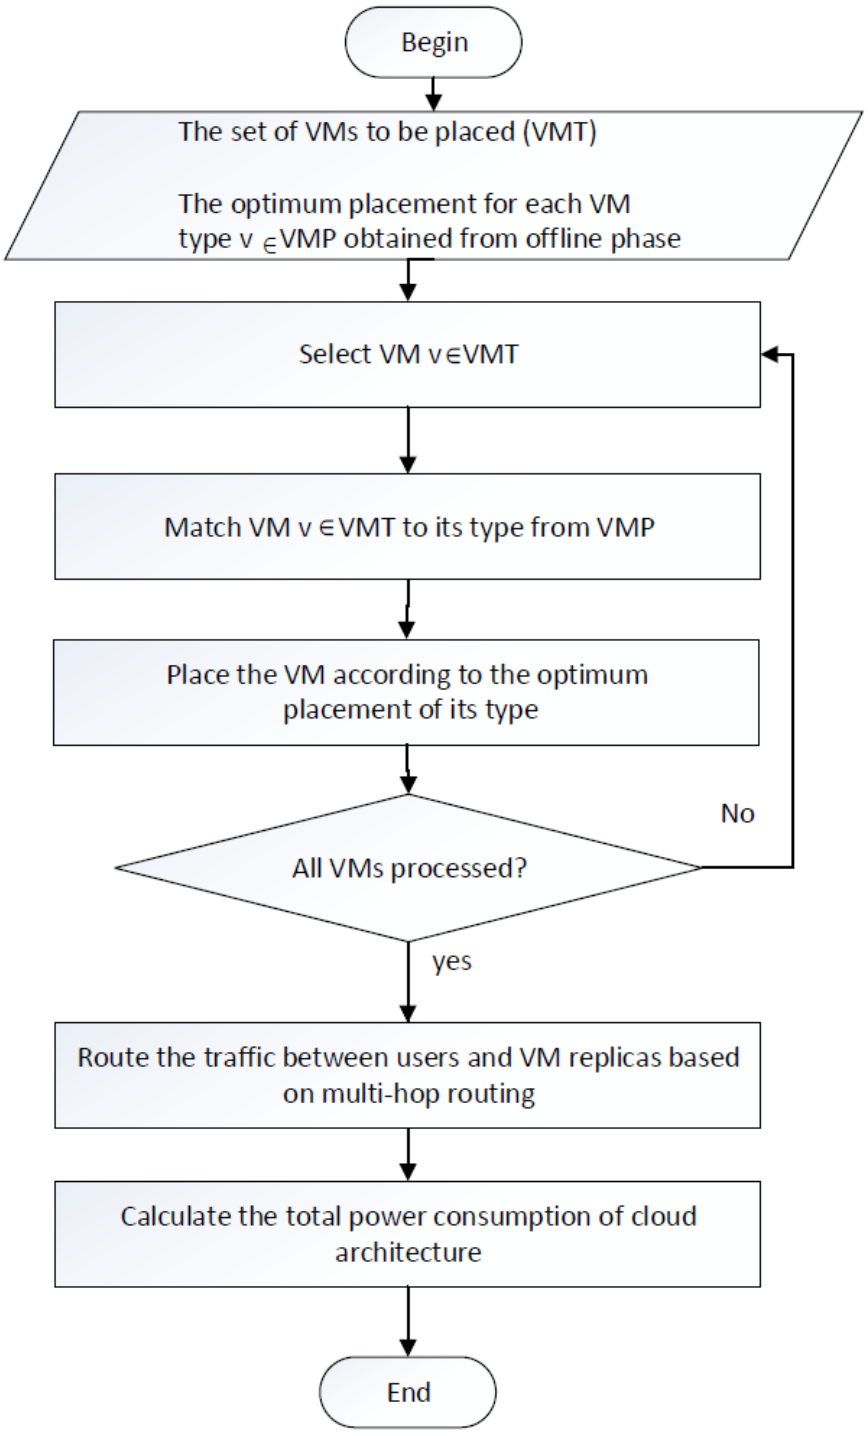

(b)

Figure 3-22: Flow chart of (a) offline phase and (b) online phase of EEVM-C heuristic algorithm.

Fig. 3-23 compares the total power consumption resulting from placing VMs using the EEVM-C heuristic with the MILP model placement considering the AT&T core network use case. Using 2.5 GHz Intel Core i7 PC with 16 GB Memory, the EEVM-C heuristic took 51 seconds to run the offline phase and 2 seconds to run the online

phase. The heuristic is evaluated under 1%, 5% and 40% workload baseline as well as 1 Mbps, 10 Mbps and 25 Mbps user data rates. In comparison, the MILP took 1-2 days to run.

Clearly, the power consumption of the MILP and the EEVM-CF are comparable. The gap between the EEVM-C and the MILP model is less than 2% of the total power consumption. This small gap is due to the multi hop non-bypass in the EEVM-C heuristic which is less efficient than the MILP model in routing traffic.

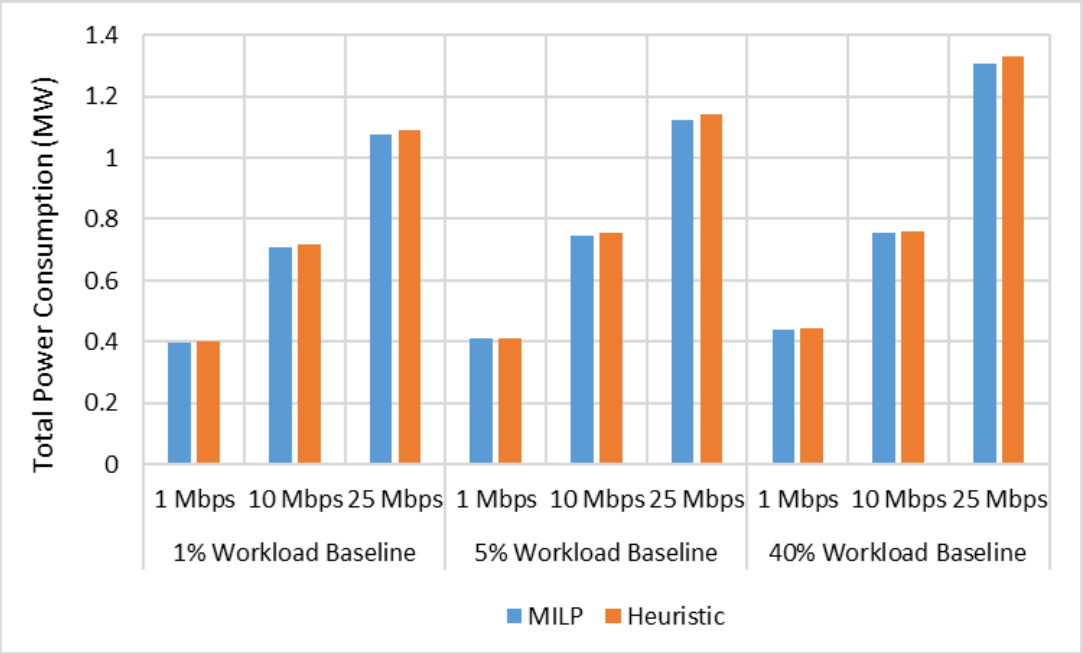

Figure 3-23: Total power consumption of MILP model compared with EEVM-C heuristic considering VMs linear workload profile with 1%, 5% and 40% CPU workload baseline and 1, 10 and 25 Mbps users data rates.

3.5 Summary

This chapter has presented a framework to find the optimum placement of VMs over distributed cloud datacentres hosted in IP over WDM networks. The optimisation is performed using a Mixed Integer Linear Programming (MILP) model considering AT&T and BT networks as use case scenarios. The results showed VMs with linear workload profile can be optimally replicated to distributed clouds as such VMs can be sliced efficiently (CPU load is linear function of number of users) to smaller replicas based on the served number of users. In contrast, replicating copies of a VM of constant linear workload results in high power consumption and therefore limits

the number of replicas. Also, it can be observed that at low users data rate, the VMs are optimally placed in one or few clouds, whereas at high data rates, multiple replicas of each VM are distributed over the distributed clouds to serve users either locally or within a single hop to reduce the network power consumption resulting from high user data rates.

In general, the results show that high traffic of highly popular VMs and VMs with high user data rates make network power consumption more important and consequently determine the placement. The results also indicate the tendency to distribute VMs with linear workload profile compared to VMs with constant workload. Furthermore, the results show that at high PUE, processing power consumption becomes more dominant, which makes it play a bigger role in VMs placement.

The total power savings achieved are up to 51% and 38%, compared to the power consumption of the traditional cloud locations in AT&T and BT core network topologies, respectively. Based on the model insights, an energy-efficient VM placement heuristic for the distributed cloud (EEVM-C) is developed with comparable energy efficiency to the MILP results.

Chapter 4: The Impact of Inter-Virtual Machines Traffic on Energy Efficient Virtual Machines Placement

4.1 Introduction

The inter-VM traffic is a major contributor to the east-west traffic (server to server traffic) which is expected to be responsible for 85% of the global cloud traffic by 2021 [62]. Designing energy efficient cloud datacentres requires the co-optimisation of both north-south traffic (user-datacentre traffic) and east-west traffic (inter datacentre traffic). For example, migrating an application VM, which has an inter-traffic with a database VM, to another datacentre in order to satisfy the increasing users demand may increase the burden on the networking infrastructure connecting datacentres.

This chapter considers the problem of optimising the placement of VMs in geo-distributed clouds in IP over WDM core networks, as seen in Fig. 4-1, so that the total power consumption is minimised. The investigation takes into consideration the cooperation traffic between different VMs and synchronisation traffic between replicas of the same VM in addition to the download traffic from VMs to users. The problem is formulised as a MILP model.

For cooperation traffic, in addition to sending traffic to users, each VM has another VM to cooperate with (e.g. traffic between an application VM and a database VM). The objective is to optimise the placement of VMs on distributed clouds and to create an optimum number of replicas that result in minimum power consumption. For synchronisation traffic, VM replicas can be created to serve distributed users, however, these replicas need to be synchronised to each other to keep the content at each location up to date (e.g. social media where a user creates a post/webpage and posts it in the nearest VM replica, then, all other replicas need to be synchronised).

A heuristic is also developed to validate the model results and provide real time solutions.

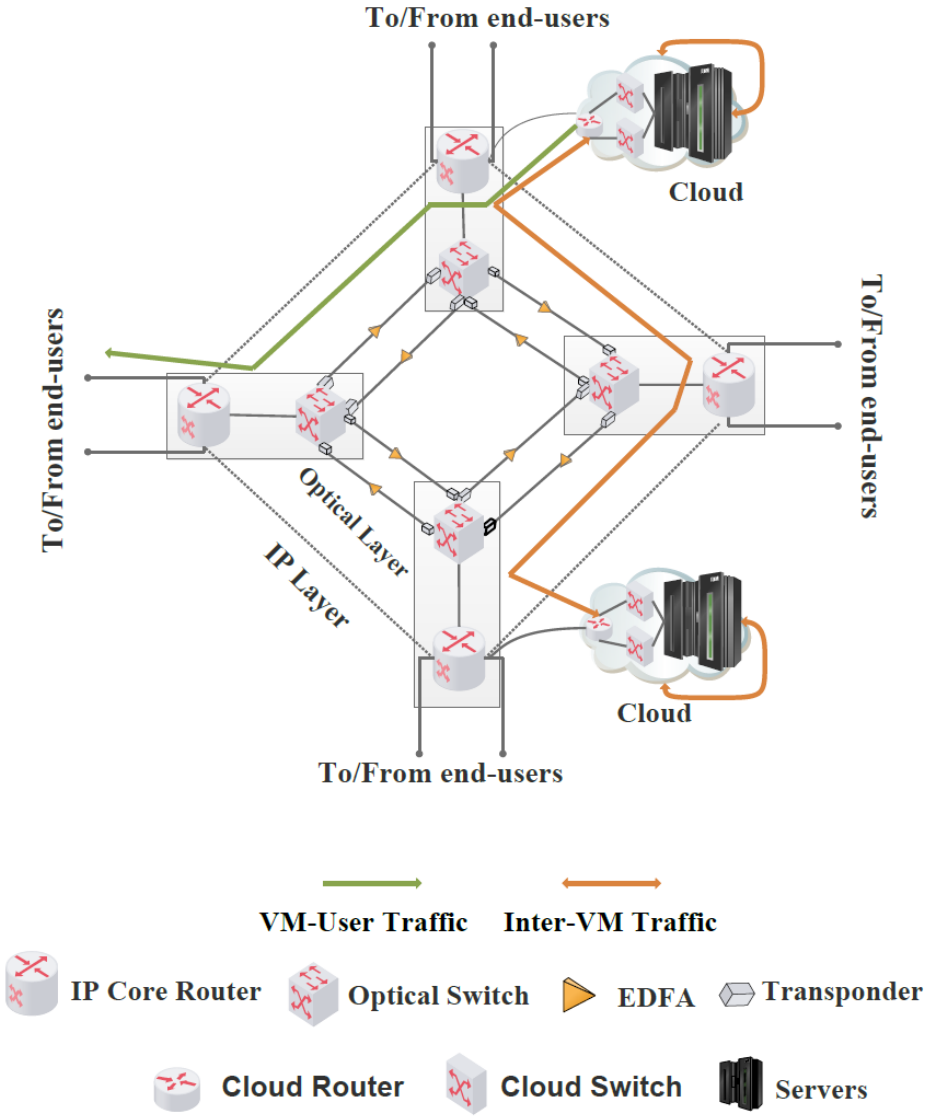

Figure 4-1: Geo-distributed clouds over an IP over WDM network with an illustration of VM-users traffic and inter-VM traffic

4.2 MILP Model

In this section, we extend the model developed in the previous chapter. The model in the previous chapter optimises the placement of VMs so the total power consumption is minimised considering download traffic between users and VMs only. In this chapter, the model is extended to study the impact of inter-VM traffic on the energy efficient placement of VMs considering two CPU workload profiles; constant profile or linear profile.

In addition to the parameters, variables and constraints introduced in Chapter 3, the parameters and variables in Tables 4-1 and Table 4-2 are introduced to model the inter-VM traffic.

Table 4-1: List of parameters used in the MILP model to model inter-VM traffic

Parameter	Description
$\boldsymbol{v}$ **and** $\boldsymbol{y}$	Indices of VMs.
$\boldsymbol{CV_{v,y}}$	$\boldsymbol{CV_{v,y}}$ =1, if there is cooperating traffic between VM $\boldsymbol{v}$ and VM $\boldsymbol{y}$, otherwise $\boldsymbol{CV_{v,y}}$ =0.
$\boldsymbol{CVM_{v,y}}$	Cooperation traffic between VM $\boldsymbol{v}$ and VM $\boldsymbol{y}$.
$\boldsymbol{SVM_{v}}$	Synchronisation traffic between VM $\boldsymbol{v}$ replicas.

Table 4-2: List of variables used in the MILP model to model inter-VM traffic

Parameter	Description
$\Psi_{\boldsymbol{d}}$	Download traffic of cloud hosted in node $\boldsymbol{d}$.
$\Phi_{\boldsymbol{v,y,s,d}}$	$\Phi_{\boldsymbol{v,y,s,d}} = \mathbf{1}$, if VM $\boldsymbol{v}$ located in node $\boldsymbol{s}$ is a candidate to cooperate with VM $\boldsymbol{y}$ located in node $\boldsymbol{d}$, otherwise $\Phi_{\boldsymbol{v,y,s,d}} = \mathbf{0}$.
$б_{\boldsymbol{v,y,d}}$	$б_{\boldsymbol{v,y,d}} = \mathbf{1}$, if cooperation traffic exists from VM $\boldsymbol{v}$ located at any node to VM $\boldsymbol{y}$ located in node $\boldsymbol{d}$, otherwise $б_{\boldsymbol{v,y,d}} = \mathbf{0}$.
$\boldsymbol{\chi_{v,y,s,d}}$	$\boldsymbol{\chi_{v,y,s,d}} = \mathbf{1}$, if cooperation traffic exists from VM $\boldsymbol{x}$ located in node $\boldsymbol{s}$ to VM $\boldsymbol{y}$ located in node $\boldsymbol{d}$, otherwise $\boldsymbol{\chi_{v,y,s,d}} = \mathbf{0}$.
$\boldsymbol{\beta_{v,y,s,d}}$ $\boldsymbol{\alpha_{v,y,s,d}}$	Binary variables set to 1 only if one or two of the following conditions are satisfied; there is a cooperating traffic from VM $\boldsymbol{v}$ to VM $\boldsymbol{y}$, VM $\boldsymbol{v}$ is located in node $\boldsymbol{s}$ or VM $\boldsymbol{y}$ is located in node $\boldsymbol{d}$, otherwise $\boldsymbol{\beta_{v,y,s,d}}$ and $\boldsymbol{\alpha_{v,y,s,d}}$ are set to 0.
$\boldsymbol{IC_{v,y,s,d}}$	Cooperating traffic from VM $\boldsymbol{v}$ to VM $\boldsymbol{y}$ located in nodes $\boldsymbol{s}$ and $\boldsymbol{d}$.
$\Theta_{\boldsymbol{v,s,d}}$	$\Theta_{\boldsymbol{v,s,d}} = \mathbf{1}$, if VM $\boldsymbol{v}$ replicas are located in nodes $\boldsymbol{s}$ and $\boldsymbol{d}$, respectively, otherwise $\Theta_{\boldsymbol{v,s,d}} = \mathbf{0}$.
$\boldsymbol{\varphi_{v,s,d}}$	$\boldsymbol{\varphi_{v,s,d}}$ = 1, if only one VM $\boldsymbol{v}$ replica is located in either node $\boldsymbol{s}$ or node $\boldsymbol{d}$, otherwise $\boldsymbol{\varphi_{v,s,d}}$= 0.
$\boldsymbol{IS_{v,s,d}}$	Synchronisation traffic between VM $\boldsymbol{v}$ replicas located in nodes $\boldsymbol{s}$ **and** $\boldsymbol{d}$, respectively.
$\boldsymbol{R_d^{(IV)}}$	Number of router ports in node $\boldsymbol{d}$ that aggregate the traffic to clouds.

Under the non-bypass approach [19], the IP over WDM network power consumption is composed of (below we re-introduce some of the equations from our model developed in Section 3.2 for completeness):

- The power consumption of router ports:

$$\boldsymbol{n}\left(\sum_{\boldsymbol{s}\in\boldsymbol{N}}\boldsymbol{R}^{(\boldsymbol{P})}\,\boldsymbol{R}_{\boldsymbol{s}}^{(\boldsymbol{AC})}+\sum_{\boldsymbol{d}\in\boldsymbol{N}}\boldsymbol{R}^{(\boldsymbol{P})}\left(\boldsymbol{R}_{\boldsymbol{d}}^{(\boldsymbol{AE})}+\boldsymbol{R}_{\boldsymbol{d}}^{(\boldsymbol{IV})}\right)+\sum_{\boldsymbol{m}\in\boldsymbol{N}}\sum_{\boldsymbol{n}\in\boldsymbol{N}\boldsymbol{m}_{\boldsymbol{m}}:\boldsymbol{n}\neq\boldsymbol{m}}\boldsymbol{R}^{(\boldsymbol{P})}\,\boldsymbol{\mathcal{W}}_{\boldsymbol{mn}}\right)\quad(\mathbf{4.1})$$

Equation (4.1) calculates the power consumption of router ports including the power consumption of router ports in each core node that aggregate the traffic from the clouds ($\sum_{s\in N}\boldsymbol{R}^{(P)}\,\boldsymbol{R}_{s}^{(AC)}$), the power consumption of router ports in each core node that aggregate the traffic from/to metro routers ($\sum_{d\in N}\boldsymbol{R}^{(P)}\boldsymbol{R}_{d}^{(AE)}$), the power consumption of router ports in each core node that aggregate the traffic to clouds due to inter-VM traffic ($\sum_{d\in N}\boldsymbol{R}^{(P)}\boldsymbol{R}_{d}^{(IV)}$) and the power consumption of intermediate router ports that transmit the traffic between the source and destination nodes ($\sum_{m\in N}\sum_{n\in N m_m:n\neq m}\boldsymbol{R}^{(P)}\,\boldsymbol{\mathcal{W}}_{mn}$).

Where total number of aggregation ports in a core node due to cloud download traffic calculated as:

$$\boldsymbol{R}_{\boldsymbol{d}}^{(\boldsymbol{IV})}=\frac{\mathbf{1}}{\boldsymbol{B}}\;\;\Psi_{\boldsymbol{d}}$$

$$\forall\,\boldsymbol{d}\in\boldsymbol{N}\qquad(\mathbf{4.2})$$

- The power consumption of transponders:

$$\boldsymbol{n}\left(\sum_{\boldsymbol{m}\in\boldsymbol{N}}\sum_{\boldsymbol{n}\in\boldsymbol{N}\boldsymbol{m}_{\boldsymbol{m}}:\boldsymbol{n}\neq\boldsymbol{m}}\boldsymbol{t}^{(\boldsymbol{P})}\,\boldsymbol{\mathcal{W}}_{\boldsymbol{mn}}\right)\qquad(\mathbf{4.3})$$

- The power consumption of EDFAs:

$$\boldsymbol{n}\left(\sum_{\boldsymbol{m}\in\boldsymbol{N}}\sum_{\boldsymbol{n}\in\boldsymbol{N}\boldsymbol{m}_{\boldsymbol{m}}:\boldsymbol{n}\neq\boldsymbol{m}}\boldsymbol{e}^{(\boldsymbol{P})}\,\boldsymbol{F}_{\boldsymbol{mn}}\,\boldsymbol{A}_{\boldsymbol{mn}}\right)\qquad(\mathbf{4.4})$$

- The power consumption of optical switches:

$$\boldsymbol{n}\left(\sum_{\boldsymbol{s}\in\boldsymbol{N}}\boldsymbol{SW}_{\boldsymbol{s}}^{(\boldsymbol{P})}\right)\qquad(\mathbf{4.5})$$

Equation (4.5) calculates the power consumption of of optical switches located in each node $\boldsymbol{s}$. Each node IP routers are connected to optical switches which are

connected by optical fibre links. The optical layer (including optical switches and fibre links) provides the large bandwidth required for communication between IP routers.

- The power consumption of regenerators:

$$n\left(\sum_{m\in N}\sum_{n\in N m_m : n\neq m} G^{(P)}\, G_{m,n}\, \mathcal{W}_{m,n}\right) \qquad (4.6)$$

The clouds power consumption is composed of:

1) Power consumption of cloud servers:

$$c\sum_{s\in N} S_s^{(C)}\, S^{(P)} \qquad (4.7)$$

2) Power consumption of cloud routers and switches:

$$c\left(\left(R^{(CP)}\frac{1}{R^{(CB)}}\left(\sum_{s\in N}\sum_{d\in N} L_{s,d}+\sum_{d\in N}\Psi_d\right)\right)\right.$$
$$+\ SW^{(R)}\, SW^{(CP)}\frac{1}{SW^{(CB)}}\left(\sum_{s\in N}\sum_{d\in N:s\neq d} L_{s,d}+\sum_{d\in N}\Psi_d\right)$$
$$\left.+\left(SW^{(R)}\, SW^{(CP)}\frac{1}{SW^{(CB)}}\sum_{s\in N}\sum_{d\in N:s=d}\sum_{v\in VM}\sum_{y\in VM} IC_{v,y,s,d}\right)\right) \qquad (4.8)$$

Where the cloud traffic is calculated as:

Cloud upload traffic:

$$L_{s,d}=\sum_{v\in VM} D_{v,s,d}^{(C)}+\sum_{v\in VM}\sum_{y\in VM} IC_{v,y,s,d}+\sum_{v\in VM} IS_{v,s,d}$$
$$\forall\, s,d\in N,\qquad s\neq d \qquad (4.9)$$

Cloud download traffic:

$$\Psi_d = \sum_{s\in N}\sum_{v\in VM}\sum_{y\in VM} IC_{v,y,s,d} + \sum_{s\in N}\sum_{v\in VM} IS_{v,s,d}$$

$$\forall\, d \in N\,, s \neq d \qquad (4.10)$$

The model is defined as follows:

The objective: Minimise total power consumption given as:

$$\begin{aligned}
& n\left(\sum_{s\in N} R^{(P)}\, R_s^{(AC)} + \sum_{d\in N} R^{(P)}\left(R_d^{(AE)} + R_d^{(IV)}\right) + \sum_{m\in N}\sum_{n\in N m_m : n\neq m} R^{(P)}\, W_{mn}\right) \\
& + n\left(\sum_{m\in N}\sum_{n\in N m_m : n\neq m} t^{(P)}\, W_{mn}\right) \\
& + n\left(\sum_{m\in N}\sum_{n\in N m_m : n\neq m} e^{(P)}\, F_{mn}\, A_{mn}\right) \\
& + n\left(\sum_{s\in N} SW_s^{(P)}\right) \\
& + n\left(\sum_{m\in N}\sum_{n\in N m_m : n\neq m} G^{(P)}\, G_{m,n}\, W_{m,n}\right) \\
& + c\sum_{s\in N} S_s^{(C)}\, S^{(P)} \\
& + c\left(\left(R^{(CP)}\frac{1}{R^{(CB)}}\left(\sum_{s\in N}\sum_{d\in N} L_{s,d} + \sum_{d\in N}\Psi_d\right)\right)\right. \\
& + SW^{(R)}\, SW^{(CP)}\frac{1}{SW^{(CB)}}\left(\sum_{s\in N}\sum_{d\in N : s\neq d} L_{s,d} + \sum_{d\in N}\Psi_d\right) \\
& \left. + \left(SW^{(R)}\, SW^{(CP)}\frac{1}{SW^{(CB)}}\sum_{s\in N}\sum_{d\in N : s=d}\sum_{v\in VM}\sum_{y\in VM} IC_{v,y,s,d}\right)\right)
\end{aligned}$$

$$(4.11)$$

Equation (4.11) calculates the total power consumption of IP over WDM networks and clouds.

Subject to: (3-17) – (3-22), (3-22) – (3-29), and the following:

Traffic flow between cooperating VMs:

$$\mathbf{3}\ \Phi_{v,y,s,d} + \boldsymbol{\alpha}_{v,y,s,d} + \boldsymbol{\beta}_{v,y,s,d} = \boldsymbol{\delta}_{v,s}^{(C)} + \boldsymbol{\delta}_{y,d}^{(C)} + \boldsymbol{CV}_{v,y}$$

$$\forall \boldsymbol{s,d} \in \boldsymbol{N}, \boldsymbol{v,y} \in \boldsymbol{VM}: \boldsymbol{v} \neq \boldsymbol{y} \qquad (\mathbf{4.12})$$

$$\sum_{\boldsymbol{s}\in\boldsymbol{N}} \Phi_{v,y,s,d} \geq б_{v,y,d}$$

$$\forall\ \boldsymbol{d} \in \boldsymbol{N}, \boldsymbol{v,y} \in \boldsymbol{VM}: \boldsymbol{v} \neq \boldsymbol{y} \qquad (\mathbf{4.13})$$

$$\sum_{\boldsymbol{s}\in\boldsymbol{N}} \Phi_{v,y,s,d} \leq \boldsymbol{L}\ б_{v,y,d}$$

$$\forall\ \boldsymbol{d} \in \boldsymbol{N}, \boldsymbol{v,y} \in \boldsymbol{VM}: \boldsymbol{v} \neq \boldsymbol{y} \qquad (\mathbf{4.14})$$

$$\sum_{\boldsymbol{s}\in\boldsymbol{N}} \boldsymbol{\chi}_{v,y,s,d} = б_{v,y,d}$$

$$\forall\ \boldsymbol{d} \in \boldsymbol{N}, \boldsymbol{v,y} \in \boldsymbol{VM}: \boldsymbol{v} \neq \boldsymbol{y} \qquad (\mathbf{4.15})$$

$$\Phi_{v,y,s,d} \geq \boldsymbol{\chi}_{v,y,s,d}$$

$$\forall\ \boldsymbol{d} \in \boldsymbol{N}, \boldsymbol{v,y} \in \boldsymbol{VM}: \boldsymbol{v} \neq \boldsymbol{y} \qquad (\mathbf{4.16})$$

$$\boldsymbol{IC}_{v,y,s,d} = \boldsymbol{\chi}_{v,y,s,d}\ \boldsymbol{CVM}_{v,y}$$

$$\forall \boldsymbol{s,d} \in \boldsymbol{N}: \boldsymbol{s} \neq \boldsymbol{d}, \boldsymbol{v,y} \in \boldsymbol{VM}: \boldsymbol{v} \neq \boldsymbol{y} \qquad (\mathbf{4.17})$$

$$\boldsymbol{IC}_{v,y,s,d} = \boldsymbol{\chi}_{v,y,s,d}\ \boldsymbol{CVM}_{v,y}$$

$$\forall \boldsymbol{s,d} \in \boldsymbol{N}: \boldsymbol{s} = \boldsymbol{d}, \qquad \boldsymbol{v,y} \in \boldsymbol{VM}: \boldsymbol{v} \neq \boldsymbol{y} \qquad (\mathbf{4.18})$$

Constraints (4.12) to (4.18) represent the traffic flow between different

cooperating VMs ($\boldsymbol{v \neq y}$). Constraint (4.12) ensures that $\Phi_{\boldsymbol{v,y,s,d}} = \mathbf{1}$ if VM $\boldsymbol{v}$ is located in node $\boldsymbol{s}$ (i.e. $\boldsymbol{\delta}_{\boldsymbol{v,s}}^{(\boldsymbol{C})}$), VM $\boldsymbol{y}$ is located in node $\boldsymbol{d}$ (i.e. $\boldsymbol{\delta}_{\boldsymbol{v,d}}^{(\boldsymbol{C})}$) and there is cooperation traffic between them (i.e. $\boldsymbol{CV_{v,y} = 1}$), otherwise $\Phi_{\boldsymbol{v,y,s,d}} = \mathbf{0}$. Constraints (4.13) and (4.14) ensure that $б_{\boldsymbol{v,y,d}} = \mathbf{1}$ if there is at least one cooperation traffic between VMs $\boldsymbol{v}$ located at any node and $\boldsymbol{y}$ located at node $\boldsymbol{d}$ ($\sum_{\boldsymbol{s \in N}} \Phi_{\boldsymbol{v,y,s,d}} = \mathbf{1}$), $б_{\boldsymbol{v,y,d}} = \mathbf{0}$ otherwise. Constraint (4.15) ensures that only one replica of VM $\boldsymbol{v}$ is selected to cooperate with VM $\boldsymbol{y}$ at node $\boldsymbol{d}$. Constraint (4.16) ensures that the node selected to provide VM $\boldsymbol{y}$ with cooperation traffic from VM $\boldsymbol{v}$ contains a replica of $\boldsymbol{v}$ which is indicated by variable $\Phi_{\boldsymbol{v,y,s,d}}$. The aim of constraints (4.13) to (4.16) is to ensure that each replica of a VM receives cooperation only from a single replica of $\boldsymbol{VM\ v}$. Constraint (4.17) calculates the cooperation traffic between VMs $\boldsymbol{v}$ and $\boldsymbol{y}$, if they are located in different nodes, whereas, Constraint (4.18) calculates the cooperation traffic between VMs, if they are located in the same node.

VM replicas synchronisation traffic:

$$\mathbf{2}\ \Theta_{\boldsymbol{v,s,d}} + \boldsymbol{\varphi}_{\boldsymbol{v,s,d}} = \boldsymbol{\delta}_{\boldsymbol{v,s}}^{(\boldsymbol{C})} + \boldsymbol{\delta}_{\boldsymbol{v,d}}^{(\boldsymbol{C})}$$

$$\forall \boldsymbol{s,d} \in \boldsymbol{N}: \boldsymbol{s \neq d}, \boldsymbol{v} \in \boldsymbol{VM} \qquad (\mathbf{4.19})$$

$$\boldsymbol{IS_{v,s,d}} = \Theta_{\boldsymbol{v,s,d}}\ \boldsymbol{SVM_v}$$

$$\forall \boldsymbol{s,d} \in \boldsymbol{N}: \boldsymbol{s \neq d}, \boldsymbol{v} \in \boldsymbol{VM} \qquad (\mathbf{4.20})$$

Constraints (4.19) and (4.20) represent the synchronisation traffic among VM $\boldsymbol{v}$ replicas. Constraint (4.19) ensures that $\Theta_{\boldsymbol{v,s,d}}$=1 if VM $\boldsymbol{v}$ replicas are located in node $\boldsymbol{s}$ (i.e. $\boldsymbol{\delta}_{\boldsymbol{v,s}}^{(\boldsymbol{C})}$) and node $\boldsymbol{d}$ (i.e. $\boldsymbol{\delta}_{\boldsymbol{v,d}}^{(\boldsymbol{C})}$), respectively, otherwise $\Theta_{\boldsymbol{v,s,d}} = \mathbf{0}$. Constraint (4.120) calculates the synchronisation traffic sent by VM $\boldsymbol{v}$ replica in node $\boldsymbol{s}$ to another replica in node $\boldsymbol{d}$.

4.3 Results and Discussions

The NSFNET network, depicted in Fig. 4-2, is considered as an example of a core network topology to optimise the placement of 1500 VMs in clouds located in its core nodes. VM users are uniformly distributed over the NSFNET 14 nodes. Each VM is considered to have 800 users. The users are considered to access the VMs with a download rate uniformly distributed between two data rates; 5 and 25 Mbps. The workload of VMs of a constant workload profile and the workload of serving the maximum number of users of VMs of a linear workload are uniformly distributed among three workloads: 10%, 50% and 100% of the server CPU capacity. VMs of a linear workload are considered to have a workload baseline of 1% or 5% of the server CPU capacity.

Each VM is considered to cooperate with 50% of the other VMs selected randomly, whereas, for VM replicas synchronisation, all VM replicas exchange traffic. The placement is studied under three inter-VM traffic scenarios: low traffic of 100 Mbps, medium traffic of 1 Gbps and high traffic of 5 Gbps [71]. VMs in cloud datacentres are hosted in IBM Power System S814 server [165], which has eight cores each with 3.72 GHz IBM power8 processor, 128 GB memory, and 1.55 TB storage while consuming 333W. Note that, the placement of VMs in a server is optimised so the I/O traffic rate is not violated. In IP over WDM networks, each router port operates at 40 Gbps while consuming 825W [75]. In the cloud datacentre network, the Juniper MX204 router [166] is considered as an aggregation router, which consumes 0.9W/Gbps (36 watt for each 40 Gbps router port). The Juniper EX4550 Ethernet switch [167] is considered as cloud LAN switch with power rating of 9W for 10GbE interface. Table 4-3 shows the input data of the evaluation scenarios. In the following results, the VMs placement and the power consumption associated with optimisation scenarios considering cooperation and synchronisation inter-VM traffic are compared with those of optimisations scenarios ignoring them.

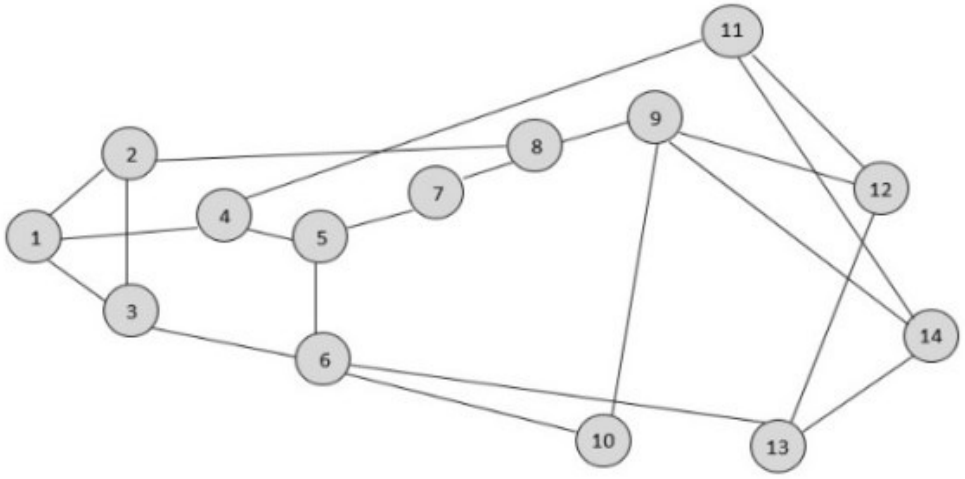

Figure 4-2: NSFNET topology.

Table 4-3: Model input parameters

40 Gbps router port power consumption ($\boldsymbol{R^{(P)}}$)	825W [75]
40 Gbps transponder power consumption ($\boldsymbol{t^{(P)}}$)	167W [75]
40 Gbps regenerator power consumption ($\boldsymbol{G^{(P)}}$)	334W, reach 2500 km [75]
EDFA power consumption ($\boldsymbol{e^{(P)}}$)	55W [75]
Optical switch power consumption ($\boldsymbol{SW^{(P)}}$)	85W [75]
Number of wavelengths in a fibre ($\boldsymbol{\mathcal{W}}$)	32 [75]
Bit rate of each wavelength ($\boldsymbol{\mathcal{W}^{(P)}}$)	40 Gbps [75]
Span distance between two EDFAs ($\boldsymbol{S}$)	80 km
Network power usage effectiveness ($\boldsymbol{n}$)	1.5 [6]
Number of VM users ($\boldsymbol{x}$)	800 users per VM
User Download rate ($\boldsymbol{r_v}$)	{5 and 25 Mbps}
Cooperation traffic from VM $\boldsymbol{v}$ to VM $\boldsymbol{y}$ ($\boldsymbol{CVM_{v,y}}$)	100 Mbps, 1 Gbps or 5 Gbps
Synchronisation traffic between VM $\boldsymbol{v}$ replicas ($\boldsymbol{SVM_v}$)	100 Mbps, 1 Gbps or 5 Gbps
Cloud router port bit rate ($\boldsymbol{R^{(CB)}}$)	40 Gbps [166]
Cloud router port power consumption ($\boldsymbol{R^{(CP)}}$)	36W [166]
Cloud switch bit rate ($\boldsymbol{SW^{(CB)}}$)	10 Gbps [167]
Cloud switch power consumption ($\boldsymbol{SW^{(CP)}}$)	9W [167]

Cloud switch redundancy ($\boldsymbol{SW^{(R)}}$)	2
Cloud power usage effectiveness ($\boldsymbol{c}$)	1.7 [143]
Number of VMs ($\boldsymbol{V}$)	1500 VMs
Server power consumption ($\boldsymbol{S^{(P)}}$)	333W [156]
Maximum server workload ($\boldsymbol{S^{(maxW)}}$)	100%
VM workload baseline in linear profile ($\boldsymbol{M}$)	1% or 5%
Set of VMs workload ($\boldsymbol{W_v}$)	{10, 50 and 100} %

4.3.1 Linear Workload (1% Workload Baseline):

Fig. 4-3 shows the optimal placement of VMs under 1% linear workload profile considering users download traffic only. VMs of 5 Mbps download rate and VMs of 25 Mbps download rate are fully replicated in all cloud locations. The optimal placement of VMs considering inter-VM traffic of 100 Mbps in addition to the user download traffic is shown in Fig. 4-4.

We separately examine the impact of cooperation traffic (Fig. 4-4 (a)) and synchronisation traffic (Fig. 4-4 (b)) and show the placement resulting from considering them jointly (Fig. 4-4 (c)). Fig. 4-4 (a) shows that taking cooperation traffic of 100 Mbps into consideration while placing VMs has no impact on the optimal VMs placement compared to the placement in Fig. 4-3 as full replication of VMs in all datacentres will confine the cooperation traffic inside the datacentre. Fig 4-4 (b) shows that the existence of synchronisation traffic has limited the number of replicas of each VM into two and four replicas under 5 and 25 Mbps download rates, respectively as traffic among VM replicas distributed to all clouds will highly burden the network. The optimum placement when considering cooperation and synchronisation traffic jointly (Fig. 4-4 (c)) is a trade-off between the placement in Fig. 4-4 (a) and (b) with the impact of synchronisation traffic dominating.

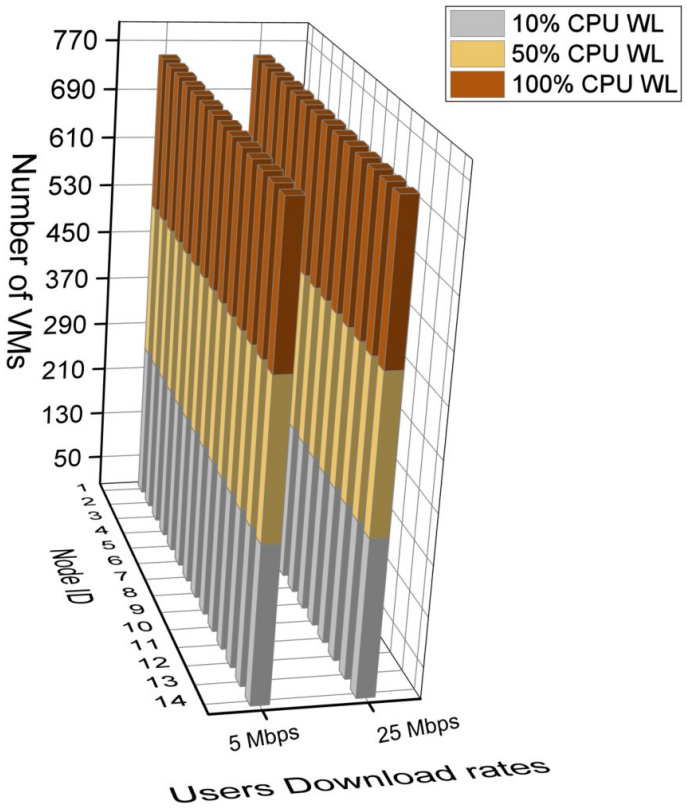

Figure 4-3: The optimal placement of VMs under 1% workload baseline considering users traffic only.

