

Received April 16, 2020, accepted May 14, 2020, date of publication May 18, 2020, date of current version June 2, 2020.

Digital Object Identifier 10.1109/ACCESS.2020.2995393

Energy Efficient Virtual Machines Placement Over Cloud-Fog Network Architecture

HATEM A. ALHARBI¹, TAISIR E. H. ELGORASHI¹,
AND JAAFAR M. H. ELMIRGHANI¹, (Senior Member, IEEE)

School of Electronic and Electrical Engineering, University of Leeds, Leeds LS2 9JT, U.K.

Corresponding author: Hatem A. Alharbi (ml13ha@leeds.ac.uk)

This work was supported by the Engineering and Physical Sciences Research Council (EPSRC), through the INTERNET Project under Grant EP/H040536/1, the STAR Project under Grant EP/K016873/1, and the TOWS Project under Grant EP/S016570/1.

ABSTRACT Fog computing is an emerging paradigm that aims to improve the efficiency and QoS of cloud computing by extending the cloud to the edge of the network. This paper develops a comprehensive energy efficiency analysis framework based on mathematical modeling and heuristics to study the offloading of virtual machine (VM) services from the cloud to the fog. The analysis addresses the impact of different factors including the traffic between the VM and its users, the VM workload, the workload versus number of users profile and the proximity of fog nodes to users. 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 case, the creation of multiple VM replicas does not increase the computing power consumption significantly (there may be a slight increase due to idle / baseline power consumption) if the number of users remains constant, however the VM replicas can be brought closer to the end users, thus reducing the transport network power consumption. In our scenario, the optimum placement of VMs over a cloud-fog architecture significantly decreased the total power consumption by 56% and 64% under high user data rates compared to optimized distributed clouds placement and placement in the existing AT&T network cloud locations, respectively.

INDEX TERMS Fog computing, cloud computing, energy efficiency, virtual machine.

I. INTRODUCTION

Cloud computing has started to transform the information and communication technology (ICT) industry by providing efficient resource-sharing solutions in an Internet-based pool of network, storage, and computational resources available to simultaneously serve many geographically distributed users. Cloud computing essentially enables the development of the emerging Internet of Things (IoT) and Big Data applications. By 2020, total cloud traffic is expected to grow to 3.7 times its level in 2015, reaching 1.2 zettabytes per month and accounting for 92% of total data center traffic [1]. This mounting traffic creates a huge burden on data centers and networks, leading to serious energy efficiency and quality-of-service (QoS) challenges [2].

Fog computing, introduced by Cisco in 2014 [3], complements central cloud services by offloading some services to geographical proximity to users at network edges for

The associate editor coordinating the review of this manuscript and approving it for publication was Rongbo Zhu¹.

more efficient service access. Research of fog computing has mainly focused on illustrating its potential advantages over cloud computing. Fog computing is proposed to provide low latency [4], conserve network bandwidth [5], and improve QoS [6] and quality of experience (QoE) [7] for various computing services. Also in [8], providing network function virtualization (NFV) in the fog computing layer achieved low delay and efficient performance in data transmission, caching, and data integrity.

The energy consumption of cloud and fog computing has received limited attention in the literature, however. In [9], the authors found that the number of hops between users and content has little impact on total energy consumption compared to the type of application running on servers and factors such as the number of downloads and updates. In [10], the authors 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. A detailed analysis of the essential service metrics regarding the cost and benefit of offloading services to the fog layer

is yet to be conducted to identify the services that the fog can efficiently host, however. Such an analysis is crucial to sustain the growth of the IoT and Big Data applications, which are proving to be pivotal to economic growth and quality of life. In [11], the authors built a theoretical model of fog computing architecture and compared it to the conventional cloud computing model. In addition to low latency, they found that offloading applications to the fog layer can significantly reduce power consumption by 41%. However, they did not consider a detailed model of the telecom network architecture.

Mathematical modelling gives a concise and accurate representation of a problem to help in understanding the problem and solving it. Mathematical optimization finds the maximum or minimum solution of a function, referred to as the objective function. Mathematical optimization, where a problem is formulated as a mixed integer linear programming (MILP) model, has been used intensively in the literature to solve network design problems. The authors of [12], [13] used MILP models to study the energy efficiency of core networks. In [14] MILP models are used to develop energy-efficient network topologies. Network carbon emissions reduction by introducing renewable energy is formulated as an MILP model in [15]. Network resilience and its impact on energy efficiency is assessed using MILP models in [16], [17]. The work in [18]–[21] introduced MILP models for energy-efficient content distribution. Network optimization for energy-efficient Big Data transport is formulated as MILP models in [22], [23]. In [24], [25] the use of Big Data analytics to optimize networks is investigated using MILP models. In [26]–[28] the problem of energy-efficient VM placement over a core network is formulated as a MILP model. Finally, [29]–[31] present a MILP model for energy-efficient network virtualization.

In this paper, we develop a comprehensive framework based on mathematical modeling and heuristics to study the offloading of VM services from the cloud to the fog layer to minimize the total power consumption of service provision. We optimize the placement of VMs over an end-to-end cloud-fog architecture that traverses the core network, metropolitan (metro) network, and access network. VM placement in the cloud at the core network allows VMs to serve users distributed across the core nodes, whereas placing VM replicas closer to the users in the fog nodes in the metro or access networks limits the traffic between users and VMs to the metro and access networks, respectively, thus eliminating the associated core network traffic (and potentially the metro traffic). This reduces the network power consumption but increases the processing power consumption due to the creation of multiple VM replicas, and therefore a trade-off exists.

The remainder of this paper is organized as follows. Section II discusses the concept of machine virtualization and VMs workload profile and introduces the MILP model for optimizing the VM placement in the cloud-fog architecture. We present the optimization model results and analyze them

in Section III. A real-time VM placement heuristic is proposed in Section IV. Finally, Section V concludes the paper.

II. ENERGY EFFICIENT PLACEMENT OF VIRTUAL MACHINES OVER CLOUD-FOG ARCHITECTURE

A. MACHINE VIRTUALIZATION

Cloud and fog processing employ VMs for efficient resource utilization. Virtualization abstracts the server resources including the CPU, RAM, hard disk and 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 [31]. Further dynamism can be achieved by migrating or replicating VMs over geo-distributed servers to achieve different features such as load balancing [32] and energy efficiency [33]. The problem of migration and replication of VMs is referred to as VMs placement. VMs placement needs to be optimized to follow variations in the VMs demands, workload of the cloud/fog resources or network status [34]. Cloud and fog processing employ VMs for efficient resource utilization. Virtualization abstracts the server resources including the CPU, RAM, hard disk and 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 [31]. Further dynamism can be achieved by migrating or replicating VMs over geo-distributed servers to achieve different features such as load balancing [32] and energy efficiency [33]. The problem of migration and replication of VMs is referred to as VMs placement. VMs placement needs to be optimized to follow variations in the VMs demands, workload of the cloud/fog resources or network status [34].

Several papers have discussed VM placement considering various factors. To reduce server load, improve QoS, and meet service-level agreements (SLAs), VMs can be migrated or replicated to another server or servers within the same data center [35] or in geographically distributed data centers [36]. Virtualized cloud architectures can also provide efficient disaster resilience in case of physical machine failure by migrating VMs into different host machines or by replicating VM content in distributed data centers [37].

Under-utilized servers can significantly increase energy consumption and consequently increase the carbon emissions and operating costs of cloud data centers. VM consolidation by bin packing into fewer servers can significantly improve energy efficiency. The authors in [38] proposed a VM placement algorithm that uses VM popularity to explore the search space, achieving up to 40% power savings and reducing the number of servers used by up to 50% compared to placement based on first fit decreasing (FFD) techniques. In [39], the authors proposed a VM placement algorithm that balances the processing and memory resources of servers resulting in reducing the total power consumption by 15% compared to

FFD techniques. In [40], the authors presented an energy-efficient approach for profile-based VM placement considering various VM profiles (e.g., central processing unit [CPU] and random access memory [RAM] requirements). Their algorithm reduced power consumption by up to 20% compared to FFD techniques. The work in [41] proposed a novel strategic formulation using an n -person cooperative game in which users request and pay for VM instances together. In most cases, however, the VM instances provided for users exceeded actual need, which led to resource over-provision and higher VM cost for users. The authors concluded that users could pay less for requested VMs if they cooperate, which led to consequent decreases in power consumption by up to 25% compared to FFD and enhanced FFD techniques.

In the fog layer, most studies of VM placement have been limited to evaluating the reduction of overall network overhead [42], optimizing the placement of physical resources in the edge network [43], and scheduling VMs to share limited fog resources to minimize SLA violations [44]. Despite the diverse factors affecting the power consumption of cloud-fog architectures, the problem of providing energy-efficient VM placement over end-to-end cloud-fog architecture has received little scholarly attention. Thus, the objective of this paper is to develop a novel framework that covers different networks and computing and optimizes the energy efficiency of VM placement. This paper compares VM placement in a cloud-fog architecture over the American Telephone and Telegraph (AT&T) network where VMs are placed in the 12 data centers built in the AT&T network [45] versus optimized VMs placement over clouds in any core nodes of the AT&T network.

Despite the diverse factors affecting the power consumption of cloud-fog architectures, the problem of providing energy-efficient VM placement over end-to-end cloud-fog architecture considering end-to-end architecture has not received any attention. Thus, the objective of this paper is to develop a novel framework that covers different networks and computing in optimizing the energy efficiency of VMs placement.

B. VM WORKLOAD PROFILE

VM power consumption is determined by hosting servers. The authors in [41] found that the CPU utilization and power consumption of a server are highly correlated. Another work [42] studied the relationship between the power consumption and CPU utilization of a server and found a linear relation between them. This paper follows the same approach but considers CPU utilization only in modeling the power consumption of VM placement.

From a CPU perspective, various studies have shown that VM workload versus the number of VM-served users mostly follows one of two profiles: constant or linear (Fig. 1). In [43], the authors presented a CPU performance benchmark study for web-application VMs serving varying numbers of users with constant CPU workloads (Fig. 1 (a)). Also, various benchmarking studies have demonstrated linear workload

profiles for database applications [44], web-based video conferencing systems [45], and multiplayer games [46] with different slope coefficients. To maintain SLAs, each VM needs a minimum workload to run an application regardless of the number of users it serves, resulting in the workload profile in Fig. 1 (b). The minimum workload required to serve a user in a VM varies from as low as 1% to 60% based on the application [44]–[46].

C. MILP MODEL

The MILP model objective is to minimize the power consumption of the end-to-end cloud-fog network architecture accounting for different networking and processing layers. The total power consumption comprises two parts: (i) the traffic-induced power consumption due to delivering VM services from the cloud (over core, metro, and access networks) or fog (over metro and access networks or access network) to users, and (ii) processing induced power consumption in clouds and fogs. The MILP model objective is subject to many constraints related to VM placement, communication network, and processing requirements and capabilities.

We approach VM placement in a cloud-fog architecture, unlike [33], where certain placement schemes were imposed on all types of VMs. We allow the MILP model to select the most energy-efficient VM placement based on VM popularity, minimum VM workload requirement, and data rate. The model aims to achieve the ideal trade-off between network power saved by replicating VMs in multiple clouds and/or fog nodes and the power consumed by those replicas. A VM replica saves power if the former power exceeds the latter. Before introducing the model, Fig. 2 defines the parameters and variables related to the different layers of the cloud-fog architecture.

1) CLOUD AND FOG NODES

A typical data center, as illustrated in Fig. 2, 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 data center.