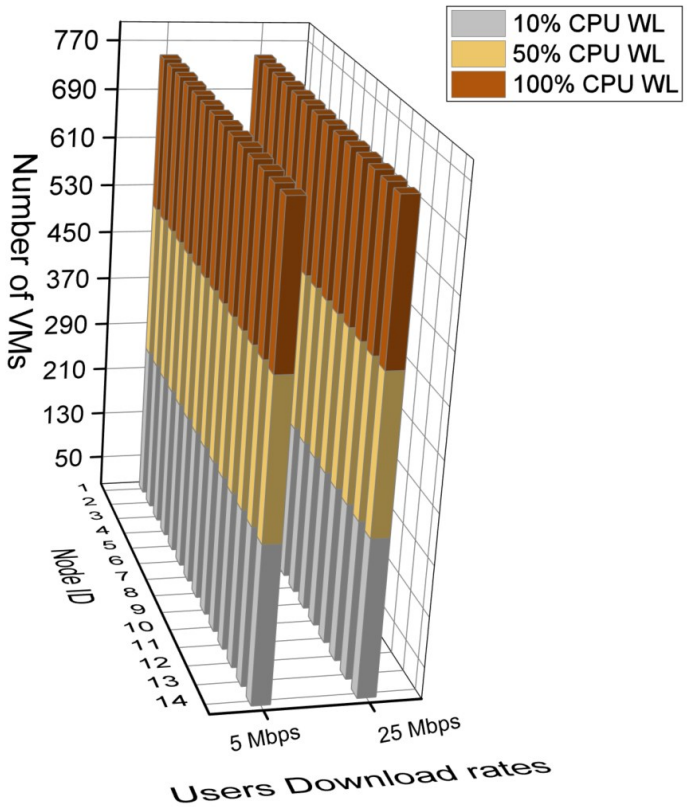

(a)

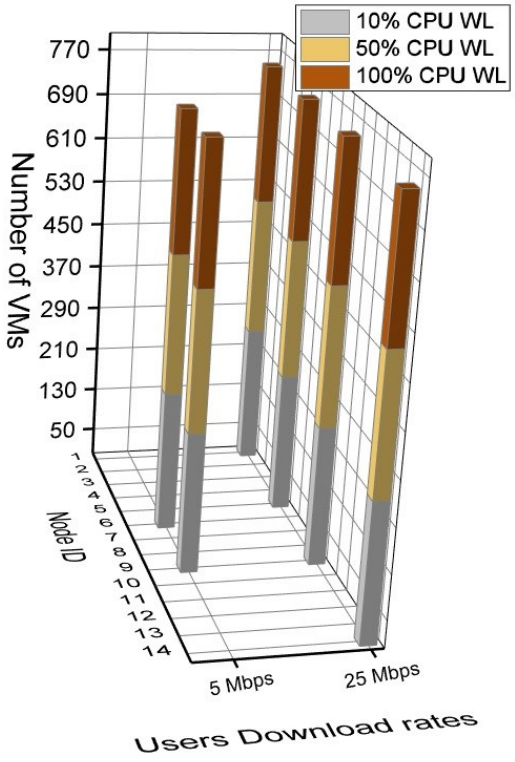

(b)

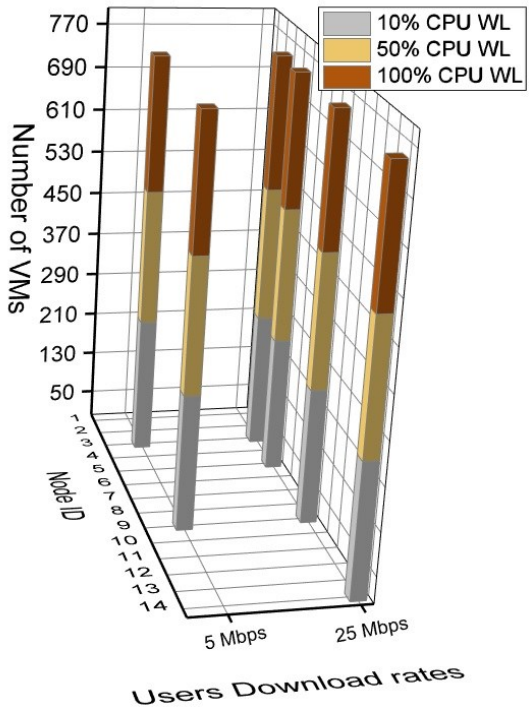

(c)

Figure 4-4: The optimal placement of VMs under 1% workload baseline considering users traffic and 100 Mbps a) cooperation traffic, b) synchronisation traffic, c) total inter-VM traffic.

In Fig. 4-5, the power consumptions of different VMs placement and inter-VM traffic scenarios are compared as follows;

- Cooperation traffic: power consumption of users' download traffic and cooperation traffic of VMs.
- Synchronisation traffic: power consumption of users' download traffic and synchronisation traffic of VMs.
- Cooperation and synchronisation traffic: power consumption of users download traffic, cooperation traffic and synchronisation traffic of VMs.

In each scenario, we study the potential increase in total power consumption resulting from optimising the placement of VMs considering users download traffic only with scenarios where cooperation and/or synchronisation traffic are optimised in additional to users download traffic. The results show very limited increase in total power consumption if the VMs are placed considering users download traffic only (as seen in Fig. 4-3) compared to a scenario where users download and cooperation traffic exist. The full replication of VMs of 5 and 25 Mbps has confined the cooperation traffic among them to the intra datacentre network and hence limited the increase in total power consumption. However, not taking synchronisation traffic into consideration when optimising VMs placement has increased the total power consumption by a factor of 1.5 in cooperation traffic scenario and cooperation and synchronisation traffic scenario. These increases are mainly caused by the synchronisation traffic among the fully distributed replicas of each VM of 5 and 25 Mbps download rate creating a full mesh traffic matrix traversing the IP over WDM network, whereas, cooperation is point-to-point traffic.

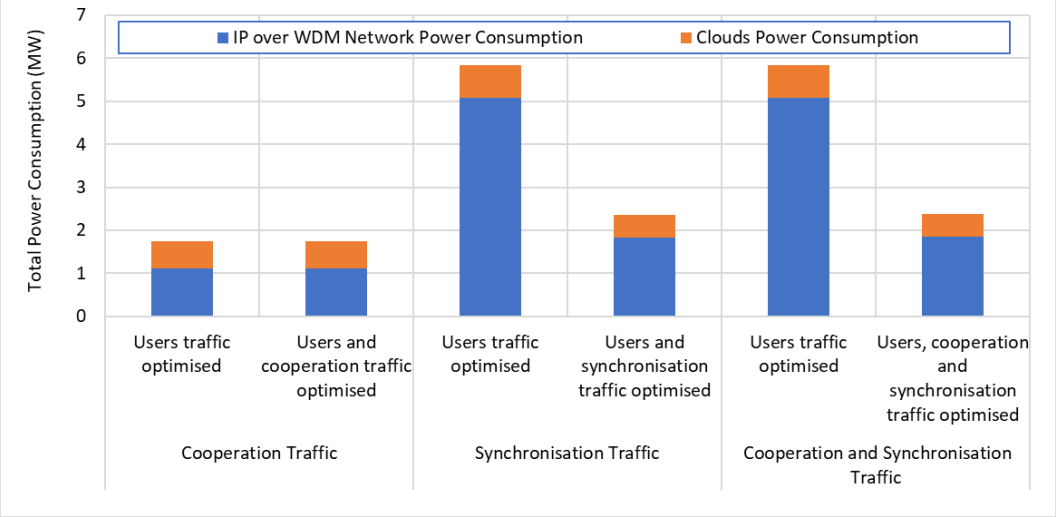

Figure 4-5: Power consumption associated with different VMs placement scenarios under 1% workload baseline considering 100 Mbps inter-VM traffic.

Fig. 4-6 shows the optimal VMs placement considering 1 Gbps inter-VM traffic. Similar to the previous scenario, the placement under 1 Gbps cooperation traffic (Fig. 4-6 (a)) has resulted in full replicas of all VMs. These replicas are collocated so cooperation traffic is kept within the datacentre. Considering VM synchronisation traffic in (Fig. 4-6 (b)) has limited the placement in a single cloud under VMs of 5 Mbps and two clouds placement under VMs of 25 Mbps. A single cloud placement is obtained from considering both inter-VM traffic jointly (Fig. 4-6 (c)). As shown in Fig. 4-7, placing VMs closer to the users without bearing in mind the existence of 1 Gbps inter-VM traffic causes the power consumption to increase by factor of 13 compared to placing them based on the existence of both traffic.

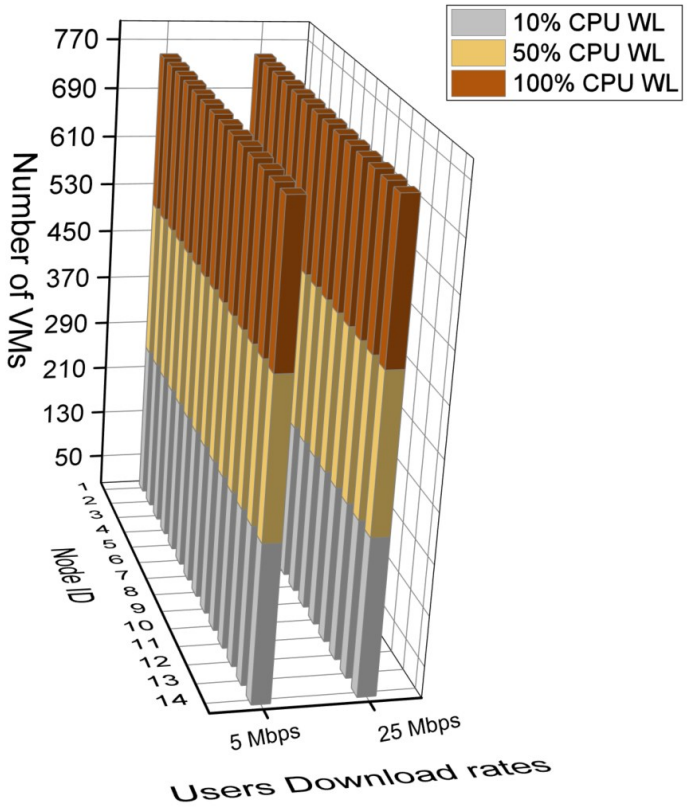

(a)

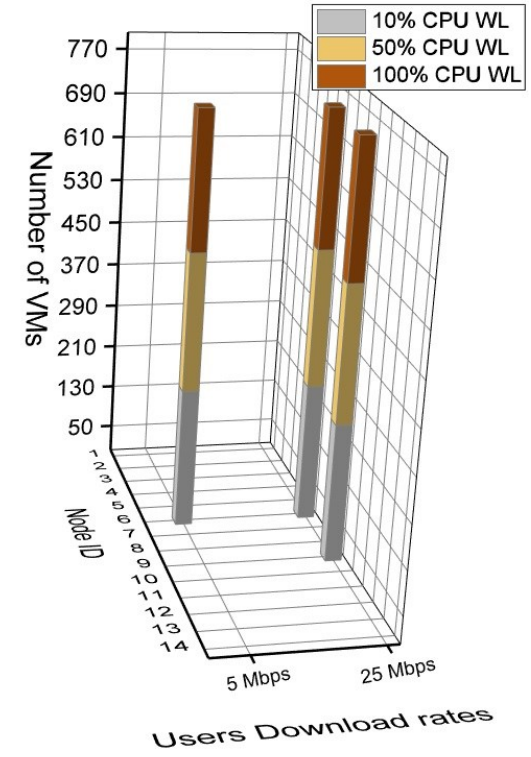

(b)

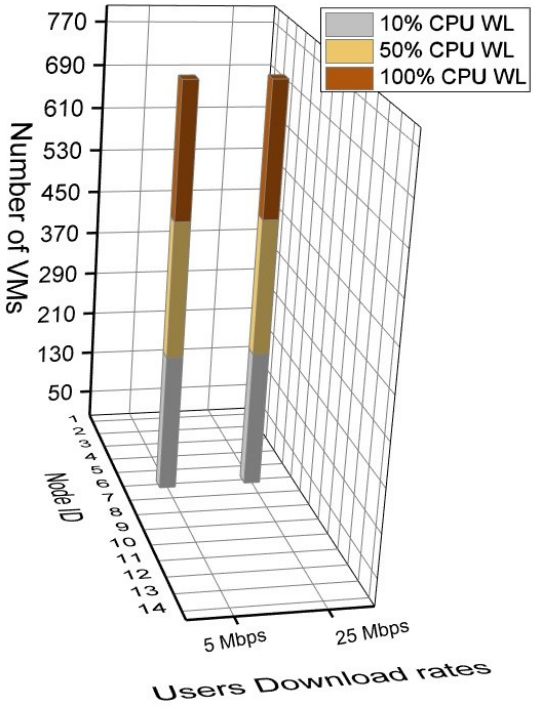

(c)

Figure 4-6: The optimal placement of VMs under 1% workload baseline considering users traffic and 1 Gbps a) cooperation traffic, b) synchronisation traffic, c) total inter-VM traffic.

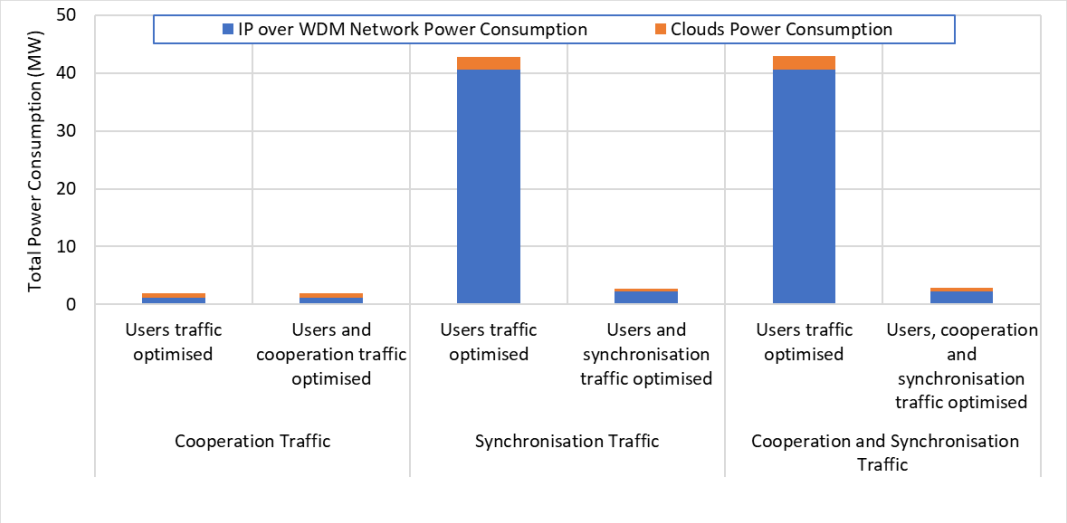

Figure 4-7: Power consumption associated with different VMs placement scenarios under 1% workload baseline considering 1 Gbps inter-VM traffic.

Fig. 4-8 shows the optimal VMs placement considering 5 Gbps inter-VM traffic. The placement under 5 Gbps cooperation traffic (Fig. 4-8 (a)) has resulted in 4 replicas of all VMs. Fig. 4-9 shows that the potential increase in total power consumption resulting from optimising the placement of VMs considering users download traffic only with scenarios where 5 Gbps cooperation and synchronisation traffic exist in additional to users download traffic can be up to a factor of 72 compared to scenarios ignoring them.

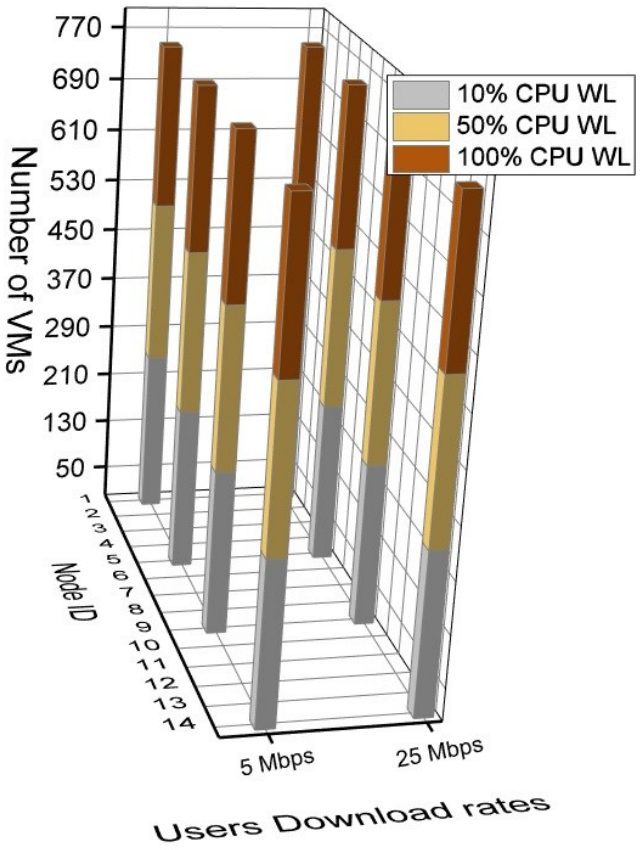
770
690
610
530
450
370
290
210
130
50
Number of VMs
10% CPU WL
50% CPU WL
100% CPU WL
Node ID
1
2
3
4
5
6
7
8
9
10
11
12
13
14
5 Mbps
25 Mbps
Users Download rates

(a)

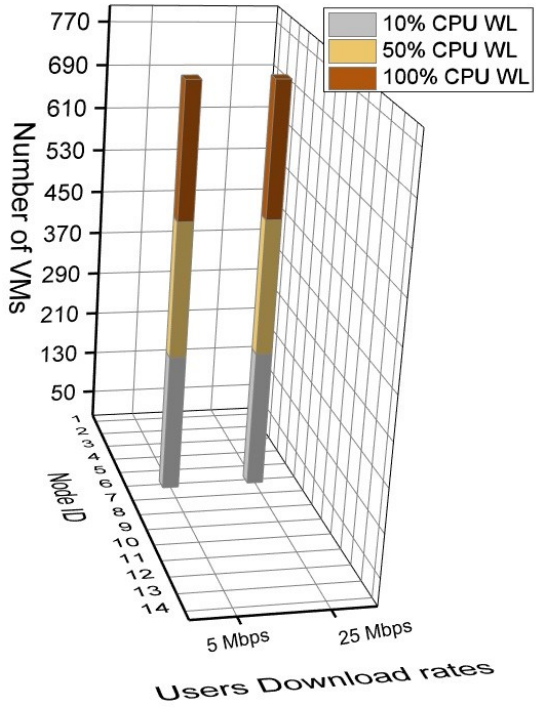
770
690
610
530
450
370
290
210
130
50
Number of VMs
10% CPU WL
50% CPU WL
100% CPU WL
Node ID
1
2
3
4
5
6
7
8
9
10
11
12
13
14
5 Mbps
25 Mbps
Users Download rates

(b)

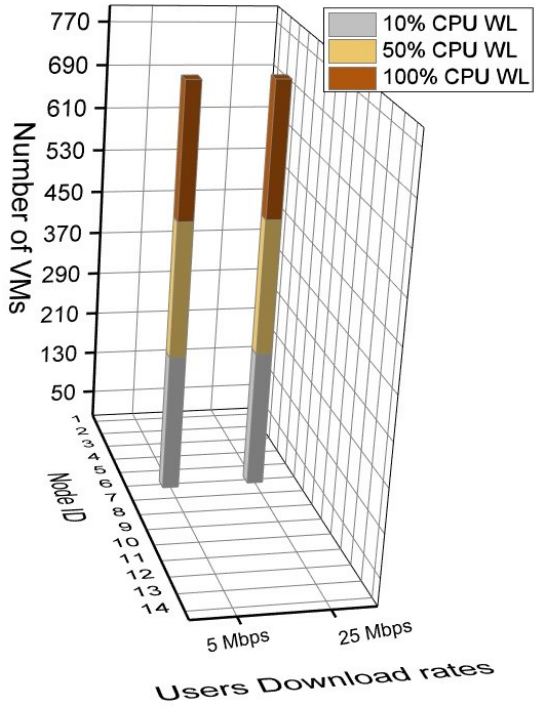

(c)

Figure 4-8: The optimal placement of VMs under 1% workload baseline considering users traffic and 5 Gbps a) cooperation traffic, b) synchronisation traffic, c) total inter-VM traffic.

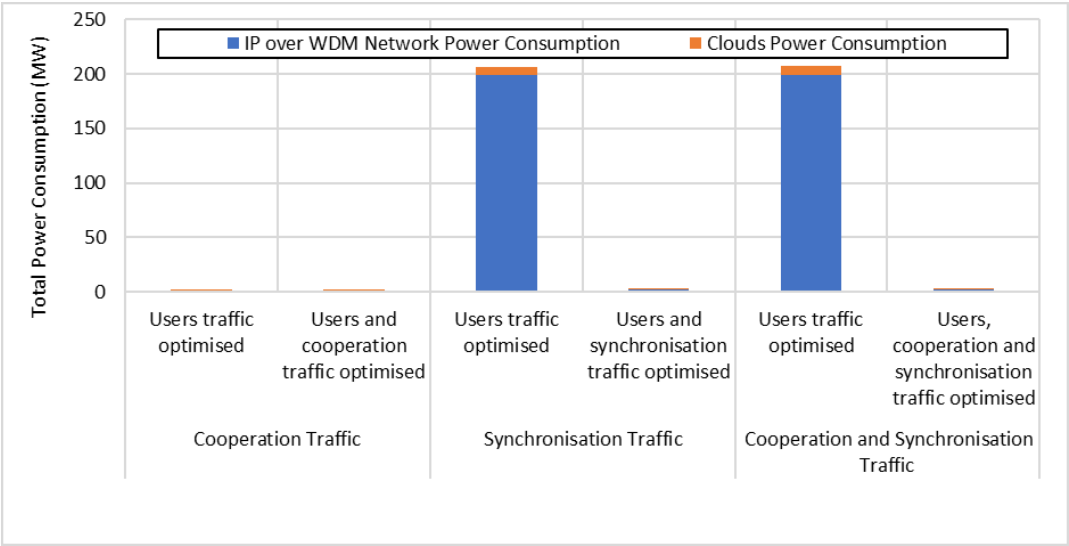

Figure 4-9: Power consumption associated with different VM placement scenarios under 1% workload baseline considering 5 Gbps inter-VM traffic.

4.3.2 Linear Workload (5% Workload Baseline):

Fig. 4-10 shows the optimum placement of VMs under 5% workload baseline considering users download traffic only. Three replicas are created of each VM of 5 Mbps download rate, while VMs with 25 Mbps download rate are fully replicated

in all cloud locations. The optimal placement of VMs under 5% workload baseline considering inter-VM traffic of 100 Mbps in addition to the user download traffic is shown in Fig. 4-11. Fig. 4-11 (a) shows that taking cooperation traffic of 100 Mbps into consideration when placing VMs has resulted in creating more replicas (four replicas) of the VMs of 5 Mbps users download rate compared to optimisation considering users traffic only (three replicas). This placement allows cooperation traffic between VMs of 25 Mbps (replicated everywhere) and VMs of 5 Mbps to traverse a maximum of a single hop in the IP over WDM network. In Fig 4-11 (b), the existence of synchronisation traffic has limited the number of replicas of each VM into two and four replicas under 5 and 25 Mbps download rates, respectively. The optimum placement considering cooperation and synchronisation traffic jointly (Fig. 4-11 (c)) is a trade-off between the placement in Fig. 4-11 (a) and (b) with the impact of synchronisation traffic dominating.

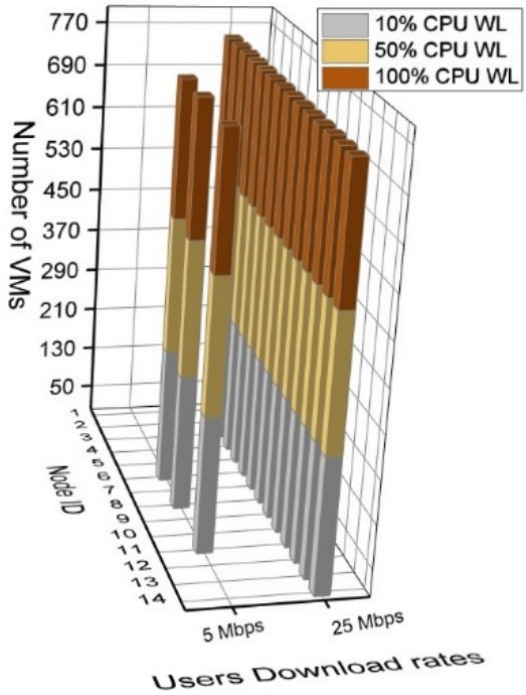

Figure 4-10: The optimal placement of VMs under 5% workload baseline considering users traffic only.

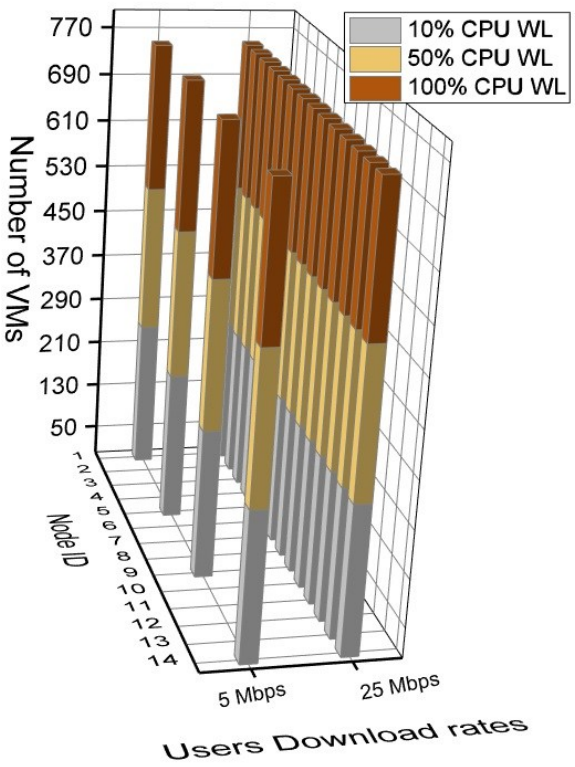
10% CPU WL
50% CPU WL
100% CPU WL
Number of VMs
770
690
610
530
450
370
290
210
130
50
Node ID
1 2 3 4 5 6 7 8 9 10 11 12 13 14
5 Mbps
25 Mbps
Users Download rates

(a)

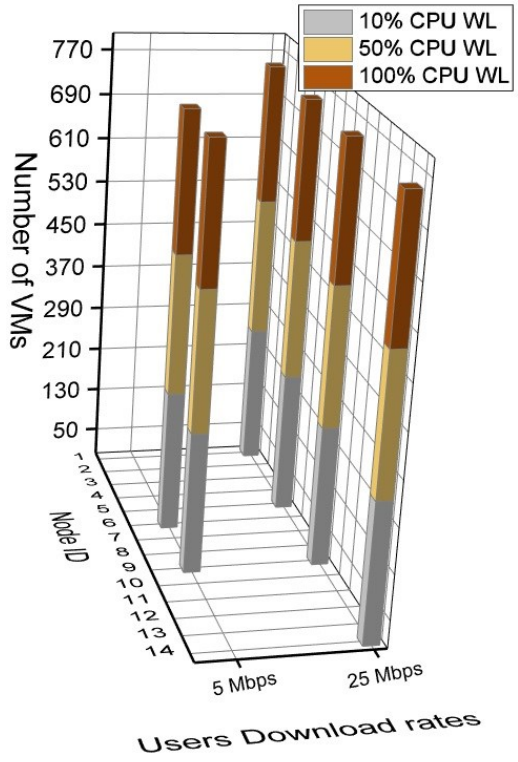
10% CPU WL
50% CPU WL
100% CPU WL
Number of VMs
770
690
610
530
450
370
290
210
130
50
Node ID
1 2 3 4 5 6 7 8 9 10 11 12 13 14
5 Mbps
25 Mbps
Users Download rates

(b)

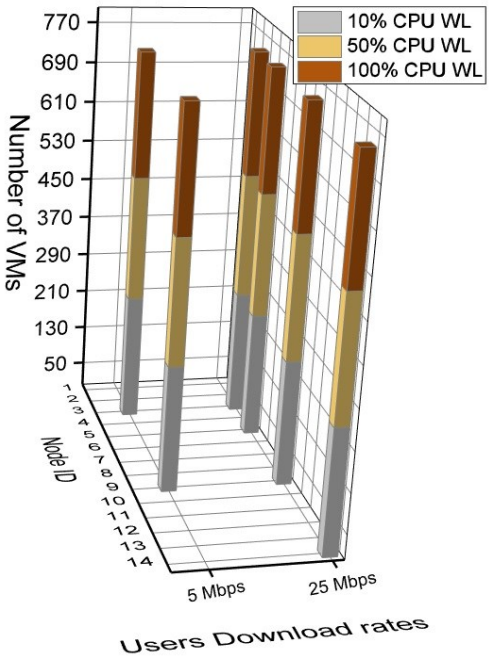

(c)

Figure 4-11: The optimal placement of VMs under 5% workload baseline considering users traffic and 100 Mbps a) cooperation traffic, b) synchronisation traffic, c) total inter-VM traffic.

Fig. 4-12 studies the potential increase in total power consumption resulting from optimising the placement of VMs considering users download traffic only with scenarios where cooperation and synchronisation traffic exist in additional to users download traffic. The results show a limited increase of 1% in total power consumption if the VMs are placed considering users download traffic only (as seen in Fig. 4-10) in a scenario where users download, and cooperation traffic exist as the difference between the two placement scenarios is limited. However, for cooperation and synchronisation traffic scenario, not taking synchronisation traffic into consideration when optimising VMs placement has increased the total power consumption by 73%. This increase is mainly caused by the full mesh

synchronisation traffic among the distributed replicas of each VM of 25 Mbps download rate.

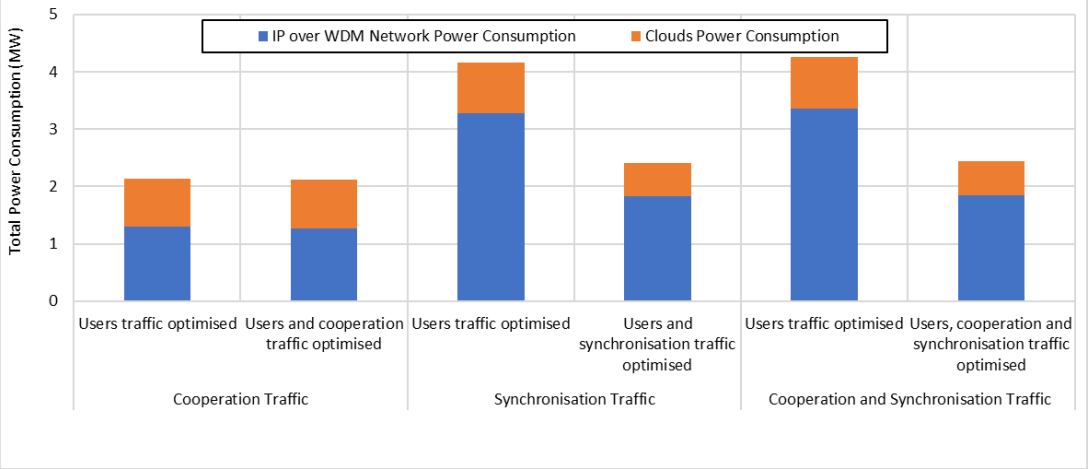

Figure 4-12: Power consumption associated with different VMs placement scenarios under 5% workload baseline considering 100 Mbps inter-VM traffic.

Fig. 4-13 shows the optimal VMs placement for VMs under 5% workload baseline considering 1 Gbps inter-VM traffic. The placement under 1 Gbps cooperation traffic (Fig. 4-13 (a)) has resulted in 4 replicas of all VMs. These replicas are collocated so cooperation traffic is kept within the datacentre. Similar to the placement of VMs under 1% workload baseline, considering 1 Gbps synchronisation traffic in Fig. 4-13 (b) has limited the number of copies to two replicas under 25 Mbps users download rate and 1 replica under 5 Mbps users download rate, whereas, considering both inter-VM traffic in Fig. 4-13 (c) has resulted in a single cloud location.

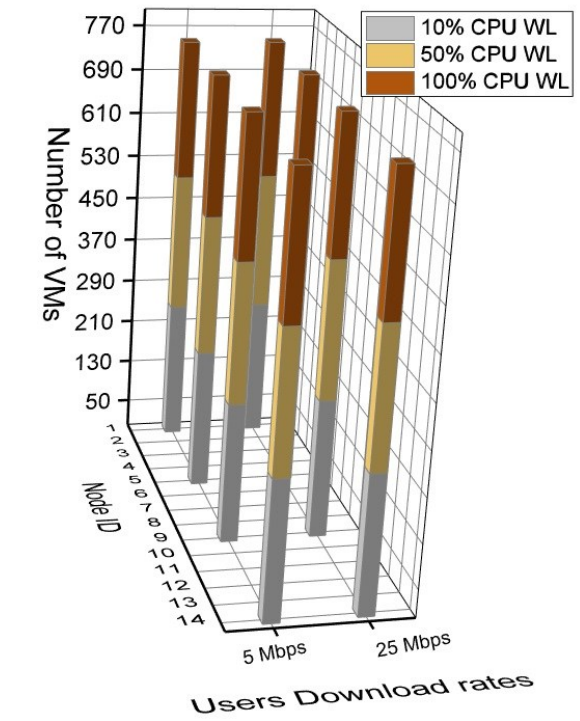
10% CPU WL
50% CPU WL
100% CPU WL
Number of VMs
770
690
610
530
450
370
290
210
130
50
Node ID
1
2
3
4
5
6
7
8
9
10
11
12
13
14
5 Mbps
25 Mbps
Users Download rates

(a)

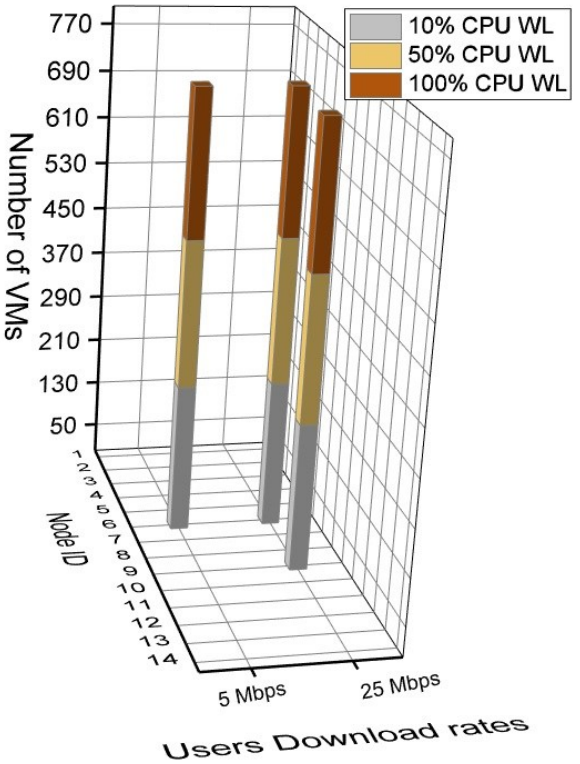
10% CPU WL
50% CPU WL
100% CPU WL
Number of VMs
770
690
610
530
450
370
290
210
130
50
Node ID
1
2
3
4
5
6
7
8
9
10
11
12
13
14
5 Mbps
25 Mbps
Users Download rates

(b)

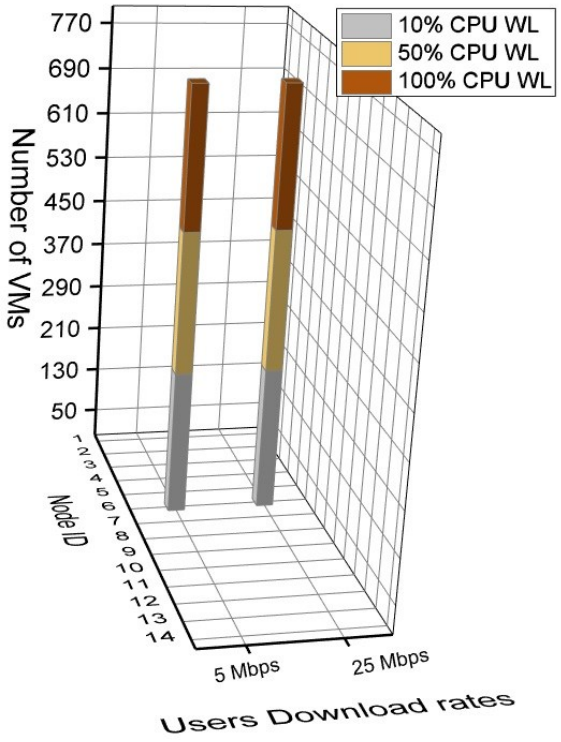

(c)

Figure 4-13: The optimal placement of VMs under 5% workload baseline considering users traffic and 1 Gbps a) cooperation traffic, b) synchronisation traffic, c) total inter-VM traffic.

Fig. 4-14 shows the power consumption of the different inter-VM scenarios, which discussed above, considering 1 Gbps inter-VM traffic and 5% workload baseline. The results show that considering only the users download traffic in the three scenarios can potentially increase the total power consumption by a factor of 0.3, 7.5 and 7.6, compared to a scenario that considers users download traffic and cooperation traffic, synchronisation traffic and both inter-VM traffic respectively if inter-VM is not taken into consideration when optimising the placement of VMs.

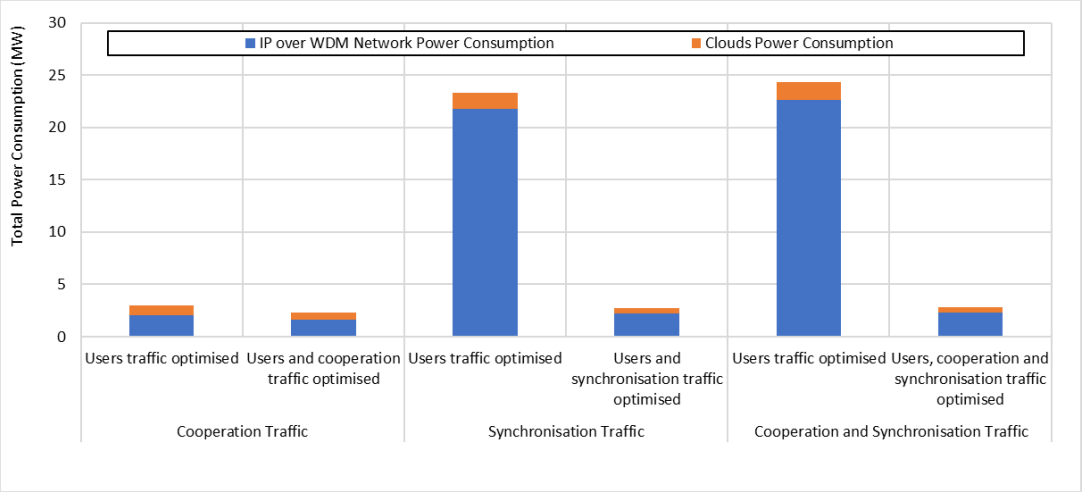

Figure 4-14: Power consumption associated with different VM placement scenarios under 5% workload baseline considering 1 Gbps inter-VM traffic.

Fig. 4-15 shows the optimal VMs placement considering 5 Gbps inter-VM traffic. The placement under 5 Gbps cooperation traffic (Fig. 4-15 (a)) has resulted in four replicas of all VMs. These replicas are placed together so cooperation traffic is kept within the datacentre. Considering VMs synchronisation traffic (Fig. 4-15(b)) has resulted in a single cloud placement as the synchronisation traffic power consumption surpasses the potential saving obtained by placing VM replicas closer to user premises. The same trend of single cloud placement is observed by considering both inter-VM traffic jointly (Fig. 4-15(c)). As shown in Fig. 4-16, placing VMs closer to the users without bearing in mind the existence of 5 Gbps inter-VM traffic causes the power consumption to increase by a factor of 39 compared to placing them based on the existence of inter-VM and users download traffic.

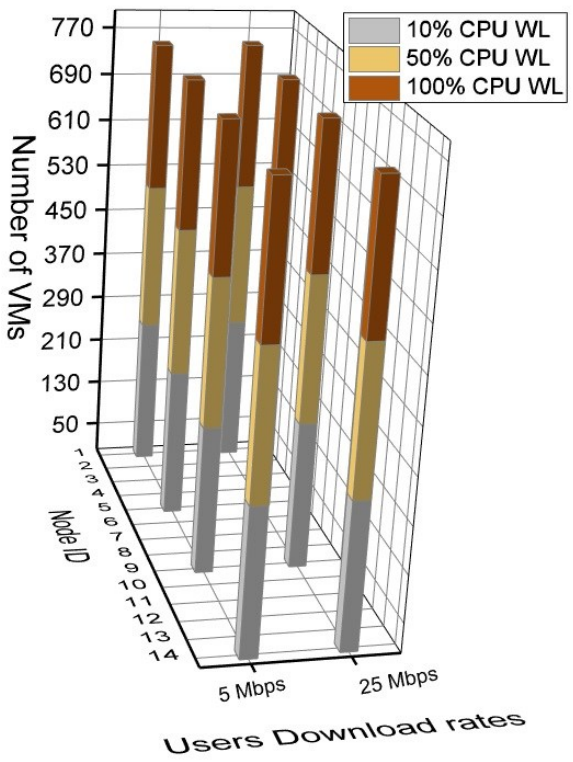
10% CPU WL
50% CPU WL
100% CPU WL
Number of VMs
770
690
610
530
450
370
290
210
130
50
Node ID
1 2 3 4 5 6 7 8 9 10 11 12 13 14
5 Mbps
25 Mbps
Users Download rates

(a)

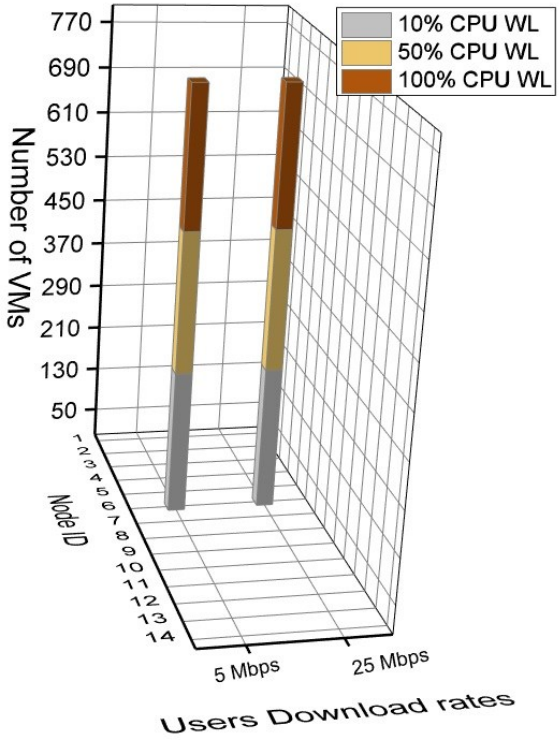
10% CPU WL
50% CPU WL
100% CPU WL
Number of VMs
770
690
610
530
450
370
290
210
130
50
Node ID
1 2 3 4 5 6 7 8 9 10 11 12 13 14
5 Mbps
25 Mbps
Users Download rates

(b)

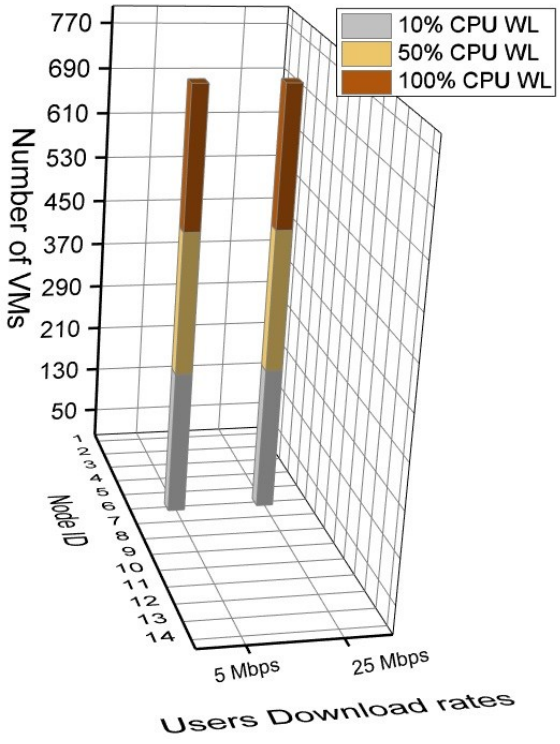

(c)

Figure 4-15: The optimal placement of VMs under 5% workload baseline considering users traffic and 5 Gbps a) cooperation traffic, b) synchronisation traffic, c) total inter-VM traffic.

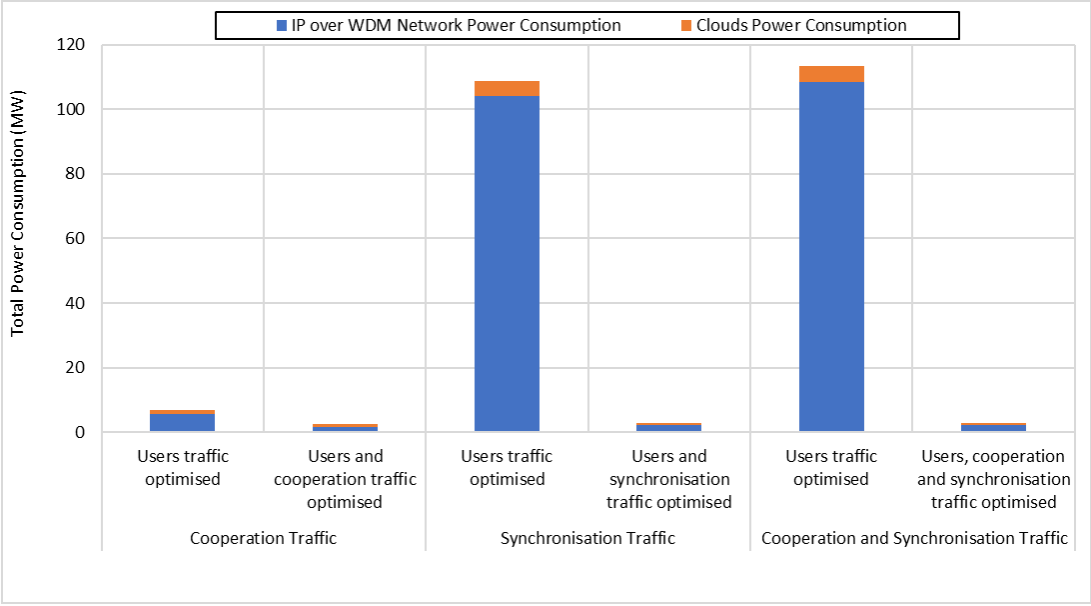

Figure 4-16: Power consumption associated with different VM placement scenarios under 5% workload baseline considering 5 Gbps inter-VM traffic.

4.3.3 Constant Workload:

Fig. 4-17 shows the results associated with energy efficient VMs placement for VMs of constant workload profile considering users download traffic only. VMs with lower workloads are distributed into multiple cloud locations to satisfy the users demand traffic and reduce the network power consumption. For VMs of a download rate of 5 Mbps, multiple replicas of VMs with 10% CPU workload are created. The increased network traffic requirement for VMs of 25 Mbps download rate has justified the creation of multiple replicas of VMs with 10% and 50% workloads.

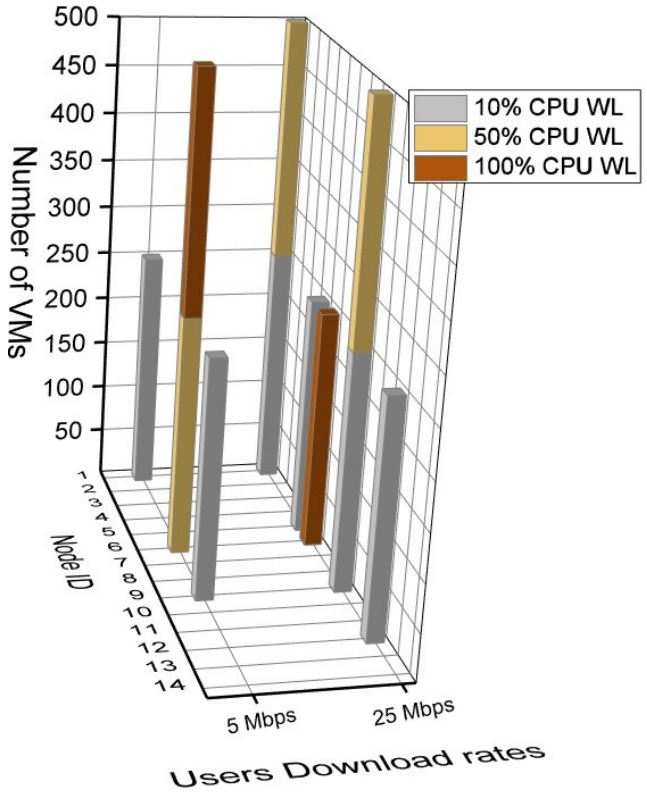

Figure 4-17: The optimal placement of VMs under constant workload profile considering users traffic only.

Fig. 4-18 shows the optimal VM placement under constant workload profile considering inter-VM traffic rate of 100 Mbps in addition to the download traffic to VM users. The existence of cooperation traffic of 100 Mbps has no major effect on the placement of VMs as seen in Fig. 4-18 (a). However, a replica of each VM is either placed in nodes 6 or 9 so cooperation traffic is kept intra cloud or within a single hop in the core network. Synchronisation traffic that sit between all replicas of a VM has resulted in fewer replicas compared to the number of replicas under cooperation traffic. A user download rate of 5 Mbps does not justify replicating VMs while only 10% and 50% workload VMs of 25 Mbps download rate have justified the creation of

multiple VM replicas. The optimal placement when considering cooperation and synchronisation traffic (Fig. 4-18(c)) is a trade-off between the placement in Fig. 4-18 (a) and (b) with the impact of synchronisation traffic dominating. Note that two of the three replicas of VMs with 25 Mbps download rate and 10% workload are moved from their optimal location considering synchronisation traffic only (nodes 8 and 11) to nodes 1 and 9 so that the hops traversed by cooperation traffic are minimised.

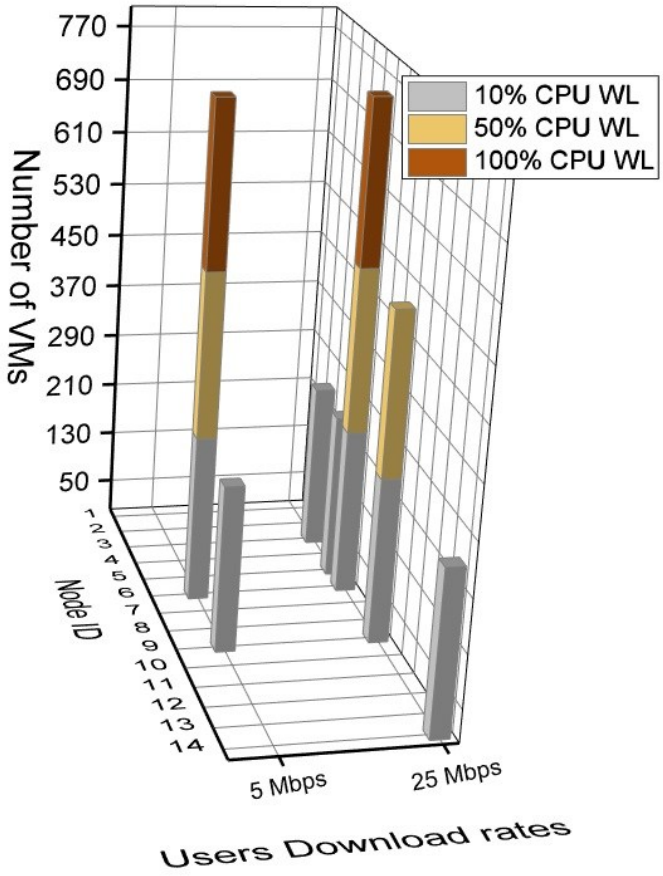

(a)

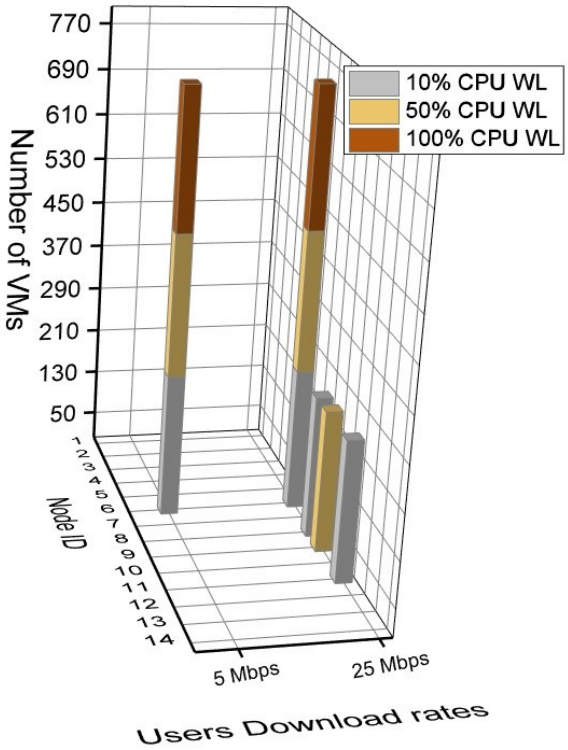

(b)

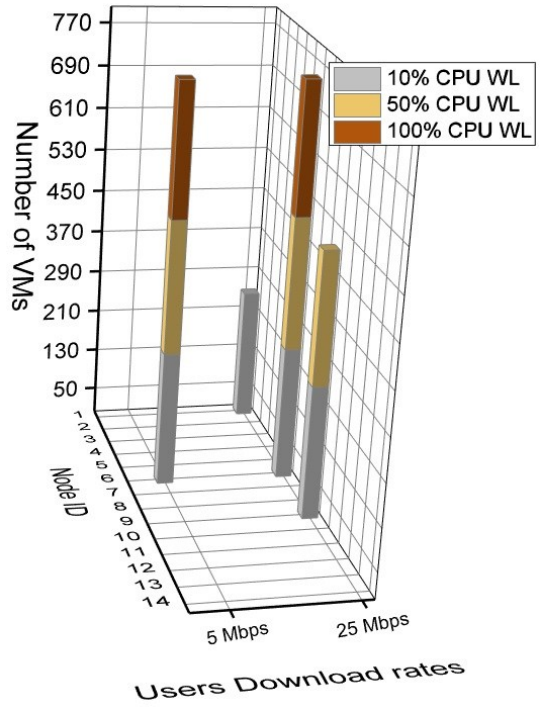

(c)

Figure 4-18: The optimal placement of VMs under constant workload profile considering users traffic and 100 Mbps a) cooperation traffic, b) synchronisation traffic, c) total inter-VM traffic.

Fig 4-19 compares the power consumption of different VM placements and inter-VM traffic scenarios. Due to the low data rate of inter-VM traffic and the constant

workload profile, a limited increase in the total power consumption (up to 3%) is achieved from considering only the users download traffic in the three scenarios compared to optimising VMs placement while considering both users and Inter-VM.

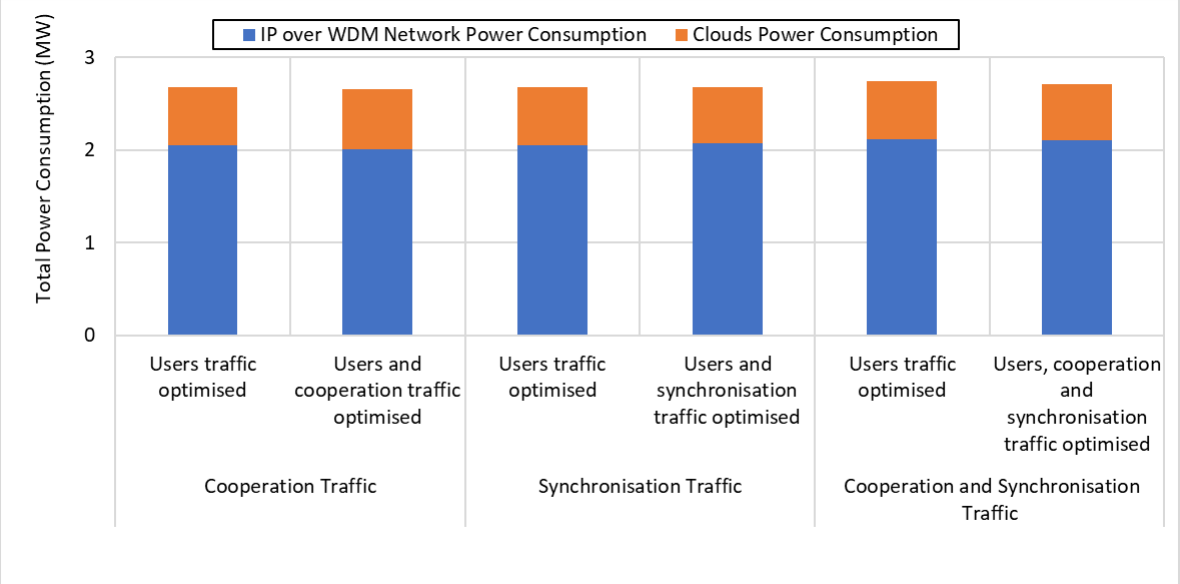

Figure 4-19: Power consumption associated with different VMs placement scenarios under constant workload profile considering 100 Mbps inter-VM traffic.

The optimal VMs placement of constant workload profile considering inter-VM traffic rate of 1 and 5 Gbps in addition to the download traffic of VMs users is illustrated in Fig. 4-20. The high inter-VM traffic has put an end to the efficiency of replicating VMs in distributed cloud considering the all inter-VM scenarios. All VMs are optimally placed in one cloud (node 6) to keep the cooperation traffic intra datacentre and eliminate the need for the synchronisation traffic.

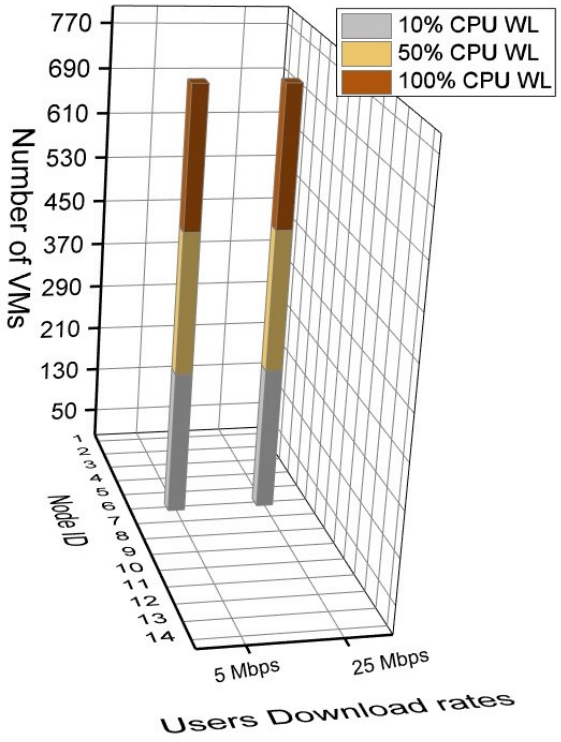

Figure 4-20: Optimal placement of VMs under constant workload profile considering users traffic and 1 or 5 Gbps a) cooperation traffic, b) synchronisation traffic, c) total inter-VM traffic.

Fig. 4-21 and Fig. 4-22 show the power consumption scenarios of 1 Gbps inter-VM traffic and 5 Gbps inter-VM, respectively, which are similar to what is discussed above. The results show that only considering users download traffic in the three scenarios can potentially increase the total power consumption by 14%, 16% and 29% under 1 Gbps inter-VM traffic and 51%, 54% and 69% under 5 Gbps inter-VM traffic, considering cooperation traffic, synchronisation traffic and both inter-VM traffic, respectively if inter-VM traffic is not taken into consideration when optimising the placement of VMs.

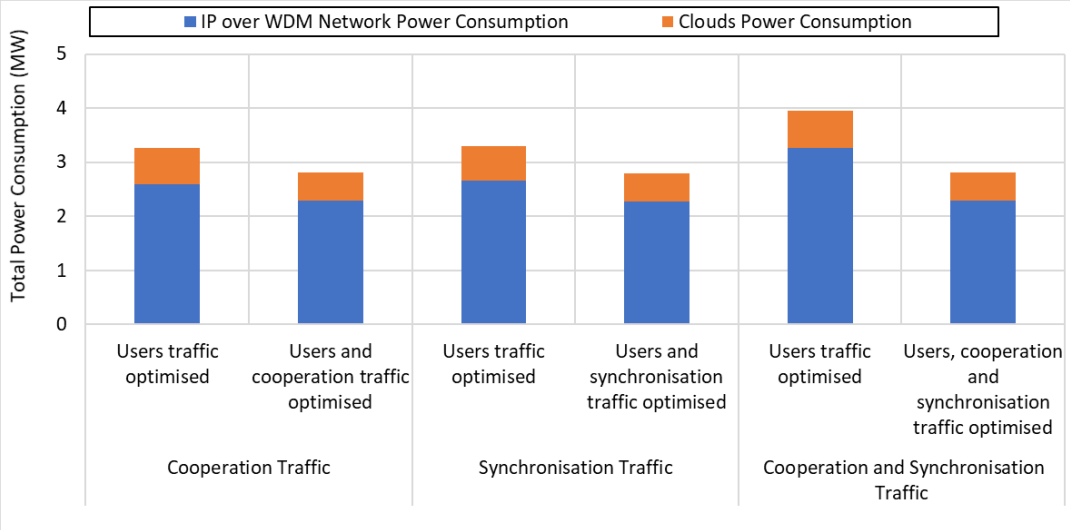

Figure 4-21: Power consumption associated with different VM placement scenarios under constant workload profile considering 1 Gbps inter-VM traffic.

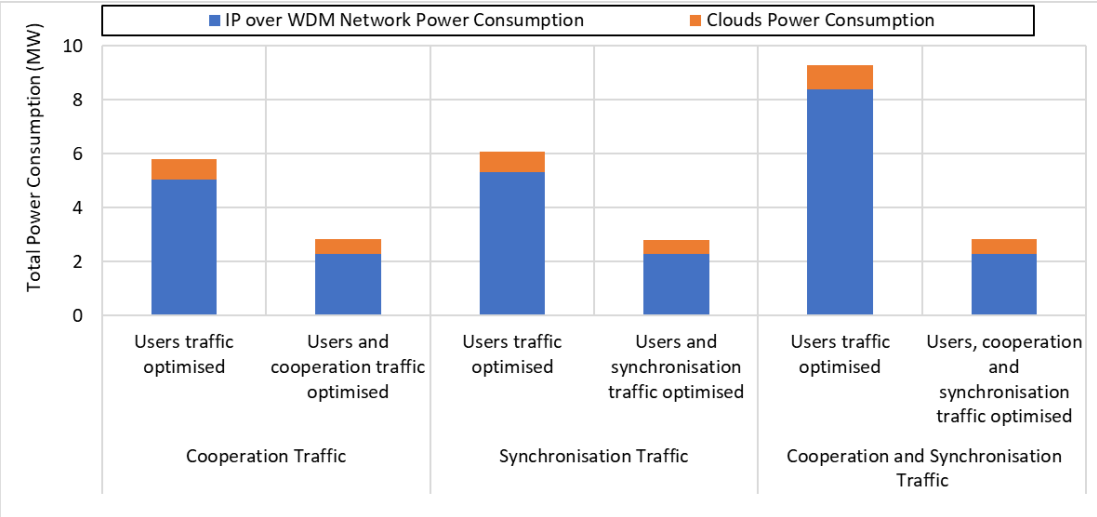

Figure 4-22: Power consumption associated with different VM placement scenarios under constant workload profile considering 5 Gbps inter-VM traffic.

Generally, this chapter shows that only considering users traffic and ignoring inter-VM traffic when placing VMs can significantly mount up the total power consumption. Thus, it is important to co-optimise the two types of traffic.

4.4 EEVM-C-IVM Heuristic

In this section, we validate the MILP operation by developing a heuristic that mimics, in real time, the behaviour of the MILP. The heuristic of energy efficient VM placement problem considering inter VMs traffic referred to Energy Efficient VMs Placement Considering Inter-VM traffic (EEVM-C-IVM). Similar to the heuristic of Chapter 3, the EEVM-C-IVM heuristic consists of two-phases: an offline phase and an online phase.

As shown in the flowchart in Fig. 4-23 (a), the offline phase starts by classifying VMs based on their workload and user data rate. Similar to the heuristic in Chapter 3, the offline phase starts by finding the most energy efficient placement for each number of replicas through an exhaustive search over all the possible placements for this number of replicas. Unlike Chapter 3 where the optimal placement for each VM type is found independently, this heuristic finds the optimum placement for all VM types jointly as the decision to place each VM type depends on the placement of the other VM types cooperating with this VM type. The optimum placement of VMs is found through an exhaustive search over all possible placement combinations considering the most energy efficient placement of 1 replica, 2 replicas…N replicas for each VM type. The size of the search space grows exponentially with the number of VM types ($\boldsymbol{N}^{|\boldsymbol{VMP}|}$), where N is the number of nodes in the network and $|\boldsymbol{VMP}|$ is the number of VMs types.

The traffic between users and VMs, traffic between cooperating VMs and synchronisation traffic (if any) resulting from the placement are routed based on minimum hop routing and the workload of VMs of a linear workload profile is calculated based on the number of users each VM serves. The network power consumption and the cloud power consumption are calculated and the optimum placement of each VM type is the placement that results in the minimum total power consumption.

In the online phase, shown in the flowchart in Fig. 4-23 (c), each VM is matched to a type and placed according to the placement obtained in the offline phase. Traffic demands resulting from this placement (user traffic and inter VM traffic) are routed on the core network using multi-hop routing [19] and the workload of clouds where the VM replicas are placed is updated. After placing all VMs, the total power consumption of distributed cloud is calculated.

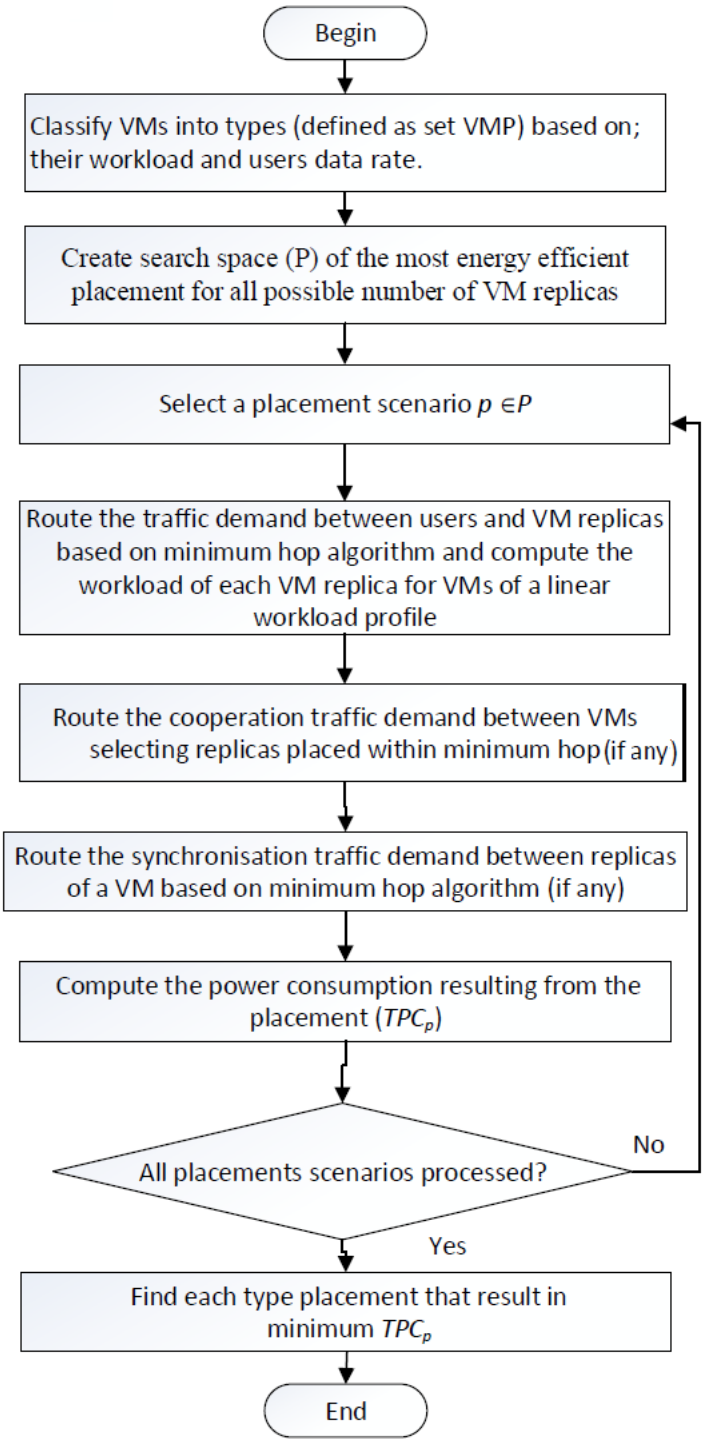
Begin
Classify VMs into types (defined as set VMP) based on; their workload and users data rate.
Create search space (P) of the most energy efficient placement for all possible number of VM replicas
Select a placement scenario p ∈P
Route the traffic demand between users and VM replicas based on minimum hop algorithm and compute the workload of each VM replica for VMs of a linear workload profile
Route the cooperation traffic demand between VMs selecting replicas placed within minimum hop (if any)
Route the synchronisation traffic demand between replicas of a VM based on minimum hop algorithm (if any)
Compute the power consumption resulting from the placement (TPC_p)
All placements scenarios processed?
No
Yes
Find each type placement that result in minimum TPC_p
End

(a)

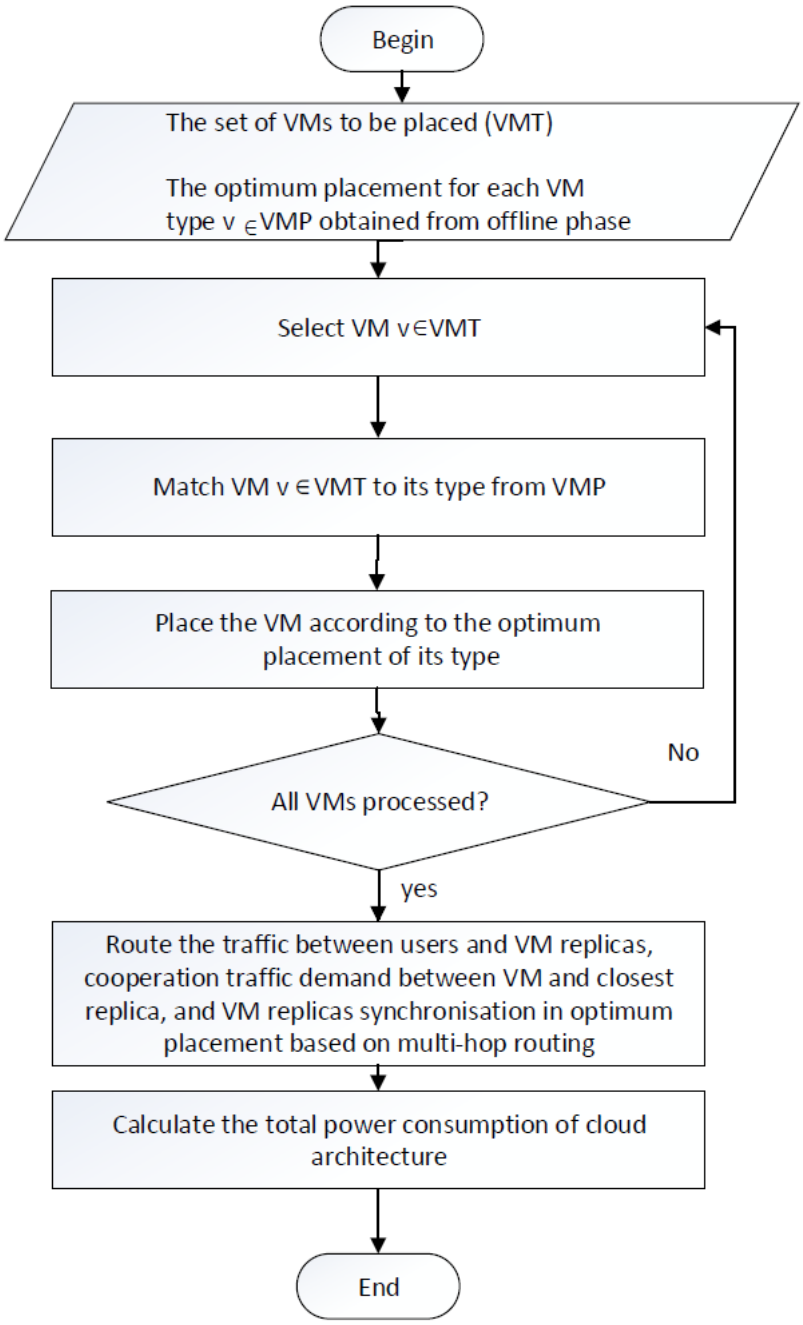

(b)

Figure 4-23: The flow chart of the EEVM-C-IVM heuristic (a) offline phase ((c) online phase.

Fig. 4-24 compares the total power consumption resulting from placing VMs using the EEVM-C-IVM heuristic with the power consumption resulting from the MILP model placement considering the NSFNET core network topology with the parameters mentioned in Section 4.3. The heuristic is evaluated for VMs under constant workload and linear workload (workload baselines of 1% and 5%), 5 and 25

Mbps users data rates and 100 Mbps and 5 Gbps inter-VM traffic. The gap between the EEVM-C-IVM heuristic and the MILP model range between 1% and 3% of the total power consumption in the synchronisation traffic scenario and cooperation and synchronisation traffic scenario, and less than 1% in the cooperation traffic scenario. The time required to run the offline phase is 155 minutes while the time required to run the online phase is 28 seconds. Running the MILP requires 1-3 days.

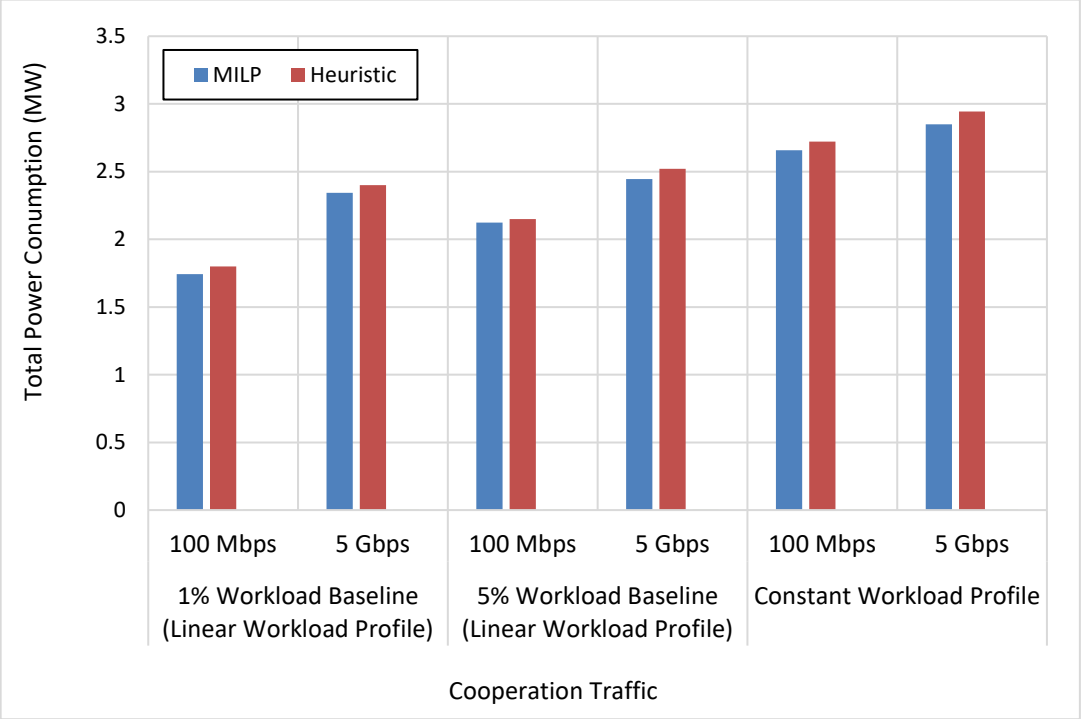

(a)

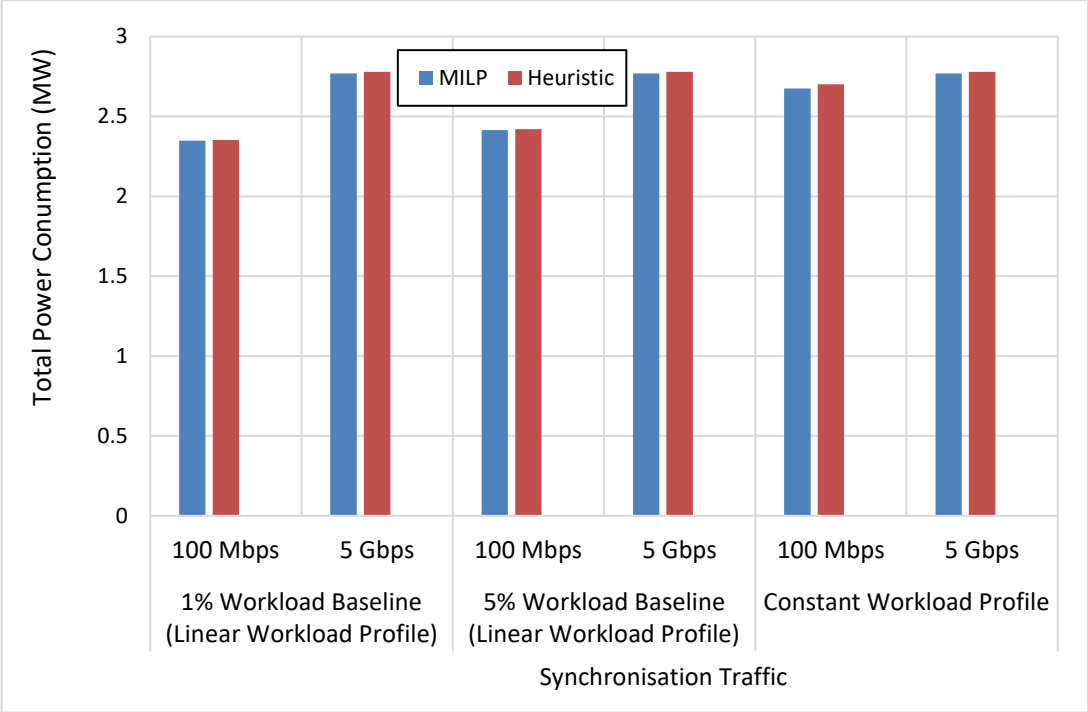

(b)

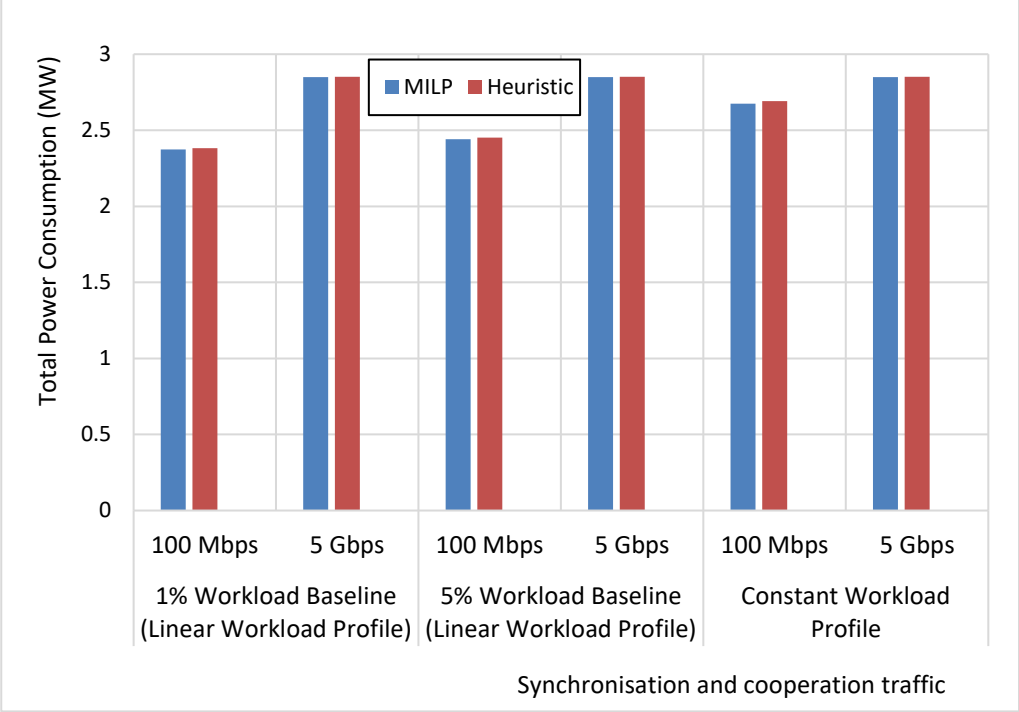

(c)

Figure 4-24: Comparison of total power consumption of (a) cooperation traffic scenario, (b) synchronisation traffic scenario, (c) cooperation and synchronisation traffic scenario, between the MILP model and EEVM-C-IVM heuristic considering VMs with constant workload and linear workload with 1% and 5% workload baseline under 100 Mbps and 5 Gbps inter-VM data rates.

4.5 Summary

In this chapter, the energy efficiency of geo-distributed VMs in IP over WDM core networks is investigated taking into consideration inter-VM cooperation traffic and synchronisation traffic between replicas of the same VM in addition to the download traffic from VMs to users. The problem was formulated as a mixed integer linear programming (MILP) model. Also, a real-time heuristic that mimics the MILP model has been developed.

For cooperation traffic, in addition to sending traffic to users, each VM has another VM to cooperate with. The objective is to find the optimum placement of VMs on distributed clouds and to create an optimum number of replicas that result in minimum power consumption. For synchronisation traffic, VM replicas can be created to serve distributed users, however, these replicas need to be synchronised to each other to keep the content at each location up to date.

Our results showed the dominating impact of synchronisation traffic on the placement of VMs, reducing the energy efficiency of replicating VMs. Neglecting inter-VM traffic when placing VMs can mount up the total power consumption by a factor of 39 for VMs with an inter-VM traffic data rate of 5 Gbps.

Chapter 5: Energy Efficient Virtual Machines Placement over Cloud-Fog Architecture

5.1 Introduction

In this chapter, the VMs placement architecture considered in Chapter 3 where users access VMs in the cloud is extended by adding fog computing layers at the metro and access networks to provision VMs in the proximity of users' premises. A comprehensive framework is developed based on mathematical modelling and heuristics to study the offloading of VMs from the cloud to the fog layers with the objective of minimising the total power consumption of providing the VMs. The placement of VMs in the cloud at the core network will allow VMs to serve users distributed across the core nodes whereas placing the VM replicas closer to the users in the fog nodes in the metro or access network will save the traffic between users and VMs from traversing the core network and therefore reduce the network power consumption, but will increase the processing power consumption due to the creation of multiple replicas of the VMs. Overall, the power consumption can be reduced if the VM users traffic is high and/or the VMs have a linear power profile. In such a linear profile, the creation of multiple VM replicas does not increase the power consumption significantly (there may be a slight increase due to idle / baseline power consumption) if the number of users remains constant.

5.2 MILP Model

In this section, a MILP model is developed to optimise the placement of VMs over the cloud-fog architecture so that the power consumption of providing the VMs is reduced. The architecture in Fig. 5-1 is considered where in addition to the clouds placed at the core network nodes, two fog layers are introduced at the metro network and the access network. The MILP model is used to select the most energy efficient

placement for each VM based on the VM popularity, workload and data rate. The model aims to achieve a trade-off between the network power saved by replicating VMs in multiple cloud and/or fog nodes and the extra power consumed by these replicas. The creation of a VM replica will result in power savings if the former power consumption exceeds the latter power consumption. The MILP model minimizes the power consumption of the end-to-end cloud-fog network architecture accounting for different communication and processing layers. The model's total power consumption comprises two main parts: telecommunication network layers (including core, metro, and access networks) and computing layers (including cloud, metro fog, and access fog). The first is the traffic-induced power consumption of communication networks generated by VM requests (placed either in the cloud or fog) and delivered to user premises in the access network. The second represents the processing-induced power consumption of placing VMs in cloud-fog data centres. The MILP model objective is subject to many constraints related to VM placement, communication network, and processing requirements and capabilities. For more clarity in MILP expressions and notations, we use superscripts to index the types of variables and parameters and subscripts for variables and parameters.