The resources at the fog nodes form mini data centers connected in a similar way to the cloud data centers. Servers, switches and routers in the cloud and fog nodes are defined by the following parameters:

Cloud and fog parameters

$SW^{(CB)}$	Cloud switch bit rate.
$SW^{(CP)}$	Cloud switch power consumption.
$SW^{(MFB)}$	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.
$SW^{(R)}$	Cloud and fog switch redundancy.
$R^{(CB)}$	Cloud router port bit rate.

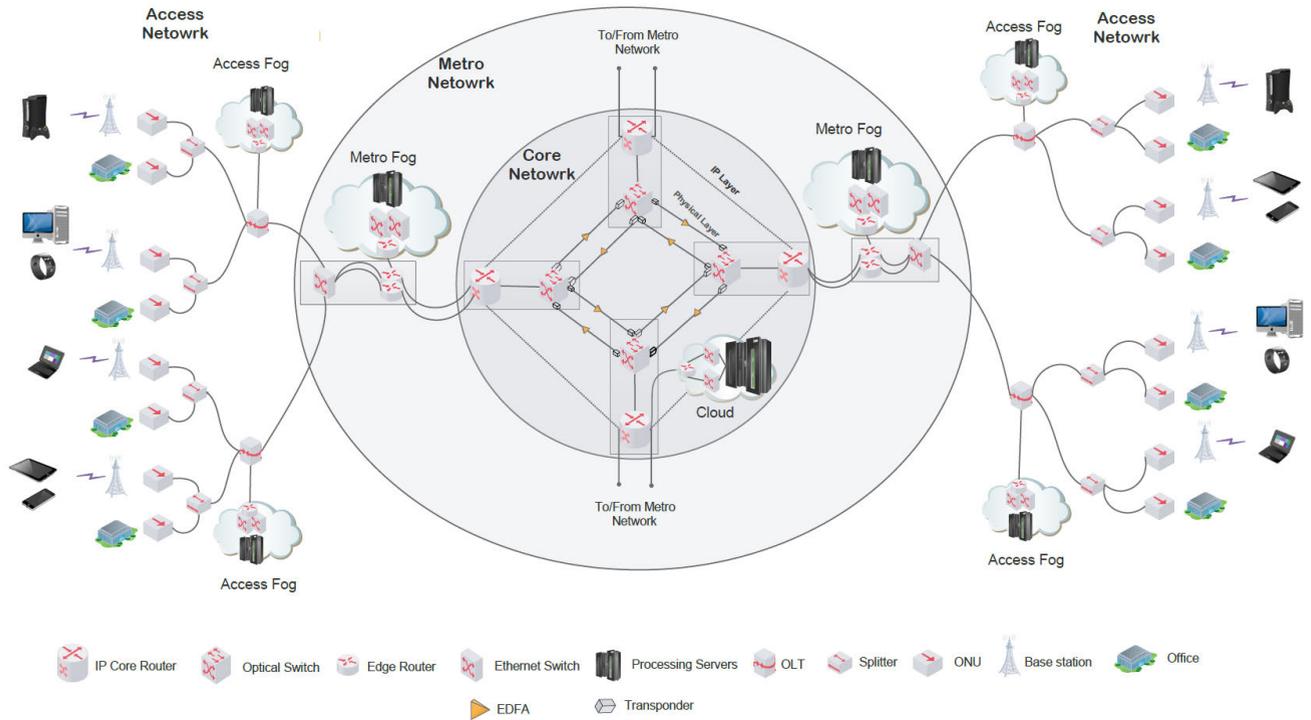


FIGURE 1. Cloud-Fog architecture.

- $R^{(CP)}$ Cloud router port 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.
- $S^{(P)}$ Power consumption of a server.
- $S^{(maxW)}$ Maximum workload of a server.
- c Cloud power usage effectiveness.
- m Metro fog power usage effectiveness.
- a Access fog power usage effectiveness.

Cloud and fog variables

- 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 .
- $F_s^{(MF)}$ $F_s^{(MF)} = 1$ 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 .

The VMs to be hosted in the cloud and/or fog and the traffic resulting from them are defined by the following parameters and variables:

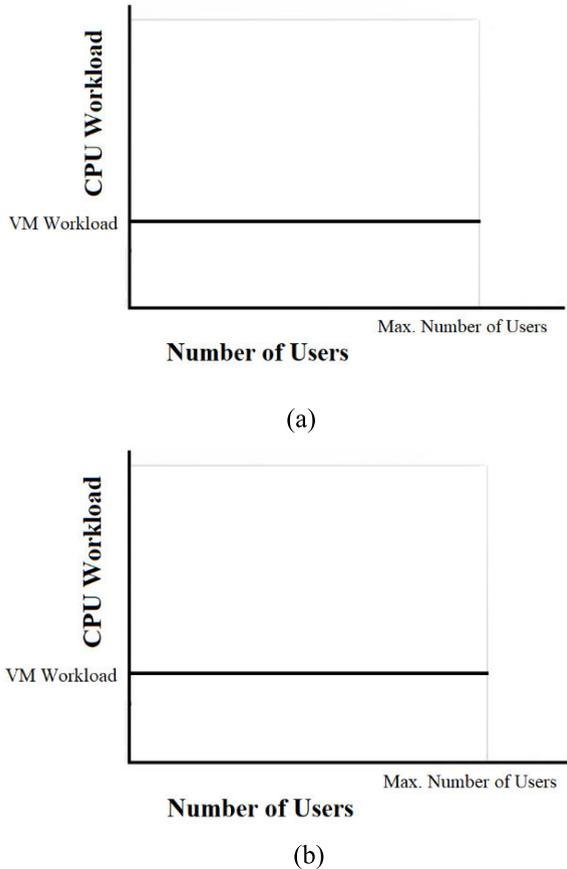


FIGURE 2. Relationship between VM workload and the number of users; (a) constant (b) linear relationship between VM workload and number of users.

VM parameters

- N Set of IP over WDM network nodes.
- VM Set of VM services.
- 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 utilization needed in the absence of load).
- T_v Traffic resulting from VM replica v serving the maximum number of users. $T_v = xr_v$
- $W_v^{(R)}$ Workload per traffic unit, $W_v^{(R)} = \frac{W_v - M}{T_v}$ evaluated for VM replica v .

VM variables

- $W_{v,s}^{(CR)}$ Workload of VM replica v hosted in cloud in nodes.
- $W_s^{(C)}$ 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 .
- $W_{v,s}^{(MFR)}$ 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 .

The clouds power consumption (*CLOUD*) is composed of:

- (i) Power consumption of cloud servers:

$$c \sum_{s \in N} S_s^{(C)} S^{(P)} \quad (1)$$

- (ii) Power consumption of cloud routers and switches:

$$c \left(\sum_{s \in N} \left((SW_s^{(C)} SW^{(R)} SW^{(CP)}) + R_s^{(C)} R^{(CP)} \right) \right) \quad (2)$$

The metro fogs (*MF*) power consumption is composed of:

- (i) Power consumption of metro fog servers:

$$m \sum_{s \in N} S_s^{(MF)} S^{(P)} \quad (3)$$

- (ii) Power consumption of metro fog switches and routers:

$$m \left(\sum_{s \in N} \left((SW_s^{(MF)} SW^{(R)} SW^{(MFP)}) + R_s^{(MF)} R^{(MFP)} \right) \right) \quad (4)$$

The access fogs power consumption (*AF*) is composed of:

- (i) Power consumption of access fog servers:

$$a \sum_{s \in N} \sum_{p \in P} S_{p,s}^{(AF)} S^{(P)} \quad (5)$$

- (ii) Power consumption of access fog switches and routers:

$$a \left(\sum_{s \in N} \sum_{p \in P} \left((SW_{p,s}^{(AF)} SW^{(R)} SW^{(AFP)}) + R_{p,s}^{(AF)} R^{(AFP)} \right) \right) \quad (6)$$

Note that, as the difference between the server idle power and full load is very small [47], we consider an on-off power profile for servers, i.e. if a server is activated, it operates at maximum power consumption.

2) ACCESS NETWORK

Passive optical networks (PONs) [48] are the selected technology for the access network in the cloud-fog architecture given in Fig. 1 due to their high bandwidth, reliability, and high data transmission compared to Ethernet access networks. At present, the gigabit PON (GPON) architecture has become the most popular solution for PON among service providers [49]. Two main active components are deployed in GPON; 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. The Optical distribution networking [50] provides a passive physical transmission between OLT and ONU. XGPON is capable of delivering data rate up to 10 Gbps over a single port. In this work, we consider 10G-PON as an example of the PON network.

The following parameters and variables are defined to represent PON networks:

Access network parameters:

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.

PON networks power consumption (PON) is composed of:

(i) Total power consumption of OLT:

$$n \left(\sum_{p \in P} \sum_{d \in N} (OLT^{(P)} OLT_{p,d}^{(N)}) \right) \quad (7)$$

(ii) Total power consumption of ONUs:

$$n \left(\sum_{p \in P} \sum_{d \in N} (ONU^{(P)} ONU_{p,d}^{(N)}) \right) \quad (8)$$

3) METRO NETWORK

A metro network [51] functions as a gateway for the access networks into the core network. Metro Ethernet is the dominant technology used in enterprise metro network. The basic components of metro Ethernet are Ethernet switch and edge routers as shown in Fig. 2. The Ethernet switch interconnects several access networks together. Also, it connects the access networks to edge routers. The best practice in ISP metro network is to use two edge routers in order to provide reliability and redundancy to the network [52]. The following parameters are defined to represent the metro network.

Metro network parameters:

$R^{(MB)}$	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.

Metro network variables:

$R_s^{(M)}$	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 ($Metro$) is composed of:

(i) Total power consumption of edge routers:

$$n \left(\sum_{s \in N} R_s^{(M)} R^{(MR)} R^{(MP)} \right) \quad (9)$$

(ii) Total power consumption of edge Ethernet switches:

$$n \left(\sum_{s \in N} SW_s^{(M)} SW^{(MP)} \right) \quad (10)$$

4) CORE NETWORK

The IP over WDM network [53] is the most commonly used architecture in core networks. The components of the IP layer and physical layer are shown on Fig. 2. In the IP layer, the core router controls the Internet traffic. It aggregates the IP traffic packets from the edge router to be sent to their destination. Optical switches make the connection between physical layer and IP layer. Optical switches are connected to fiber links. In each switching node, the transponder provides

optical-electronic-optical (OEO) conversion for full wavelength conversion. In addition, for long distance transmission, erbium-doped fiber amplifiers (EDFAs) are used to amplify the optical signal in each fiber [53]. Regenerators are used to re-amplify, re-shape and re-time (3R) the optical signal in long-haul transmission [54]. The IP over WDM network can be implemented using either the non-bypass approach or the lightpath bypass approach. Under the non-bypass approach, the packets are processed by the IP layer of every intermediate node during their journey from the source to destination. On the other hand, under the bypass approach, the intermediate nodes introduce a shortcut by bypassing the IP layer (of intermediate nodes) on the way to the destination node.

The following parameters and variables are defined to represent the IP over WDM core network:

Core network parameters:

- 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 nodes.
- $G^{(P)}$ 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.