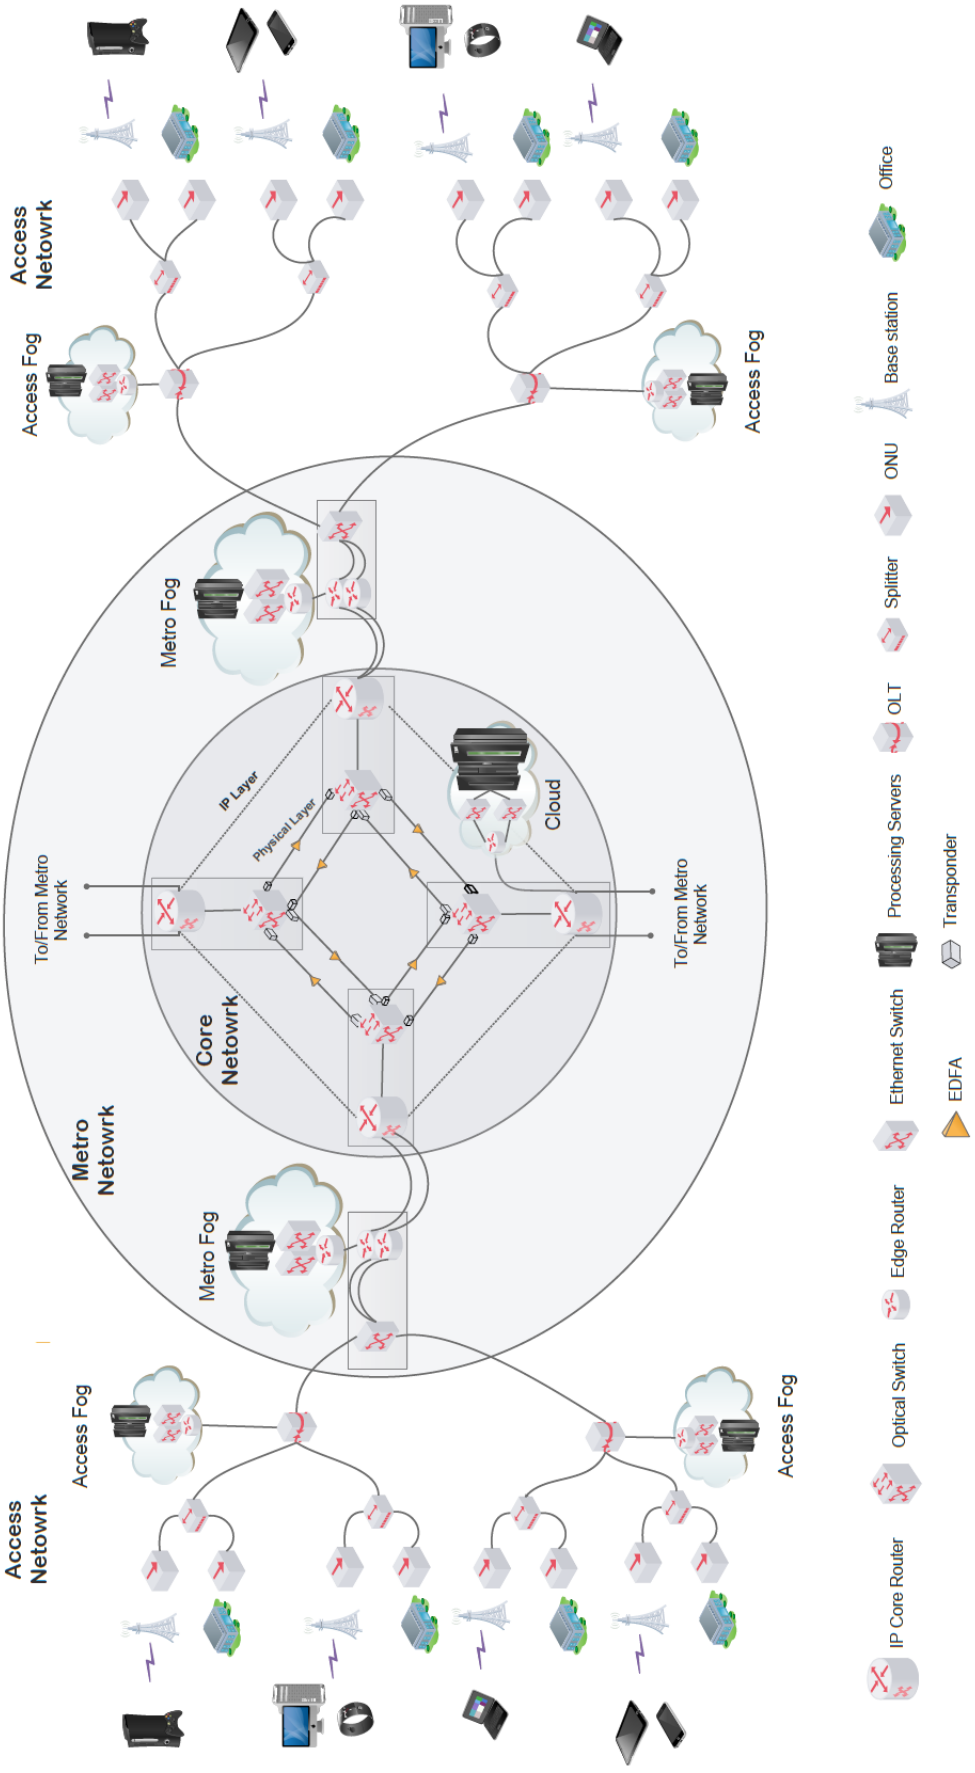

Figure 5-1: Cloud-Fog architecture

The resources at the fog nodes form a mini datacentre connected in a similar way to the cloud datacentres (as shown in Fig. 5.1). In addition to the cloud parameters and variables defined in Chapter 3, Table 5-1 and Table 5-2 define the parameters and variables of fog nodes.

Table 5-1: List of fog input parameters used in the MILP model

Parameter	Description
$\boldsymbol{SW^{(MFB)}}$	Metro fog switch bit rate.
$\boldsymbol{SW^{(MFP)}}$	Metro fog switch power consumption.
$\boldsymbol{SW^{(AFB)}}$	Access fog switch bit rate.
$\boldsymbol{SW^{(AFP)}}$	Access fog switch power consumption.
$\boldsymbol{R^{(MFB)}}$	Metro fog router port bit rate.
$\boldsymbol{R^{(MFP)}}$	Metro fog router port power consumption.
$\boldsymbol{R^{(AFB)}}$	Access fog router port bit rate.
$\boldsymbol{R^{(AFP)}}$	Access fog router port power consumption.
$\boldsymbol{m}$	Metro fog power usage effectiveness.
$\boldsymbol{a}$	Access fog power usage effectiveness.

Table 5-2: List of fog variables used in the MILP model

Variable	Description
$\boldsymbol{F_s^{(MF)}}$	$\boldsymbol{F_s^{(MF)} = 1}$ if a fog processing node is hosted in the metro network connected to core node $\boldsymbol{s}$, otherwise $\boldsymbol{F_s^{(MF)} = 0}$.
$\boldsymbol{\delta_{v,s}^{(MF)}}$	$\boldsymbol{\delta_{v,s}^{(MF)} = 1}$ if the fog processing node hosted in the metro network connected to node $\boldsymbol{s}$ hosts a replica of VM $\boldsymbol{v}$, otherwise $\boldsymbol{\delta_{v,s}^{(MF)} = 0}$.
$\boldsymbol{R_s^{(MF)}}$	Number of router ports used in the fog processing node hosted in the metro network connected to node $\boldsymbol{s}$.
$\boldsymbol{SW_s^{(MF)}}$	Number of switches used in the fog processing node hosted in the metro network connected to node $\boldsymbol{s}$.
$\boldsymbol{S_s^{(MF)}}$	Number of processing servers in the fog processing node hosted in the metro network connected to node $\boldsymbol{s}$.
$\boldsymbol{F_{p,s}^{(AF)}}$	$\boldsymbol{F_{p,s}^{(AF)} = 1}$ if a fog processing node is built in access network $\boldsymbol{p}$ connected to core node $\boldsymbol{s}$, otherwise $\boldsymbol{F_{p,s}^{(AF)} = 0}$.
$\boldsymbol{\delta_{v,p,s}^{(AF)}}$	$\boldsymbol{\delta_{v,p,s}^{(AF)} = 1}$ if the fog processing node in access network $\boldsymbol{p}$ connected to core node $\boldsymbol{s}$, hosts a replica of VM $\boldsymbol{v}$, otherwise $\boldsymbol{\delta_{v,p,s}^{(AF)} = 0}$.
$\boldsymbol{R_{p,s}^{(AF)}}$	Number of router ports used in the fog processing node located in the access network $\boldsymbol{p}$ connected to core node $\boldsymbol{s}$.
$\boldsymbol{SW_{p,s}^{(AF)}}$	Number of switches used in the fog processing node located in access network $\boldsymbol{p}$ connected to core node $\boldsymbol{s}$.

$S_{p,s}^{(AF)}$	Number of processing servers in the fog processing node located in the access network p connected to core node s.

In addition to the VMs parameters and variables defined in Chapter 3, the VMs to be hosted in the fog and the traffic resulting from them are defined by the variables in Table 5-3:

Table 5-3: List of VM variables used in the MILP model

Variable	Description
$W_{v,s}^{(MFR)}$	Workload of the VM replica v hosted in the fog processing node located in the metro network connected to node s.
$W_{s}^{(MF)}$	Total workload of the metro fog processing node located in core node s.
$D_{v,s}^{(MF)}$	Traffic from the VM replica v hosted in the fog processing node of the metro network connected to core node s.
$W_{v,p,s}^{(AFR)}$	Workload of the VM replica v hosted in the fog processing node located in the access network p connected to core node s.
$W_{p,s}^{(AF)}$	Total workload of the fog processing node located in the access network p connected to core node s.
$D_{v,p,s}^{(AF)}$	Traffic flow from the VM replica v hosted in the fog processing node located in the access network p connected to core node s.

In addition to the clouds, IP over WDM, metro and access power networks consumption equations (3.1) -(3.11) defined in Chapter 3, the metro fog and access fog nodes power consumptions are defined as follows:

The metro fog nodes power consumption is composed of:

- Power consumption of metro fog servers:

$$m \sum_{s \in N} S_{s}^{(MF)} S^{(P)} \qquad (5.1)$$

- Power consumption of metro fog switches and router ports:

$$m\left(\sum_{s \in N}\left(\left(SW_{s}^{(MF)} SW^{(R)} SW^{(MFP)}\right) + R_{s}^{(MF)} R^{(MFP)}\right)\right) \qquad (5.2)$$

The access fog nodes power consumption is composed of:

- Power consumption of access fog servers:

$$a \sum_{s \in N} \sum_{p \in P} S_{p,s}^{(AF)} \, S^{(P)} \tag{5.3}$$

- Power consumption of access fog switches and router ports:

$$a \left(\sum_{s \in N} \sum_{p \in P} \left(\left(SW_{p,s}^{(AF)} \, SW^{(R)} \, SW^{(AFP)} \right) + R_{p,s}^{(AF)} R^{(AFP)} \right) \right) \tag{5.4}$$

Note that similar to the cloud servers we consider an on-off power profile for the fog servers, i.e. if a server is activated, it will operate at maximum power consumption.

The model is defined as follows:

The objective: Minimise the total power consumption given as:

$$\begin{aligned}
& m \sum_{s \in N} S_s^{(MF)} \, S^{(P)} + m \left(\sum_{s \in N} \left(\left(SW_s^{(MF)} \, SW^{(R)} \, SW^{(MFP)} \right) + R_s^{(MF)} \, R^{(MFP)} \right) \right) \\
& + a \sum_{s \in N} \sum_{p \in P} S_{p,s}^{(AF)} \, S^{(P)} \\
& + a \left(\sum_{s \in N} \sum_{p \in P} \left(\left(SW_{p,s}^{(AF)} \, SW^{(R)} \, SW^{(AFP)} \right) + R_{p,s}^{(AF)} R^{(AFP)} \right) \right) \\
& + c \sum_{s \in N} S_s^{(C)} \, S^{(P)} + c \left(\sum_{s \in N} \left(\left(SW_s^{(C)} \, SW^{(R)} \, SW^{(CP)} \right) + R_s^{(C)} \, R^{(CP)} \right) \right) \\
& + n \left(\sum_{p \in P} \sum_{d \in N} \left(OLT^{(P)} \, OLT_{p,d}^{(N)} \right) \right) + n \left(\sum_{p \in P} \sum_{d \in N} \left(ONU^{(P)} \, ONU_{p,d}^{(N)} \right) \right) \\
& + n \left(\sum_{s \in N} R_s^{(M)} \, R^{(MR)} R^{(MP)} \right) + n \left(\sum_{s \in N} SW_s^{(M)} \, SW^{(MP)} \right) \\
& + n \left(\sum_{s \in N} R^{(P)} \, R_s^{(AC)} + \sum_{d \in N} R^{(P)} R_d^{(AE)} + \sum_{m \in N} \sum_{n \in N m_m : n \neq m} R^{(P)} \, W_{mn} \right) \\
& + n \left(\sum_{m \in N} \sum_{n \in N m_m : n \neq m} t^{(P)} \, W_{mn} \right) + n \left(\sum_{m \in N} \sum_{n \in N m_m : n \neq m} e^{(P)} \, F_{mn} \, A_{mn} \right)
\end{aligned}$$

$$+ n\left(\sum_{s\in N} SW_s^{(P)}\right) + n\left(\sum_{m\in N}\sum_{n\in N m_m : n\neq m} G^{(P)}\, G_{m,n}\, W_{m,n}\right) \tag{5.5}$$

Equation (5.5) gives the total power consumption as the sum of the power consumption of the access fogs, the metro fogs, the clouds, the access network, the metro network and the IP over WDM core network.

Subject to:

In addition to the constraints (3-14) - (3-29), the model is subject to the following constraints:

Serving VM demands:

$$\sum_{p\in P}\sum_{d\in N} D_{v,p,d} = \sum_{s\in N}\sum_{d\in N} D_{v,s,d}^{(C)} + \sum_{s\in N} D_{v,s}^{(MF)} + \sum_{p\in P}\sum_{s\in N} D_{v,p,s}^{(AF)}$$

$$\forall\ v \in VM \tag{5.6}$$

Constraint (5.6) ensures that all the user demands for a VM are served by the clouds, the metro fogs or the access fogs.

Placing VMs in metro fog constraints:

$$D_{v,s}^{(MF)} \geq \delta_{v,s}^{(MF)}$$

$$\forall\ s \in N, v \in VM \tag{5.7}$$

$$D_{v,s}^{(MF)} \leq L\, \delta_{v,s}^{(MF)}$$

$$\forall\ s \in N, v \in VM \tag{5.8}$$

Constraints (5.7) and (5.8) relate the binary variable that indicates whether a VM is hosted in a metro fog or not ($\delta_{v,s}^{(MF)}$) to the traffic between users of this VM and the metro fog ($D_{v,s}^{(MF)}$) by setting $\delta_{v,s}^{(MF)} = 1$ if $D_{v,s}^{(MF)} > 0$ and $\delta_{v,s}^{(MF)} = 0$ otherwise.

Placing VM in access fog constraints:

$$D_{v,p,s}^{(AF)} \geq \delta_{v,p,s}^{(AF)}$$

$$\forall s \in N, v \in VM, p \in P \quad (5.9)$$

$$D_{v,p,s}^{(AF)} \leq L\, \delta_{v,s}^{(AF)}$$

$$\forall s \in N, v \in VM, p \in P \quad (5.10)$$

Constraint (5.9) and (5.10) relate the binary variable that indicates whether a VM is hosted in an access fog or not ($A\delta_{v,p,s}^{(AF)}$) to the traffic between users of this VM and the cloud ($D_{v,p,s}^{(AF)}$), by setting $AF\delta_{vsp} = 1$ if $D_{v,p,s}^{(AF)} > 0$ and $\delta_{v,s}^{(AF)} = 0$ otherwise.

Metro fog location constraints:

$$\sum_{v \in VM} \delta_{v,s}^{(MF)} \geq F_s^{(MF)}$$

$$\forall s \in N \quad (5.11)$$

$$\sum_{v \epsilon VM} \delta_{v,s}^{(MF)} \leq L\, F_s^{(MF)}$$

$$\forall s \in N \quad (5.12)$$

Constraints (5.11) and (5.12) ensure that a metro fog is built in metro nodes selected to host VMs by setting $F_s^{(MF)} = 1$ if $\sum_{v \epsilon VM} \delta_{v,s}^{(MF)} > 0$ and $F_s^{(MF)} = 0$ otherwise.

Access fog location constraints:

$$\sum_{v \epsilon VM} \delta_{v,p,s}^{(AF)} \geq F_{p,s}^{(AF)}$$

$$\forall s \in N \quad (5.13)$$

$$\sum_{v \epsilon VM} \delta_{v,p,s}^{(AF)} \leq L\, F_{p,s}^{(AF)}$$

$$\forall s \in N \quad (5.14)$$

Constraint (5.13) and (5.14) ensure that an access fog is built in access nodes selected to host VMs by setting $F_{p,s}^{(AF)} = 1$ if $\sum_{v \epsilon VM} \delta_{v,p,s}^{(AF)} > 0$ and $F_{p,s}^{(AF)} = 0$ otherwise.

Fog workload:

$$W_{v,s}^{(MFR)} = \delta_{v,s}^{(MF)} \; W_v \qquad (\textbf{Constant Workload Profile})$$
$$\forall \; v \in VM, s \in N \tag{5.15}$$

$$W_{v,s}^{(MFR)} = \left(\frac{D_{v,s}^{(MF)}}{r_v \; x} \; M \; \delta_{v,s}^{(MF)} \right)$$
$$+ \left(W_v^{(R)} \sum_{d \in N} D_{v,s,d}^{(MF)} \right) \quad (\textbf{Linear Workload Profile})$$
$$\forall \; v \in VM, s \in N \tag{5.16}$$

$$W_s^{(MF)} = \sum_{v \epsilon VM} W_{v,s}^{(MFR)}$$
$$\forall \, s \in N \tag{5.17}$$

$$W_{v,p,s}^{(AFR)} = \delta_{v,p,s}^{(AF)} \; W_v \qquad (\textbf{Constant Workload Profile})$$
$$\forall \; v \in VM, s \in N, p \in P \tag{5.18}$$

$$W_{v,p,s}^{(AFR)} = \left(\frac{D_{v,p,s}^{(AF)}}{r_v \; x} \; M \; \delta_{v,p,s}^{(AF)} \right)$$
$$+ \left(W_v^{(R)} D_{v,p,s}^{(AF)} \right) \qquad (\textbf{Linear Workload Profile})$$
$$\forall \; v \in VM, s \in N, p \in P \tag{5.19}$$

$$W_{p,s}^{(AF)} = \sum_{v \epsilon VM} W_{v,p,s}^{(AFR)}$$
$$\forall \, s \in N \tag{5.20}$$

Constraints (5.15) and (5.18) calculate the workload of a VM replica under a constant workload profile in a metro fog and an access fog, respectively. Constraints (5.16) and (5.19) calculate the workload of a VM replica in a metro fog and an access fog, respectively as a linear function of the traffic resulting from serving users of the replica with a minimum CPU usage. Note that in the model, a VM workload is calculated as a function of the traffic associated with the users served by it. Constraints (5.17) and (5.20) calculate the total workload of a metro fog and an access fog, respectively by summing the workload of VMs hosted in it.

Number of servers in fog:

$$S_s^{(MF)} \geq \frac{W_s^{(MF)}}{S^{(maxW)}}$$

$$\forall s \in N \quad (5.21)$$

$$S_{p,s}^{(AF)} \geq \frac{W_{p,s}^{(AF)}}{S^{(maxW)}}$$

$$\forall s \in N, p \in P \quad (5.22)$$

Constraints (5.21) and (5.22) calculate the number of servers in each metro fog and access fog, respectively, based on the CPU utilisation.

Number of router ports and switches in fog:

$$R_s^{(MF)} \geq \frac{\sum_{v \in VM} D_{v,s}^{(MF)}}{R^{(MFB)}}$$

$$\forall s \in N \quad (5.23)$$

$$SW_s^{(MF)} \geq \frac{\sum_{v \in VM} D_{v,s}^{(MF)}}{SW^{(MFB)}}$$

$$\forall s \in N \quad (5.24)$$

$$R_{p,s}^{(AF)} \geq \frac{\sum_{v \in VM} D_{v,p,s}^{(AF)}}{R^{(AFB)}}$$

$$\forall s \in N, p \in P \quad (5.25)$$

$$SW_{p,s}^{(AF)} \geq \frac{\sum_{v \in VM} D_{v,p,s}^{(AF)}}{SW^{(AFB)}}$$

$$\forall s \in N, p \in P \quad (5.26)$$

Constraints (5.23) -(5.26) calculate the number of router ports and switches in each metro fog and access fog.

Number of metro routers ports and ethernet switches in metro network:

$$R_s^{(M)} \geq \frac{\sum_{v \in VM} \sum_{s \in N} D_{v,s,d}^{(C)} + \sum_{v \in VM} D_{v,s}^{(MF)}}{R^{(MB)}}$$

$$\forall s \in N \quad (5.27)$$

$$SW_s^{(M)} \geq \frac{\sum_{v \in VM} \sum_{s \in N} D_{v,s,d}^{(C)} + \sum_{v \in VM} D_{v,s}^{(MF)}}{SW^{(MB)}}$$

$$\forall \boldsymbol{s} \in \boldsymbol{N} \qquad (\mathbf{5.28})$$

Constraints (5.27) and (5.28) calculate the number of routers ports and switches, respectively, in each metro network.

5.3 Results and Discussions

In the following, the problem of VMs placement in the cloud-fog architecture considering the US's AT&T network and the UK's BT network is investigated.

5.3.1 AT&T Network Use Case:

In this section, the optimal VMs placement over AT&T cloud-fog architecture is investigated as the first use case (AT&T core networks topology is illustrated in Fig. 5-2 [144]). We start by optimising a single VM as the simplest representative problem. Then we consider optimisation in a realistic scenario with multiple VMs.

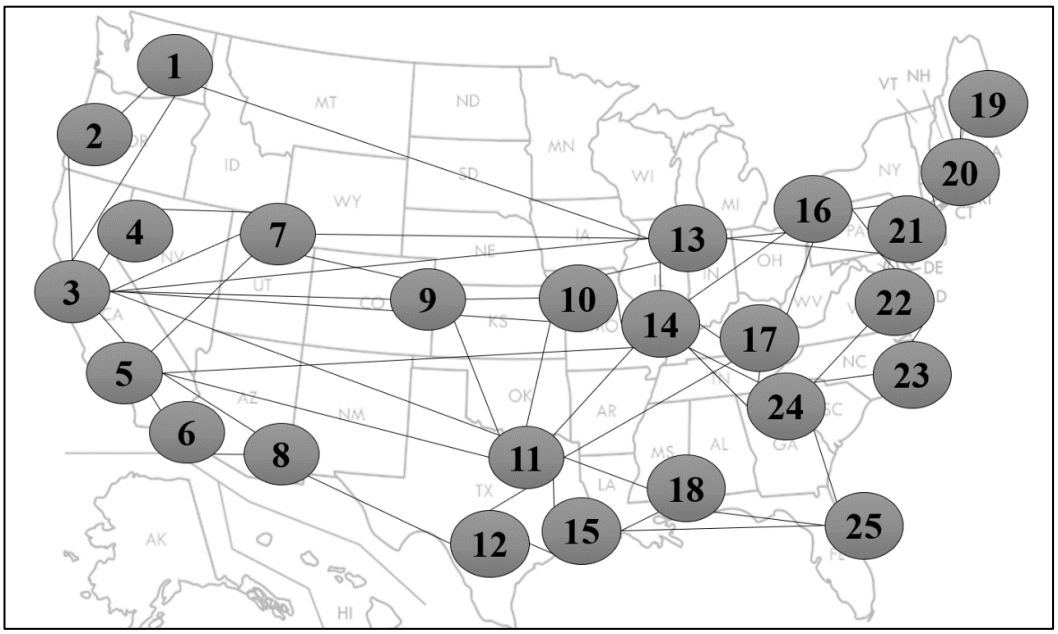

Figure 5-2: AT&T core network topology.

5.3.1.1 Simple Representative Scenario:

This section investigates how the energy efficient placement of a single VM over cloud-fog architecture will vary based on three factors; the CPU requirements, download traffic and PUE values.

Similar to the analysis carried out in Chapter 3, the impact of the VM workload on the VM placement is examined by considering the constant and linear workload profiles considering the same input parameters. The VM is considered to has 800 users. The workload of the constant workload profile and the workload of serving the maximum number of users of a VM with linear workload are considered to have one of three workloads: 10%, 50% and 100% of the server CPU capacity. VM of linear profile is considered to have no baseline. The users are considered to access the VM with one of following download rates; 0.1 Mbps, 1 Mbps, 10 Mbps, 20 Mbps, 50 Mbps, 100 Mbps or 200 Mbps.

Based on the US data centre energy usage [146], the PUE varies based on the datacentre size as more efficient cooling technologies are used in larger datacentres. For best practice datacentres, PUE of clouds, metro fogs and access fogs take the values of 1.3, 1.4 and 1.5, receptively [146]. For datacentres from 2014, the PUE values considered are 1.7, 1.9 and 2.5, respectively [146] . In network infrastructures, a typical telecom office PUE value is 1.5 [75].

In this scenario, the model takes into consideration the Cisco CRS-1 router [148] as a core network IP router which consumes 638W per 40 Gbps port and the Cisco NCS 5502 router [153] as the cloud and metro networks router which consumes 30W per 40 Gbps port. In the metro and fog datacentre, Cisco NCS 5501 [153] is considered with a power consumption of 13W per 40 Gbps port. Furthermore, the Cisco Nexus 93180YC-EX [154] switch is considered as metro, cloud and metro fog LAN Ethernet switch with upload capacity of 600 Gbps and power rating at 470W. In access fog, the Cisco Nexus 93180YC-EX [154] switch is considered with capacity of 240 Gbps while consuming 210W.

Table 5-4, Table 5-5, Table 5-6 and Table 5-7 define the IP over WDM network, metro network, access network and cloud-fog input parameters, respectively.

Table 5-4: IP Over WDM core network input parameters used in the model

40 Gbps router port power consumption ($\boldsymbol{R^{(P)}}$)	638W [148]
40 Gbps transponder power consumption ($\boldsymbol{t^{(P)}}$)	129W [149]

40 Gbps regenerator power consumption ($\boldsymbol{G^{(P)}}$)	114W, reach 2000 km [150]
EDFA power consumption ($\boldsymbol{e^{(P)}}$)	11W [151]
Optical switch power consumption ($\boldsymbol{SW^{(P)}}$)	85W [152]
Number of wavelengths in a fibre ($\boldsymbol{\mathcal{W}}$)	32 [20]
Bit rate of each wavelength ($\boldsymbol{\mathcal{W}^{(P)}}$)	40 Gbps [20]
Span distance between two EDFAs ($\boldsymbol{S}$)	80 km [151]
Network power usage effectiveness ($\boldsymbol{n}$)	1.5 [6]

Table 5-5: Metro network input parameters used in the model

Metro router redundancy ($\boldsymbol{R^{(MR)}}$)	2
Metro router port bit rate ($\boldsymbol{R^{(MB)}}$)	40 Gbps
Metro router port power consumption ($\boldsymbol{R^{(MP)}}$)	30W [153]
Metro ethernet switch bit rate ($\boldsymbol{SW^{(MB)}}$)	600 Gbps [154]
Metro ethernet switch power consumption ($\boldsymbol{SW^{(MP)}}$)	470W [154]

Table 5-6: Access network input parameters used in the model

Number of PON networks in a node ($\boldsymbol{P}$)	2
Maximum number of single VM users ($\boldsymbol{x}$)	800 concurrent users
Number of ONU devices in a PON network ($\boldsymbol{ONU_{p,d}^{(N)}}$)	512
Power consumption of ONU device ($\boldsymbol{ONU^{(P)}}$)	5W [155]
Number of OLTs in a PON network ($\boldsymbol{OLT_{p,d}^{(N)}}$)	1
OLT Capacity ($\boldsymbol{OLT^{(B)}}$)	1280 Gbps [145]

Table 5-7: Clouds and fog input parameters used in the model

Number of VMs ($\boldsymbol{V}$)	1
User download rate ($\boldsymbol{r_v}$)	{0.1, 1, 10, 20, 50, 100 or 200 Mbps}
Maximum workload of VM ($\boldsymbol{W_v}$)	10%, 50% and 100%
Server power consumption ($\boldsymbol{S^{(P)}}$)	333 Watt [156]
Maximum server workload ($\boldsymbol{S^{(maxW)}}$)	100%

Cloud and metro fog switch bit rate ($\boldsymbol{SW^{(CB)}}$, $\boldsymbol{SW^{(MFB)}}$)	600 Gbps [154]
Cloud and metro fog switch power consumption ($\boldsymbol{SW^{(CP)}}$, $\boldsymbol{SW^{(MFP)}}$)	470 Watt [154]
Access fog switch bit rate ($\boldsymbol{SW^{(AFB)}}$)	240 Gbps [154]
Access fog switch power consumption ($\boldsymbol{SW^{(MFP)}}$).	210 Watt [154]
Cloud and fog switch redundancy ($\boldsymbol{SW^{(R)}}$)	2
Cloud and fog router port bit rate ($\boldsymbol{R^{(CB)}}$, $\boldsymbol{R^{(MFB)}}$, $\boldsymbol{R^{(AFB)}}$)	40 Gbps
Cloud router port power consumption ($\boldsymbol{R^{(CP)}}$)	30 Watt [153]
Metro and access fog router port power consumption ($\boldsymbol{R^{(MFP)}}$, $\boldsymbol{R^{(AFP)}}$)	13 Watt [153]
Cloud power usage effectiveness ($\boldsymbol{c}$)	1.3 or 1.7 [146]
Metro fog power usage effectiveness ($\boldsymbol{m}$)	1.4 or 1.9 [146]
Access fog power usage effectiveness ($\boldsymbol{a}$)	1.5 or 2.5 [146]

Fig. 5-3 (a), (b) and (c) show the optimal placement of VM of 10%, 50% and 100% CPU requirements, respectively, considering the best practice PUE values. In each figure, the x-axis is the VM workload profile, the y-axis is the data rates which range from 0.1 Mbps to 200 Mbps and the z-axis is the percentage of VM replicas in each location over the cloud-fog architecture. The placement of VMs in the cloud at the core network will allow VMs to serve users distributed across the core nodes whereas placing the VM replicas closer to the users in the fog nodes in the metro or access network will save the traffic between users and VMs from traversing the core network and therefore reduce the network power consumption, but will increase the processing power consumption due to the creation of multiple replicas of the VMs.

As discussed in Chapter 3, the placement of VM with linear workload profile is not affected by the VM workload as serving users will consume the same power whether centralised in a single VM or distributed among multiple replicas with smaller workloads. However, the higher PUE of fog nodes compared to the cloud, results in a situation where distributing replicas into fogs processing nodes incurs additional power consumption as the PUE value of fog nodes is higher than clouds. Hence, there is a trade-off between network power saved by replicating VMs into fogs

and the additional power consumed by these replicas. The creation of a VM replica will result in power savings if the former power exceeds the latter power. At data rates of 1 Mbps and higher, VMs of 10%, 50% and 100% workloads are offloaded to access fog processing nodes considering a linear workload profile.

For constant workload profile, replicas are less energy efficient, therefore, offloading VMs to fog nodes decreases as the VM workload increases. While VMs of 10% workload and 20 Mbps are fully offloaded to metro fogs, 50% and 100% workload VMs are replicated only to clouds. Also, users of VM of 50% workload at 100 Mbps data rate as well as VMs of 100% workload at 200 Mbps data rate are served by clouds and metro fogs. A VM replica is offloaded to 14 metro fogs (in core nodes 1, 2, 4, 6, 7, 8, 13, 16, 19, 20, 21, 22, 23, 25) while users from other nodes are served by the replica placed in the cloud in core node 11 which they can access by traversing a single hop in the core network. These 14 metro fog nodes are selected to host replicas of the VM as the traffic flows will traverse more than a single hop in the IP over WDM network to access the VM placed in the cloud hosted in node 11 and therefore increase the needs for IP router ports (the most power consuming device in the IP over WDM network).

The results also show that VM with higher data rates justify the creation of more replicas closer to users premises in the fog layer, thus, the power consumption of networks, which is the major contributor to the power consumption in the cloud-fog architecture, is reduced. For example, VMs of 10% workload under the linear workload profile, are fully replicated to clouds and offloaded to access fogs for VMs of 0.1 Mbps, and $\geq$ 1 Mbps user data rates, respectively.

Placing VMs in cloud architecture with higher PUE (2014 PUE), as in Fig. 5-4, increases the replicas power consumption and therefore limits offloading VMs into the fog processing nodes, e.g., VM of constant workload profile with 100% workload and 200 Mbps data rate, that are fully offloaded to metro fogs considering clouds of best practice PUE, are limited to clouds with 2014 PUE.

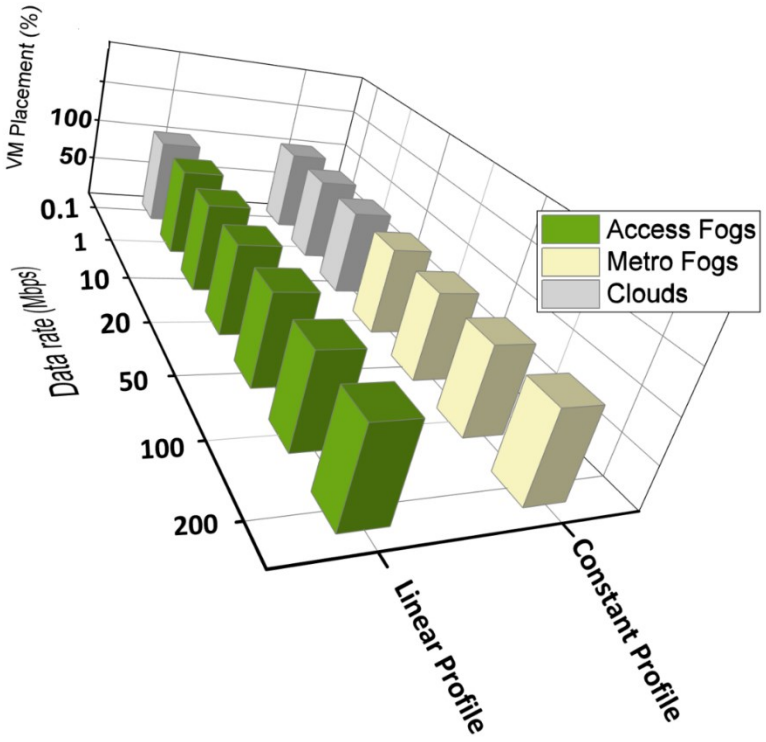
VM Placement (%)
100
50
0.1
1
10
20
50
100
200
Data rate (Mbps)
Linear Profile
Constant Profile
Access Fogs
Metro Fogs
Clouds

(a)

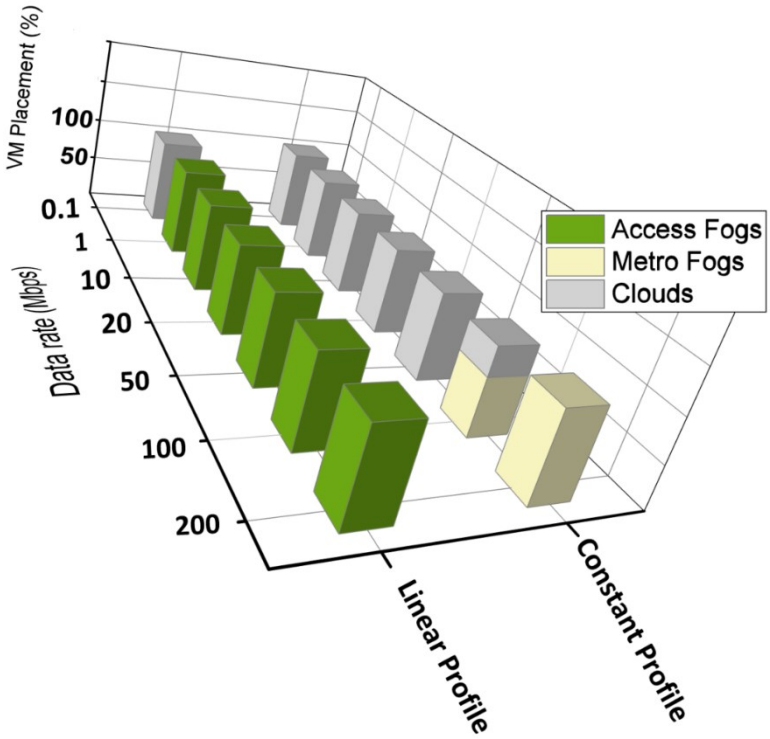

(b)

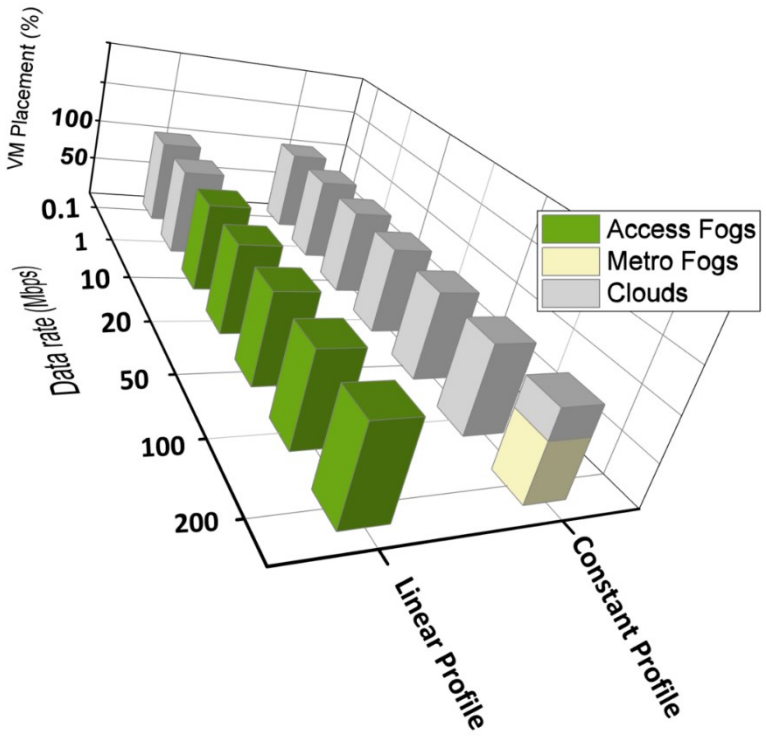

(c)

Figure 5-3: Optimal VM placement of (a) constant profile at 10% of CPU and linear profile with peak utilisation at 10%, (b) 50% case, (c) 100% case at different data rates considering best practice PUE value ($\boldsymbol{c}$=1.3, $\boldsymbol{m = 1.4, a = 1.5}$).

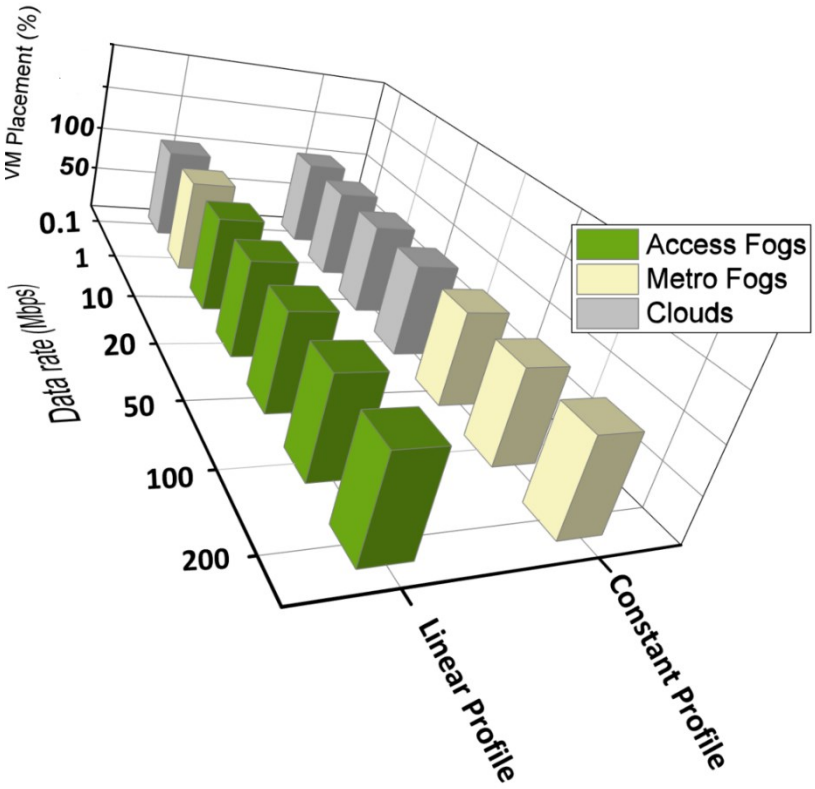
VM Placement (%)
100
50
0.1
1
10
20
50
100
200
Data rate (Mbps)
Linear Profile
Constant Profile
Access Fogs
Metro Fogs
Clouds

(a)

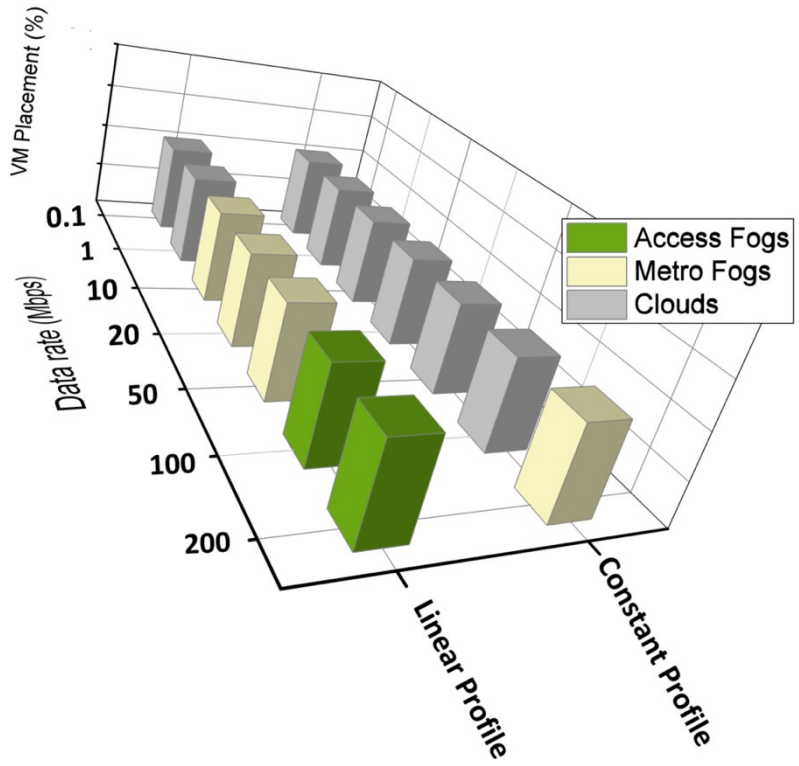
VM Placement (%)
0.1
1
10
20
50
100
200
Data rate (Mbps)
Linear Profile
Constant Profile
Access Fogs
Metro Fogs
Clouds

(b)

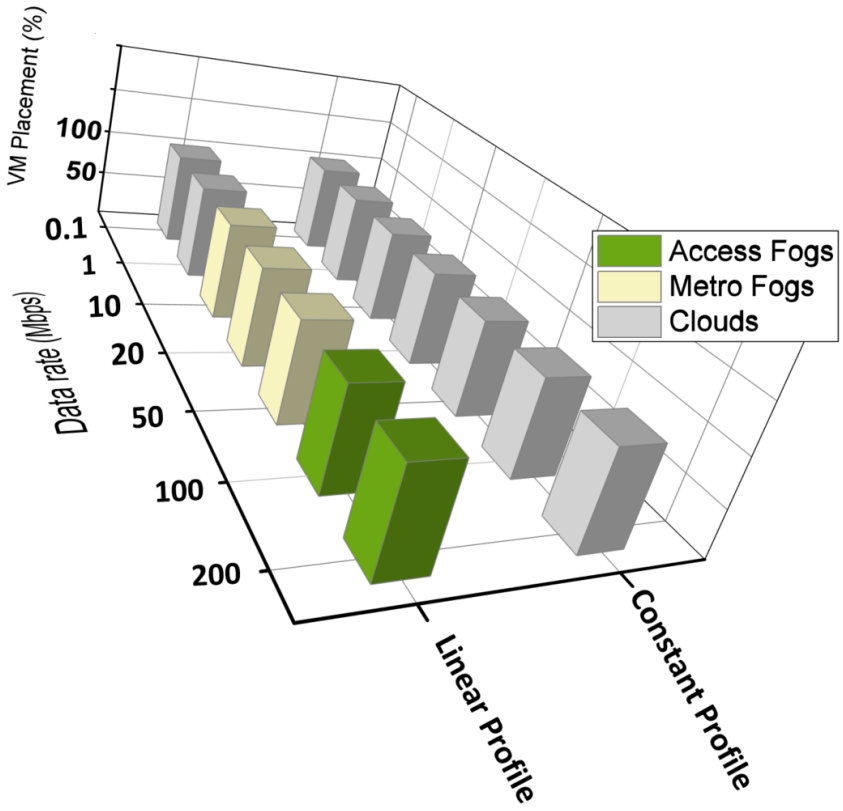

(c)

Figure 5-4: Optimal VM placement of (a) constant profile at 10% of CPU and linear profile with peak utilisation at 10%, (b) 50% case, (c) 100% case at different data rates considering 2014 PUE value ($\boldsymbol{c}$=1.7, $\boldsymbol{m = 1.9, a = 2.5}$).