Core network variables:

- $C_{i,j}$ 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 (*Core*) is composed of [53]:

(i) The power consumption of router ports:

$$n \left(\sum_{s \in N} R^{(P)} R_s^{(AC)} + \sum_{d \in N} R^{(P)} R_d^{(AE)} + \sum_{m \in N} \sum_{n \in Nm_m, n \neq m} R^{(P)} W_{m,n} \right) \quad (11)$$

(ii) The power consumption of transponders:

$$n \left(\sum_{m \in N} \sum_{n \in Nm_m, n \neq m} t^{(P)} W_{m,n} \right) \quad (12)$$

(iii) The power consumption of EDFAs:

$$n \left(\sum_{m \in N} \sum_{n \in Nm_m, n \neq m} e^{(P)} F_{m,n} A_{m,n} \right) \quad (13)$$

(iv) The power consumption of optical switches:

$$n \left(\sum_{s \in N} SW_s^{(P)} \right) \quad (14)$$

(v) The power consumption of regenerator:

$$n \left(\sum_{m \in N} \sum_{n \in Nm_m, n \neq m} G^{(P)} G_{m,n} W_{m,n} \right) \quad (15)$$

The model is defined as follows:

The objective: *Minimize total power consumption given as the sum of the power consumptions:*

$$Core + Metro + PON + CLOUD + MF + AF \quad (16)$$

Expression (16) gives the total power consumption as the sum of the power consumption of the IP over WDM core network, the metro network, the PON access network, clouds, metro fogs and access fogs.

Subject to:

Serving VM demand constraints:

$$\sum_{p \in P} \sum_{d \in N} D_{v,p,d} = \sum_{s \in N} \sum_{d \in N} D_{v,s,d}^{(C)} + \sum_{s \in N} D_{v,s}^{(MF)} + \sum_{p \in P} \sum_{s \in N} D_{v,p,s}^{(AF)} \quad \forall v \in VM \quad (17)$$

Constraint (17) ensures that the users demand for a VM ($\sum_{p \in P} \sum_{d \in N} D_{v,p,d}$) is satisfied by VMs placed at the clouds and/or the metro fogs and/or the access fogs.

Placing VM in cloud constraints:

$$L \sum_{d \in N} D_{v,s,d}^{(C)} \geq \delta_{v,s}^{(C)} \quad \forall s \in N, v \in VM \quad (18)$$

$$\sum_{d \in N} D_{v,s,d}^{(C)} \leq L \delta_{v,s}^{(C)} \quad \forall s \in N, v \in VM \quad (19)$$

Constraints (18) and (19) relate the binary variable that indicates whether a VM is hosted in a cloud or not ($\delta_{v,s}^{(C)}$) to the

traffic between users of this VM and the cloud ($\sum_{d \in N} D_{v,s,d}^{(C)}$) by setting $\delta_{v,s}^{(C)} = 1$ if $\sum_{d \in N} D_{v,s,d}^{(C)} > 0$ and $\delta_{v,s}^{(C)} = 0$ otherwise.

Placing VM in metro fog constraints:

$$D_{v,s}^{(MF)} \geq \delta_{v,s}^{(MF)} \quad \forall s \in N, v \in VM \quad (20)$$

$$D_{v,s}^{(MF)} \leq L\delta_{v,s}^{(MF)} \quad \forall s \in N, v \in VM \quad (21)$$

Constraints (20) and (21) relate the binary variable that indicates whether a VM is hosted in a 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:

$$D_{v,p,s}^{(AF)} \geq \delta_{v,p,s}^{(AF)} \quad \forall s \in N, v \in VM, p \in P \quad (22)$$

$$D_{v,p,s}^{(AF)} \leq L\delta_{v,p,s}^{(AF)} \quad \forall s \in N, v \in VM, p \in P \quad (23)$$

Constraints (22) and (23) relate the binary variable that indicates whether a VM is hosted in an access fog or not ($AF\delta_{v,sp}$) to the traffic between users of this VM and the cloud ($D_{v,p,s}^{(AF)}$), by setting $AF\delta_{v,sp} = 1$ if $D_{v,p,s}^{(AF)} > 0$ and $\delta_{v,s}^{(AF)} = 0$ otherwise.

Clouds locations constraints:

$$\sum_{v \in VM} \delta_{v,s}^{(C)} \geq C_s \quad \forall s \in N \quad (24)$$

$$\sum_{v \in VM} \delta_{v,s}^{(C)} \leq LC_s \quad \forall s \in N \quad (25)$$

Constraints (24) and (25) ensure that a cloud is built in core nodes selected to host VMs by setting $C_s = 1$ if $\sum_{v \in VM} \delta_{v,s}^{(C)} > 0$ and $C_s = 0$ otherwise.

Metro fogs location constraints:

$$\sum_{v \in VM} \delta_{v,s}^{(MF)} \geq F_s^{(MF)} \quad \forall s \in N \quad (26)$$

$$\sum_{v \in VM} \delta_{v,s}^{(MF)} \leq LF_s^{(MF)} \quad \forall s \in N \quad (27)$$

Constraints (26) and (27) ensure that metro fogs are built in metro nodes selected to host VMs are by setting $MFog_s = 1$ if $\sum_{v \in VM} \delta_{v,s}^{(MF)} > 0$ and $MFog_s = 0$ otherwise.

Access fog location constraints:

$$\sum_{v \in VM} \delta_{v,p,s}^{(AF)} \geq F_{p,s}^{(AF)} \quad \forall s \in N \quad (28)$$

$$\sum_{v \in VM} \delta_{v,p,s}^{(AF)} \leq LF_{p,s}^{(AF)} \quad \forall s \in N \quad (29)$$

Constraints (28) and (29) ensure that an access fog is built in access nodes selected to host VMs by setting $F_{p,s}^{(AF)} = 1$ if $\sum_{v \in VM} \delta_{v,p,s}^{(AF)} > 0$ and $F_{p,s}^{(AF)} = 0$ otherwise.

Cloud and fog workload constraints:

$$W_{v,s}^{(CR)} = \delta_{v,s}^{(C)} W_v \quad \forall v \in VM, s \in N \quad (30)$$

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

$$\forall v \in VM, s \in N \quad (31)$$

$$W_s^{(C)} = \sum_{v \in VM} W_{v,s}^{(CR)} \quad \forall s \in N \quad (32)$$

$$W_{v,s}^{(MFR)} = \delta_{v,s}^{(MF)} W_v \quad \forall v \in VM, s \in N \quad (33)$$

$$W_{v,s}^{(MFR)} = \left(\frac{D_{v,s}^{(MF)}}{r_v x} M \delta_{v,s}^{(MF)} \right) + \left(W_v^{(R)} D_{v,s}^{(MF)} \right) \quad \forall v \in VM, s \in N \quad (34)$$

$$W_s^{(MF)} = \sum_{v \in VM} W_{v,s}^{(MFR)} \quad \forall s \in N \quad (35)$$

$$W_{v,p,s}^{(AFR)} = \delta_{v,p,s}^{(AF)} W_v \quad \forall v \in VM, s \in N, p \in P \quad (36)$$

$$W_{v,p,s}^{(AFR)} = \left(\frac{D_{v,p,s}^{(AF)}}{r_v x} M \delta_{v,p,s}^{(AF)} \right) + \left(W_v^{(R)} D_{v,p,s}^{(AF)} \right) \quad \forall v \in VM, s \in N, p \in P \quad (37)$$

$$W_{p,s}^{(AF)} = \sum_{v \in VM} W_{v,p,s}^{(AFR)} \quad \forall s \in N \quad (38)$$

Constraints (30), (33), and (36) calculate VM replica workload in a cloud, a metro fog, and an access fog, respectively, under a constant workload profile. Constraints (31), (34), and (37) calculate the VM replica workload in a cloud, a metro fog, and an access fog, respectively, as a linear function of the traffic resulting from serving users of the replicas plus the workload baseline. Constraints (32), (35), and (38) calculate the total workload of a cloud, a metro fog, and an access fog, respectively, by summing the workloads of VMs hosted in them. Constraints (30)–(38) ensure that the VM CPU workload v satisfies user requirements to maintain QoS. For instance, in the constant workload profile, the VM replica is placed in locations that satisfy the full workload W_v . Also, in the linear workload profile, the QoS is maintained by ensuring that the VM replica workload v satisfies user demand $\left(W_v^{(R)} \sum_{d \in N} D_{v,s,d}^{(C)} \right)$ and the workload baseline $\left(\frac{\sum_{d \in N} D_{v,s,d}^{(C)}}{r_v x} M \right)$, which is the minimum baseline required to run the VM.

Number of servers in cloud and fog constraints:

$$S_s^{(C)} \geq \frac{W_s^{(C)}}{S^{(maxW)}} \quad \forall s \in N \quad (39)$$

$$S_s^{(MF)} \geq \frac{W_s^{(MFR)}}{S^{(maxW)}} \quad \forall s \in N \quad (40)$$

$$S_{p,s}^{(AF)} \geq \frac{W_{p,s}^{(AFR)}}{S^{(maxW)}} \quad \forall s \in N, p \in P \quad (41)$$

Constraints (39) - (41) calculate the number of servers in each cloud, metro fog and access fog, respectively based on the CPU utilization as the CPU draws the largest proportion of the server power consumption [55].

Number of router ports and switches in cloud and fog:

$$R_s^{(C)} \geq \frac{\sum_{v \in VM} \sum_{d \in N} D_{v,s,d}^{(C)}}{R^{(CB)}} \quad \forall s \in N \quad (42)$$

$$SW_s^{(C)} \geq \frac{\sum_{v \in VM} \sum_{d \in N} D_{v,s,d}^{(C)}}{SW^{(CB)}} \quad \forall s \in N \quad (43)$$

$$R_s^{(MF)} \geq \frac{\sum_{v \in VM} D_{v,s}^{(MF)}}{R^{(MFB)}} \quad \forall s \in N \quad (44)$$

$$SW_s^{(MF)} \geq \frac{\sum_{v \in VM} D_{v,s}^{(MF)}}{SW^{(MFB)}} \quad \forall s \in N \quad (45)$$

$$R_{p,s}^{(AF)} \geq \frac{\sum_{v \in VM} D_{v,p,s}^{(AF)}}{R^{(AFB)}} \quad \forall s \in N, p \in P \quad (46)$$

$$SW_{p,s}^{(AF)} \geq \frac{\sum_{v \in VM} D_{v,p,s}^{(AF)}}{SW^{(AFB)}} \quad \forall s \in N, p \in P \quad (47)$$

Constraints (42) - (47) calculate the number of routers ports and switches in each cloud, metro fog and access fog, respectively.

Number of metro router ports and ethernet switches in metro network constraints:

$$R_s^{(M)} \geq \frac{\sum_{v \in VM} \sum_{s \in N} D_{v,s,d}^{(C)} + \sum_{v \in VM} D_{v,s}^{(MF)}}{R^{(MB)}} \quad \forall s \in N \quad (48)$$

$$SW_s^{(M)} \geq \frac{\sum_{v \in VM} \sum_{s \in N} D_{v,s,d}^{(C)} + \sum_{v \in VM} D_{v,s}^{(MF)}}{SW^{(MB)}} \quad \forall s \in N \quad (49)$$

Constraints (48) and (49) calculate the number of routers ports and switches, respectively, in each metro network.

Traffic demand on IP over WDM core network constraint:

$$L_{s,d} = \sum_{v \in VM} D_{v,s,d}^{(C)} \quad \forall s, d \in N \quad (50)$$

Constraint (50) 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:

$$\sum_{j \in N: i \neq j} L_{i,j}^{s,d} - \sum_{j \in N: i \neq j} L_{i,j}^{s,d} = \begin{cases} L_{s,d} & i = s \\ -L_{s,d} & i = d \\ 0 & \text{otherwise} \end{cases} \quad \forall s, d, i \in N : s \neq d \quad (51)$$

Constraint (51) represents the flow conservation for IP layer on the IP over WDM network. It ensures that the total incoming traffic equal the total outgoing traffic in all nodes; excluding the source and destination nodes.

Virtual link capacity constraint:

$$\sum_{s \in N} \sum_{d \in N: s \neq d} L_{i,j}^{s,d} \leq C_{i,j} \mathcal{W}^{(B)} \quad \forall i, j \in N : s \neq d \quad (52)$$

Constraint (52) ensures that the traffic transmitted through a virtual link does not exceed its maximum capacity.

Flow conservation constraint in the optical layer:

$$\sum_{n \in Nm_m} \dot{W}_{m,n}^{i,j} - \sum_{n \in Nm_m} \mathcal{W}_{m,n}^{i,j} = \begin{cases} C_{i,j} & m = i \\ -C_{i,j} & m = j \\ 0 & \text{otherwise} \end{cases} \quad \forall i, j, m \in N : i \neq j \quad (53)$$

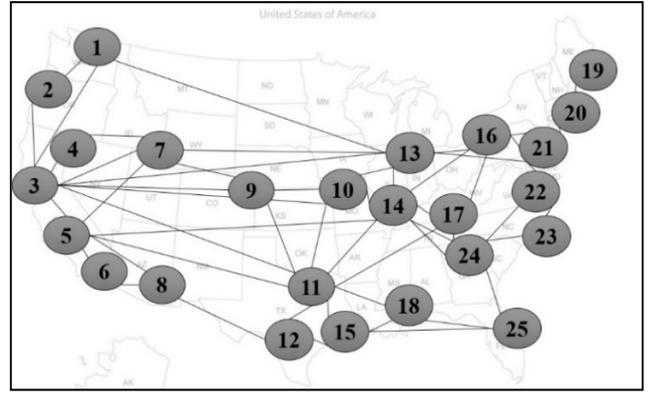


FIGURE 3. AT&T core network topology.

Constraint (53) 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:

$$\sum_{i \in N} \sum_{j \in N: i \neq j} \dot{W}_{m,n}^{i,j} \leq \dot{W} F_{m,n} \quad \forall m, n \in N \quad (54)$$

$$\dot{W}_{mn} = \sum_{i \in N} \sum_{j \in N: i \neq j} \dot{W}_{m,n}^{i,j} \quad \forall m, n \in N \quad (55)$$

Constraints (54) and (55) represent the physical link capacity limit. Constraint (54) ensures that the number of wavelengths in virtual links traversing a physical link does not exceed the maximum capacity of fibers in the physical link. Constraint (55) calculates the number of wavelengths in a physical link as the sum of wavelength channels in virtual links traversing the physical link.

Total number of router ports in a core node:

$$R_s^{(AC)} = \frac{1}{\mathcal{W}^{(B)}} \sum_{d \in N} L_{s,d} \quad \forall s \in N \quad (56)$$

$$R_s^{(AE)} = R^{(MR)} \left(\frac{1}{\mathcal{W}^{(B)}} \sum_{s \in N} L_{s,d} \right) \quad \forall d \in N \quad (57)$$

Constraint (56) calculates the total number of router ports in each core node that aggregate the traffic from/to the clouds. Constraint (57) calculates the total number of router ports in each core node that aggregate the traffic from/to edge routers.

III. CLOUD-FOG ARCHITECTURE MILP MODEL RESULTS

This section investigates optimal VM placement over a distributed AT&T cloud architecture. Fig. 3 illustrates the core AT&T network topology [56], which consists of 25 nodes and 54 bidirectional links [56]. We consider an architecture in which each core node is connected to two PON networks through a metro network consisting of a single ethernet switch and two metro routers (Fig. 3.1). The PON access network connects 512 locations. The total capacity of each OLT is 1,280 Gbps [57].

We start by considering the optimization of a single VM placement as the simplest representative problem. We then consider optimization in a realistic scenario with multiple VMs.

A. SIMPLE REPRESENTATIVE SCENARIO

We investigate how the energy-efficient placement of a single VM over cloud-fog architecture varies based on three factors: CPU requirements, download traffic, and power usage effectiveness (PUE) values. The impact of the VM workload profile on 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 a VM with a constant workload profile and the workload of a VM with a linear workload profile that serves the maximum number of users are both considered. Three workloads—10%, 50%, and 100% of the server CPU capacity—are considered. The users are assumed to access VMs 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 assumed to have 800 users. The PUE is a metric used to determine the total energy consumption of the facility hosting the clouds, fog nodes, or network nodes, which includes the power consumption of computing and communication hardware, IT cooling, lighting, etc. PUE is the ratio of this total power consumption to the IT (computing and communication) infrastructure power consumption. Based on United States (US) data center energy usage [58], PUE varies based on data center size, as larger data centers tend to use more efficient cooling technologies. For best-practice data centers, the PUE of clouds, metro fogs, and access fogs take values of 1.3, 1.4, and 1.5, respectively [58]. For data centers from 2014, the PUE values are 1.7, 1.9, and 2.5, respectively [58]. In network infrastructures, a typical telecom office PUE value is 1.5 [35].

The Cisco Carrier Routing System 1 (CRS-1) [59] 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. Also, the Cisco NCS 5502 router [60] is considered as the cloud and metro networks router which consumes 30W per 40 Gbps port. In the metro and fog datacentre, Cisco NCS 5501 [60] is considered with a power consumption of 13W per 40 Gbps port. Furthermore, the Cisco Nexus 93180YC-EX [61] 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 [61] switch is considered with capacity of 240 Gbps while consuming 210W. Tables 1-3 show the IP over WDM, metro and access network parameters and Table 4 shows the Clouds and fogs parameters. 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-core Intel 2.6 GHz Sandy Bridge E5-2670 processors [62].

TABLE 1. IP Over WDM core network input parameters of the model.

Router port power consumption ($R^{(P)}$)	638 Watt [59]
Transponder power consumption ($t^{(P)}$)	129 Watt [63]
Regenerator power consumption ($G^{(P)}$)	114 Watt, reach 2000 km [64]
EDFA power consumption ($e^{(P)}$)	11 Watt [65]
Optical switch power consumption ($SW^{(P)}$)	85 Watt [66]
Number of wavelengths in a fiber (\mathcal{W})	32 [67]
Bit rate of each wavelength ($\mathcal{W}^{(P)}$)	40 Gbps [67]
Span distance between two EDFAs (S)	80 km [65]
Network power usage effectiveness (n)	1.5 [33]1.5 [33]

TABLE 2. Metro network input parameters of the model.

Metro router redundancy ($R^{(MR)}$)	2
Metro edge router port bit rate ($R^{(MB)}$)	40 Gbps
Metro edge router port power consumption ($R^{(MP)}$)	30 Watt [60]
Metro ethernet switch bit rate ($SW^{(MB)}$)	600 Gbps [61]
Metro ethernet switch power consumption ($W^{(MP)}$)	470 Watt [61]

TABLE 3. Access network input parameters of the model.

Number of PON networks in a node (P)	2
Number of users of VM services in each PON based on VMs popularity groups ($U_{v,p,d}$)	13,000 users in each PON, six VMs popularity groups; 16%, 5%, 2%, 1%, 0.5% and 0.05%.
Maximum number of users of a single VM (x)	800 concurrent users
Number of ONU devices in a PON network ($ONU_{p,d}^{(N)}$)	512
Power consumption of ONU device ($ONU^{(P)}$)	5 Watt [68]
Number of OLTs in a PON network ($OLT_{p,d}^{(N)}$)	1
OLT Capacity ($OLT_{p,d}^{(B)}$)	1280 Gbps [57]
OLT Power consumption ($OLT^{(P)}$)	1842 W [57]

Fig. 4 (a), (b), and (c) show the optimal placement of VMs with 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 rate, which ranges from 0.1 Mbps to 200 Mbps, and the z-axis

TABLE 4. Cloud and fog input parameters of the model.

Number of VMs (V)	1
User download rate (r_v)	{0.1, 1, 10, 20, 50, 100 or 200 Mbps}
Maximum workload of VM (W_v)	10%, 50% and 100%
Server power consumption ($S^{(P)}$)	333 Watt [69]
Maximum server workload ($S^{(maxW)}$)	100%
Cloud and metro fog switch bit rate ($SW^{(CB)}, SW^{(MFB)}$)	600 Gbps [61]
Cloud and metro fog switch power consumption ($SW^{(CP)}, SW^{(MFP)}$)	470 Watt [61]
Access fog switch bit rate ($SW^{(AFB)}$)	240 Gbps [61]
Access fog switch power consumption ($SW^{(MFP)}$)	210 Watt [61]
Cloud and fog switch redundancy ($SW^{(R)}$)	2
Cloud and fog router port bit rate ($R^{(CB)}, R^{(MFB)}, R^{(AFB)}$)	40 Gbps
Cloud router port power consumption ($R^{(CP)}$)	30 Watt [60]
Metro and access fog router port power consumption ($R^{(MFP)}, R^{(AFP)}$)	13 Watt [60]
Cloud power usage effectiveness (c)	1.3 or 1.7 [58]
Metro fog power usage effectiveness (m)	1.4 or 1.9 [58]
Access fog power usage effectiveness (a)	1.5 or 2.5 [58]

is the percentage of VM replicas in each location over the cloud-fog architecture.

The placement of VMs with linear workload profiles is not affected by VM workload, as serving users consumes the same power whether centralized in a single VM or distributed among multiple replicas with smaller workloads. However, the higher PUE of fog nodes compared to the cloud entails that distributing replicas into fog processing nodes incurs additional power consumption, as the PUE values of fog nodes are higher than that of clouds. Hence, a trade-off exists between the network power saved by replicating VMs into fog nodes and the additional power consumed by these replicas. The creation of a VM replica results in power savings if the former power exceeds the latter. At data rates of 1 Mbps and higher, VMs with 10%, 50%, and 100% workloads are offloaded to access fog processing nodes considering linear workload profiles.

For constant workload profiles, replicas are less energy-efficient, so offloading VMs to fog nodes decreases as VM workload increases. While VMs with 10% workload and 20 Mbps are fully offloaded to metro fogs, 50% and 100% workload VMs are replicated only to clouds. Also, users of VMs with 50% workload at 100 Mbps as well as VMs of 100% workload at 200 Mbps data rate are served by clouds and metro fog nodes. A single VM replica is offloaded to