5.3.1.2 Realistic Scenario

In this scenario, analysis based on realistic number of users and VM popularity is studied. The number of users, number of VMs, VM popularity, VM workload, users' data rate are considered to be the same as in Chapter 3. The PUE values are considered to follow the best practice PUE, where the PUE values for cloud, metro fog and access fog are 1.3, 1.4 and 1.5, respectively.

The optimised VMs placement over the cloud-fog architecture, referred to as Optimised clouds and fogs placements (OC&F) approach, is compared to the OC approach (the optimised VMs placement over distributed clouds) and AT&T clouds investigated in Chapter 3.

Two network components scenarios are considered to show how the energy efficiency of network components is a key factor in the placement of VMs across the cloud-fog architecture.

a) Network Components with Low Energy Efficiency:

In the following, the OC&F approach is compared to the AT&T clouds and OC approach considering the different linear workload baseline scenarios 1%, 5% and 40% discussed above. In addition to the parameters in Table 5-4 to Table 5-7, Table 5-8 shows the additional/modified parameters considered for the following results.

Table 5-8: Input parameters used in the model

Number of VM users in each PON based on VM popularity groups ($\boldsymbol{U_{v,p,d}}$)	13,000 users in each PON, six VMs popularity groups; 16%, 5%, 2%, 1%, 0.5% and 0.05%
Number of VMs ($\boldsymbol{V}$)	300
User download rate ($\boldsymbol{r_v}$)	{1, 10 or 25 Mbps}
Maximum workload of VM ($\boldsymbol{W_v}$)	50%
Cloud power usage effectiveness ($\boldsymbol{c}$)	1.3 [146]
Metro fog power usage effectiveness ($\boldsymbol{m}$)	1.4 [146]
Access fog power usage effectiveness ($\boldsymbol{a}$)	1.5 [146]

i. Linear Workload (1% Workload Baseline):

Fig. 5-5 shows the power consumption resulting from placing VMs of 1% workload baseline considering the different placement approaches under 1, 10 and 25 Mbps user data rates. The efficiency of VMs with 1% workload baseline allows the creation of more efficient VM replicas as the workload is proportional to the number of users served by the VM with a limited workload baseline required by each VM. Under 1 Mbps data rate, the OC&F approach achieves 6% reduction in the total power consumption compared to the AT&T clouds. The total reductions widen more to 40% under 10 Mbps data rate and 64% under 25 Mbps data rate. The savings achieved by the OC&F approach compared to the OC approach are 4%, 31% and 48% under the low, medium and high data rates, respectively.

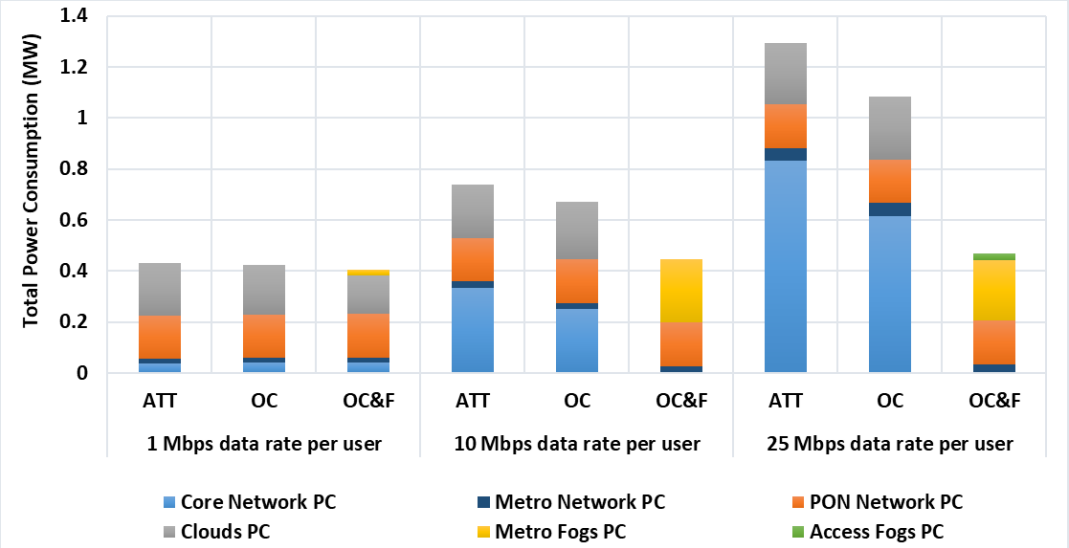

Figure 5-5: The power consumption of different VMs placement approaches considering VMs of 1% workload baseline.

In Fig. 5-6, the OC&F placement approach is further investigated by looking at how VMs of different data rates and popularity groups are placed considering the low, medium and high data rates. Note that the different colours show if VM of a certain popularity is placed in this location or not, i.e. it does not represent the number of replicas.

Figure 5-6(a) shows that VMs with a low user data rate of 1 Mbps have only justified creating three metro fogs in nodes 6, 8, and 19 as the traffic flows from these nodes traverses more than a single hop in the IP over WDM network to access the replicas optimally placed in the distributed clouds built in nodes 3, 11, 20, and 24. Thus, these fog nodes are built to serve the users demand locally, and consequently, eliminate the needs for IP router ports. However, VMs with the lowest popularity (0.05%) have only justified the creation of two replicas only in nodes 11 and 20.

VMs with 10 Mbps data rate are fully offloaded to every metro fog as shown in Fig. 5-6(b). Note that, VMs users are uniformly distributed across the metro and access networks, thus, the placement of a VM is consistent across all the metro fogs. In Fig. 5-7(c), VMs with high data rate of 25 Mbps show that in addition to a full replication in metro fogs, VMs with 16% popularity group justified creating VM replicas in some access fogs. Although, we are able to reduce the traffic traversing the metro network and consequently reduce the total power consumption, however, VMs with 16% popularity are not fully replicated to access fogs. There are a number of replicas offloaded to metro fogs. The reason for that is the on-off power consumption profile

of fog and network resources. Thus, before creating a new fog node in the access network, VMs are consolidated into the available resources remained from the placement of other VMs that share the same architecture.

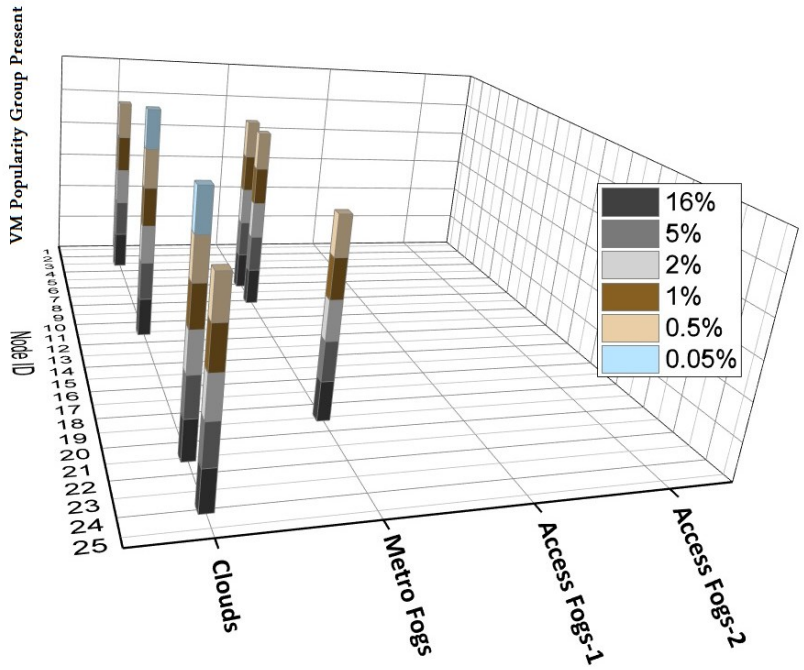

(a)

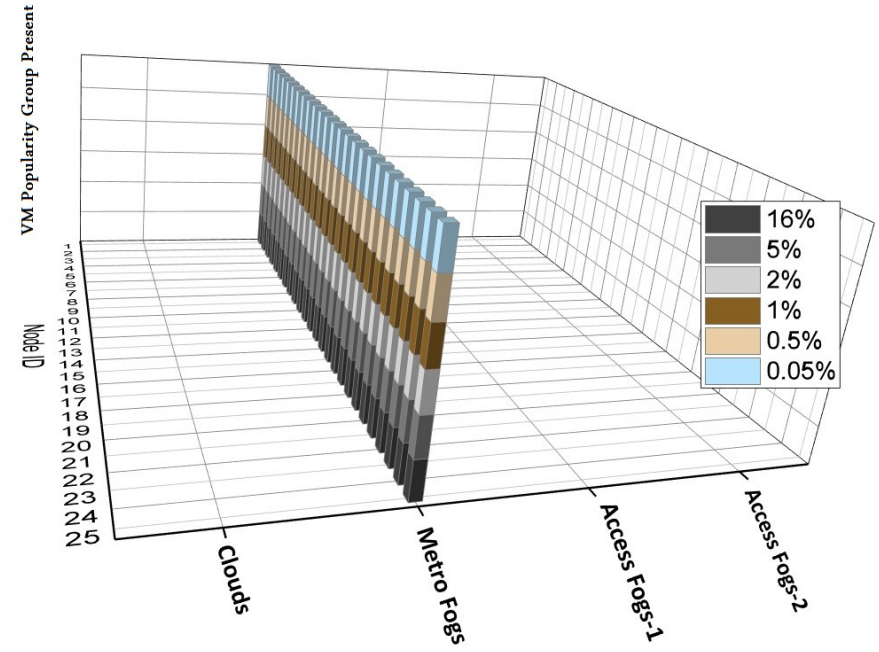

(b)

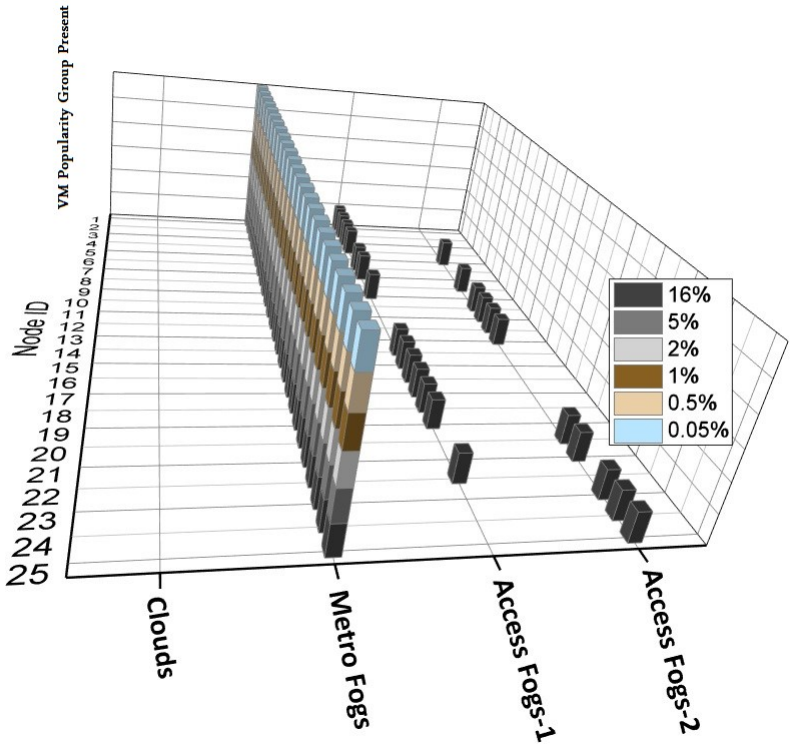

(c)

Figure 5-6: Optimal placement of different VMs popularity groups of 1% workload baseline under the OC&F approach with (a) 1 Mbps data rate per user, (b) 10 Mbps data rate per user and (c) 25 Mbps data rate per user.

In Fig. 5-7, OC&F1 and OC&F2 placement approaches are introduced. The former represents the optimal placement considering clouds and metro fogs only and the latter shows the optimal placement considering the three computing layers; clouds, metro

and access fogs. These two approaches show how introducing fog nodes in access network (OC&F2), in addition to metro fogs, are able to save total power consumption compared to an approach that considers only fog nodes connected to metro network (OC&F1). Under a 25 Mbps user data rate, it can be observed that the OC&F2 approach achieves 6% extra power saving compared to the OC&F1 approach.

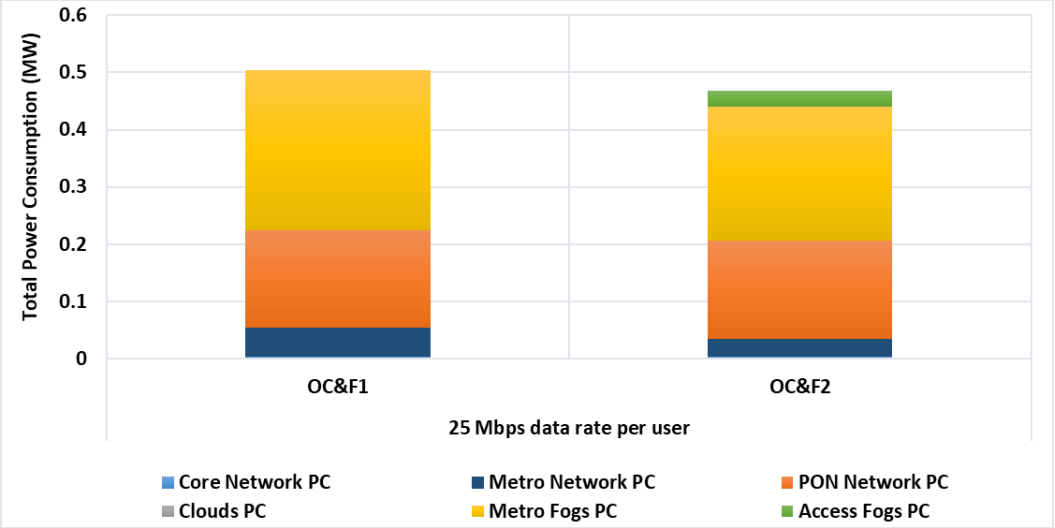

Figure 5-7: The power consumption considering OC&F1 and OC&F2 placement approaches. OC&F1 represents the optimal placement considering clouds and metro fogs only and OC&F2 represents the optimal placement considering clouds, metro and access fogs.

Fig. 5-8 shows the number of servers required to host VM replicas under the OC&F approach. The number of servers is a function of the number of VM replicas hosted and their workload. For instance, in the OC&F approach under 25 Mbps user data rate (Fig. 5-9(c)), two servers are activated in the access nodes selected to host replicas of VMs of 16% popularity. Also, 18 servers are activated in each metro fog in order to serve users of VMs with ≤ 5% popularity as well users of VMs with 16% popularity who are users not served by the replica created in the access fogs. Such number of servers can be practically attached to the metro edge routers to create the metro fog layer and to OLT in the access network to create the access fog layer.

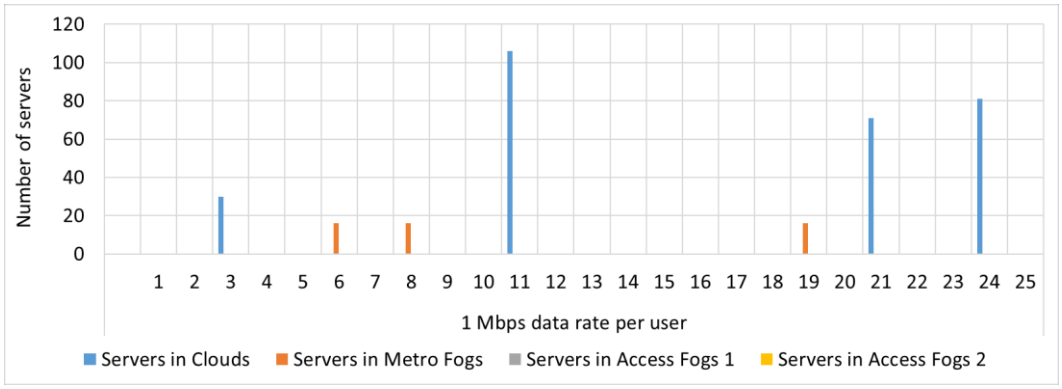

(a)

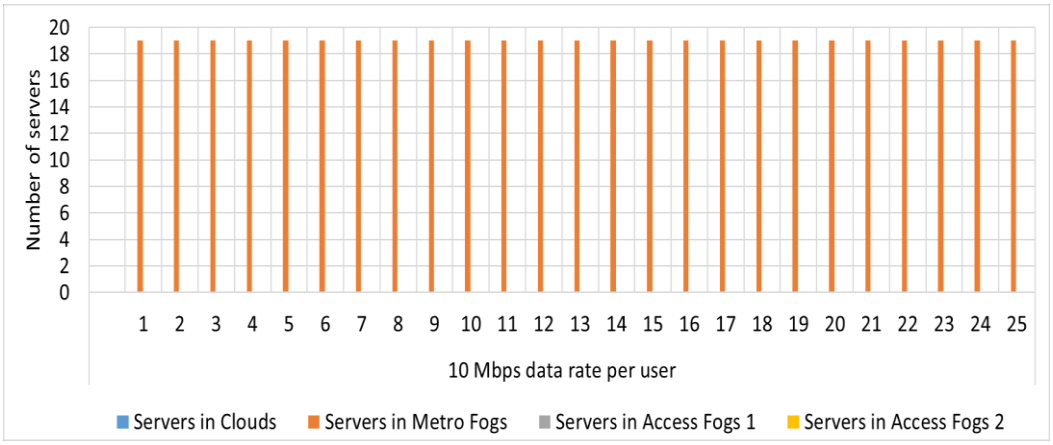

(b)

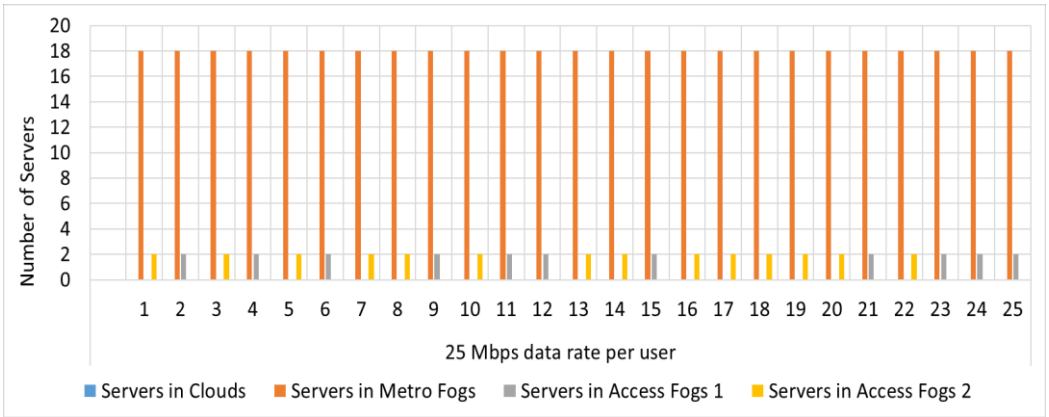

(c)

Figure 5-8: Number of servers in OC&F approach required to host VMs of 1% workload baseline with (a) 1 Mbps data rate per user (b) 10 Mbps data rate per user (c) 25 Mbps data rate per user.

ii. Linear Workload (5% Workload Baseline):

Fig. 5-9 shows the power savings achieved under VMs of linear workload profile of 5% minimum CPU utilisation. Increasing the minimum CPU utilisation of the VM workload profile to 5% reduces the efficiency of creating more VMs replicas. The

total savings achieved under the OC&F approach compared to the AT&T cloud are 12%, 35% and 55% under the low, medium and high data rates, respectively. Compared to the OC approach, there is no extra power saving achieved under low data rate, as the total traffic has not justified replicating any VM into fogs. Under the medium and high user data rates, the power savings achieved are 28% and 47%, respectively.

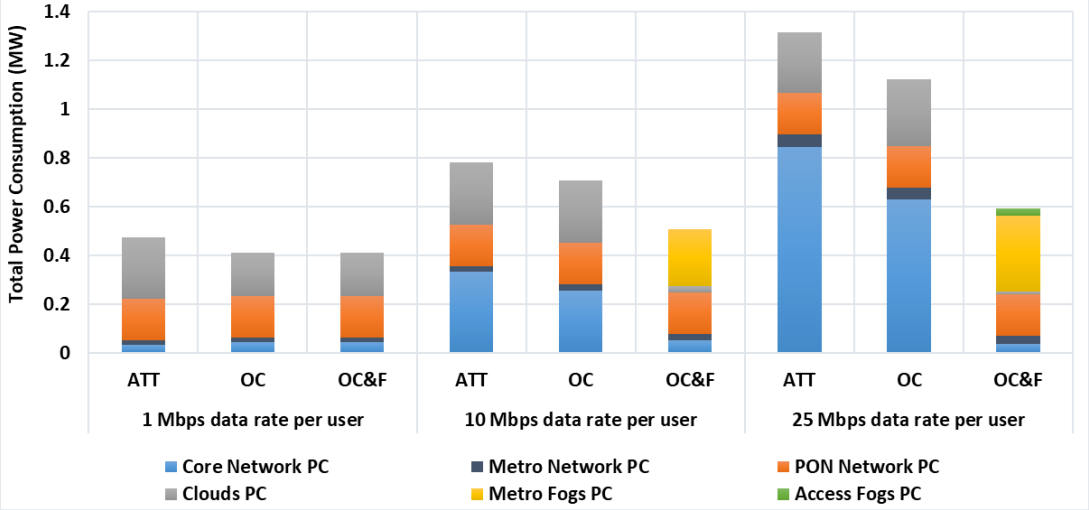

Figure 5-9: The power consumption of different VMs placement approaches considering VMs of 5% minimum CPU workload.

Fig. 5-10(a) and (b) illustrate the placements of the VMs of 5% minimum CPU utilisation considering the OC&F placement approach under low and high user data rates, respectively. VMs with low user data rates are distributed among distributed clouds. The low user data rates have not justified offloading VMs to any fog node. VMs of $\geq \mathbf{1}\%$ popularity have justified the creation of five cloud locations. VMs with 0.5% and 0.05% popularity groups have only justified the creation of three and two replicas, respectively. Under the high user data rates, , it can be observed that VMs with $\geq \mathbf{0.5}\%$ and $\leq \mathbf{5}\%$ popularity groups are fully offloaded to the metro fogs. In addition, VMs with 16% popularity have justified the creation of replicas in some access fogs. Whereas, VMs with 0.05% popularity group have only justified the creation of two replicas in nodes 3 and 14.

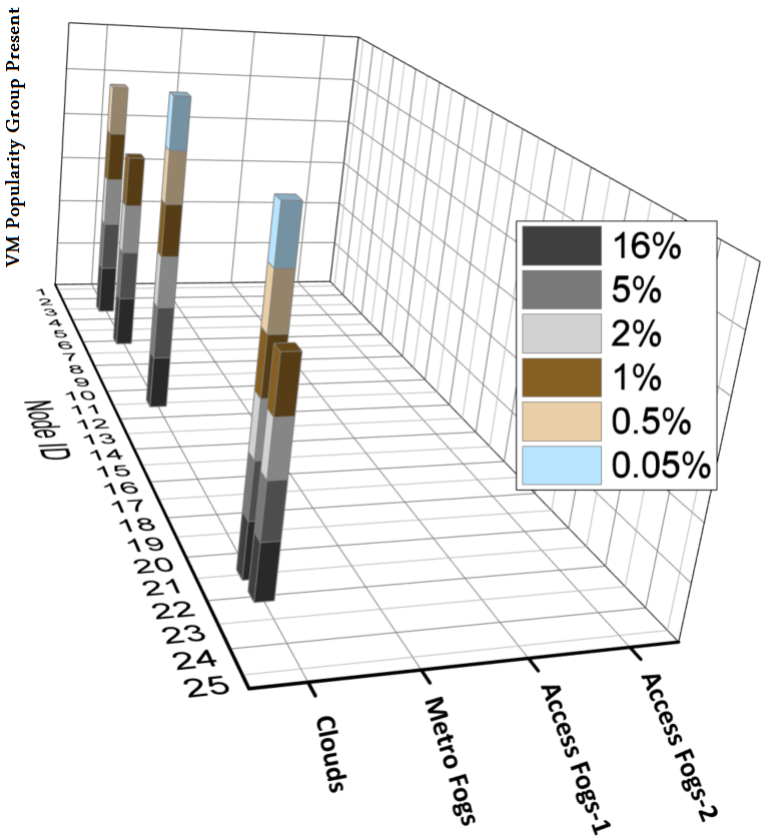

(a)

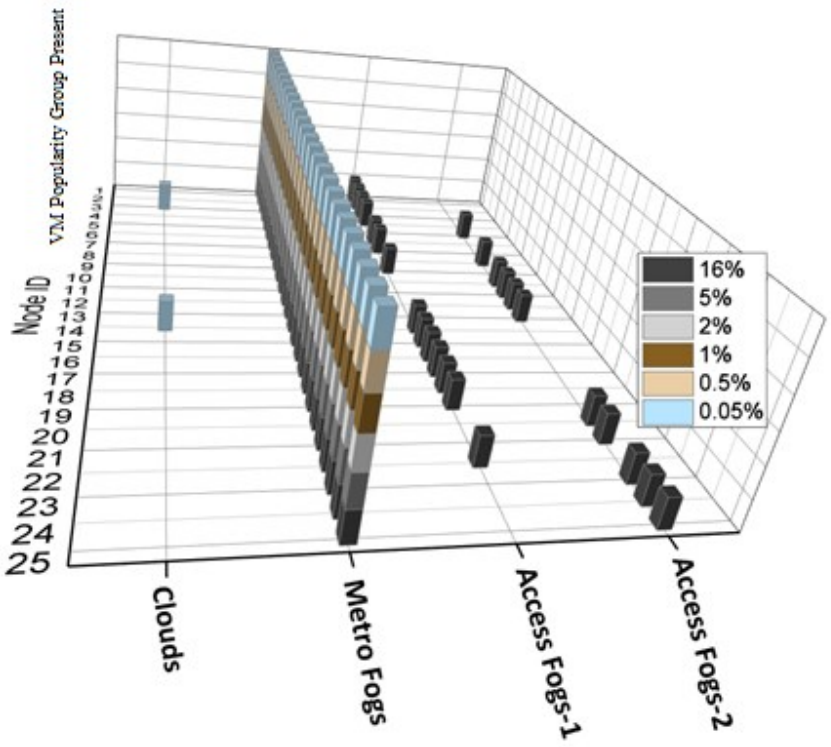

(b)

Figure 5-10: Optimal placement of different VMs popularity groups of 5% workload baseline under the OC&F approach with (a) 1 Mbps data rate per user and (b) 25 Mbps data rate per user.

iii. Linear Workload (40% Workload Baseline):

Fig. 5-11 shows the power savings achieved under VMs of 40% minimum CPU utilisation. The total savings achieved under the OC&F approach compared to the AT&T clouds are 53%, 44% and 48% under the low, medium and high users data rates, respectively. Compared to the OC, there is no extra power saving achieved under the low user data rates, as the total traffic has not justified replication of any VM into any fog node. Under the medium and high user data rates, the power savings achieved are 12% and 31%, respectively.

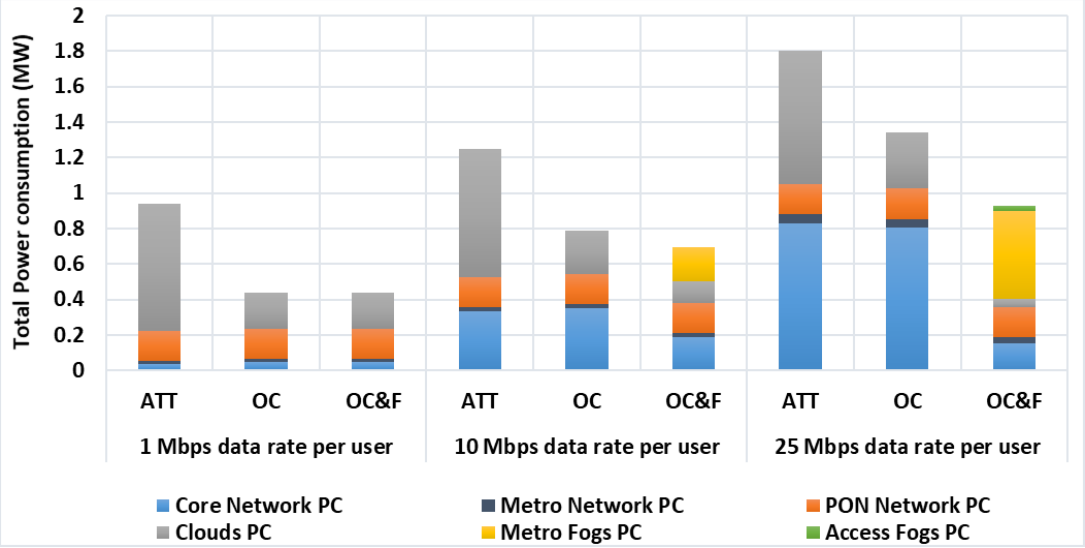

Figure 5-11: The power consumption of different VMs placement approaches considering VMs of 40% minimum CPU workload.

Fig. 5-12 (a), (b) and (c) illustrate the optimal VMs placement under low, medium and high user data rates, respectively, considering the OC&F approach. It can be observed that that increasing the minimum CPU utilisation of VM workload to 40% reduces the efficiency of creating more replicas of VMs with a low popularity across distributed cloud and fog nodes, compared to VMs with 1% or 5% minimum CPU utilisation. VMs with data rate of 1 Mbps are replicated among distributed clouds. The low user data rates have not justified offloading VMs to any fog node. VMs of $\geq$ **1%** popularity have justified the creation of five cloud locations. However, VMs with 0.5% and 0.05% popularity groups have only justified the creation of three and one replicas, respectively. Under medium user data rates, VMs with $\geq$ **1%** popularity are offloaded to metro fogs whereas other popularity groups are optimally placed in clouds. Under high user data rates, despite the high workload baseline, VMs with high popularity of 16% justified the creation of VM replicas in some access fogs. VMs

with $\geq \mathbf{0.5}\%$ and $\leq \mathbf{5}\%$ popularity groups are fully offloaded to metro fogs whereas VMs with 0.05% popularity group have not justified the creation of multiple replicas. Only a single replica is optimally placed in node 11 to serve its distributed users.

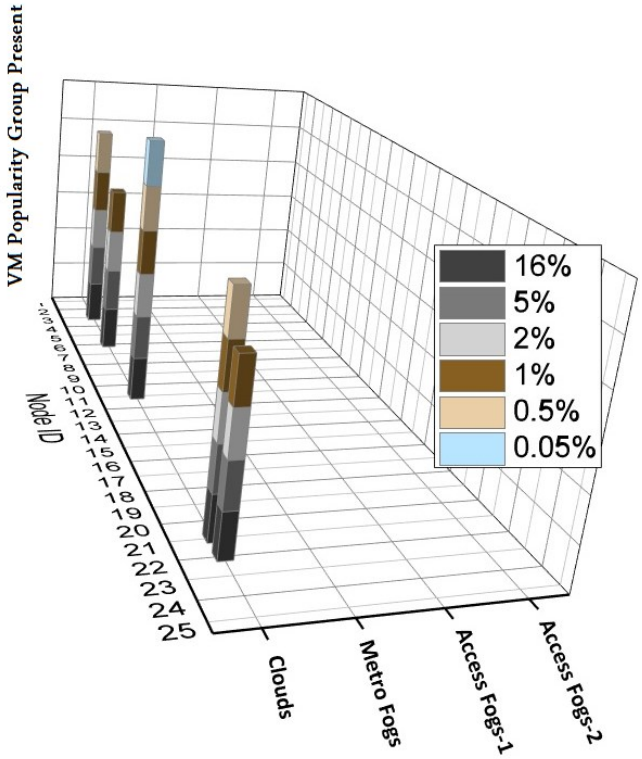

(a)

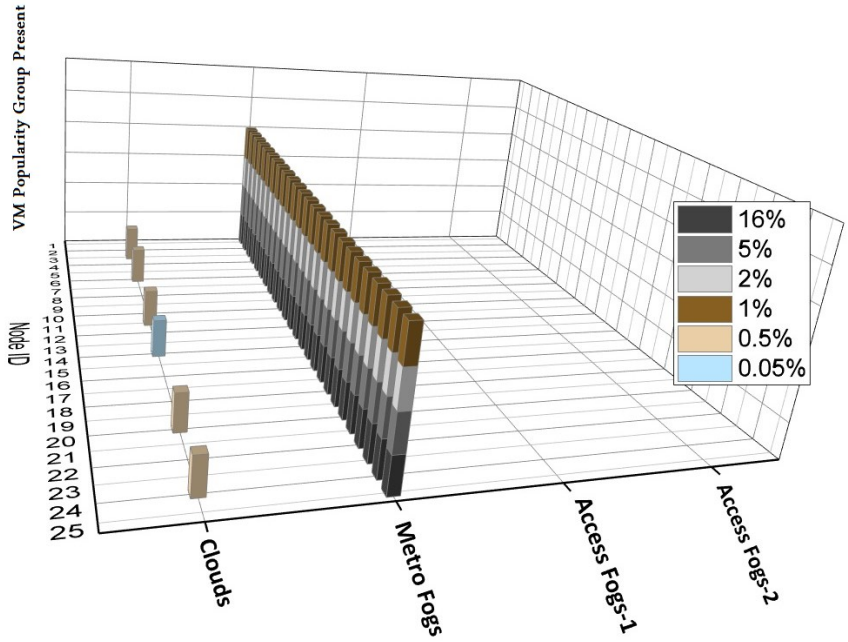

(b)

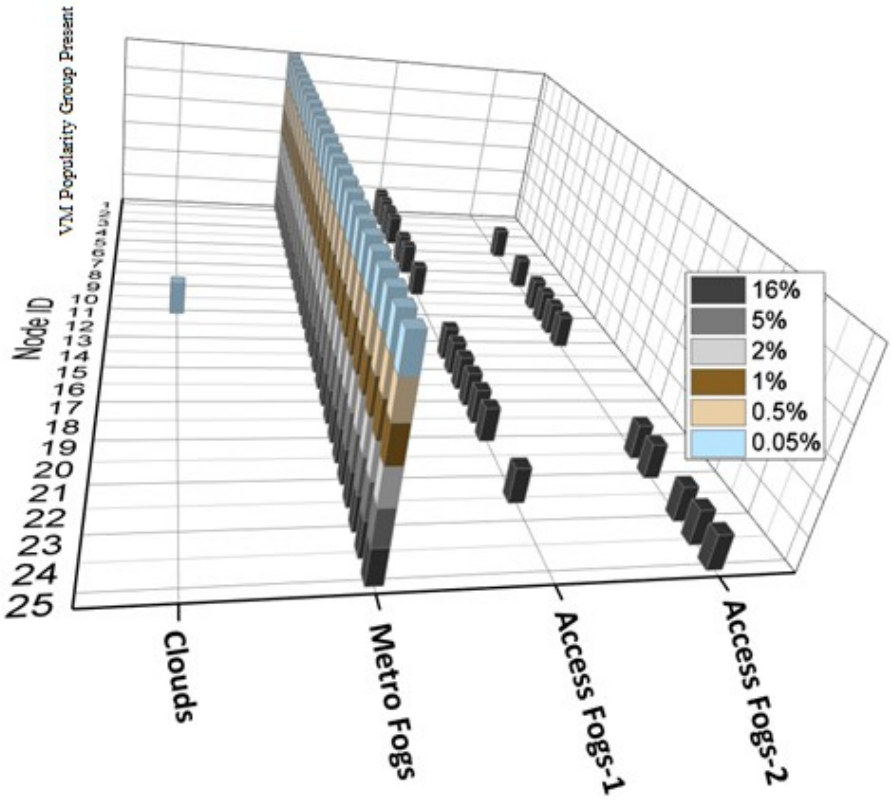

(c)

Figure 5-12: Optimal placement of different VMs popularity groups of 40% workload baseline under the OC&F approach with (a) 1 Mbps data rate per user, (b) 10 Mbps data rate per user and (c) 25 Mbps data rate per user.

b) Network Components with High Energy Efficiency:

The previous scenario has considered Cisco CRS-1 as a core network IP router which consumes 638W per 40 Gbps port and Cisco NCS 5502 router as cloud and metro networks which consumes 30W per 40 Gbps port. In this part of the study, in metro and fog datacentres, the Cisco NCS 5501 router is considered which consumes 13W per 40 Gbps port. However, this is not the only option available in the market. In this scenario, Juniper's PTX1000 and Juniper's MX series 5G universal routing platform are considered to evaluate the cloud-fog architecture model. In core network, PTX1000 provides a 40 Gbps router ports while only consuming 20W [168]. Also, MX series 5G universal routing platform [169] provides a 40 Gbps router ports for clouds while consuming 20W . It also provides a 40 Gbps router ports for metro, metro fog and access fog while consuming 36W. In addition to cloud-fog architecture parameters in Table 5-4 Table to 5-7, Table 5-8 shows the considered IP router ports input parameters of the model.

Table 5-9: IP routers input parameters used in the model

40 Gbps core network router port power consumption ($\boldsymbol{R}^{(\boldsymbol{P})}$)	20W [168]
40 Gbps cloud router port power consumption ($\boldsymbol{R}^{(\boldsymbol{CP})}$)	20W [169]
40 Gbps metro, metro fog and access fog router port power consumption ($\boldsymbol{R}^{(\boldsymbol{MP})}, \boldsymbol{R}^{(\boldsymbol{MFP})}, \boldsymbol{R}^{(\boldsymbol{AFP})}$)	36W [169]

Fig. 5-13 shows that the total saving achieved in the OC approach is 5% compared to the AT&T clouds and no extra saving is achieved by introducing OC&F approach, therefore, no replicas are offloaded to any fog node, as shown in Fig 5-16. Offloading VM replicas to clouds becomes more energy efficient as traversing IP routers becomes more energy efficient than creating multiple fog nodes.

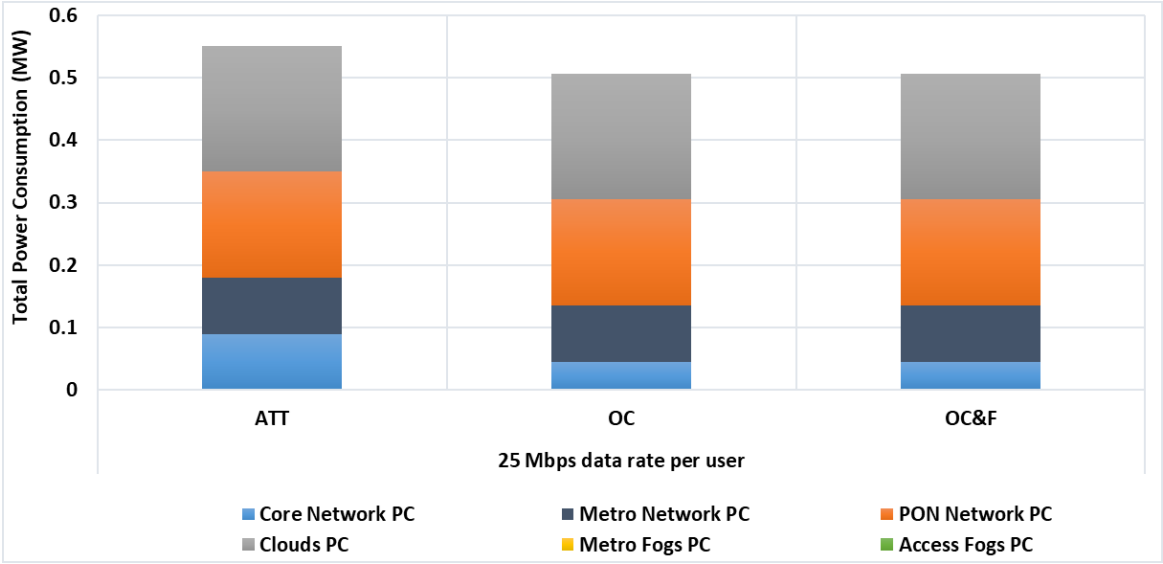

Figure 5-13: The power consumption of different VMs placement approaches considering VMs of 1% workload baseline.