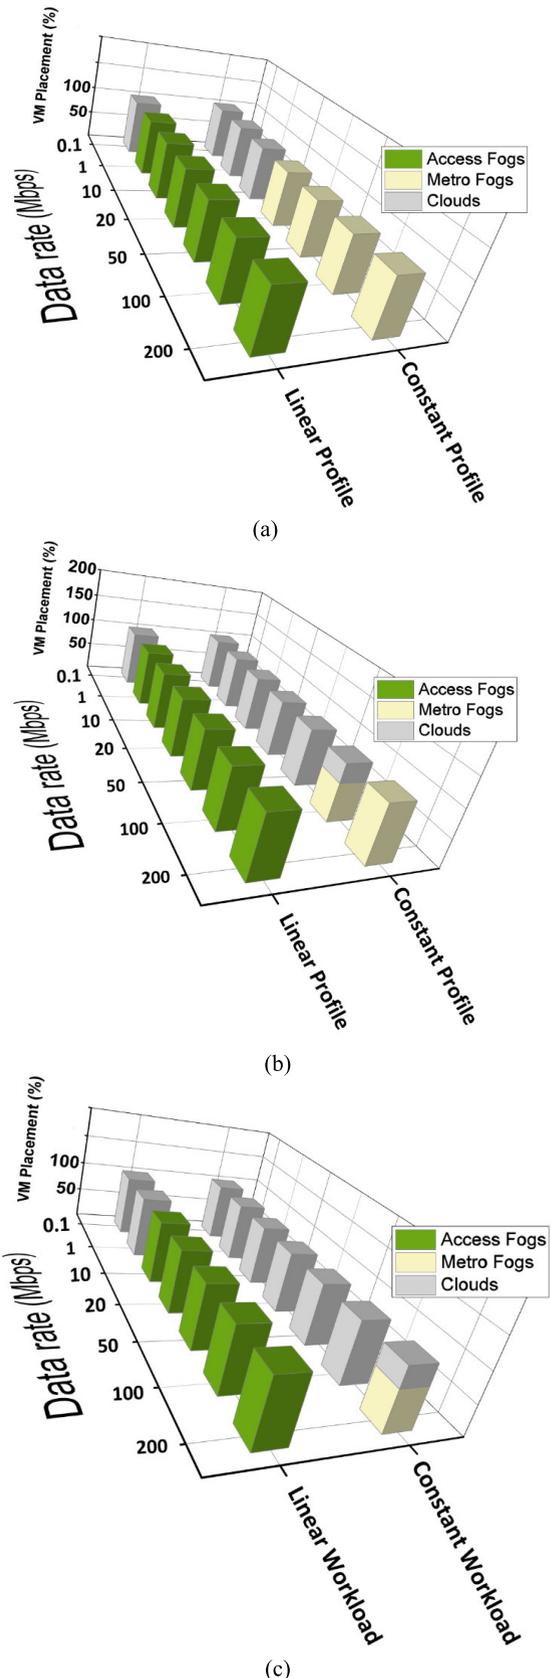


FIGURE 4. Optimal VM placement of (a) constant profile at 10% of CPU and linear profile with peak utilization at 10%, (b) 50% case, (c) 100% case at different data rates considering best practice PUE value ($c = 1.3$, $m = 1.4$, $a = 1.5$).

14 metro fog nodes (in core nodes 1, 2, 4, 6, 7, 8, 13, 16, 19, 20, 21, 22, 23, and 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 VM replicas because the traffic flows traverse more than a single hop in the Internet protocol (IP) over a wavelength-division multiplexing (WDM) network to access the VM placed in the cloud in node 11, which increases the need for IP router ports (the most power-consuming devices in the IP over WDM network).

The results also show that VMs with higher data rates justify the creation of more replicas closer to user premises in the fog layer. Thus, the power consumption of the network, which is the greatest contributor to the power consumption of the cloud-fog architecture, is reduced. For example, VMs with 10% workload under the linear workload profile are fully replicated to clouds and offloaded to access fog nodes for VMs of 0.1 Mbps and ≥ 1 Mbps user data rates, respectively.

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

B. REALISTIC SCENARIO

In this scenario, a realistic number of users and VM popularity is studied. According to the Cisco Visual Network Index (VNI) [70], in 2016, the average US broadband data rate was 36 Mbps. Therefore, each OLT is assumed to be able to serve $\sim 35,000$ connections (or users). The Cisco VNI also reported that 76% of all Internet traffic crossed clouds in 2016. SimilarWeb [71], an online tool that provides Internet traffic statistics and analytics, shows that the top 300 applications or websites have a 50% share of all traffic. Accordingly, 13,000 users are assumed in each PON ($\sim 50\%$ of cloud traffic, i.e., 38% of total traffic) to access the VMs (placed either in the cloud or fog) hosting the top 300 applications or websites. The popularity of these VMs is assumed to follow a Zipf distribution [72]. To simplify the analysis, VM popularity is divided into six groups: 16%, 5%, 2%, 1%, 0.5%, and 0.05% of 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 to serve 800 users. Based on the literature [44]–[46], [73], [74], in such a case, a VM can serve 800 users at a low error rate. VMs with linear workloads are assumed to have workload baselines of 1%, 5%, or 40% of total server CPU capacity based on the CPU requirements of state-of-the-art applications [44], [46], [73] (i.e., a 1% workload baseline for database applications, 5% for website applications, and 40% for video games and web conference applications). The users are assumed to access the VMs at one of the following

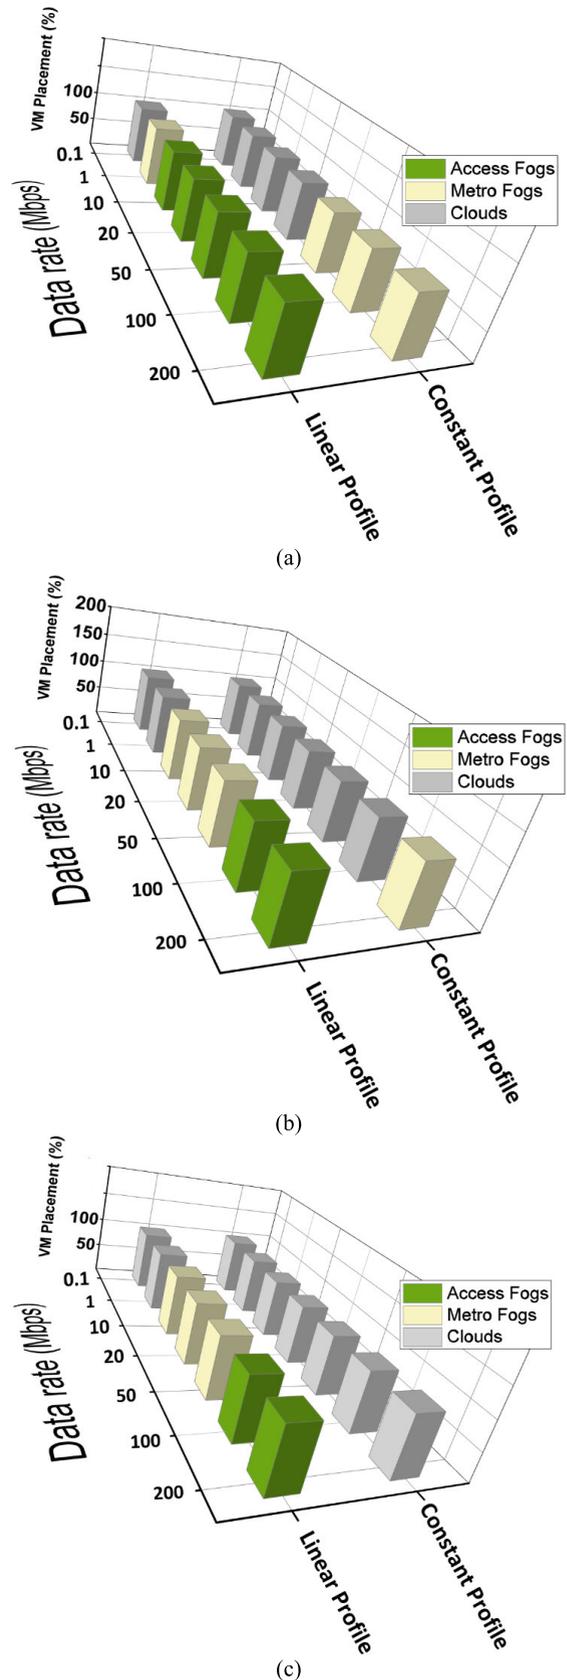


FIGURE 5. Optimal VM placement of (a) constant profile at 10% of CPU and linear profile with peak utilization at 10%, (b) 50% case, (c) 100% case at different data rates considering 2014 PUE value ($c = 1.7$, $m = 1.9$, $\alpha = 2.5$).

TABLE 5. Input parameters used in the model.

Number of VMs (V)	300
User download rate (r_u)	{1 Mbps, 10 Mbps, 25 Mbps}
Maximum workload of VM (W_v)	50%
Cloud power usage effectiveness (c)	1.3 [58]
Metro fog power usage effectiveness (m)	1.4 [58]
Access fog power usage effectiveness (a)	1.5 [58]

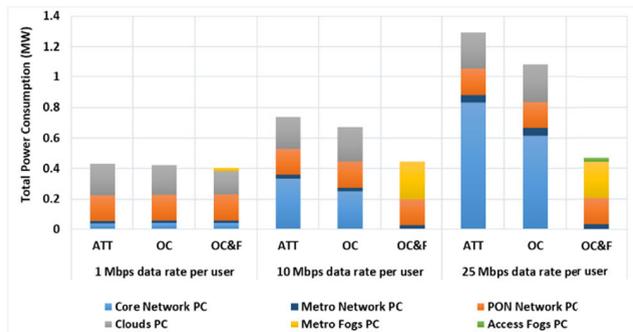


FIGURE 6. The power consumption of different VMs placement approaches considering VMs of 1% minimum CPU workload.

data rates: 1 Mbps (low), 10 Mbps (medium), or 25 Mbps (high), which represent the recommended download speeds to access the content of state-of-the-art applications (i.e., 1 Mbps for light web browsing [75] (emails, Google Docs [76], and websites with lower definition video content [77]), 10 Mbps for applications processing high-definition video quality [78] and online multiplayer games [79], and 25 Mbps for applications processing ultra-high video quality [80]).

Optimized VM placement over a cloud-fog architecture, referred to as the optimized cloud and fog placements (OC&F) approach, is compared to the optimized cloud (OC) approach, in which VMs are optimally placed in clouds distributed over the core network, and the AT&T cloud (ATT), in which the VMs are placed in nodes 1, 3, 5, 6, 8, 11, 13, 17, 19, 20, 22, and 25 according to AT&T data center map [56].

In addition to the parameters in Table 1 to Table 4, Table 5 shows the additional/modified parameters considered for the following results.

1) LINEAR WORKLOAD PROFILE (1% WORKLOAD BASELINE)

Fig. 6 shows the power consumption resulting from placing VMs with 1% minimum CPU workload considering the various placement approaches at 1, 10, and 25 Mbps user data rates. The efficiency of VMs with 1% minimum CPU workload allows the creation of more efficient VM replicas, as the workload is proportional to the number of users served by the VM with a trivial minimum workload required by each VM. At a 1 Mbps data rate, the OC&F approach achieves a 6% reduction of total power consumption compared to

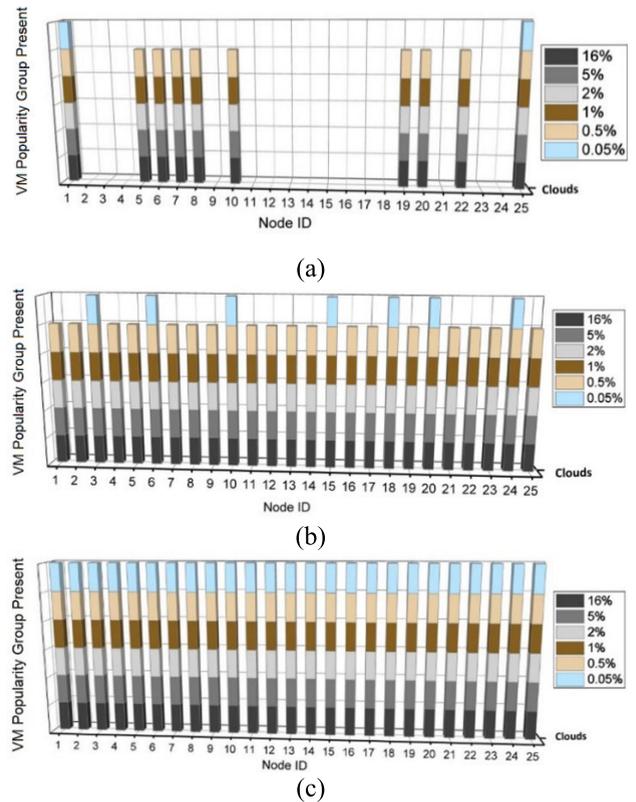


FIGURE 7. Optimal placement of different VMs popularity groups of 1% workload baseline under the OC approach with (a) 1 Mbps data rate per user, (b) 10 Mbps data rate per user and (c) 25 Mbps data rate per user.

the ATT approach. The total reductions amount to 40% at a 10 Mbps data rate and 64% at a 25 Mbps data rate. Compared to the OC approach, the savings achieved by the OC&F approach are 4%, 31%, and 48% at the low, medium and high data rates, respectively; compared to the ATT approach, the savings achieved by the OC approach are 2%, 9%, and 16% at the low, medium and high data rates, respectively.