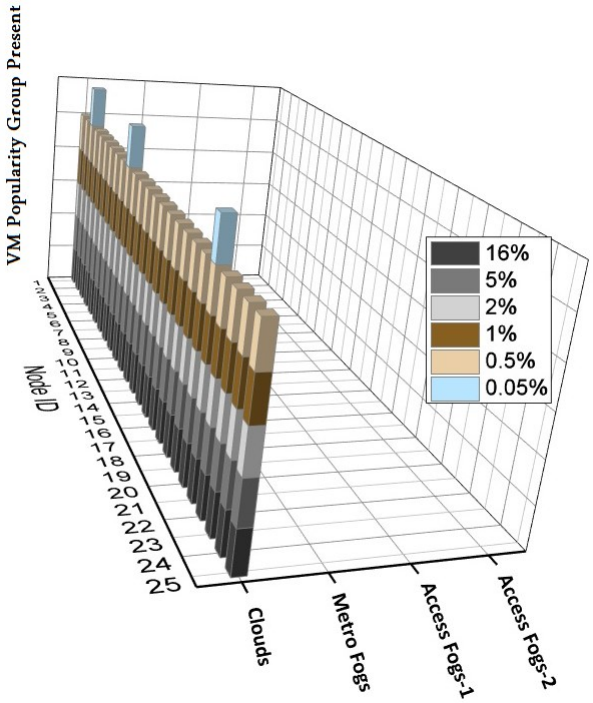

Figure 5-14: Optimal placement of different VM popularity groups of 1% workload baseline under OC&F approach with 25 Mbps data rate per user.

5.3.2 BT networks Use Case:

In this use case, the BT network topology [144] is considered as a core network topology example (illustrated in Fig. 5-15). Similar to Section 3.1.2, 18.8k users are considered in each PON to access the 300 applications/websites hosted in VMs. The popularity of these applications / websites VMs is considered to follow a Zipf distribution with the same popularity groups defined in Table 5.8.

The datacentres and networks components parameters of the cloud-fog architecture in the BT topology use case are considered to take the same values considered in Section 5.3.1.

The OC&F approach is compared to the OC approach for the BT topology obtained in Chapter 3. It is also compared to SC where the VMs are placed in node 6 (City of London).

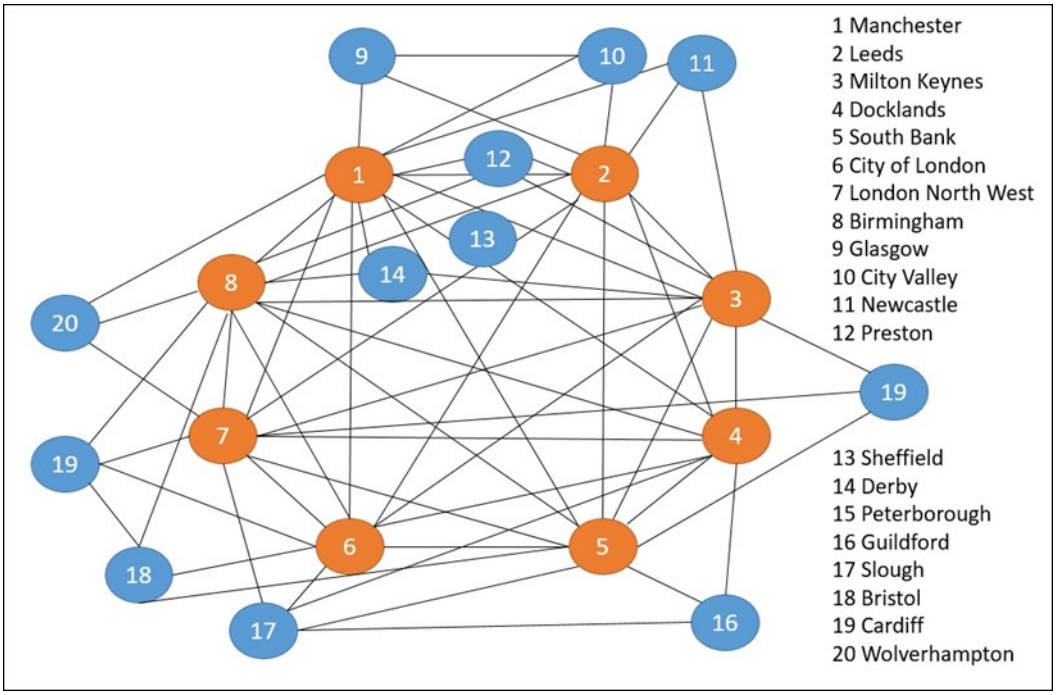

Figure 5-15: BT core network topology.

Fig. 5-16 shows the power consumption resulting from placing the VMs considering the different placement approaches under different users' data rates. Compared to the OC approach, the total savings achieved under the OC&F approach are 2%, 37% and 59%, under 1 Mbps, 10 Mbps and 25 Mbps user data rates, respectively. The total power savings achieved under the OC&F approach compared to the SC approach are 5%, 48% and 75%, under 1 Mbps, 10 Mbps and 25 Mbps user data rates, respectively.

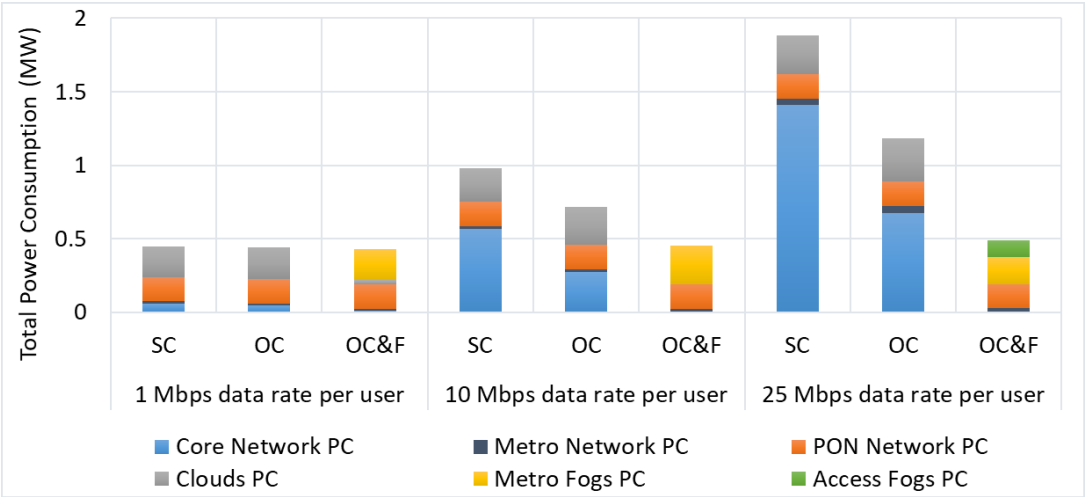

Figure 5-16: The power consumption of different VMs placement approaches considering VMs of 1% workload baseline.

Fig. 5-17 shows the optimal VMs placement under the OC&F approach. Fig. 5-20(a) shows that the cloud in node 2 is selected to serve distributed users of VMs of 1 Mbps users data rate and 0.05% popularity. VMs of 1 Mbps and a popularity of ≥ 0.5% are replicated into every metro fog location except in node 2. Users in node 2 access VMs of 0.5%-16% popularity in the cloud located in node 2 to avoid setting another computing location in node 2 as on-off power profile is considered across cloud-fog architecture. Under the 10 Mbps user data rate, VMs of every popularity group are fully offloaded to metro fog nodes as shown in Fig. 5-20(b). Under the 25 Mbps user data rate, VMs of ≥ **2**% popularity have justified creating a replica in each access fog node. Other VMs of **1**% popularity and less are offloaded to metro fog nodes, as shown in Fig. 5-20(c).

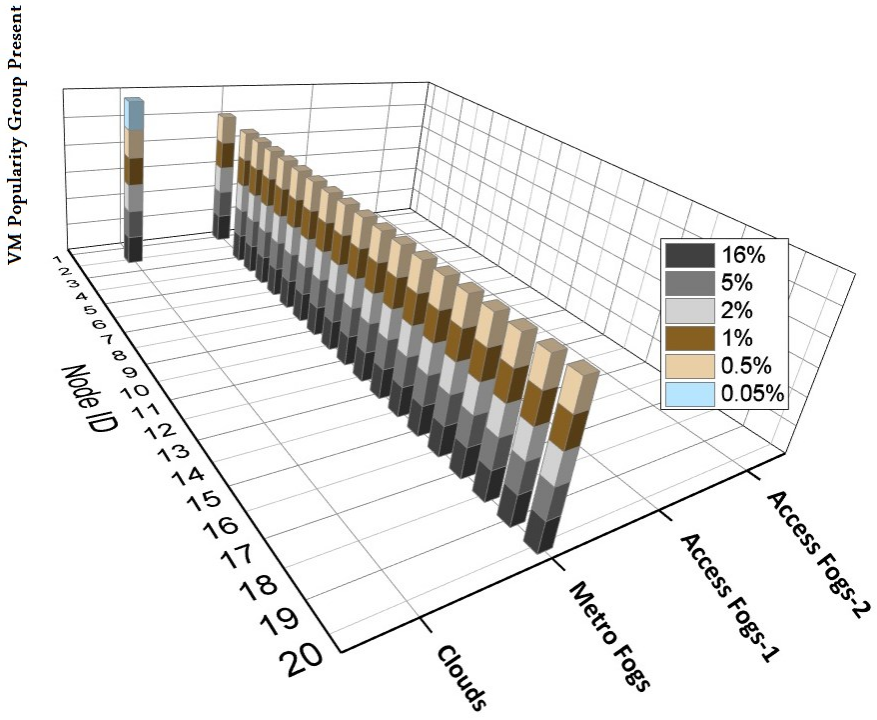

(a)

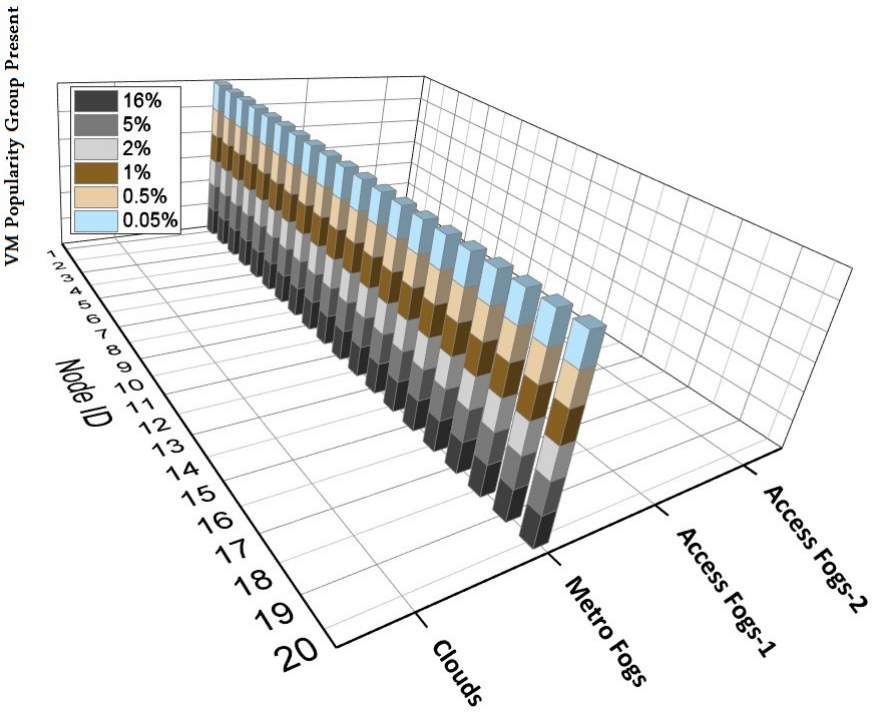

(b)

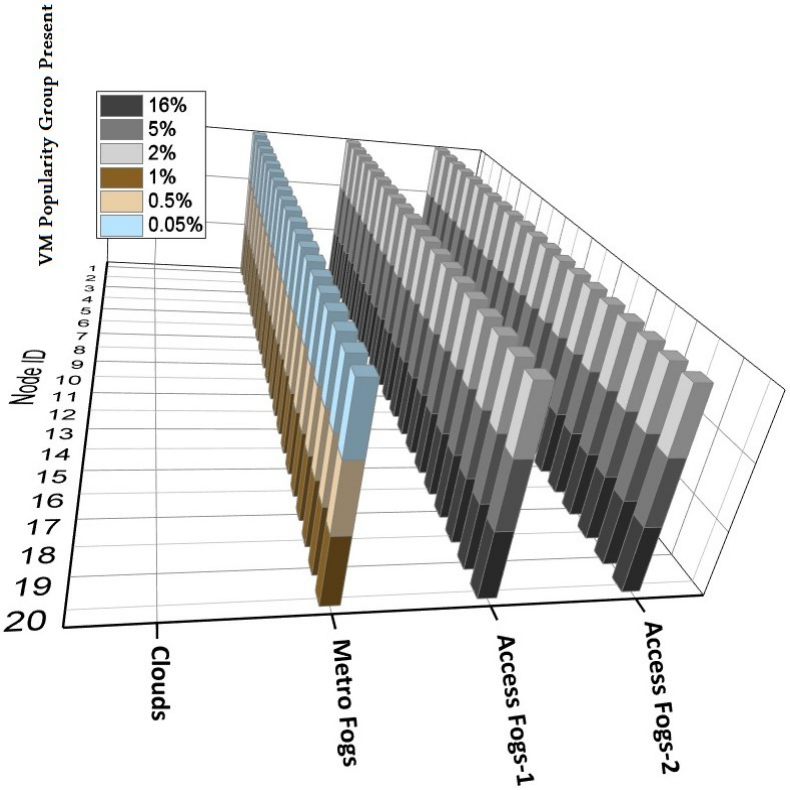

(c)

Figure 5-17: The optimal placement of different VMs popularity groups with 1% workload baseline under the OC&F approach with (a) 1 Mbps data rate per user, (b) 10 Mbps data rate per user and (c) 25 Mbps data rate per user.

We can observe that the VMs placement in the BT network architecture tends to offload more replicas to the access fog compared to the AT&T network. The reason behind this trend is that the top 300 VMs considered in the evaluation scenario are of higher popularity among users in the BT network compared to the AT&T network.

5.4 EEVM-CF Heuristic

As described in Chapter 2, the VM placement problem over cloud-fog architecture is NP-Hard. Therefore, it is not practical to apply the MILP model in a real time implementation. The optimal solutions obtained from the MILP model can offer a benchmark for determining the performance of developed heuristics. Thus, to mimic the MILP model, a real time heuristic is implemented. The heuristic in Chapter 3 is extended to optimise the placement over cloud-fog architecture. We refer to the extended heuristic as Energy Efficient VMs Placement Heuristic for the Cloud-Fog architecture (EEVM-CF).

Similar to the heuristic in Chapter 3, in the offline phase, as shown in the flowchart in Fig. 5-18(a), VMs are classified into different types and the optimal placement of different VMs types are found. The search space, $\boldsymbol{P}$, to find the optimum placement for each VM type includes the most energy efficient placement to place 1 replica, 2 replicas… up to $\boldsymbol{N}$ replicas in the clouds, where $\boldsymbol{N}$ is the number of clouds. For fog nodes, there are two placement scenarios. In one scenario, VMs are replicated to the metro fog and in all core nodes and in the other scenario VMs are replicated to the two access fogs in all core nodes. The traffic resulting from replicating the VMs in clouds and fogs and the workload of VMs of a linear workload profile are calculated based on the number of users each VM serves. The networks, the cloud and the fog power consumptions are calculated and the optimum placement of a VM type is the placement that results in the minimum total power consumption.

Then, VMs are matched to their type in real time (online phase), which is shown in the flowchart in Fig. 5-18(b) and placed according to the placement obtained in the offline phase. Then, the traffic resulting from replicating the VM in the cloud is routed on core network based on multi-hop routing [19] and the workload of clouds where the VM replicas are placed is updated. After placing all VMs, the total power

consumption of the distributed cloud is calculated. After placing all VMs, the total power consumption of cloud-fog architecture is calculated.

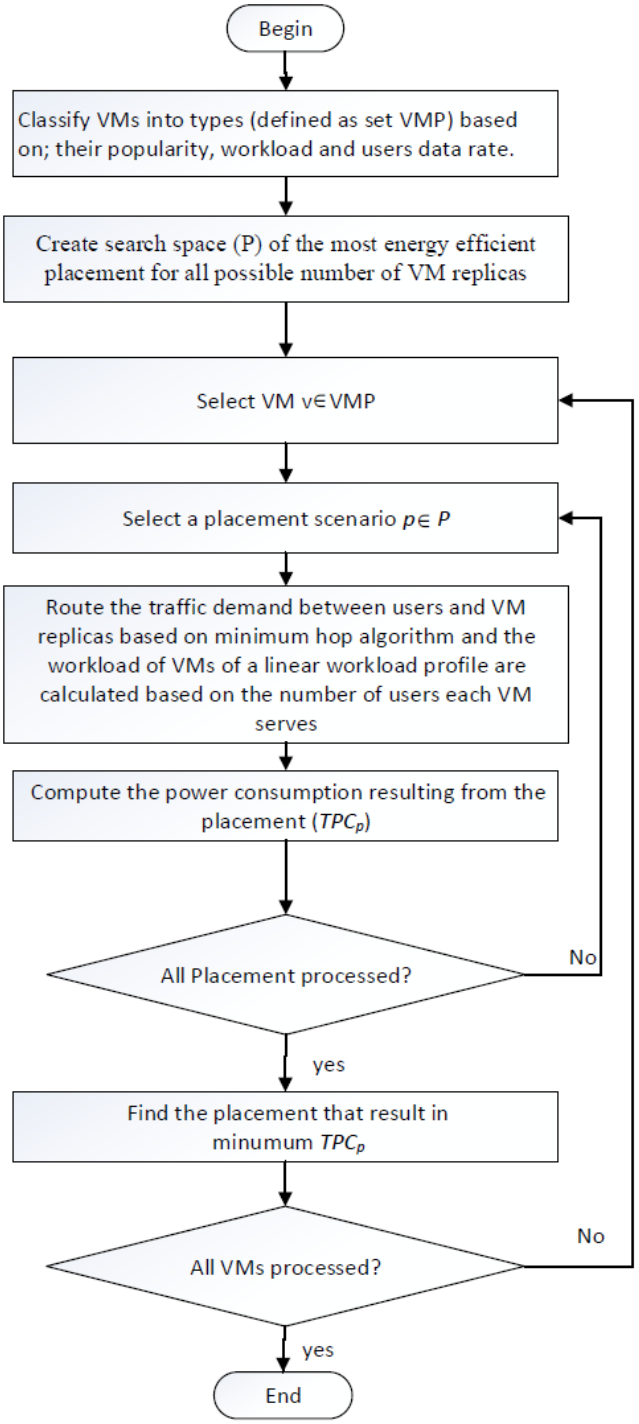

(a)

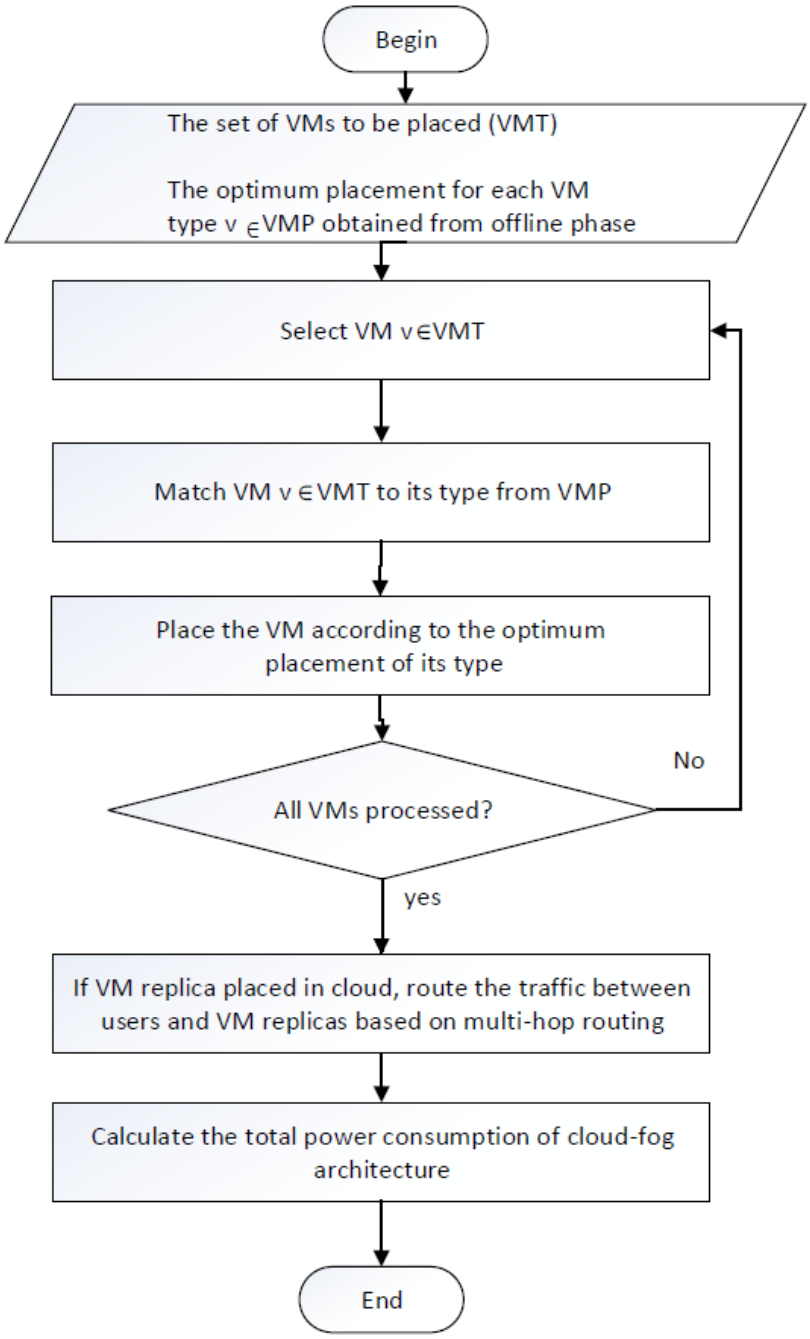

(b)

Figure 5-18: Flowchart of (a) the offline phase and (b) the online phase of EEVM-CF heuristic.

The heuristics are examined by considering the AT&T network as a network example. The EEVM-CF heuristic took 55 seconds to evaluate the offline phase and 2 seconds to evaluate the online phase running on an Intel i-7 core machine with 16 GB of RAM. Fig. 5-19 compares the total power consumption of the EEVM-CF to the that of the MILP model considering the network, cloud, and fog parameters discussed in Section 5.2. The heuristic is evaluated at 1%, 5%, and 40% workload baselines at 1 Mbps, 10 Mbps, and 25 Mbps user data rates. Clearly, the power

consumption of the MILP and the EEVM-CF are comparable; the gap between them ranges from 1–2% of total power consumption due to almost identical VM placement by the model and the heuristic.

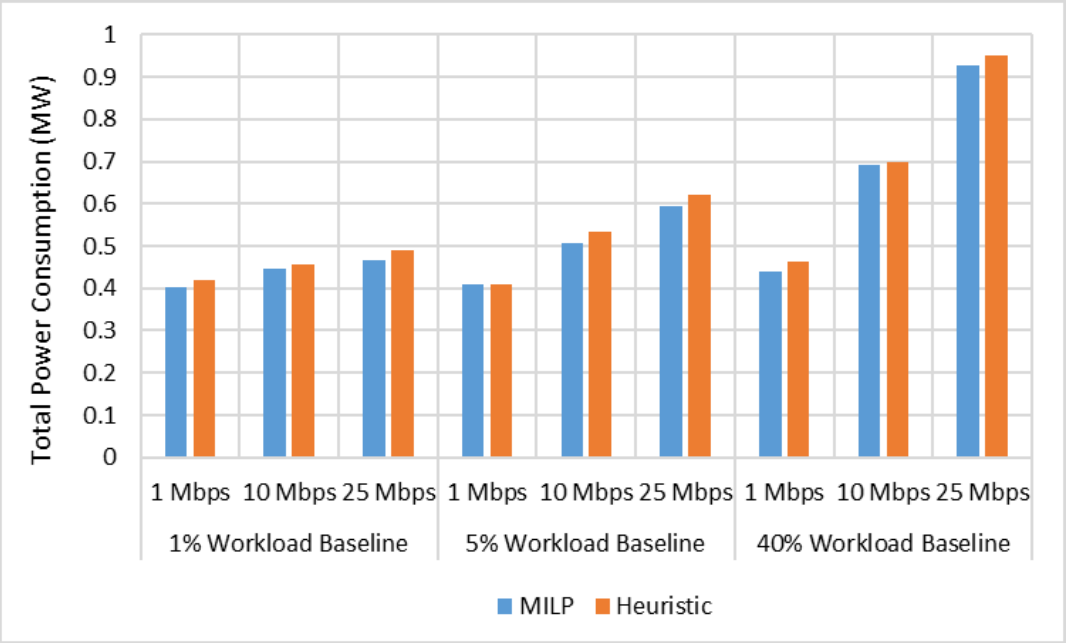

Figure 5-19: Total power consumption of the MILP model compared with EEVM-CF heuristics considering VMs with 1%, 5% and 40% CPU workload baseline.

5.5 Summary

In this chapter, the placement of VMs over a cloud-fog architecture is investigated with the aim of minimising the total power consumption. The optimisation is performed using a Mixed Integer Linear Programming (MILP) model considering AT&T and BT networks as use case scenarios. The MILP model is used to analyse the impact of different factors including VM popularity, the traffic between the VM and its users, the VM workload, the workload vs. number of users profile, the proximity of fog nodes and the PUE.

The decision to serve users from fog nodes is driven by the trade-off between network power saved by placing VMs in fog nodes and the increase in processing power that results from replicating VMs to the fog. Our results demonstrate that VMs placement in fog computing might lead to power saving depending on many factors;

workload and network bandwidth requirements of VMs, VMs popularity among users and energy efficiency of distributed clouds. Overall, the power consumption can be reduced if the VM users traffic is high and/or the VMs have a linear power profile. In such a linear profile, the creation of multiple VM replicas does not increase the power consumption significantly (there may be a slight increase due to idle / baseline power consumption) if the number of users remains constant.

The results also show that (in AT&T networks) the processing requirements efficiency of VMs of a linear workload profile with high data rate and minimum CPU utilisation of 1% allows offloading VMs with 16% popularity to the access fogs. Other VMs are optimally replicated to metro fog nodes. Significant power savings of 48% compared to optimised placement in distributed clouds and 64% compared to a placement considering traditional cloud locations, have resulted from this offloading. VMs with linear workload of a minimum CPU utilisation of 40% tend to offload fewer replicas into fog nodes as the high workload baseline means that VM consolidation in fewer locations is the most efficient approach.

Furthermore, we have developed a heuristic based on an offline exhaustive search, referred to as energy efficient VM placement heuristic for the cloud-fog architecture (EEVM-CF) to place VMs over the cloud-fog architectures in real-time. The heuristic results closely approach those of the MILP model.

Chapter 6: Impact of the Repeal of Net Neutrality on Communication Networks

6.1 Introduction

Network neutrality (net neutrality) is the principle of equally treating all Internet traffic regardless of its source, destination, content, application… etc. Under net neutrality, Internet Service Providers (ISPs) are compelled to charge all Content Providers (CPs) the same per Gbps rate despite the growing profit achieved by CPs. Any discrimination on the internet traffic violates the concept of net neutrality.

In this chapter, a techno-economic Mixed Integer Linear Programming (MILP) model is built to study the potential profit ISPs can achieve by a differentiated pricing scheme under the repeal of net neutrality. We consider that the ISP offers the CP service classes, which represent different data rate requirements. The model optimises the pricing scheme of differentiated service classes to maximise the ISP profit based on price elasticity of demand (PED). The MILP model finds the resulting equilibrium pricing, core network power consumption and traffic.

6.2 Pricing Model:

In economics, the relationship between users demand and price is referred to as price elasticity of demand (PED) [170]. PED measures the percentage change in demand resulting from one percent change in price. To decide pricing strategy of a product, the seller looks at different sensitivities to various factors that may affect their decision to purchase a product. The dominant factor in determining PED is the users' ability and willingness at any given price. Many factors have an effect on users' behaviour such as substitution availability, market competition, frequency of purchase, Necessity of the product, and how much the product price represents in users income. The PED is calculated as follows:

$$\boldsymbol{PED} = \frac{\%\ \boldsymbol{Change\ in\ Demand}}{\%\ \boldsymbol{Change\ in\ Price}}\ . \qquad (\mathbf{6.1})$$

Equation (6.1) measures the sensitivity of demand after a change in a product price. For example, if the price increases by 2% and demand decreases by 2%, then the elasticity at the PED = $\frac{-2\%}{2\%} = -1$.

In telecommunications, it is not easy task to estimate an exact value of PED for various Internet applications as the factors that affect the elasticity change from area to another e.g. wealth, popularity of an application, quality of services provided by ISPs/CPs or competition between different CPs. However, PED for broadband subscriptions in Organisation for Economic Co-operation and Development (OECD) countries has been analysed in [171] by studying the relationship between price, income and broadband adoption. Additional factors have been included in [172], which are age and education to study PED for broadband subscriptions in Latin America and the Caribbean countries. They found that 1% decrease in price would lead to 0.43% and 2.2% increase in demand, respectively, over the two selected areas.

6.3 Profit-Driven MILP Model:

We develop a profit-driven MILP model where the objective is to maximise the total profit of an ISP offering core network infrastructure to CPs to deliver content from distributed clouds and/or fog nodes to their users.

In terms of computational complexity, we do not need to run the model in real-time as the optimization of pricing is an offline problem solved at the service planning and service update phases. However, it is important to ensure that we can obtain solutions for networks with large number of nodes. The above MILP optimization model has a total of $\boldsymbol{O}(\boldsymbol{N^4})$ variables and $\boldsymbol{O}(\boldsymbol{N^3})$ constraints where $\boldsymbol{N}$ is the number of nodes. For a network of $\boldsymbol{N}$=25, there is a total of about $\mathbf{25^4}$ variables and $\mathbf{25^3}$ constraints. Using a 2.5 GHz Intel core i7 with 16 GB memory, the model runs for a maximum of 7 minutes to obtain the optimum pricing for each price elasticity of demand (PED) scenario.

We consider a monopolist ISP who owns the network backbone, i.e. CPs have to subscribe to the monopolist ISP to reach their customers. According to the FCC, 40% of total US Internet subscribers have only a single ISP option in their area [173]. The ISP has the power to control the pricing scheme. Under the net neutrality repeal, the ISP can deliver CPs content of different data rate requirements at a varying price per bit rate. We consider three classes to represent different data rate requirements of CPs services;

- Class A for high data rate content (i.e. UHD video service).
- Class B for medium data rate content (i.e. HD video service).
- Class C for low data rate content (i.e. SD video service).

The ISP needs to optimise the price of the three classes to maximise its profit. We consider content with higher data rate, which causes extra burden on core network, to be priced higher per bit rate than content with a lower data rate. End-users will perceive varied video definitions from CPs based on their CP subscribed class. We assume that CPs will transfer the ISP new prices to their users to maintain their profit margin. Therefore, for the sake of simplicity we consider CPs to offer the same classes to their users. We assume a certain number of users to initially subscribe to each class under net neutrality. As the ISP and consequently the CPs vary the per bit rate charges for the different classes, users can choose to upgrade, downgrade or unsubscribe to the service. The number of users subscribing to each class depends on the PED. We assume that users leaving class A will join class B, users leaving class B will join class C and users leaving class C will unsubscribe to the service.

Before introducing the model, in addition to the parameters and variables defined in Chapter 3, Table 6-1 and Table 6-2 define the parameters and variables used in the model:

Table 6-1: List of parameters used in the profit-driven MILP model

Parameter	Description
$\boldsymbol{\alpha}$	Set of service classes.
$\boldsymbol{B}$	Wavelength data rate.
$\boldsymbol{CN}$	Number of clouds hosted in core network.
$\boldsymbol{u}$	Total number of users in net neutrality scenario (i.e. before net neutrality is repealed).
$\boldsymbol{LB}$	Minimum percentage of users served by CP to be maintained by the pricing scheme.

$\boldsymbol{d_i}$	Download rate of class $\boldsymbol{i}$.
ℭ	The cost in US$ of provisioning a Gbps of IP over WDM network bandwidth per month.
ℑ	The cost in US$ of provisioning a Gbps of metro and access network bandwidth per month.
$\boldsymbol{PS}$	The net neutrality selling price in US$ of a Gbps of network bandwidth per month.
$\boldsymbol{E_i}$	Price elasticity of demand of class $\boldsymbol{i}$.
$\boldsymbol{N_{d,i}}$	Number of users of class $\boldsymbol{i}$ located in node $\boldsymbol{d}$ under net neutrality scenario.
$\boldsymbol{\delta_s}$	$\boldsymbol{\delta_s = 1}$, if a cloud datacentre is hosted in node $\boldsymbol{s}$, otherwise $\boldsymbol{\delta_s = 0}$.
$\boldsymbol{F_d}$	$\boldsymbol{F_d = 1}$, if there is no fog datacentre hosted in node $\boldsymbol{d}$, otherwise $\boldsymbol{F_d = 0}$.
ℨ	Set of all possible solutions.
$\boldsymbol{\rho_{s,i}}$	The price of class $\boldsymbol{i}$ under solution $\boldsymbol{s}$ and class $\boldsymbol{i}$.
$\boldsymbol{yn_{s,d,i}}$	The number of users in solution $\boldsymbol{s}$ subscribing to class $\boldsymbol{i}$ in node $\boldsymbol{d}$ as a result of its PED, where $\frac{\boldsymbol{PS}}{\boldsymbol{\rho_{s,i} - PS}}\,\boldsymbol{E_i} = \sum_{\boldsymbol{d \in N}} \left(\frac{\boldsymbol{yn_{s,d,i}} - \boldsymbol{N_{d,i}}}{\boldsymbol{N_{d,i}}} \right)$ $\forall\ \boldsymbol{i} \in \boldsymbol{\alpha}\,,\ \ \boldsymbol{s} \in$ ℨ.

Table 6-2: List of variables used in the profit-driven MILP model

Variable	Description
$\boldsymbol{r_i}$	ISP's revenue achieved by delivering traffic of class $\boldsymbol{i}$ to CP users.
$\boldsymbol{R}$	Total ISP's revenue in US$ of delivering networking services to CPs content.
$\boldsymbol{C}$	Total ISP cost in US$ of provisioning core network.
$\boldsymbol{P_i}$	The price in US$ per Gbps of network bandwidth per month charged to the class $\boldsymbol{i}$.
$\boldsymbol{U_{d,i}}$	Number of users who subscribe to class $\boldsymbol{i}$ located in node $\boldsymbol{d}$.
$\boldsymbol{CD_{i,d}}$	Cloud flow from users in node $\boldsymbol{d}$ subscribed to class $\boldsymbol{i}$.
$\boldsymbol{\mathcal{Z}_{s,i}}$	$\boldsymbol{\mathcal{Z}_{s,i} = 1}$, if solution $\boldsymbol{s}$ is selected for class $\boldsymbol{i}$, otherwise $\boldsymbol{\mathcal{Z}_{s,i} = 0}$.
$\boldsymbol{ys_{s,d,i}}$	The number of users in solution $\boldsymbol{s}$ subscribing to class $\boldsymbol{i}$ in node $\boldsymbol{d}$, $\boldsymbol{ys_{s,d,i} > 0}$ if solution $\boldsymbol{s}$ is selected for class $\boldsymbol{i}$, otherwise $\boldsymbol{ys_{s,d,i} = 0}$.

Total ISP's cost and revenue of delivering CP contents are calculated as follows:

Cost of provisioning core, metro and access networks infrastructure:

$$\sum_{\boldsymbol{s \in N}} \boldsymbol{R_s^{(AC)}}\ \boldsymbol{B}\ \ ℭ + \sum_{\boldsymbol{i \in \alpha}} \sum_{\boldsymbol{d \in N}} \boldsymbol{U_{d,i}}\, ℑ\ \boldsymbol{d_i} \qquad (\boldsymbol{6.2})$$

Revenue of delivering networking services to CP users:

$$\sum_{i\in\alpha} r_i \qquad (6.3)$$

The model is defined as follows:

The objective:

Maximise total profit given as:

$$\sum_{i\in\alpha} r_i \quad - \quad \left(\sum_{s\in N} R_s^{(AC)} B \quad \text{C} + \sum_{i\in\alpha}\sum_{d\in N} U_{d,i}\, \text{Ɔ}\ d_i\right) \qquad (6.4)$$

Equation (6.4) gives the total profit in US dollar.

Subject to:

Subject to: (3-23) – (3-29), and the following constraints:

Revenue of each class:

$$r_i = U_{d,i}\, d_i\, P_i \quad \forall\, i \in \alpha \qquad (6.5)$$

Constraint (6.5) calculates the revenue the ISP achieves by delivering a service class by multiplying the class price by the total traffic in each class. Note that, the total revenue is obtained by multiplying two variables ($U_{d,i}$ *and* P_i) which is a non-linear process. A look up table of solutions under different PED values defined by parameters ρ_{si},$ys_{s,d,i}$, $yn_{s,d,i}$ is used for linearization. Constraints (6.6)-(6.10) select the optimum number of users and price for each class and the calculate the resulting revenue.

$$ys_{s,d,i}\begin{cases} = (yn_{s,d,i}\, Z_{s,i}) & if\ \ i = \text{A} \\ \leq (yn_{s,d,i} + N_{d,\text{A}}) Z_{s,i} & if\ \ i = \text{B} \\ \leq (yn_{s,d,i} + N_{d,\text{B}}) Z_{s,i} & if\ \ i = \text{C} \end{cases}$$

$$\forall\ i \in \alpha,\ \ s \in \text{Ƨ}\ , d \in N \qquad (6.6)$$

$$P_i = \sum_{s\in\text{Ƨ}} (\rho_{s,i}\, Z_{s,i}) \qquad \forall\, i \in \alpha \qquad (6.7)$$

$$\sum_{s\in\text{Ƨ}} Z_{s,i} = 1 \qquad \forall\, i \in \alpha \qquad (6.8)$$

$$\boldsymbol{U_{d,i} = \sum_{s \in \mathcal{S}} ys_{s,d,i} \qquad \forall\ d \in N, \qquad i \in \alpha} \tag{6.9}$$

$$\boldsymbol{r_i = \sum_{s \in \mathcal{S}} \left(\sum_{d \in N} ys_{s,d,i}\ d_i\ \rho_{s,i} \right) \quad \forall\ i \in \alpha} \tag{6.10}$$

Constraint (6.6) calculates the number of users located in node $\boldsymbol{d}$ of class $\boldsymbol{i}$ in solution $\boldsymbol{s}$. The number of users in class A is the number of users subscribing to the class as a result of its PED (from the look up table described in the parameter $\boldsymbol{yn_{s,d,i}}$ above). In the case of class B, the number of users available to class B includes all users subscribing to the class B as a result of its PED plus any users downgrading their subscription from class A to class B. In the case of class C, the number of users available to class C includes users subscribing to class C as a result of its PED plus any users downgrading their subscription from class B to class C. Constraint (6.7) gives the price of each class based on the solution selected from the lookup table. Constraint (6.8) ensures that only one solution is selected. Constraint (6.9) calculates the number of users of class $\boldsymbol{i}$ in node $\boldsymbol{d}$. Constraint (6.10) calculates the revenue the ISP achieves by delivering a service class by multiplying the class price by the total traffic in each class.

Constraints on number of users and prices:

$$\boldsymbol{\sum_{d \in N} \sum_{i \in \alpha} U_{d,i} \geq u\ LB} \tag{6.11}$$

$$\boldsymbol{P_1 \geq P_2 \geq P_3} \tag{6.12}$$

$$\boldsymbol{\frac{\sum_{d \in N} U_{d,i}}{\sum_{d \in N} \sum_{i \in \alpha} U_{d,i}} = \frac{U_{d,i}}{\sum_{i \in \alpha} U_{d,i}}}$$

$$\boldsymbol{\forall\ i \in \alpha, d \in N} \tag{6.13}$$

Constraint (6.11) defines the minimum user percentage the CP service needs to maintain. Constraint (6.12) ensures that the price of a lower class does not exceed the price of upper classes, i.e. the price of class C does not exceed the price of class B and the price of class B does not exceed the price of class A. Constraint (6.13) ensures that the ratio of users in different nodes is identical.

Core network traffic:

$$CD_{d,i} = U_{d,i}\, F_d\, d_i$$

$$\forall\, d \in N, \quad i \in \alpha \qquad (6.14)$$

$$\sum_{s \in N} L_{s,d} = \sum_{i \in \alpha} \sum_{d \in N} CD_{d,i}$$

$$\forall\, d \in N \qquad (6.15)$$

Constraint (6.14) ensures that nodes with a fog built in its proximity are not served by a cloud. Constraint (6.15) calculates the download traffic from CP cloud to users in different nodes.

User demands can be used to decide on datacentre locations as follows:

$$L \sum_{d \in N} L_{s,d} \geq \delta_s \qquad \forall\, s \in N \qquad (6.16)$$

$$\sum_{d \in N} L_{s,d} \leq L\, \delta_s \qquad \forall\, s \in N \qquad (6.17)$$

Constraints (6.16) and (6.17) relate the binary parameter that indicates whether there is a datacentre built in node s or not (δ_s) to the traffic between users in node d and datacentre in node s.

The mathematical model given above maximises the total profit of an ISP. To calculate the core network power consumption that results from the profit-driven model, network components defined in Chapter 3 (Table 3-9 and Table 3-10) are considered.

The IP over WDM network power consumption is composed of;

$$n\left(\sum_{s \in N} R^{(P)}\, R_s^{(AC)} + \sum_{d \in N} R^{(P)} R_d^{(AE)} + \sum_{m \in N} \sum_{n \in N m_m : n \neq m} R^{(P)}\, W_{mn}\right)$$

$$+ n\left(\sum_{m \in N} \sum_{n \in N m_m : n \neq m} t^{(P)}\, W_{m,n}\right) + n\left(\sum_{m \in N} \sum_{n \in N m_m : n \neq m} e^{(P)}\, F_{m,n}\, A_{m,n}\right)$$

$$+ n\left(\sum_{s \in N} SW_s^{(P)}\right) + n\left(\sum_{m \in N} \sum_{n \in N m_m : n \neq m} G^{(P)}\, G_{m,n}\, W_{m,n}\right)$$

$$(6.18)$$

Equation (6.18) calculates the total IP over WDM network power consumption.