In Fig. 7 and Fig. 8, we further investigate the OC and OC&F placement approaches by examining the placement of VMs at different data rates and in different popularity groups. Fig. 7 shows the optimal VM placement with the OC approach. Note that the different colors indicate the creation of VM replicas in the cloud, not the number of replicas. VM efficiency has allowed the creation of multiple replicas, as the workload is proportional to the number of users served by a VM with a limited workload baseline. Efficient VM workload profiles justify the replication of VMs with popularity greater than 0.5% into 10 clouds at a 1 Mbps data rate and into 25 clouds (full replication) at a 10 Mbps data rate. VMs of 0.05% popularity are only replicated into two clouds. The high traffic of VMs at a 25 Mbps data rate allows full replication for the different popularity groups across all clouds.

Fig. 8 (a) shows that VMs at a low user data rate of 1 Mbps only justify creating three metro fogs in nodes 6, 8, and 19, as the traffic flows from these nodes traverse more than a single

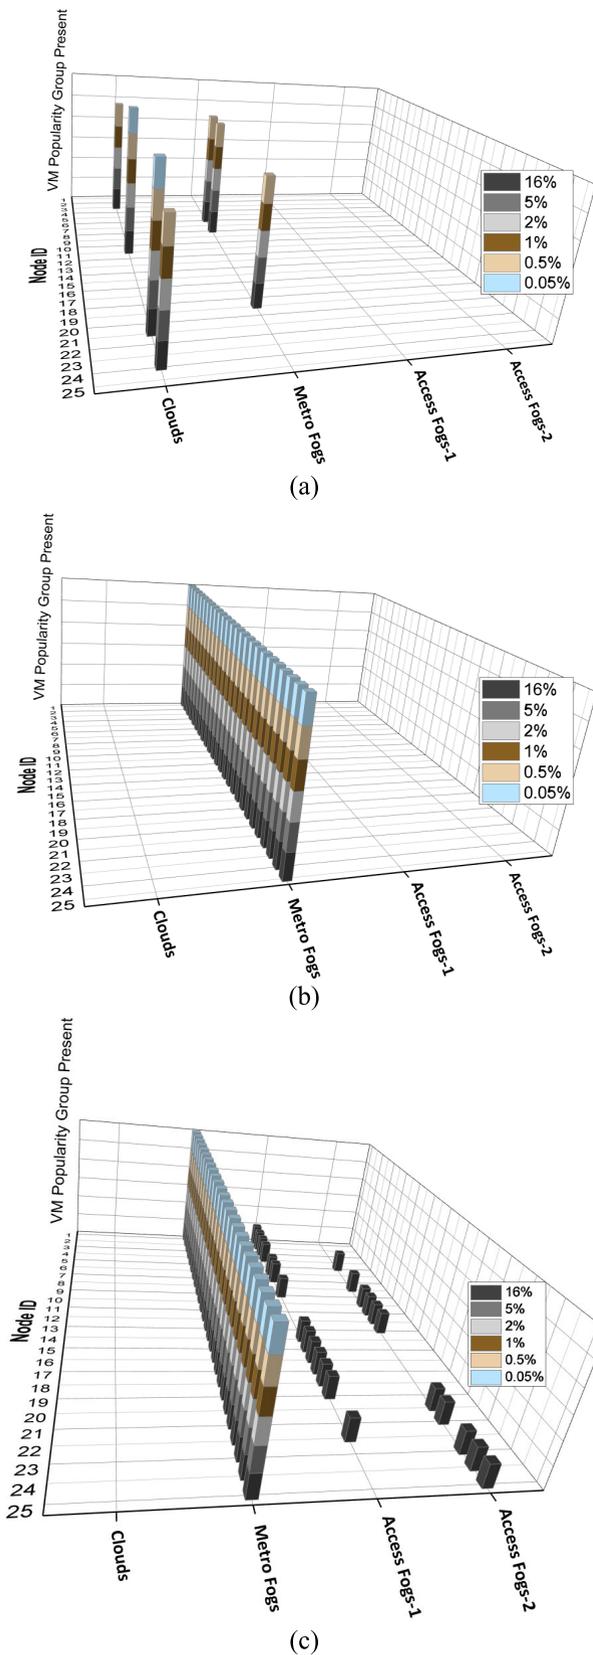


FIGURE 8. Optimal placement of different VMs popularity groups of 1% workload baseline under the OC&F approach with (a) 1 Mbps data rate per user, (b) 10 Mbps data rate per user and (c) 25 Mbps data rate per user.

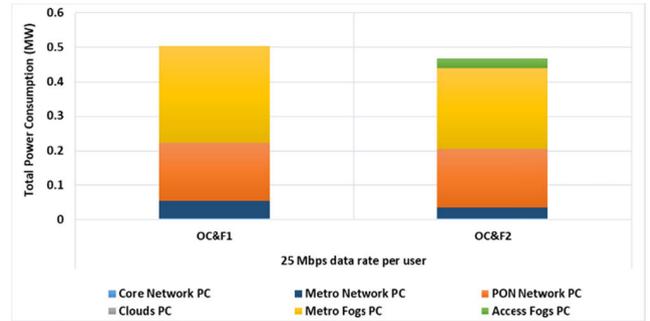


FIGURE 9. The power consumption considering OC&F1 and OC&F2 placement approaches. OC&F1 represents the optimal placement considering clouds and metro fogs only and OC&F2 represents the optimal placement considering clouds, metro and access fogs.

hop in the IP over the WDM network to access the replicas optimally placed in the distributed clouds in nodes 3, 11, 20, and 24. These fog nodes are thus built to serve user demand locally and consequently eliminate the need for IP router ports. However, VMs with the lowest popularity (0.05%) only justify the creation of two replicas in nodes 11 and 20.

VMs at a 10 Mbps data rate are fully offloaded to every metro fog, as shown in Fig. 8 (b). VM users are uniformly distributed across the metro and access networks, so VM placement is consistent across all the metro fog nodes. In Fig. 8 (c), VMs at a high data rate of 25 Mbps show full replication in metro fog nodes, and VMs with 16% popularity justify creating VM replicas in some access fog nodes. Although we can reduce the traffic traversing the metro network and consequently reduce the total power consumption, VMs with 16% popularity are not fully replicated to access fog nodes. A number of replicas are offloaded to metro fog nodes because of the on-off power consumption profiles of fog and network resources. Thus, before creating a new fog node in the access network, VMs are consolidated into the available resources that remain from the placement of other VMs that share the same architecture.

Fig. 9 introduces the OC&F1 and OC&F2 placement approaches. The former represents the optimal placement considering clouds and metro fog nodes only, and the latter represents the optimal placement considering all three computing layers: cloud, metro fog, and access fog. These two approaches show how introducing fog nodes in the access network (OC&F2) in addition to the metro fog can reduce total power consumption compared to an approach that considers only fog nodes connected to a metro network (OC&F1). At a 25 Mbps user data rate, the OC&F2 approach saves 6% more power than the OC&F1 approach.

Fig. 10 and Fig. 11 show the number of servers required to host VM replicas under the OC and OC&F approaches, respectively. The number of servers is a function of the number of VM replicas hosted and their workloads. For instance, the OC&F approach at a 25 Mbps user data rate (Fig. 9 (c)) requires 18 servers in each metro fog and two servers in

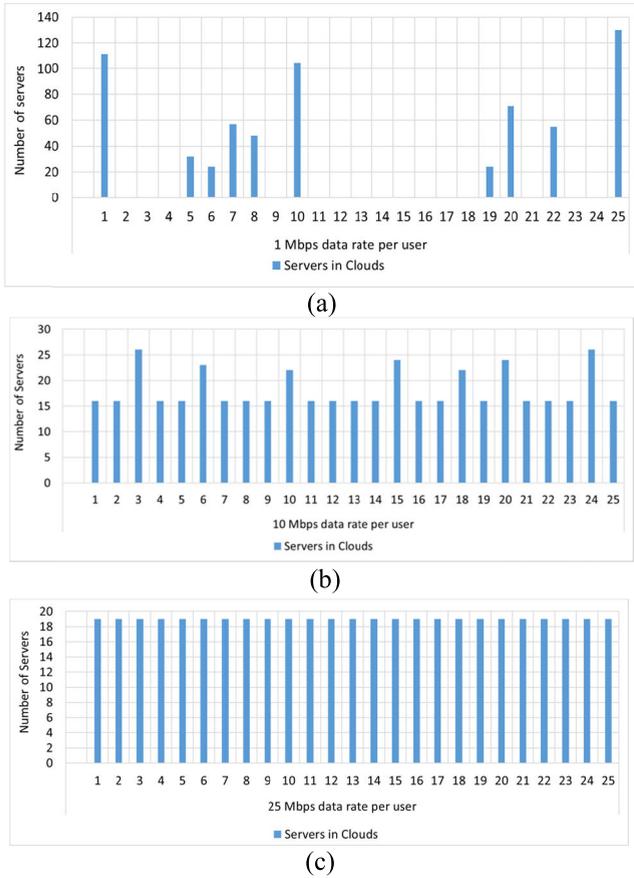


FIGURE 10. Number of servers in OC approach required to host VMs of 1% workload baseline with (a) 1 Mbps data rate per user (b) 10 Mbps data rate per user (c) 25 Mbps data rate per user.

access fogs to host VM replicas. Such a number of servers can be practically attached to the metro edge routers to create the metro fog layer and to the OLT in the access network to create the access fog layer.

2) LINEAR WORKLOAD PROFILE (5% WORKLOAD BASELINE)

Fig. 12 shows the power savings achieved by VMs with linear workload profiles and 5% minimum CPU utilization. Increasing the minimum CPU utilization of the VM workload profile to the current value of 5% reduces the efficiency of creating more VM replicas. The total savings achieved by the OC&F approach compared to the ATT approach are 12%, 35%, and 55% at the low, medium, and high data rates, respectively. Compared to the OC approach, no extra power saving is achieved at the low data rate, as the total traffic does not justify replicating any VMs into fogs. At the medium and high user data rates, the power savings are 28% and 47%, respectively.

Fig. 12 (a) and (b) illustrate the placement of VMs with 5% minimum CPU utilization considering the OC&F placement approach at low and high user data rates, respectively. VMs

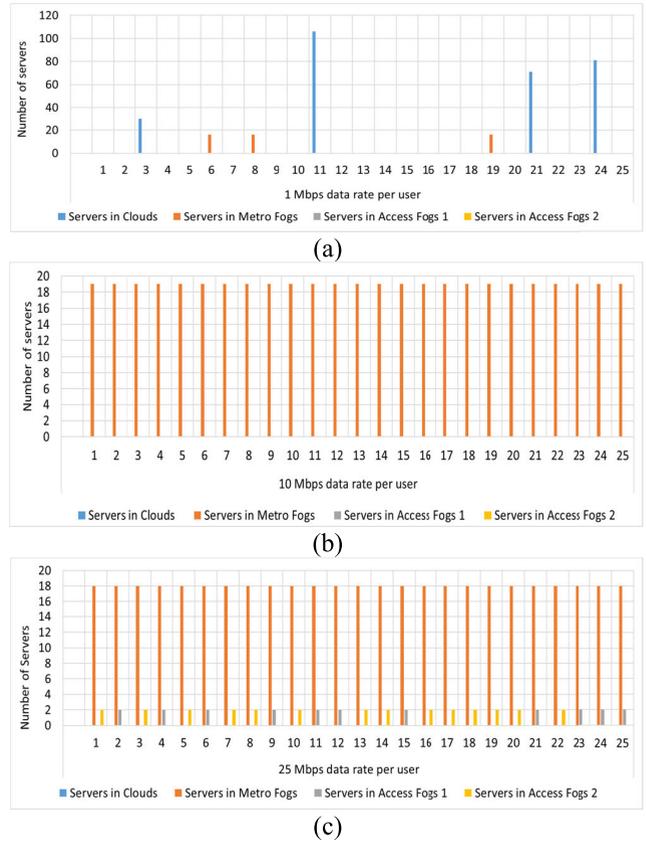


FIGURE 11. Number of servers in OC&F approach required to host VMs of 1% workload baseline with (a) 1 Mbps data rate per user (b) 10 Mbps data rate per user (c) 25 Mbps data rate per user.

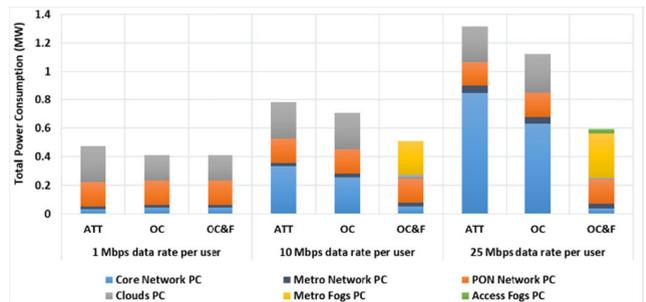
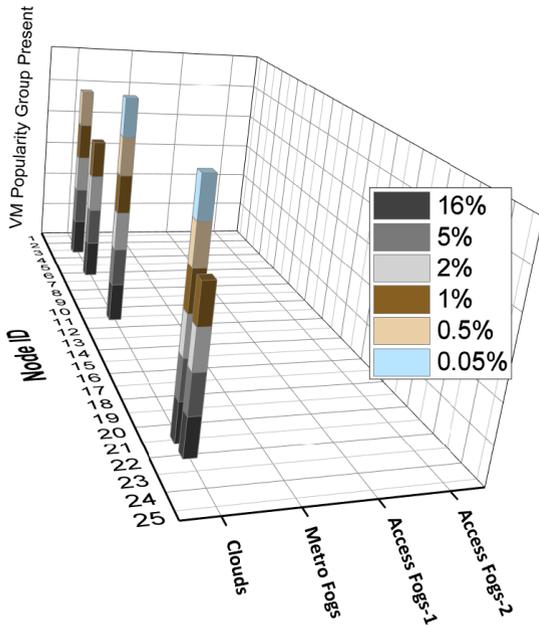
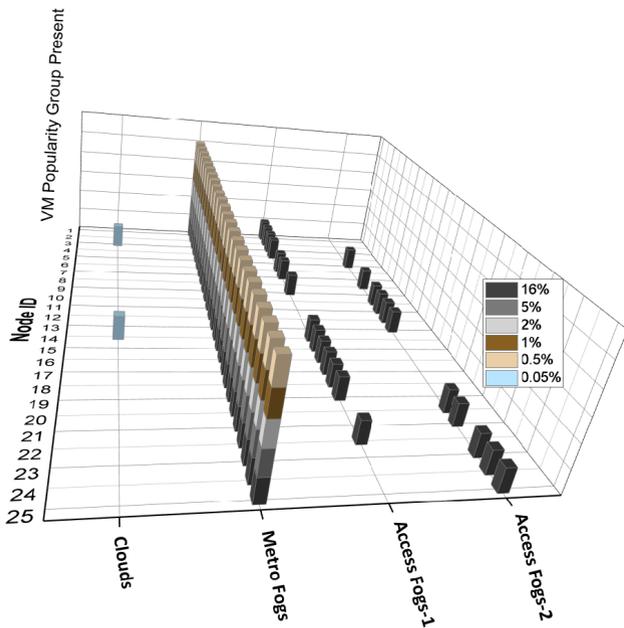


FIGURE 12. The power consumption of different VMs placement approaches considering VMs of 5% workload baseline.

with low user data rates are dispersed among distributed clouds. The low user data rates do not justify offloading VMs to any fog nodes. VMs of $\geq 1\%$ popularity justify the creation of five cloud locations. VMs with 0.5% and 0.05% popularity only justify the creation of three and two replicas, respectively. At the high user data rates, VMs with $\geq 0.5\%$ and $\leq 5\%$ popularity are fully offloaded to the metro fogs. In addition, VMs with 16% popularity justify the creation of replicas in some access fogs, whereas VMs with 0.05% popularity only justify the creation of two replicas in nodes 3 and 14.



(a)



(b)

FIGURE 13. The optimal placement of different VMs popularity groups of 5% minimum CPU workload under the OC&F approach with (a) 1 Mbps and (b) 25 Mbps data rate per user.

3) LINEAR WORKLOAD PROFILE (40% WORKLOAD BASELINE)

Fig. 14 shows the power savings achieved by VMs with 40% minimum CPU utilization. The total savings achieved by the OC&F approach compared to the ATT approach are 53%, 44%, and 48% at the low, medium, and high user data rates, respectively. Compared to the OC, no extra power saving is achieved at the low user data rates, as the total traffic does not justify replication of any VM into any fog node. At the

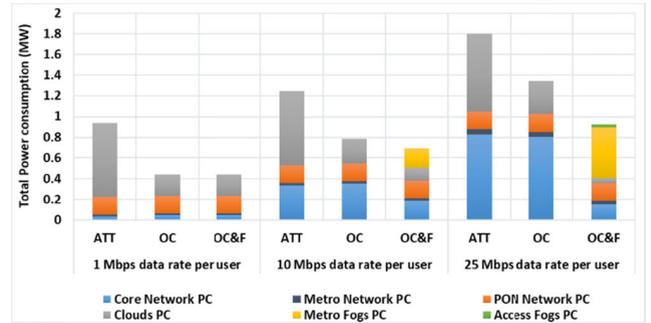
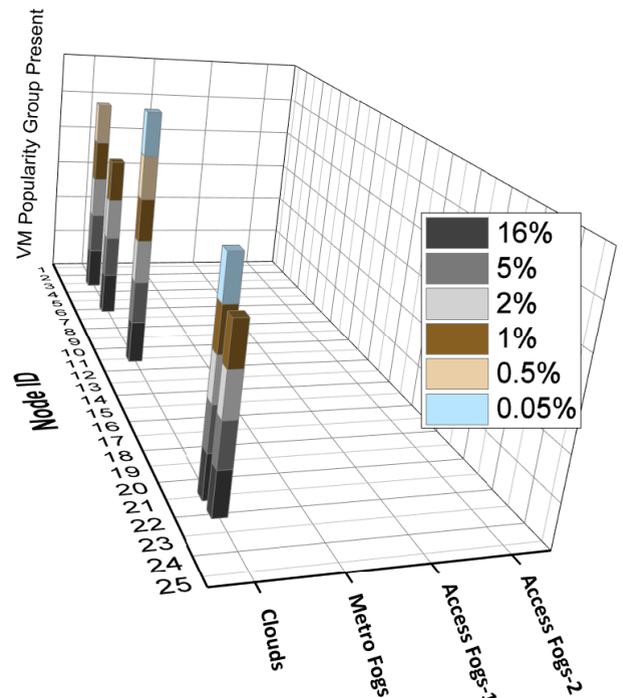


FIGURE 14. The power consumption of different VMs placement approaches considering VMs of 40% minimum CPU workload.

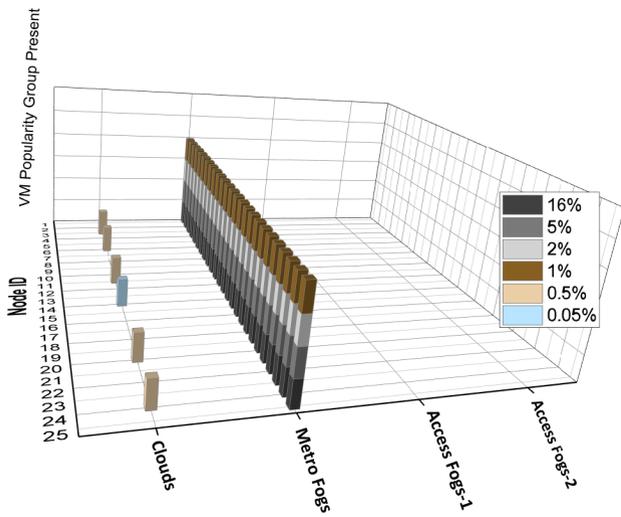
medium and high user data rates, the power savings achieved are 12% and 31%, respectively.

Fig. 15 (a), (b), and (c) illustrate optimal VM placements at low, medium, and high user data rates, respectively, with the OC&F approach. Increasing the minimum CPU utilization of VM workload to 40% reduces the efficiency of creating more replicas of VMs with low popularity across distributed cloud and fog nodes compared to VMs with 1% or 5% minimum CPU utilization. VMs with a data rate of 1 Mbps are replicated among distributed clouds. The low user data rate does not justify offloading VMs to any fog nodes. VMs with $\geq 1\%$ popularity justify the creation of five cloud locations. However, VMs with 0.5% and 0.05% popularity only

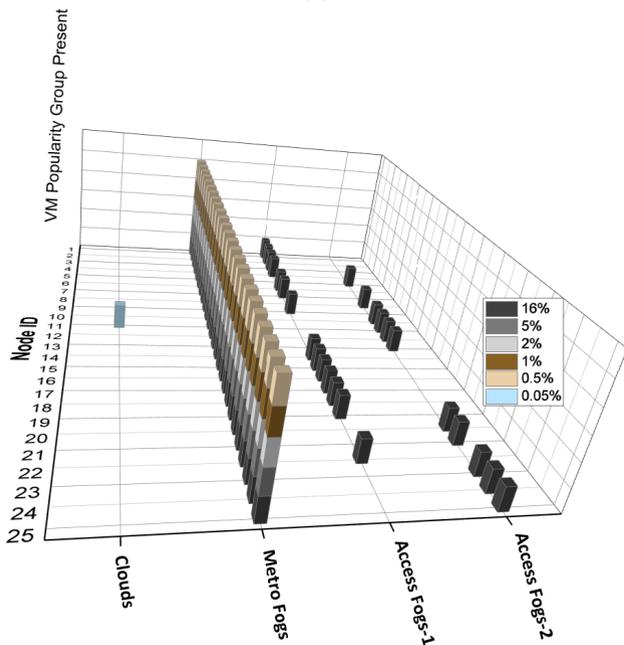


(a)

FIGURE 15. 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)



(c)

FIGURE 15. (Continued.) 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.

justify the creation of three and one replicas, respectively. At medium user data rates, VMs with $\geq 1\%$ popularity are offloaded to metro fogs, whereas other popularity groups are optimally placed in clouds. At high user data rates, despite the high workload baseline, VMs with high popularity of 16% justify the creation of VM replicas in some access fog nodes. VMs with $\geq 0.5\%$ and $\leq 5\%$ popularity are fully offloaded to metro fogs, whereas VMs with 0.05% popularity do not justify the creation of multiple replicas. Only a single replica is optimally placed in node 11 to serve its distributed users.