The total traffic carried over the core physical links is given as:

$$\sum_{m \in N} \sum_{n \in N m_m : n \neq m} W_{m,n} B \qquad (6.19)$$

6.4 Results and Discussions:

In this section, we evaluate the increase in ISP profit and the reduction in network traffic and subsequently power consumption resulting from the optimised pricing scheme under the repeal of net neutrality. We define the three services classes as follows;

- Class A; for UHD video service; 18 Mbps download rate.
- Class B; for HD video service; 7.2 Mbps download rate.
- Class C; for SD video service; 2 Mbps download rate.

We investigate CP's end users' choices of service classes based on different PED. We show how users behaviour under the different PED; 0.2, 0.4, 0.6, 0.8, 1 or 2 affect the equilibrium price of each class the ISP will charge the CP for delivering its content.

As discussed above, we assume that CP will transfer the price increase to their customers at the same rate (if the CP absorbs some of the increase in prices, then this may represent a different PED). As a benchmark, we consider users to be distributed among classes according to the Cisco forecast report [174], where UHD, HD, and SD users distribution are 19%, 56% and 25% respectively. We consider 1.8 million users active simultaneously in the network. This figure is obtained as follows: The number of users is 44 million users in Netflix in the US and the average user spent around 1 hour per day watching movies in 2015 [175]. Therefore, the average number of users during one hour of the day is 1.8 million users, which is an average number that does not consider the popularity of different viewing times in the day. The concentration of users at any node in AT&T network is based on the population of the state where the node is located (see Fig. 6.1).

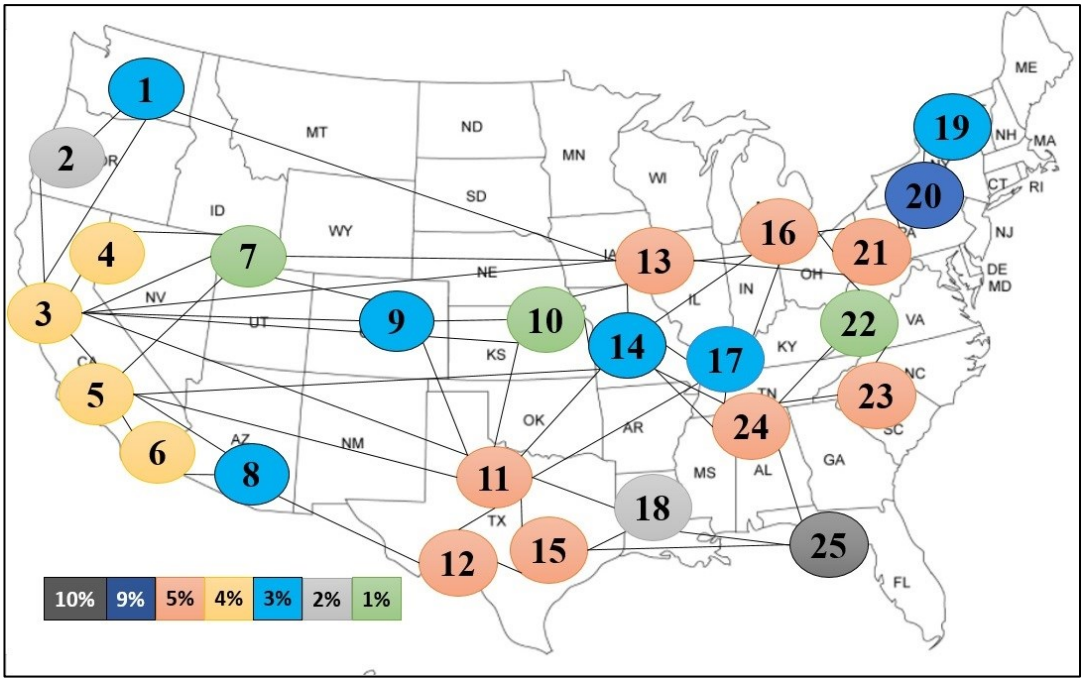

Figure 6-1: AT&T core network with percentage of population in each node.

In the following, the connectivity selling price and the cost of provisioning different telecom network layers are explained. we consider the BT network connectivity selling price as the net neutrality price of the three classes where 10 Gbps connectivity is priced at £12,600 ($15,750) per year [168], i.e. $131 per 1 Gbps link per month. The actual cost of provisioning ISP core network infrastructure is sensitive information and not usually shared by ISPs. However, we estimate the cost of provisioning 1 Gbps of network as $118 considering 10% as the ISP profit margin (the profit margin average for AT&T [121] and Comcast [176] were approximately 9% and 12%, respectively between 2013-2018). We divided the cost among the three network layers; core, metro and access network based on their power consumption percentage; 24%, 6% and 70%, respectively [177] which corresponds to $28, $7, $83, respectively. The cost of $28 per Gbps in the core network is associated with a single hop. For the AT&T architecture the average hop count between clouds and other nodes is 1.

As shown in Fig. 6.1, we choose AT&T core network (a primary core network topology in the US) as a core network topology example. This core network consists of 25 nodes and 54 bidirectional links. AT&T hosts data centres in nodes 1, 3, 5, 6, 8, 11, 13, 17, 19, 20, 22, and 25 [144]. These nodes are used to host datacentres to serve distributed CPs users. The input parameters used are given in Table 6.3.

Table 6-3: Input parameters of profit-driven MILP model

40 Gbps router port power consumption ($\boldsymbol{R^{(P)}}$)	638W [148]
40 Gbps transponder power consumption ($\boldsymbol{t^{(P)}}$)	129W [149]
40 Gbps regenerator power consumption ($\boldsymbol{G^{(P)}}$)	114W, reach 2000 km [150]
EDFA power consumption ($\boldsymbol{e^{(P)}}$)	11W [151]
Optical switch power consumption ($\boldsymbol{SW^{(P)}}$)	85W [152]
Number of wavelengths in a fibre ($\boldsymbol{\mathcal{W}}$)	32 [20]
Bit rate of each wavelength ($\boldsymbol{\mathcal{W}^{(B)}}$)	40 Gbps [20]
Span distance between two EDFAs ($\boldsymbol{S}$)	80 km [151]
Network power usage effectiveness ($\boldsymbol{n}$)	1.5 [6]
Total users (u)	1.8 million users [175]
The cost of provisioning 1 Gbps of core network bandwidth per month (Ⓒ)	$28
The cost of provisioning 1 Gbps of metro and access network bandwidth per month (Ↄ).	$90
The net neutrality selling price of downloading 1 Gbps of network bandwidth per month (Ps)	$131 [168]
Set of classes ($\boldsymbol{\alpha}$)	3 classes; A, B and C
Number of users of class $\boldsymbol{i}$ located in node $\boldsymbol{d}$ under net neutrality scenario ($\boldsymbol{N_{d,i}}$)	19% of total users for class A, 56% for class B, and 28% for class C [174]. Number of users in each node is based on the population of the state where the node is located (see Fig. 6.1).

Download rate of class $\boldsymbol{i}$ $(\boldsymbol{d_i})$	18 Mbps for class A, 7.2 Mbps for class B, and 2 Mbps for class C [174]
Price elasticity of demand $(\boldsymbol{E_i})$	0.2, 0.4, 0.6, 0.8, 1 or 2
Minimum percentage of users served by CP to be maintained by the pricing scheme ($\boldsymbol{LB}$).	0 or 100

In the following subsections, we evaluate two scenarios; equal PED for all classes and different PED for different classes. Under each scenario we study three scenarios of delivering CPs contents to users; a cloud-based delivery and a cloud-fog based deliver and fog-based delivery.

6.4.1 Equal PED among classes:

In the following, we study three scenarios of delivering CPs contents to users; a cloud-based delivery and a cloud-fog based deliver and fog-based delivery.

6.4.1.1 Cloud based delivery:

Fig. 6-2 to 6-4 show the profit-driven model results for AT&T core network where content is delivered from the 12 datacentres in the AT&T topology [144]. The number of users and the corresponding price of each class under different PED are illustrated in Fig. 6-2. The primary y-axis shows price per Gbps per month of each class in US dollar. These prices represent the equilibrium point of users' willingness to follow the price increase which result in maximum profit for the ISP. The secondary y-axis corresponds to the percentage of users subscribed to each class. The x-axis shows different PED scenarios from 2 to 0.2. The former represents highest sensitivity to the price change, whereas, the latter represents the contrary. PED values are shown along with the case of net neutrality where the price of different classes is fixed to 113$ and the percentage of users in each class follows Cisco forecast report [174] as discussed above. For each PED value we consider two cases; a case where the optimised pricing scheme should maintain 100% of the users that existed under net neutrality ($\boldsymbol{LB} \geq$

$\boldsymbol{100}$) and another case where the pricing scheme can result in users leaving the service ($\boldsymbol{LB \geq 0}$).

Fig. 6-3 is a plot of the monthly profit of ISP considering different PED values as well as net neutrality scenario. Total traffic of core network and the power consumption from this traffic under different PED scenarios and the net neutrality scenario are plotted in Figure 5.

In case of content with $\boldsymbol{PED = 2}$, under $\boldsymbol{LB \geq 100}$ or $\boldsymbol{LB \geq 0}$, Fig. 6-2 shows that repealing net neutrality has increased class C users to 48% of the total number of users compared to 14% only under the net neutrality pricing scheme. This increase is a result of some users of class B downgrading to class C as the class B price increased slightly by 18% (the number of users in class B reduced to 36%) and new users joining the service (the total number of users increased to 102%) attracted by the 1% decrease in class C price. Note that LB defines the minimum user percentage the CP service needs to maintain. The users of class A are reduced to 18% of the total number of users as a result of the slight increase in price by 19%. This pricing scheme and distribution of users have resulted in an increase in the total profit by 54% compared to the net neutrality scenario as seen in Fig. 6-3. For a less sensitive content with $\boldsymbol{PED = 0.2}$ under $\boldsymbol{LB \geq 0}$, the equilibrium pricing scheme resulted in 28% of the users leaving the service as the increase in the classes price resulted in increasing the profit by a factor of 8.3 compared to the net neutrality scenario. Maintaining all the users of the service ($\boldsymbol{LB \geq 100}$) has slightly reduced the profit by 10%.

In addition to growing ISP profit, we also observe in Fig. 6-4 a decline in the core network traffic by up to 55% under $\boldsymbol{PED = 0.2}$, $\boldsymbol{LB \geq 0}$ and a consequent reduction in power consumption by 49%. This reduction in core network traffic and power consumption occurred for two reasons; 1) some cloud service users leaving classes A and B to subscribe to class C as the charges per Gb/s of the classes A and B increase. 2) the total cloud service subscribers diminished due to the increase in class C price (in case of $\boldsymbol{LB \geq 0}$).

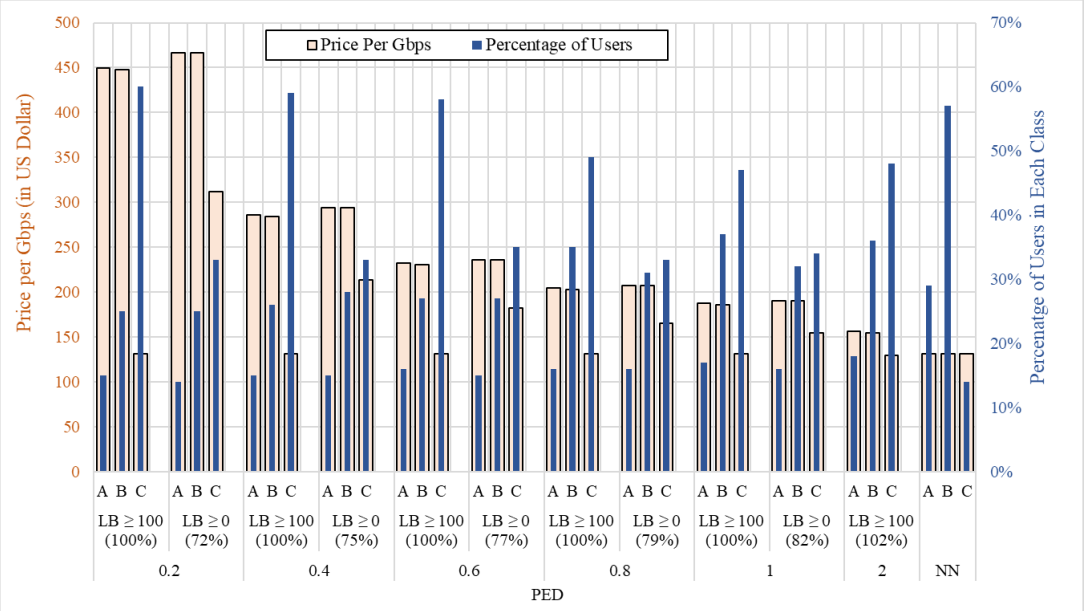

Figure 6-2: Price per Gbps per month and the corresponding number of users in each class based on different PED after removing net neutrality (cloud-based delivery).

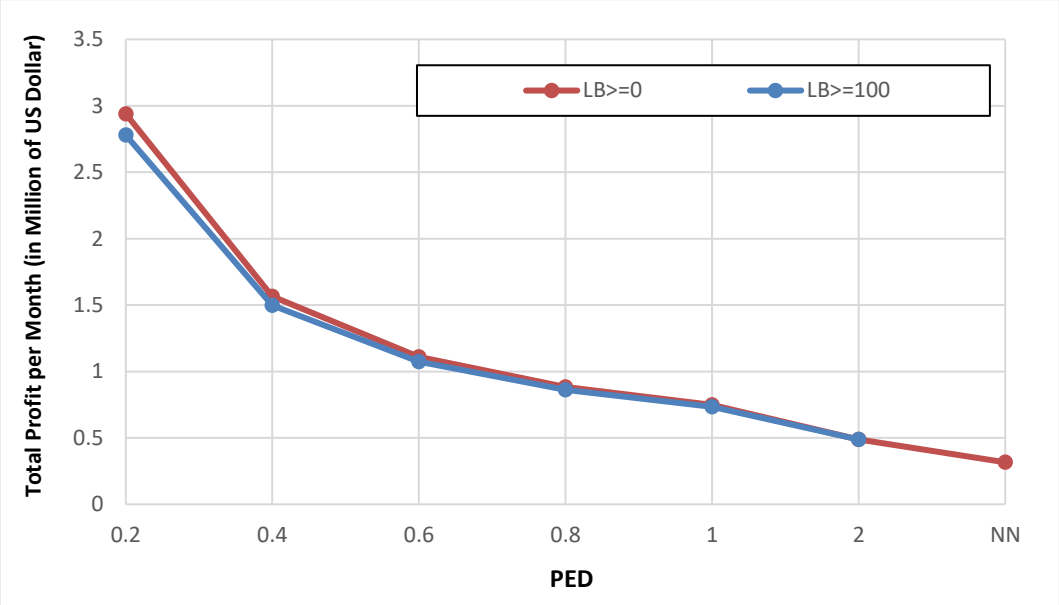

Figure 6-3: Total profit per month under different PED scenarios for cloud-based delivery.

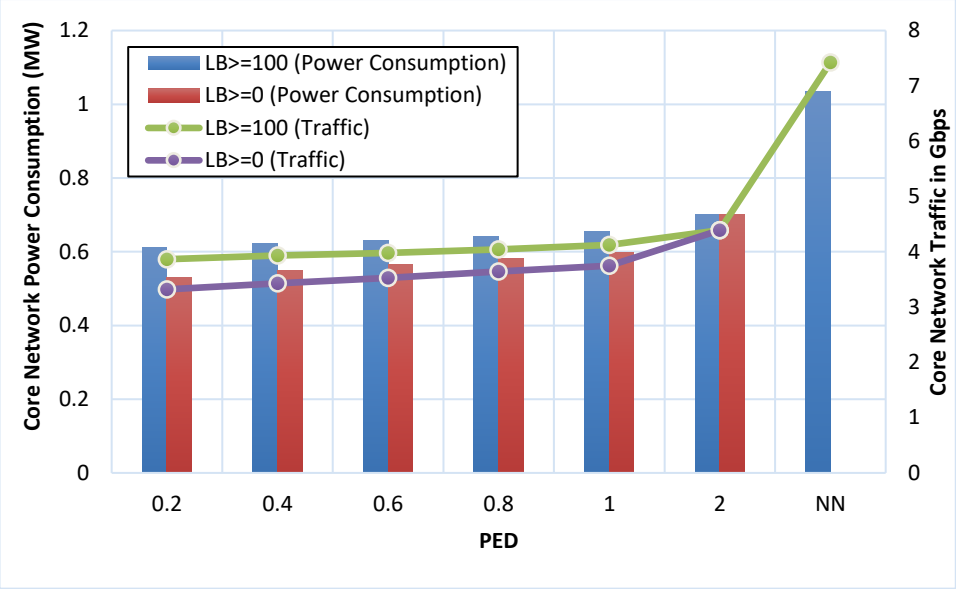

Figure 6-4: Total core network traffic and power consumption under different PED scenarios for cloud based delivery.

6.4.1.2 Cloud-Fog based delivery:

Next, we introduce 10 fog nodes in addition to the 12 datacentre locations. These fogs are assumed to be built in the proximity of nodes with the highest population in AT&T core network, so no core network cost (Ⓒ) is incurred by serving demands of these nodes. Fig. 6-5 shows that the prices per Gbps per month under different PED that are less than the previous case (cloud-based delivery) as we reduced the cost of the core network by introducing the fog nodes. Under $\boldsymbol{PED = 2}$, the prices compared to the net neutrality case in class A and B are increased by 12% and 11%, respectively, while the price of class C dropped by 1% as opposed to 19%, 18% and 1% with cloud-based delivery. The reduced prices attracted more users resulting in increasing the profit by 18% compared to the net neutrality case as seen in Fig. 6-6 as opposed to a 54% increase in profit with cloud-based delivery. Fig. 6-7 shows a reduction in core network traffic (40%) and power consumption (35%) by repealing net neutrality in the cloud-fog architecture.

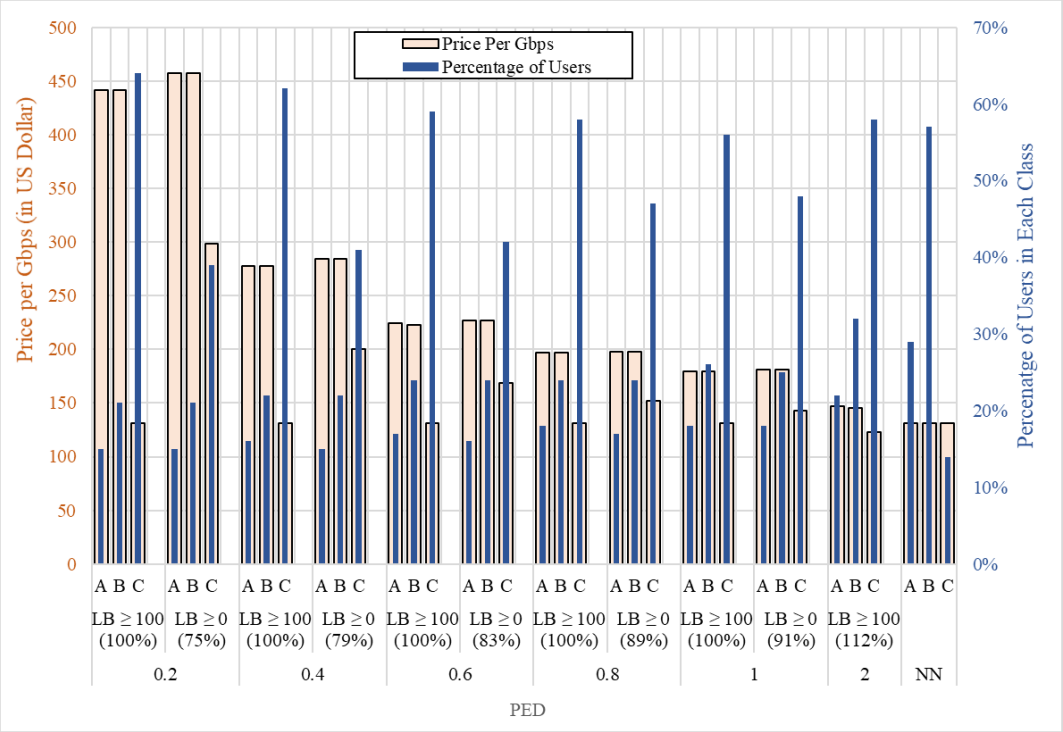

Figure 6-5: Price per Gbps per month and the corresponding number of users in each class based on different PED after removing net neutrality (cloud-fog based delivery).

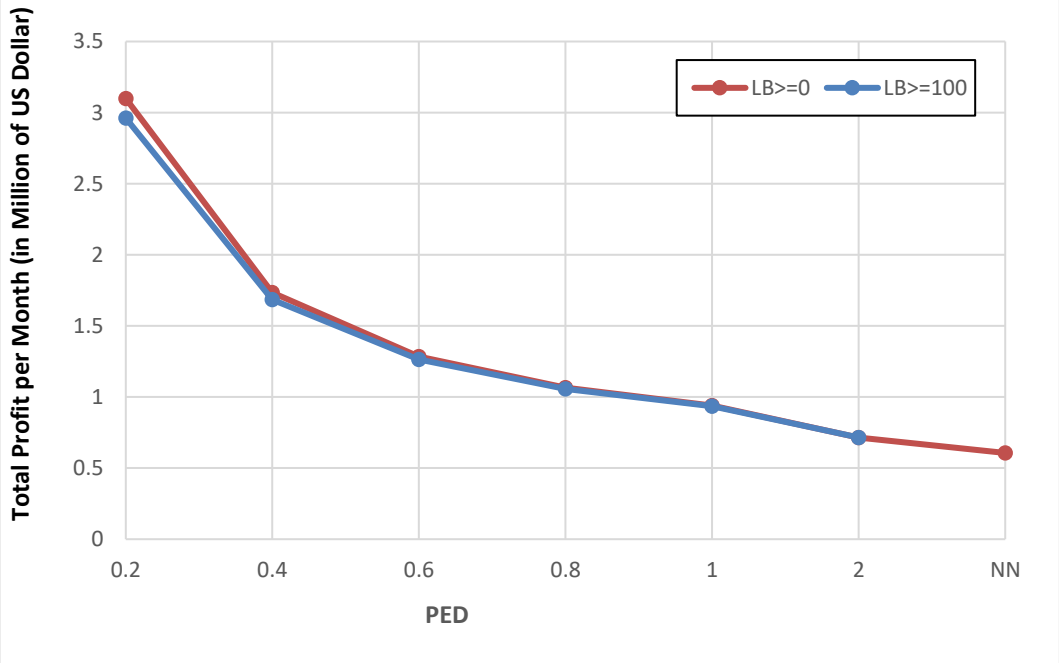

Figure 6-6: Total profit per month of profit-driven model under different PED (cloud-fog based delivery).

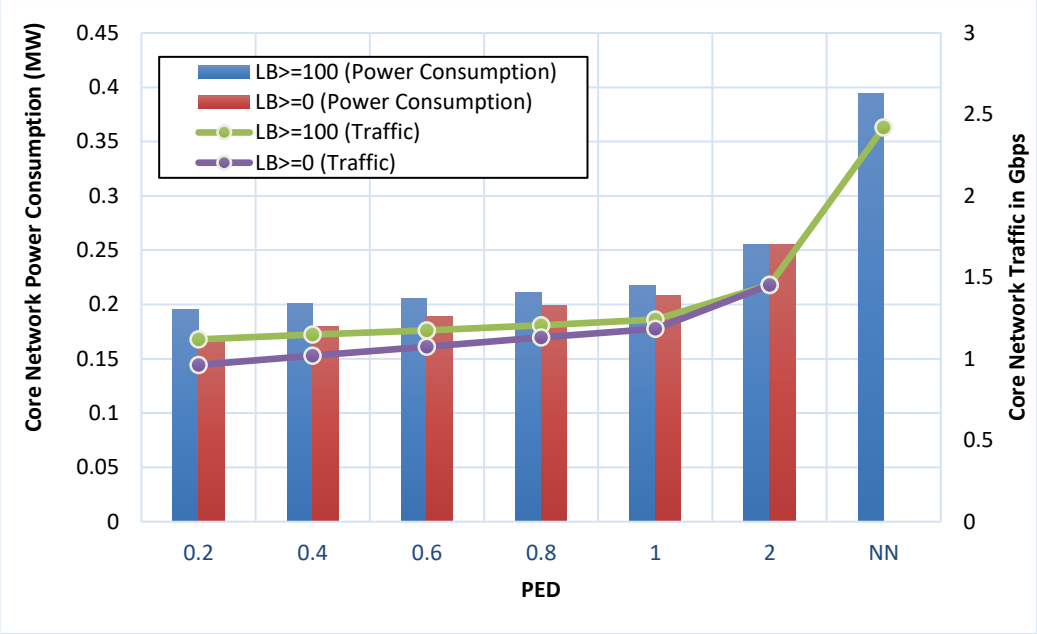

Figure 6-7: Total core network power consumption and bandwidth of profit- driven model under different PED (cloud-fog based delivery).

6.4.1.3 Fog based delivery:

Here, we consider a scenario in which all users access CP contents from a local fog node. Although, deploying a fog node locally to serve CP customers increases the capital expenditure (CAPEX) and operating expenses (OPEX) of provisioning multiple locations (i.e. 25 fog nodes in AT&T network), it reduces the communication network transit cost burden to the minimum. However, fog nodes are not always an option due to the finite capacity of processing and storage. The results show that the prices are further reduced under fog-based delivery (Figs. 6-8) as no core network cost (c) is incurred by serving demands. For instance, under $\boldsymbol{PED = 2}$, the prices compared to the net neutrality case in class A and B is increased by 9% while the price of class C is decreased by 11% resulting in increasing the profit by 6% compared to the net neutrality scenario as seen in Fig. 6-9.

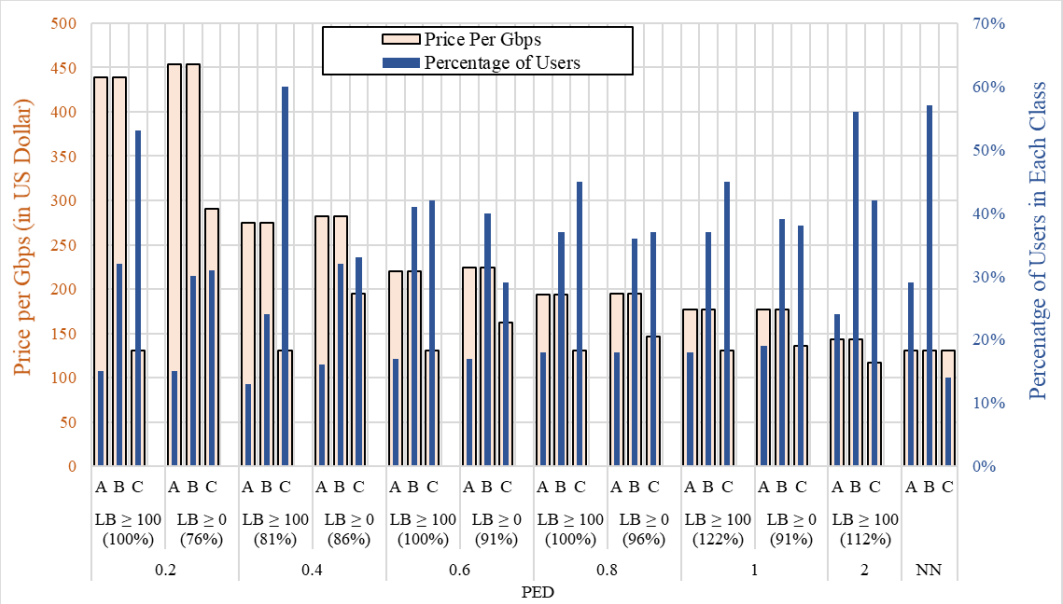

Figure 6-8: Price per Gbps per month and the corresponding number of users in each class based on different PED after removing net neutrality (fog-based delivery).

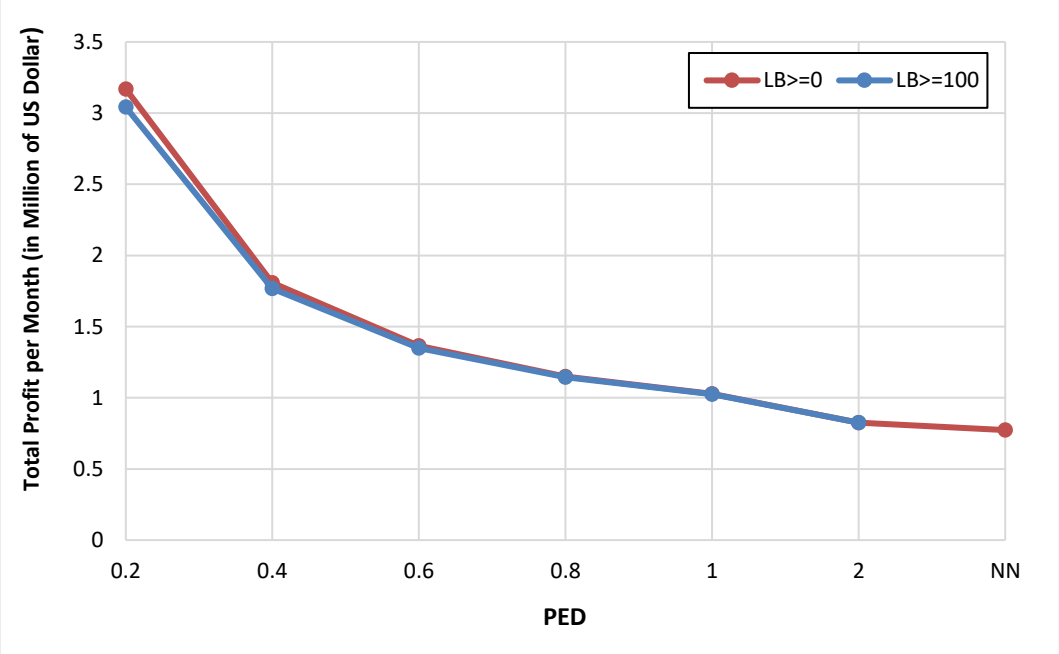

Figure 6-9: Total profit per month of profit-driven model under different PED (fog-based delivery).

6.4.2 Different PED Among Classes:

In this section, we consider a scenario where elasticity of demand varies among the different classes of service. We consider class C to be less sensitive to price change than class B. Also, we considered class B to be less sensitive than class A. The elasticity of demand for classes A, B and C are considered to be 2, 0.8 and 0.2,

respectively. Thus, 1% increase in price of classes A, B, and C will result in 2%, 0.8% and 0.2% decrease in the number of users, respectively. Hence, the users in class C will highly try to follow the increase in the price, compared to users of classes A and B in order to stay in the service. Fig. 6-10 shows the price per Gbps for classes A and B is the same under different scenarios and delivery schemes as a result of the high PED of class A. Class C is priced at the same level of classes A and B for $\boldsymbol{LB \geq 0}$ as the low PED of class C limits the number of users leaving the services as a result of increasing the price. Fig. 6-11 shows an increase in profit by up to 88%, 29% and 16% under cloud-based delivery, cloud-fog based delivery and fog-based delivery, respectively, compared to the net neutrality scenario. Fig 6-12 shows a decrease in core network traffic by up to 43% and 30% under cloud-based delivery, cloud-fog based delivery, respectively, compared to the net neutrality scenario. Also, the total reductions of core network power consumptions (as shown in Fig 6-13) are up to 40% and 32% under cloud-based delivery, and cloud-fog based delivery, respectively.

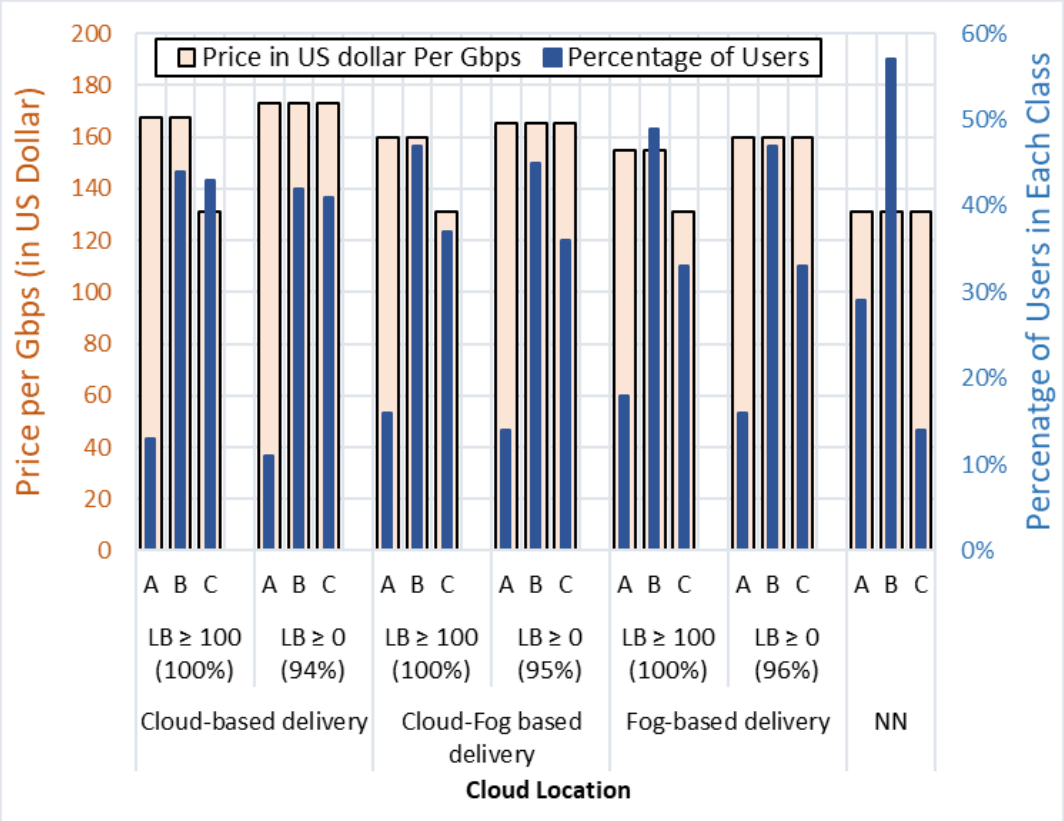

Figure 6-10: Price per Gbps per month and the corresponding number of users in each class for different CP delivery scenarios where PED values of different classes A, B and C are 2, 0.8 and 0.2, respectively.

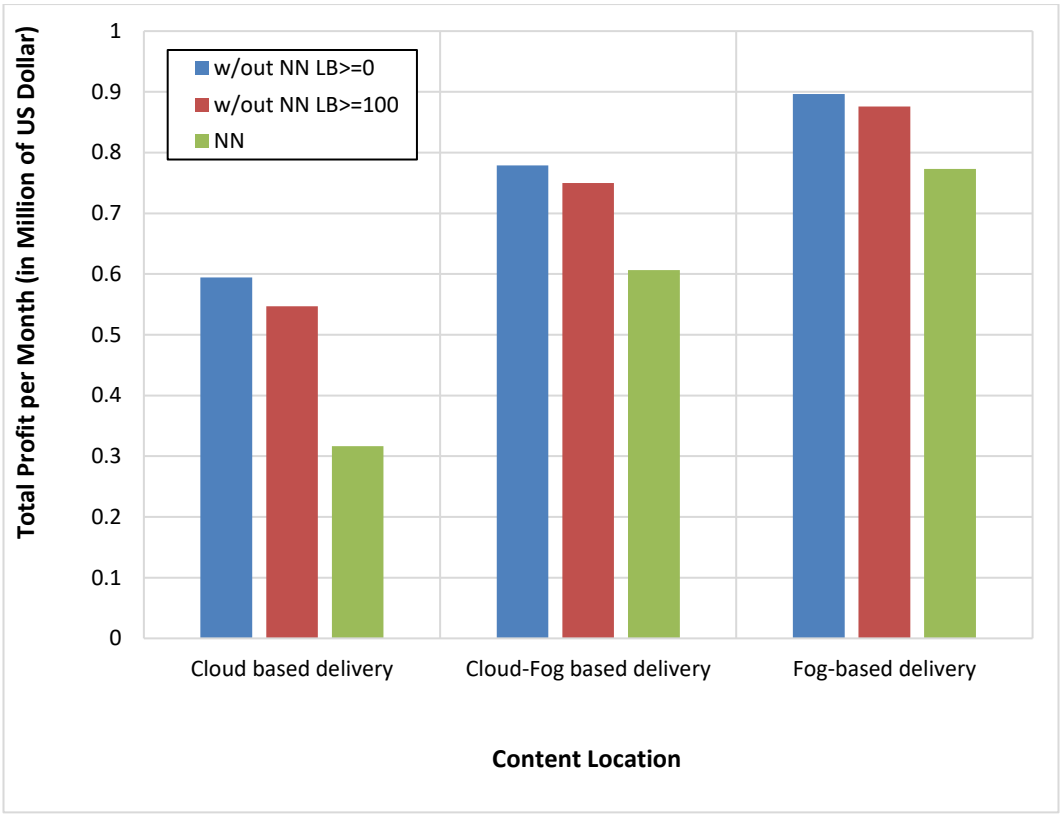

Figure 6-11: Total profit per month of profit-driven model for different CP delivery scenarios where PED values of different classes A, B and C are 2, 0.8 and 0.2, respectively.

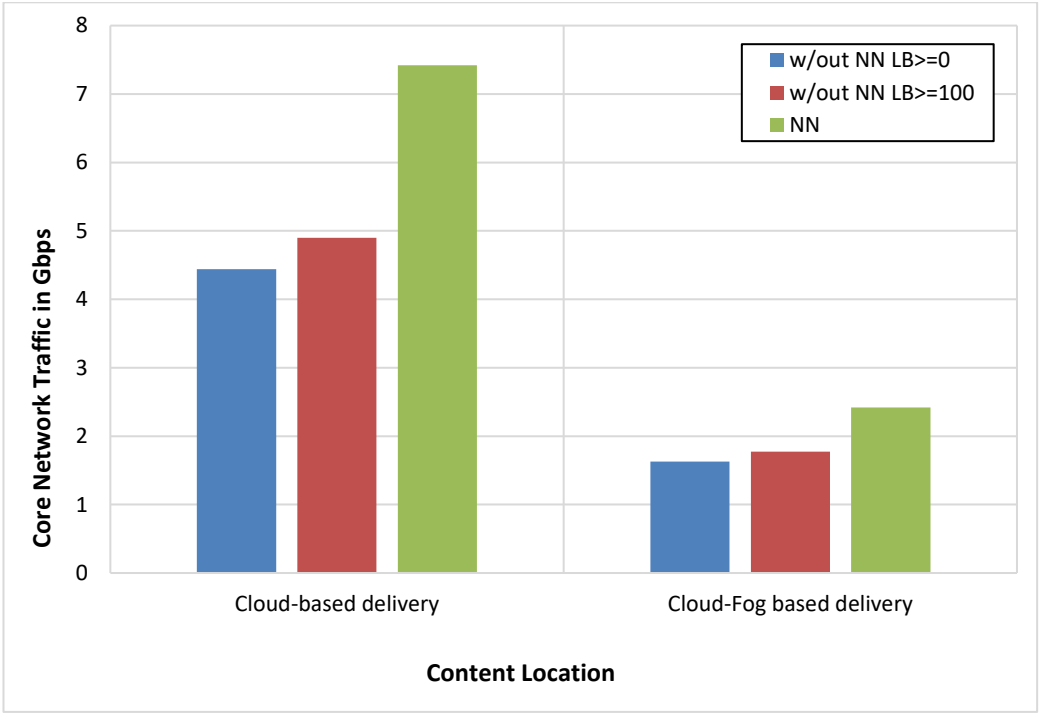

Figure 6-12: Total traffic resulting from profit-driven model for different CP delivery scenarios where PED values of different classes A, B and C are 2, 0.8 and 0.2, respectively.

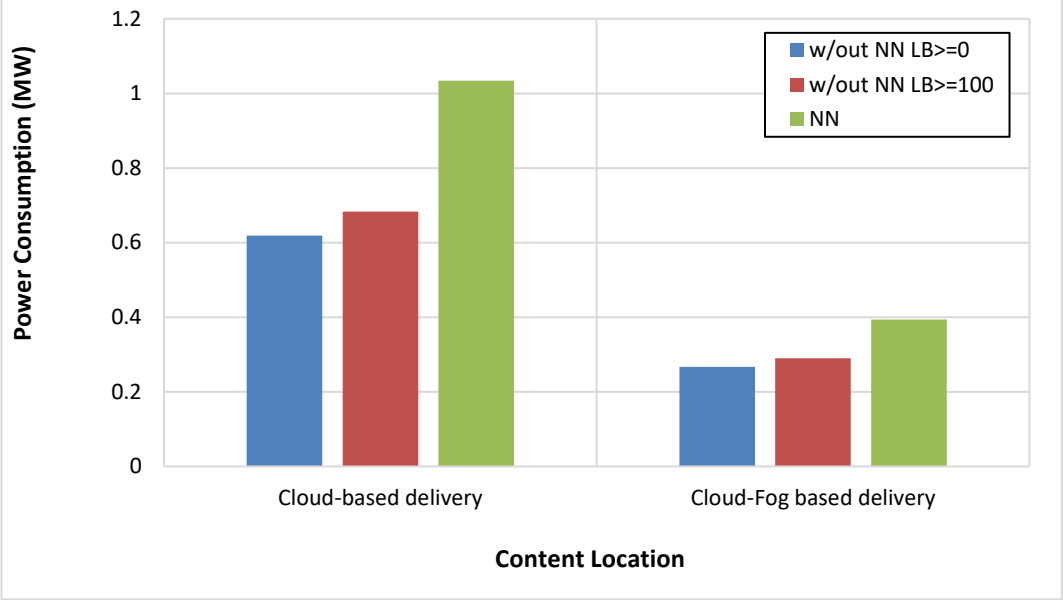

Figure 6-13: total core network power consumption resulting from profit-driven model for different CP delivery scenarios where PED values of different classes A, B and C are 2, 0.8 and 0.2, respectively.

6.5 Summary

In this chapter, we developed a MILP model to optimise the pricing scheme used by ISPs to charge CPs for delivering their video content under the repeal of net neutrality where ISPs can apply discrimination treatment on data intensive traffic. A techno-economic Mixed Integer Linear Programming (MILP) model is developed to maximise the ISP profit by optimising the ISP pricing scheme to charge different classes of service differently subject to PED. We considered three classes of service that represent different data rate requirements of video content. The analysis addressed three CP delivery scenarios; cloud-based delivery, cloud-fog based delivery and fog-based delivery. The results show that the discriminatory pricing scheme on video contents, which has low sensitivity to the change of price, can increase the ISPs profit by a factor of 8. The results also show that by influencing the way end-users consume data-intensive content, the core network traffic and consequently power consumption are reduced by up to 49% and 55%, respectively, compared to the net neutrality scenario.

Chapter 7: Conclusions and Future work

In this chapter, the thesis is summarised highlighting the original contributions and conclusions of each chapter. Directions for potential future work are also outlined.

7.1 Conclusions

The thesis has addressed the energy efficiency and profitability of cloud-fog architectures through the following contributions:

In Chapter 3, energy efficient VMs placement was investigated considering the problem of joint optimisation of energy consumption of optical core networks represented by an IP over WDM network and distributed clouds. The analysis addressed the impact of different factors including VM popularity, the traffic between the VM and its users, the VM workload, the workload versus number of users and their profile and power usage effectiveness (PUE). To evaluate the potential power savings, the problem is formulated as a MILP model and a real-time heuristic was developed to find the optimal VMs placement. The results show that popular VMs and VMs with high user data rate make network power consumption a vital factor in determining the VM placement. Also, the results indicate VMs with linear workload profile should be distributed to multiple clouds compared to VMs with constant workload. In addition, the results show that at high PUE, VM processing power consumption becomes a dominant factor in deciding VM placement. The results showed that the potential power savings achieved from the energy efficient VM placement over distributed clouds in AT&T and BT core networks considering the state-of-the-art applications are 51% and 38%, respectively, compared to a placement in conventional cloud locations.

In Chapter 4, the problem of energy efficient VMs placement over distributed clouds was further investigated by studying the impact of inter-VM traffic, in addition to users' traffic. The investigation considered the impact of two types of inter-VM

traffic; synchronisation and cooperation traffic. Cooperation traffic exists between different VMs to support the completion of their processing jobs. Synchronisation traffic exists among replicas of the same VM to keep the content of each replica up to date. The problem was formulated as a MILP model and a real-time heuristic was also developed. The results show that neglecting inter-VM traffic when placing VMs can increase the total power consumption by a factor of 39 for VMs with an inter-VM traffic data rate of 5 Gbps.

In Chapter 5, a framework based on MILP mathematical modelling and a heuristic was developed to study the energy efficient VMs placement problem over a cloud-fog architecture including the three layers of telecommunication networks; core, metro and access networks. The framework jointly optimised the use of networking and computing resources from the core to access networks, taking into consideration minimising the total power consumption of providing the service. It also addressed the impact of different factors including the VM popularity, traffic between the VM and its users, the VM workload, the workload versus number of users profile, proximity of fog nodes and the PUE. The optimal VMs placement and total power consumption savings were investigated considering AT&T and BT network topologies. The placement of VMs in cloud allows them to serve geo-distributed users across the core nodes whereas placing the replicas of a VM in the fog nodes eliminates core network crossing for the traffic between VMs and users and therefore reduces the network power consumption. However, placing VM replicas in distributed fogs increases the processing power consumption due to the creation of multiple replicas of the VMs. The creation of a VM replicas therefore results in power savings if the former power consumption exceeds the latter power consumption. Significant power savings of 48% compared to optimised placement in distributed clouds and 64% compared to a placement considering traditional cloud locations, have resulted from offloading VMs to the fog.

In Chapter 6, the potential impact of the repeal of net neutrality on communication networks has been explored. A techno-economic MILP model is developed with an objective to maximise ISPs profit by optimising the pricing scheme used by an ISP to charge CPs for delivering their video content. ISPs are considered to differentially charge classes of service representing data rate requirements of different video content quality. The analysis addressed three CP delivery scenarios; cloud-based delivery, cloud-fog based delivery and fog-based delivery. The results show that the proposed

pricing scheme can increase the ISPs profit by a factor of 8. The results also show that by applying a discriminatory pricing scheme on data-intensive content, the network traffic and consequently power consumption can be reduced by up to 49% and 55%, respectively, compared to the net neutrality scenario.

7.2 Future Work

In the following, potential future research directions in energy efficient cloud-fog architectures are proposed:

7.2.1 Concave Workload Profile

The work in this thesis has considered two different workload profiles; linear and constant VM workload profiles. However, there are other workload profiles that can be tackled when studying the relationship between users and VM workload such as concave workload profile [160]. This requires non-linear programming modelling.

7.2.2 Minimising Carbon Emission Intensity of Cloud-Fog Architecture

The VMs placement study can be extended to consider minimising the total carbon emission intensity of the cloud-fog architecture. Carbon emission intensity is the carbon emission per unit of energy consumed. In the UK grid system, the emission value varies from region to another based on the fuel type for each region. For instance, the carbon emission intensity in North Scotland is much lower than the carbon emission intensity in London [178]. However, London has a greater population than North Scotland. Thus, there is a trade-off between cloud emission reduction by placing VMs in North Scotland and the network emission increase incurred due to the need to route the users' traffic to London.

7.2.3 Energy Efficient Content Distribution over Cloud-Fog architecture

The study of energy efficient VMs placement can be extended to study energy efficient content distribution over cloud-fog architectures. The locations to store content can be optimised to minimise energy taking into consideration the requirements of the state-of-the-art content delivery applications (e.g. storage requirements, data rates, CPU requirements for video compression … etc).

7.2.4 Net Neutrality in A Competitive ISP Market

The net neutrality work presented in this thesis considered a monopoly ISP market, where a single ISP exists. However, more than a single ISP can be offering their services creating a competitive ISP market. The problem of the competitive ISP market can be mathematically solved by game theory, where the strategic interaction in between rational decision-makers is considered. In game theory, each player chooses the best decision taking into account the other players' decisions.

7.2.5 Minimising Latency in Cloud-Fog Architectures

In this thesis, the focus has been on investigating the cloud-fog architecture considering power consumption and profit as the performance metrics. However, the demand for delay-sensitive applications such as autonomous cars, augmented and virtual reality, real-time data analytics, and video surveillance applications is growing. One of the research directions that can be pursued is energy efficient and profitable placement in cloud-fog architectures where delay-sensitive applications have the priority to access fog nodes over less critical applications.

References

[1] Cisco Systems, “White paper: Cisco Visual Networking Index: Forecast and Trends, 2017–2022,” 2018.

[2] H. F. Atlam, R. J. Walters, and and G. B. Wills, “Fog Computing and the Internet of Things : A Review,” *Big Data Cogn. Comput*, pp. 1–18, 2018.

[3] W. Paper, “Cisco Fog Computing Solutions : Unleash the Power of the Internet of Things,” 2015.

[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] Microsoft, “Virtual Machine Load Balancing overview,” 2018. [Online]. Available: https://docs.microsoft.com/en-us/windows-server/failover-clustering/vm-load-balancing-overview. [Accessed: 24-Feb-2018].

[6] 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.

[7] I. Pietri and R. Sakellariou, “Mapping Virtual Machines onto Physical Machines in Cloud Computing : A Survey,” *J. ACM Comput. Surv.*, vol. 49, no. 3, 2016.

[8] Climate Change News, “‘Tsunami of data’ could consume one fifth of global electricity by 2025.” [Online]. Available: https://www.climatechangenews.com/2017/12/11/tsunami-data-consume-one-fifth-global-electricity-2025/. [Accessed: 13-Aug-2019].

[9] Y. Wu, M. Tornatore, S. Ferdousi, and B. Mukherjee, “Green Data Center Placement in Optical Cloud Networks,” *IEEE Trans. Green Commun. Netw.*, vol. 1, no. 3, pp. 1–1, 2017.

[10] H. A. Alharbi, T. E. H. El-gorashi, A. Q. Lawey, and J. M. H. Elmirghani, “Energy Efficient Virtual Machines Placement in IP over WDM Networks,”

in *19th International Conference on Transparent Optical Networks*, 2017, pp. 1–4.

[11] L. Zhang, T. Han, and N. Ansari, “Energy-Aware Virtual Machine Management in Inter-datacenter Networks over Elastic Optical Infrastructure,” *IEEE Trans. Green Commun. Netw.*, vol. 2400, no. XX, pp. 1–1, 2017.

[12] A. Choudhary, S. Rana, and K. J. Matahai, “A Critical Analysis of Energy Efficient Virtual Machine Placement Techniques and its Optimization in a Cloud Computing Environment,” *Phys. Procedia*, vol. 78, no. December 2015, pp. 132–138, 2016.

[13] F. Farahnakian *et al.*, “Using Ant Colony System to Consolidate VMs for Green Cloud Computing,” *IEEE Trans. Serv. Comput.*, vol. 8, no. 2, pp. 187–198, 2015.

[14] B. Kantarci, L. Foschini, A. Corradi, and H. T. Mouftah, “Inter-and-intra data center VM-placement for energy-efficient large-Scale cloud systems,” in *IEEE Globecom’12 Workshops: First International workshop on Management and Security technologies for Cloud Computing*, 2012, pp. 708–713.

[15] D. Puthal *et al.*, “Energy-efficient VM-placement in cloud data center,” *Sustain. Comput. Informatics Syst.*, vol. 20, pp. 48–55, 2018.

[16] A. Manzalini, R. Minerva, and T. Italia, “Clouds of Virtual Machines in Edge Networks,” *IEEE Commun. Mag.*, no. July, pp. 63–70, 2013.

[17] J. Baliga, R. Ayre, K. Hinton, W. V. Sorin, and R. S. Tucker, “Energy consumption in optical IP networks,” *J. Light. Technol.*, vol. 27, no. 13, pp. 2391–2403, 2009.

[18] Y. Zhang, P. Chowdhury, M. Tornatore, and B. Mukherjee, “Energy Efficiency in Telecom Optical Networks,” *IEEE Commun. Surv. TUTORIALS*, vol. 12, no. 4, pp. 441–458, 2010.

[19] G. S. G. Shen and R. S. Tucker, “Energy-Minimized Design for IP Over WDM Networks,” *IEEE/OSA J. Opt. Commun. Netw.*, vol. 1, no. 1, pp. 176–186, 2009.

[20] “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.

[21] Jun Zheng; Hussein T. Mouftah, *Optical WDM Networks:Concepts and Design Principles*, 1st ed. Wiley-IEEE Press, 2004.

[22] K. Zhu, S. Member, and B. Mukherjee, “Traffic Grooming in an Optical WDM Mesh Network,” vol. 20, no. 1, pp. 122–133, 2002.

[23] D. Neilson, “Photonics for Switching and Routing,” *IEEE J. Sel. Top. QUANTUM Electron. VOL. 12, NO. 4, JULY/AUGUST 2006*, vol. 12, no. 4, pp. 669–678, 2006.

[24] M. Forzati *et al.*, “Next-Generation Optical Access Seamless Evolution: Concluding Results of the European FP7 Project OASE,” *J. Opt. Commun. Netw.*, vol. 7, no. 2, pp. 109–123, 2015.

[25] Stephen B. Weinstein; Yuanqiu Luo; Ting Wang, “Recent advances and looking to the future,” in *The ComSoc Guide to Passive Optical Networks:Enhancing the Last Mile Access*, 2012, pp. 87–114.

[26] Huawei, “Next-Generation PON Evolution,” 2010. [Online]. Available: www.huawei.com/ilink/en/download/HW_077443. [Accessed: 07-Mar-2017].

[27] ITU-T, “G.987, 10-Gigabit-capable passive optical network (XG-PON) systems: Definitions, abbreviations and acronyms,” pp. 1–19, 2012.

[28] P. Mell and T. Grance, “The NIST definition of cloud computing,” *Natl. Inst. Stand. andTechnology*, 2011.

[29] Google, “Google App Engine.” [Online]. Available: https://cloud.google.com/appengine/. [Accessed: 23-Feb-2019].

[30] Microsoft, “Microsoft Azure.” .

[31] Microsoft, “What is middleware?” [Online]. Available: https://azure.microsoft.com/en-gb/overview/what-is-middleware/. [Accessed: 23-Feb-2019].

[32] Amazon, “Amazon EC2.” [Online]. Available: https://aws.amazon.com/ec2. [Accessed: 23-Feb-2019].

[33] S. Antonio, “Cisco Delivers Vision of Fog Computing to Accelerate Value from Billions of Connected Devices,” *Cisco*, 2014. [Online]. Available: https://newsroom.cisco.com/press-release-content?articleId=1334100. [Accessed: 27-Sep-2017].

[34] “OpenFog Consortium.” [Online]. Available: https://www.openfogconsortium.org/. [Accessed: 19-Jan-2019].

[35] S. Misra and S. Sarkar, “Theoretical modelling of fog computing: a green computing paradigm to support IoT applications,” *IET Networks*, vol. 5, no. 2, pp. 23–29, 2016.

[36] I. Stojmenovic and S. Wen, “The Fog Computing Paradigm : Scenarios and Security Issues,” in *Federated Conference on Computer Science and Information Systems*, 2014, vol. 2, pp. 1–8.

[37] J. Zhu, D. S. Chan, M. S. Prabhu, P. Natarajan, H. Hu, and F. Bonomi, “Improving Web Sites Performance Using Edge Servers in Fog Computing Architecture,” in *IEEE Seventh International Symposium on Service-Oriented System Engineering Improving*, 2013.

[38] M. Chiang and T. Zhang, “Fog and IoT: An Overview of Research Opportunities,” *IEEE Internet Things J.*, vol. 3, no. 6, pp. 854–864, 2016.

[39] A. Yousefpour, G. Ishigaki, and J. P. Jue, “Fog Computing : Towards Minimizing Delay in the Internet of Things,” in *IEEE 1st International Conference on Edge Computing*, 2017, pp. 17–24.

[40] Cisco Systems, “Fog Computing and the Internet of Things: Extend the Cloud to Where the Things Are,” *Www.Cisco.Com*, p. 6, 2016.

[41] X. Xu, Q. Liu, L. Qi, Y. Yuan, W. Dou, and A. X. Liu, “A Heuristic Virtual Machine Scheduling Method for Load Balancing in Fog-Cloud Computing,” *Proc. - 4th IEEE Int. Conf. Big Data Secur. Cloud, BigDataSecurity 2018, 4th IEEE Int. Conf. High Perform. Smart Comput. HPSC 2018 3rd IEEE Int. Conf. Intell. Data Secur.*, pp. 83–88, 2018.

[42] I. Acm, O. Skarlat, M. Nardelli, S. Schulte, and S. Dustdar, “Towards QoS-aware Fog Service Placement,” in *IEEE 1st International Conference on Fog and Edge Computing (ICFEC)*, 2017, pp. 89–96.

[43] O. F. R. and M. P. L. F. Bittencourt, J. Diaz-Montes, R. Buyya, “Mobility-Aware Application Scheduling in Fog Computing,” *IEEE CLOUD Comput.*, vol. 4, no. 2, pp. 26–35, 2017.

[44] S. K. Mishra, D. Puthal, J. J. P. C. Rodrigues, B. Sahoo, and E. Dutkiewicz, “Sustainable Service Allocation Using a Metaheuristic Technique in a Fog Server for Industrial Applications,” *IEEE Trans. Ind. Informatics*, vol. 14, no. 10, pp. 4497–4506, 2018.

[45] N. Iotti, M. Picone, and S. Cirani, “Improving Quality of Experience in Future Wireless Access Networks through Fog Computing,” *IEEE Internet Comput.*, vol. 21, no. 2, pp. 26–31, 2017.

[46] F. Jalali, K. Hinton, R. Ayre, T. Alpcan, and R. S. Tucker, “Fog computing may help to save energy in cloud computing,” *IEEE J. Sel. Areas Commun.*, vol. 34, no. 5, pp. 1728–1739, 2016.

[47] S. Sarkar, S. Chatterjee, and S. Misra, “Assessment of the Suitability of Fog Computing in the Context of Internet of Things,” *IEEE Trans. Cloud Comput.*, vol. 6, no. 1, pp. 46–59, 2018.

[48] A. G.-J. R. XAVI MASIP-BRUIN, EVA MARÍN-TORDERA, GHAZAL TASHAKOR, ADMELA JUKAN, “Foggy Clouds and Cloudy Fogs : A Real Need for Coordinated Management of Fog-to-Cloud Computing Systems,” *IEEE Wirel. Commun.*, no. October, pp. 120–128, 2016.

[49] H. Liu *et al.*, “Mobile Edge Cloud System: Architectures, Challenges, and Approaches,” *IEEE Syst. J.*, pp. 1–14, 2017.

[50] R. Deng, R. Lu, C. Lai, and T. H. Luan, “Towards power consumption-delay tradeoff by workload allocation in cloud-fog computing,” *IEEE Int. Conf. Commun.*, vol. 2015-Septe, pp. 3909–3914, 2015.

[51] R. J. Creasy, “The Origin of the VM/370 Time-Sharing System,” *IBM J. Res. Dev.*, vol. 25, no. 5, pp. 483–490, 1981.

[52] S. Nanda and T.-C. Chiueh, “A Survey on Virtualization Technologies,” *Rpe Rep.*, vol. 179, pp. 1–42, 2005.

[53] M. Masdari, S. S. Nabavi, and V. Ahmadi, “An overview of virtual machine placement schemes in cloud computing,” *J. Netw. Comput. Appl.*, vol. 66, pp.

106–127, 2016.

[54] C. P. Bezemer, A. Zaidman, B. Platzbeecker, T. Hurkmans, and A. Hart, “Enabling multi-tenancy: An industrial experience report,” in *IEEE International Conference on Software Maintenance, ICSM*, 2010.

[55] R. Wasim *et al.*, “A survey on virtual machine migration and server consolidation frameworks for cloud data centers,” *J. Netw. Comput. Appl.*, vol. 52, pp. 11–25, 2015.

[56] D. Kliazovich, P. Bouvry, and S. Ullah, “DENS : data center energy-efficient network-aware scheduling,” *Cluster Comput.*, vol. 16, no. 1, pp. 65–75, 2013.

[57] M. H. Ferdaus, M. Murshed, R. N. Calheiros, and R. Buyya, *Network-Aware Virtual Machine Placement and Migration in Cloud Data Centers*. 2015.

[58] Jarrod, “Linux Web Server Performance Benchmark – 2016 Results.” [Online]. Available: https://www.rootusers.com/linux-web-server-performance-benchmark-2016-results/. [Accessed: 24-Nov-2017].

[59] P. Padala, X. Zhu, Z. Wang, S. Singhal, and K. G. Shin, “Performance Evaluation of Virtualization Technologies for Server Consolidation,” 2007.

[60] P. R. M. Vasconcelos, G. A. D. A. Freitas, G. A. De Araujo, and T. G. Marques, “Virtualization technologies in web conferencing systems: A performance overview,” *2016 11th Int. Conf. Internet Technol. Secur. Trans. ICITST 2016*, pp. 376–383, 2017.

[61] A. Bharambe, J. Pang, and S. Seshan, “Colyseus : A Distributed Architecture for Online Multiplayer Games,” in *USENIX Association - NSDI ’06: 3rd Symposium on Networked Systems Design & Implementation 155*, 2006, pp. 155–168.

[62] Cisco, “Cisco Global Cloud Index : Forecast and Methodology , 2016–2021,” 2018.

[63] Google, “Google Help - Network Check,” 2018. [Online]. Available: https://support.google.com/wifi/answer/6246634?hl=en-GB. [Accessed: 02-Feb-2018].

[64] YouTube, “Video Quality Report,” 2018. [Online]. Available: https://www.google.com/get/videoqualityreport/#methodology.

[65] YouTube, “Live encoder settings, bitrates, and resolutions,” *YouTube Help*, 2017. [Online]. Available: https://support.google.com/youtube/answer/2853702.

[66] A. Bharambe *et al.*, “Donnybrook: Enabling Large-Scale, High-Speed, Peer-to-Peer Games,” in *SIGCOMM*, 2008, pp. 389–400.

[67] Netflix, “Internet Connection Speed Recommendations,” *Netflix Help Center*, 2014. [Online]. Available: https://help.netflix.com/en/node/306. [Accessed: 27-Jan-2018].

[68] VMware, “vSphere Replication,” 2018. [Online]. Available: https://www.vmware.com/uk/products/vsphere/replication.html. [Accessed: 27-Feb-2018].

[69] C. Gebhardt and A. Tomlinson, “Challenges for Inter Virtual Machine Communication,” 2010.

[70] X. Meng, V. Pappas, and L. Zhang, “Improving the scalability of data center networks with traffic-aware virtual machine placement,” *Proc. - IEEE INFOCOM*, 2010.

[71] K. LaCurts and S. Deng, “Choreo: network-aware task placement for cloud applications,” in *Conference on Internet measurement conference*, 2013, pp. 191–204.

[72] F. Farahnakian, T. Pahikkala, P. Liljeberg, J. Plosila, N. T. Hieu, and H. Tenhunen, “Energy-aware VM Consolidation in Cloud Data Centers Using Utilization Prediction Model,” *IEEE Trans. Cloud Comput.*, vol. 7, no. 2, pp. 1–1, 2019.

[73] M. C. Silva Filho, C. C. Monteiro, P. R. M. Inácio, and M. M. Freire, “Approaches for optimizing virtual machine placement and migration in cloud environments: A survey,” *J. Parallel Distrib. Comput.*, vol. 111, pp. 222–250, 2018.

[74] Z. Usmani and S. Singh, “A Survey of Virtual Machine Placement Techniques in a Cloud Data Center,” *Phys. Procedia*, vol. 78, no. December 2015, pp. 491–498, 2016.

[75] J. M. H. Elmirghani *et al.*, “GreenTouch GreenMeter Core Network Energy-

Efficiency Improvement Measures and Optimization," *J. Opt. Commun. Netw.*, vol. 10, no. 2, p. A250, 2018.

[76] T. Wood, E. Cecchet, K. K. Ramakrishnan, P. Shenoy, J. Van Der Merwe, and A. Venkataramani, "Disaster Recovery as a Cloud Service : Economic Benefits & Deployment Challenges University of Massachusetts Amherst," *Econ. Benefits Deploy. Challenges. HotCloud,* pp. 8–15, 2010.

[77] A. Bianco, L. Giraudo, and D. Hay, "Optimal Resource Allocation for Disaster Recovery," *2010 IEEE Glob. Telecommun. Conf. GLOBECOM 2010*, pp. 1–5, 2010.

[78] B. Sahoo, Z. Xue, S. Sharma, D. Puthal, A. Y. Zomaya, and S. K. Mishra, "Energy-Efficient Deployment of Edge Dataenters for Mobile Clouds in Sustainable IoT," *IEEE Access*, vol. 6, pp. 56587–56597, 2018.

[79] A. Horri, M. S. Mozafari, and G. Dastghaibyfard, "Novel resource allocation algorithms to performance and energy efficiency in cloud computing," *J. Supercomput.*, vol. 69, no. 3, pp. 1445–1461, 2014.

[80] M. Pedram and I. Hwang, "Power and performance modeling in a virtualized server system," *Proc. Int. Conf. Parallel Process. Work.*, pp. 520–526, 2010.

[81] J. Wang, C. Huang, K. He, X. Wang, X. Chen, and K. Qin, "An energy-aware resource allocation heuristics for VM scheduling in cloud," *Proc. - 2013 IEEE Int. Conf. High Perform. Comput. Commun. HPCC 2013 2013 IEEE Int. Conf. Embed. Ubiquitous Comput. EUC 2013*, pp. 587–594, 2014.

[82] Q. Zheng *et al.*, "Virtual machine consolidated placement based on multi-objective biogeography-based optimization," *Futur. Gener. Comput. Syst.*, vol. 54, no. March, pp. 95–122, 2016.

[83] X. Li, Z. Qian, S. Lu, and J. Wu, "Energy efficient virtual machine placement algorithm with balanced and improved resource utilization in a data center," *Math. Comput. Model.*, vol. 58, no. 5–6, pp. 1222–1235, 2013.

[84] M. Dabbagh, B. Hamdaoui, M. Guizani, and A. Rayes, "Efficient datacenter resource utilization through cloud resource overcommitment," *Proc. - IEEE INFOCOM*, vol. 2015-Augus, pp. 330–335, 2015.

[85] H. Shen, S. Member, and L. Chen, "Management in Datacenters Using Finite-

Markov Decision Process," *IEEE/ACM Trans. Netw.*, vol. 25, no. 6, pp. 3836–3849, 2017.

[86] N. Fernando, S. W. Loke, and W. Rahayu, "Mobile cloud computing: A survey," *Futur. Gener. Comput. Syst.*, vol. 29, no. 1, pp. 84–106, 2013.

[87] Y. Luo, "Network I/O Virtualization for Cloud Computing," *IT Prof.*, vol. 12, no. 5, pp. 36–41, 2010.

[88] I. Takouna, R. Rojas-Cessa, K. Sachs, and C. Meinel, "Communication-aware and energy-efficient scheduling for parallel applications in virtualized data centers," *Proc. - 2013 IEEE/ACM 6th Int. Conf. Util. Cloud Comput. UCC 2013*, pp. 251–255, 2013.

[89] B. Zhang, Z. Qian, W. Huang, X. Li, and S. Lu, "Minimizing communication traffic in data centers with power-aware VM placement," *Proc. - 6th Int. Conf. Innov. Mob. Internet Serv. Ubiquitous Comput. IMIS 2012*, pp. 280–285, 2012.

[90] S. Wang, A. Zhou, C. H. Hsu, X. Xiao, and F. Yang, "Provision of Data-Intensive Services Through Energy-and QoS-Aware Virtual Machine Placement in National Cloud Data Centers," *IEEE Trans. Emerg. Top. Comput.*, vol. 4, no. 2, pp. 290–300, 2016.

[91] V. Shrivastava, P. Zerfos, K.-W. Lee, H. Jamjoom, Y.-H. Liu, and S. Banerjee, "Application-aware virtual machine migration in data centers," *INFOCOM, 2011 Proc. IEEE*, pp. 66–70, 2011.

[92] D. Z. and Z. Y. Weiwei Fang, Zishuo Wang, Jaime LIoret, "Optimising data placement and traffic routing for energy saving in Backbone Networks," *Trans. Emerg. Tel. Tech.*, vol. 25, no. 2014, pp. 914–925, 2013.

[93] A. Varasteh and M. Goudarzi, "Server consolidation techniques in virtualized data centers: A survey," *IEEE Syst. J.*, vol. PP, no. 99, pp. 1–12, 2015.

[94] O. Ayoub, F. Musumeci, M. Tornatore, and A. Pattavina, "Efficient Routing and Bandwidth Assignment for Inter-Data-Center Live Virtual-Machine Migrations," *J. Opt. Commun. Netw.*, vol. 9, no. 3, p. B12, 2017.

[95] M. Kurtadikar, A. Patil, P. Toshniwal, and J. Abraham, "An Inter-VM Communication Model Supporting Live Migration," *Int. Conf. Cloud*

Ubiquitous Comput. Emerg. Technol., pp. 63–68, 2013.

[96] N. Tziritas *et al.*, “Data Replication and Virtual Machine Migrations to Mitigate Network Overhead in Edge Computing Systems,” *IEEE Trans. Sustain. Comput.*, vol. PP, no. 99, pp. 1–14, 2017.

[97] M. Alicherry and T. V. Lakshman, “Network aware resource allocation in distributed clouds,” *Proc. - IEEE INFOCOM*, pp. 963–971, 2012.

[98] T. K. ZA. Tzanakaki, M. P. Anastasopoulos, S. Peng, B. Rofoee, Y. Yan D. Simeonidou, G. Landi, G. Bernini, N. Ciulli, J.F. Riera, E. Escalona, J. A. Garcia-Espin, K. Katsalis, “A Converged Network Architecture for Energy Efficient Mobile Cloud Computing,” *2015 3rd IEEE Int. Conf. Mob. Cloud Comput. Serv. Eng.*, vol. 1233, no. c, pp. 120–125, 2014.

[99] A. Ceselli, M. Premoli, S. Secci, and A. Mislove, “Mobile Edge Cloud Network Design Optimization,” *IEEE/ACM Trans. Netw.*, vol. 25, no. 3, pp. 1818–1831, 2017.

[100] 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.

[101] M. Dayarathna, Y. Wen, A. Mislove, and R. Fan, “Data Center Energy Consumption Modeling : A Survey,” *IEEE Commun. Surv. TUTORIALS*, vol. 18, no. 1, pp. 732–794, 2016.

[102] G. Katsaros, J. Subirats, J. O. Fitó, J. Guitart, P. Gilet, and D. Espling, “A service framework for energy-aware monitoring and VM management in Clouds,” *Futur. Gener. Comput. Syst.*, vol. 29, no. 8, pp. 2077–2091, 2013.

[103] Y. Gao, H. Guan, Z. Qi, B. Wang, and L. Liu, “Quality of service aware power management for virtualized data centers,” *J. Syst. Archit.*, vol. 59, no. 4–5, pp. 245–259, 2013.

[104] A. Beloglazov and R. Buyya, “Optimal online deterministic algorithms and adaptive heuristics for energy and performance efficient dynamic consolidation of virtual machines in Cloud data centers,” *Concurr. Comput. Pract. Exp.*, vol. 24, no. 13, pp. 1397–1420, 2012.

[105] D. Dong and J. Herbert, “Energy efficient VM placement supported by data

analytic service," *Proc. - 13th IEEE/ACM Int. Symp. Clust. Cloud, Grid Comput. CCGrid 2013*, pp. 648–655, 2013.

[106] N. Trung Hieu, M. Di Francesco, and A. Ylä-Jääski, "Virtual Machine Consolidation with Usage Prediction for Energy-Efficient Cloud Data Centers," in *8th International Conference on Cloud Computing*, 2015, pp. 750–757.

[107] "IBM System Energy Estimator," *S814-8286-41A*. [Online]. Available: http://see.au-syd.mybluemix.net/see/EnergyEstimator. [Accessed: 15-Mar-2017].

[108] W. Dargie, "A stochastic model for estimating the power consumption of a processor," *IEEE Trans. Comput.*, vol. 64, no. 5, pp. 1311–1322, 2015.

[109] 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.

[110] W. Fang, X. Liang, S. Li, L. Chiaraviglio, and N. Xiong, "VMPlanner: Optimizing virtual machine placement and traffic flow routing to reduce network power costs in cloud data centers," *Comput. Networks*, vol. 57, no. 1, pp. 179–196, 2013.

[111] B. Yu, Y. Han, X. Wen, X. Chen, and Z. Xu, "An energy-aware algorithm for optimizing resource allocation in software defined network," *2016 IEEE Glob. Commun. Conf. GLOBECOM 2016 - Proc.*, pp. 1–7, 2016.

[112] B. Zhang, Z. Qian, W. Huang, X. Li, and S. Lu, "Minimizing communication traffic in data centers with power-aware VM placement," *Proc. - 6th Int. Conf. Innov. Mob. Internet Serv. Ubiquitous Comput. IMIS 2012*, pp. 280–285, 2012.

[113] N. Tziritas, C. Z. Xu, T. Loukopoulos, S. U. Khan, and Z. Yu, "Application-Aware workload consolidation to minimize both energy consumption and network load in cloud environments," *Proc. Int. Conf. Parallel Process.*, pp. 449–457, 2013.

[114] B. Yu, Y. Han, X. Wen, X. Chen, and Z. Xu, "An energy-aware algorithm for optimizing resource allocation in software defined network," *2016 IEEE Glob. Commun. Conf. GLOBECOM 2016 - Proc.*, pp. 1–7, 2016.

[115] L. Tsai and W. Liao, “Cost-aware workload consolidation in green cloud datacenter,” *2012 1st IEEE Int. Conf. Cloud Networking, CLOUDNET 2012 - Proc.*, pp. 29–34, 2012.

[116] P. Arroba, J. M. Moya, J. L. Ayala, and R. Buyya, “Dynamic Voltage and Frequency Scaling-aware dynamic consolidation of virtual machines for energy efficient cloud data centers,” *Concurr. Comput.* , vol. 29, no. 10, 2017.

[117] A. Beloglazov and R. Buyya, “OpenStack Neat: a framework for dynamic and energy-efficient consolidation of virtual machines in OpenStack clouds,” *Concurr. Comput. Pract. Exp.*, 2014.

[118] E. Feller, C. Rohr, D. Margery, and C. Morin, “Energy management in IaaS clouds: A holistic approach,” *Proc. - 2012 IEEE 5th Int. Conf. Cloud Comput. CLOUD 2012*, pp. 204–212, 2012.

[119] H. N. Van, F. D. Tran, and J. M. Menaud, “Performance and power management for cloud infrastructures,” *Proc. - 2010 IEEE 3rd Int. Conf. Cloud Comput. CLOUD 2010*, pp. 329–336, 2010.

[120] Cisco, “The Zettabyte Era: Trends and Analysis,” 2017.

[121] Ycharts.com, “AT&T Profit Margin (Quarterly),” 2018. [Online]. Available: https://ycharts.com/companies/T/profit_margin. [Accessed: 01-Aug-2018].

[122] Ycharts, “Netflix Profit Margin (Quarterly),” 2018. [Online]. Available: https://ycharts.com/companies/NFLX/profit_margin. [Accessed: 01-Aug-2018].

[123] T. Garrett, L. E. Setenareski, L. M. Peres, L. C. E. Bona, E. P. D. Jr, and A. Mislove, “Monitoring Network Neutrality : A Survey on Traffic Differentiation Detection,” *IEEE Commun. Surv. Tutorials*, no. c, pp. 1–32, 2018.

[124] A. M. Kakhki, D. Choffnes, A. Mislove, and E. Katz-bassett, “BingeOn Under the Microscope : Understanding T-Mobile ’ s Zero-Rating Implementation,” *Proc. 2016 Work. QoE-based Anal. Manag. Data Commun. Networks (Internet-QoE ’16) ACM*, pp. 43–48, 2016.

[125] J. S. Huang and B. Zhu, “Video Traffic Detection Method for Deep Packet Inspection,” *Adv. CSIE*, vol. 2, pp. 135–140, 2012.

[126] Cisco Systems, “Traffic Profiling,” pp. 1–12.

[127] D. Tsilimantos, T. Karagkioules, A. Nogales-g, and S. Valentin, “Traffic profiling for mobile video streaming,” *arXiv*, pp. 1–7, 2017.

[128] W. Dai, J. W. Baek, and S. Jordan, “Neutrality between a vertically integrated cable provider and an over-the-top video provider,” *J. Commun. Networks*, vol. 18, no. 6, pp. 962–974, 2016.

[129] Amazon, “Netflix Case Study - Amazon Web Services (AWS),” 2018. [Online]. Available: https://aws.amazon.com/solutions/case-studies/netflix/. [Accessed: 15-Aug-2018].

[130] AT&T, “AT&T Content Delivery Network (CDN) Services,” 2018. [Online]. Available: https://www.business.att.com/solutions/Family/cloud/content-delivery-network/. [Accessed: 15-Aug-2018].

[131] Comcast, “CDN - Comcast Technology Solutions,” 2018. [Online]. Available: https://www.comcasttechnologysolutions.com/resources/comcast-cdn-one-sheet. [Accessed: 15-Aug-2018].

[132] “AT & T Switched Ethernet Classic (ASE Classic) and On-Demand (ASEoD).”

[133] D. Grunwald, “The Internet Ecosystem: The Potential for Discrimination,” *Fed. Commun. Law J.*, vol. 63, no. 2, p. 5, 2011.

[134] A. Odlyzko, “Network Neutrality, Search Neutrality, and the Never-ending Conflict between Efficiency and Fairness in Markets,” *Rev. Netw. Econ.*, vol. 8, no. 1, pp. 40–60, 2009.

[135] P. Nooren, A. Leurdijk, and N. van Eijk, “Net neutrality and the value chain for video,” *Info*, vol. 14, no. 6, pp. 45–58, 2012.

[136] R. T. B. Ma, J. Wang, and D. M. Chiu, “Paid prioritization and its impact on net neutrality,” *IEEE J. Sel. Areas Commun.*, vol. 35, no. 2, pp. 367–379, 2017.

[137] P. Mailé, G. Simon, and B. Tuffin, “Toward a net neutrality debate that conforms to the 2010s,” *IEEE Commun. Mag.*, vol. 54, no. 3, pp. 94–99, 2016.

[138] “IBM ILOG CPLEX Optimization.” [Online]. Available:

www.ibm.com/Data_Science/CPLEX%0A. [Accessed: 01-Mar-2019].

[139] L. Nonde, T. E. H. Elgorashi, and J. M. H. Elmirghani, “Cloud virtual network embedding: profit, power and acceptance,” *2015 IEEE Glob. Commun. Conf. GLOBECOM 2015*, 2015.

[140] AMPL, “A Mathematical Programming Language - AMPL.” [Online]. Available: https://www.ampl.com/REFS/amplmod.pdf. [Accessed: 01-Mar-2019].

[141] “lpsolve.” [Online]. Available: http://lpsolve.sourceforge.net/5.1/AMPL.htm. [Accessed: 01-Mar-2019].

[142] M. Pioro and D. Medhi, *Routing , Flow and Capacity Design in Communication and Computer Networks*. 2003.

[143] Google Data Centers, “Efficiency: How we do it,” 2018. [Online]. Available: https://www.google.com/about/datacenters/efficiency/. [Accessed: 27-Jan-2018].

[144] AT&T, “AT&T’s 38 Global Internet Data Centers.” [Online]. Available: https://www.beatriceco.com/pdf/AB-0119-09_eb_idcmap.pdf. [Accessed: 15-Aug-2018].

[145] ZTE CORPORATION, “ZXA10 C300 Hardware Description,” 2013. [Online]. Available: http://enterprise.zte.com.cn/en/products/network_lnfrastructure/broadband_access/xpon_onu/201312/t20131209_414454.html. [Accessed: 15-Mar-2017].

[146] A. Shehabi *et al.*, “United States Data Center Energy Usage Report,” 2016.

[147] “Facilities | N8HPC,” 2016. [Online]. Available: http://n8hpc.org.uk/facilities/. [Accessed: 08-Feb-2016].

[148] Cisco System, “Cisco CRS-1 4-Slot Single-Shelf System,” 2014.

[149] Cisco System, “Cisco ONS 15454 40 Gbps CP-DQPSK Full C-Band Tuneable Transponder Card.”

[150] Oclaro, “OTS-4400 Regenerator,” 2014.

[151] MRV, “EDFA Optical Amplifiers,” 2012.

[152] “Intelligent Optical System 600,” *Glimmerglass*, 2013. [Online]. Available:

http://www.glimmerglass.com/ products/intelligent-optical-systems/. [Accessed: 16-Mar-2017].

[153] Cisco, “Cisco Network Convergence System 5500 Series Data Sheet,” pp. 1–11, 2018.

[154] Cisco Systems, “Cisco Nexus 9300-EX and 9300-FX Platform Switches,” 2018.

[155] C. Gray, R. Ayre, K. Hinton, and R. S. Tucker, “Power consumption of IoT access network technologies,” in *2015 IEEE International Conference on Communication Workshop (ICCW)*, 2015, pp. 2818–2823.

[156] IBM, “IBM Systems Energy Estimator,” 2018.

[157] Cisco, “United States - 2021 Forecast Highlights,” 2018.

[158] “Similar Web,” 2018. [Online]. Available: https://www.similarweb.com/. [Accessed: 10-May-2019].

[159] L. A. Adamic and B. A. Huberman, “Zipf’s law and the Internet,” *Glottometrics*, vol. 3, no. 1, pp. 143–150, 2002.

[160] A. Venkatraman, V. Pandey, B. Plale, and S. Shei, “Benchmarking Effort of Virtual Machines on Multicore Machines,” 2007.

[161] E. S. Mi, Ningfang, N Giuliano Casale, Ludmila Cherkasova, “Injecting Realistic Burstiness to a Traditional Client-Server Benchmark,” in *ICAC’09*, 2009, pp. 15–19.

[162] Microsoft, “Azure locations.” [Online]. Available: https://azure.microsoft.com/en-us/global-infrastructure/locations/. [Accessed: 23-Jan-2019].

[163] Amazon, “London in the Cloud.” [Online]. Available: https://aws.amazon.com/london/. [Accessed: 23-Jan-2019].

[164] Google Cloud, “Cloud locations.” [Online]. Available: https://cloud.google.com/about/locations/. [Accessed: 23-Jan-2019].

[165] IBM, “Power System S814LC,” *Data sheet*, 2015. [Online]. Available: https://www.ibm.com/. [Accessed: 28-Jan-2018].

[166] “Datasheet Mx Series 3D Universal Edge,” *Juniper Networks*, 2017. [Online].

Available: https://www.juniper.net/assets/us/en/local/pdf/datasheets/1000597-en.pdf. [Accessed: 23-Nov-2017].

[167] "Data Sheet Ex4550 Ethernet Switch," *Juniper Networks*, 2017. [Online]. Available: https://www.juniper.net/assets/uk/en/local/pdf/datasheets/1000414-en.pdf. [Accessed: 23-Nov-2017].

[168] J. Networks, "PTX1000, PTX10001, PTX10002, AND PTX10003 FIXEDCONFIGURATION PACKET TRANSPORT ROUTERS," 2019.

[169] Juniper, "MX Series 5G Universal Routing Platforms," 2018.

[170] R. S. Pindyck and D. L. Rubinfeld, *Microeconomics*. 2005.

[171] R. Cadman, C. Dineen, and R. Cadman, "Price and Income Elasticity of Demand for Broadband Subscriptions : A Cross-Sectional Model of OECD Countries," *SPC Network, 19*, pp. 3–8, 2009.

[172] H. Galperin and C. A. Ruzzier, "Price elasticity of demand for broadband : Evidence from Latin America and the Caribbean," *Telecomm. Policy*, vol. 37, pp. 429–438, 2013.

[173] USAToday, "Interent bill too high," 2018.

[174] Cisco, "United States - 2021 Forecast Highlights," pp. 1–6, 2016.

[175] B. Insider, "Netflix subscribers over the years," 2019. [Online]. Available: https://www.businessinsider.com/netflix-subscribers-chart-2017-1?r=US&IR=T. [Accessed: 01-Jul-2018].

[176] Ycharts.com, "Comcast Corp Historical Profit Margin (Quarterly) Data," 2018. [Online]. Available: https://ycharts.com/companies/CMCSA/profit_margin. [Accessed: 23-Aug-2018].

[177] J. Baliga, K. Hinton, and R. S. Tucker, "Energy Consumption of the Internet," in *COIN-ACOFT 2007 - Joint International Conference on the Optical Internet and the 32nd Australian Conference on Optical Fibre Technology*, 2007, pp. 1–3.

[178] National Grid, "Carbon Intensity Forecast Methodology," 2019. [Online]. Available: http://carbonintensity.org.uk/. [Accessed: 12-Aug-2018].