Fig. 16 shows the number of servers required to host VM replicas with the OC&F approach at a 25 Mbps data rate

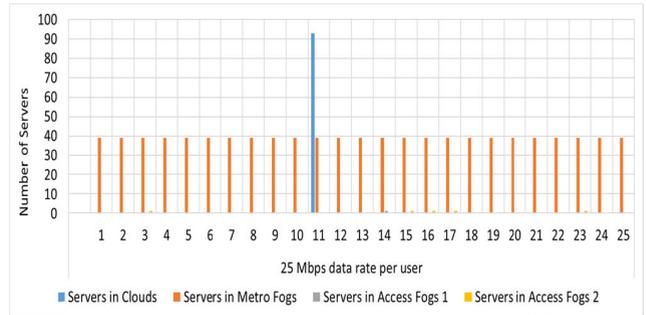
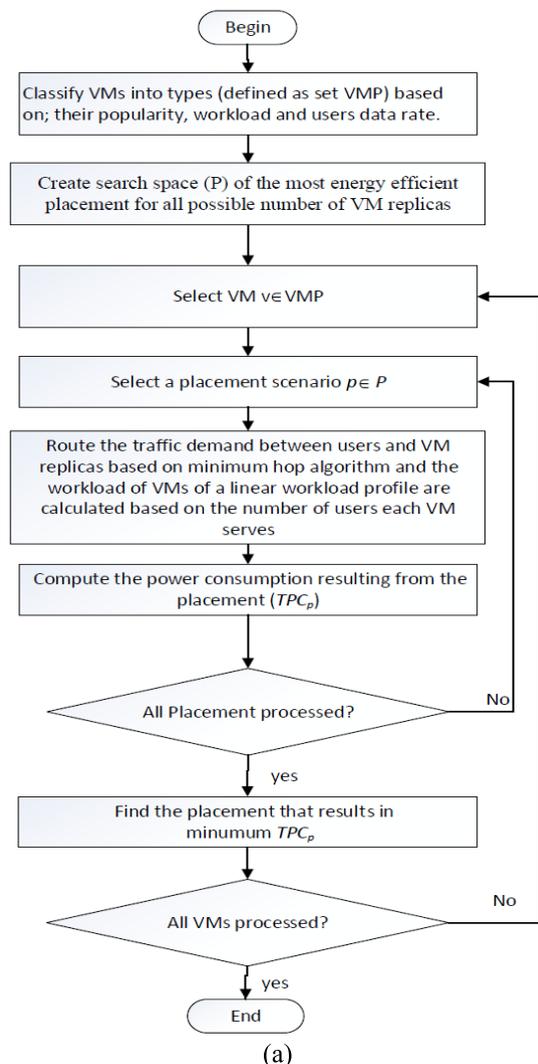


FIGURE 16. Number of servers required to host VMs of 40% minimum CPU workload under the OC&F approach with 25 Mbps data rate per user.

per user. The number of servers is a function of the number of VM replicas hosted and their workloads.

IV. ENERGY EFFICIENT VIRTUAL MACHINES PLACEMENT HEURISTIC FOR CLOUD AND FOG ARCHITECTURE

The VM placement problem over cloud-fog architecture is a nondeterministic polynomial (NP)-hard problem. For



(a)

FIGURE 17. Flowchart of (a) the offline phase and (b) the online phase of EEVM-CF heuristic.

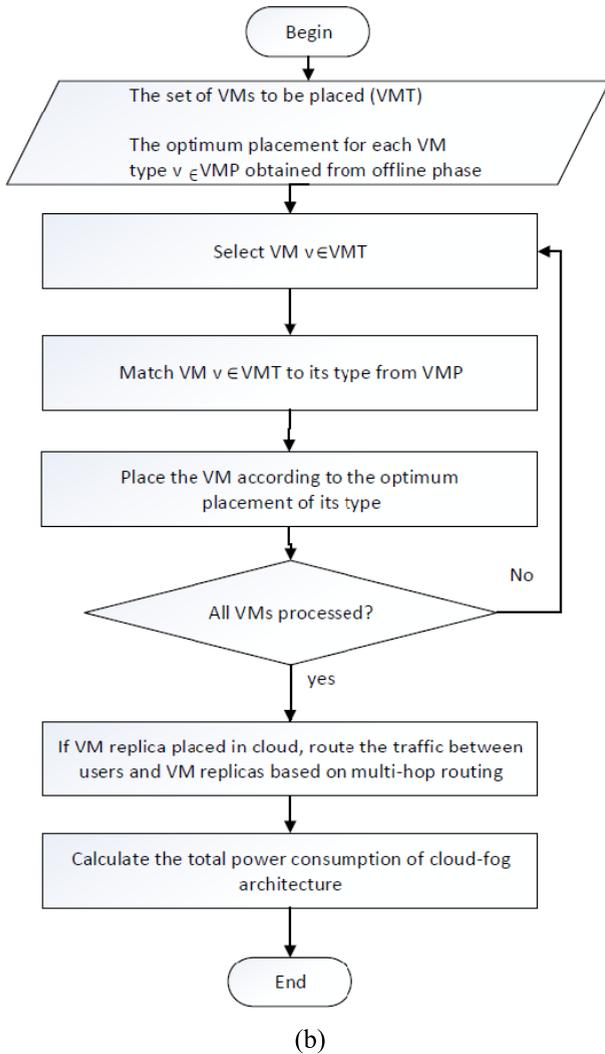


FIGURE 17. (Continued.) Flowchart of (a) the offline phase and (b) the online phase of EEVM-CF heuristic.

example, if v is the number of VMs and s is the number of servers, then the number of possible VM placements in different servers is vs . In the case of replicating VMs into multiple data centers (N), an exhaustive search of distributed data center locations requires the evaluation of placement combinations to find the optimal number and locations of VM copies.

In general, MILP solutions of allocation problems is known to be NP-hard (non-deterministic polynomial-time) i.e. no algorithm can be identified that can find a solution in polynomial time [29]. For our MILP model, the placement of VMs in a cloud-fog architecture over the AT&T network with $N = 100$ placement locations, (25 cloud locations, 25 metro fog locations, and 50 access fog locations), requires examining a number of solutions given as $\left(\sum_{i=1}^N \frac{N!}{(N-i)!}\right)$. Therefore, it is not practical to apply the MILP model in a real time large implementation. Heuristics can provide simple and fast operation in real time with performance that may approach that of the optimal MILP

solution. The optimal solutions obtained from the MILP model (potentially in a small network) can thus offer a benchmark for determining the performance of the heuristics developed.

Therefore, it is not practical to apply the MILP model in a large real-time implementation. Heuristics can provide fast, simple operations in real time that may approach that of the optimal MILP solution. The optimal solutions obtained from the MILP model can thus offer a benchmark for determining the performance of the developed heuristics. A supervised learning algorithm—a branch of machine learning in which an input is matched to an output based on a sample of input-output pairs—is adopted here to develop a heuristic solution. VMs are classified into different types based on user download rates, VM workloads, and VM popularity. The optimum placement of different VM types is found in an offline phase. VMs are matched to their types in real time (an online phase) and placed according to the placement obtained in the offline phase. In this section, we develop a real-time implementation of the MILP model called an *energy-efficient VM placement heuristic for the cloud-fog architecture* (EEVM-CF) to mimic the MILP model. The EEVM-CF heuristic consists of two phases: offline and online.

Fig. 17 (a) shows the flowchart of the offline phase of the heuristic. The offline phase starts by classifying VMs into multiple types based on their popularity, CPU usage, and user data rates (i.e., type 1 has 16% popularity, 1% CPU baseline usage, and a 1 Mbps user data rate; type 2 has 16% popularity, 1% CPU baseline usage, and 10 Mbps, etc.). The offline heuristic then checks all the candidate nodes hosting VMs in the cloud-fog architecture. The most energy-efficient placement for each possible number of replicas is found through an exhaustive search over all the possible placements for this number of replicas, i.e. the most energy-efficient placement for one replica, two replicas, etc., up to N replicas (N being the number of cloud and fog nodes in the network). For the cloud nodes, VMs can be placed in any number and combination of clouds. For fog nodes, there are two placement scenarios. In the first, VMs are replicated to the metro fog and in all nodes. In the second, VMs are replicated to the two access fog nodes in all nodes. The most energy efficient placements considering all possible number of replicas are thus created from the search space P to find the optimum placements for different VM types. For each VM type, each placement in P is examined, and the total cloud-fog power consumption is calculated. The traffic resulting from replicating the VMs in clouds and fogs and the workloads of VMs with linear workload profiles is calculated based on the number of users each VM serves. After checking all the candidate nodes, the optimum placement of a VM type is the placement that results in the minimum network, cloud, and fog power consumptions.

In the online phase (real-time placement of VMs), VMs are then matched to their types (the online phase) (Fig. 17 (b)) and placed according to the placements obtained in the offline phase. The traffic resulting from replicating VMs in the cloud

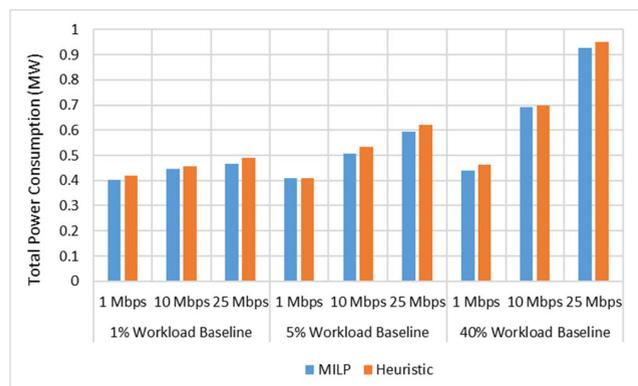


FIGURE 18. Total power consumption of the MILP model compared with EEVM-CF heuristics considering VMs with 1%, 5% and 40% CPU workload baseline.

is then routed over the core network based on minimum hop routing [30], and the workload of the cloud-fog where the VM replicas are placed is updated. After placing all VMs, the total power consumption of the cloud-fog architecture is calculated.

The heuristics are examined by considering the AT&T network as an 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. 18 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 III. 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% to 2% of the total power consumption due to almost identical VM placements obtained by the MILP model and the heuristic.

V. CONCLUSIONS

In this paper, the placement of VMs over a cloud-fog architecture is investigated with the aim of minimizing the total power consumption. The optimization 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 analyze the impact of different factors including VM popularity, the traffic between the VM and its users, the VM workload, the profile of the workload versus number of users, the proximity of fog nodes and the PUE.

The decision to serve users from fog nodes is driven by the trade-off between the network power saved by placing VMs in fog nodes close to end users, and the increase in processing power that results from replicating VMs to the fog. Our results demonstrate that VM placement in fog computing might lead to power saving depending on many factors which include workload and network bandwidth requirements of VMs, VMs popularity among users and the energy efficiency of distributed clouds.

The results evaluate a range of boundary and typical scenarios. For example, the processing power consumption of VMs of a linear workload profile with high data rate and minimum CPU utilization of 1% allows offloading VMs with 16% popularity to the access fog nodes. Other VMs are optimally replicated to metro fog nodes. Significant power savings of 48% compared to optimized placement in distributed clouds and 64% compared to a placement considering traditional cloud locations, have resulted from this offloading. VMs with linear workload and a minimum CPU utilization 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.

ACKNOWLEDGMENT

The first author would like to acknowledge the Government of Saudi Arabia for funding his Ph.D. scholarship. All data are provided in full in the results section of this paper.

REFERENCES

- [1] Cisco. (2016). *Cisco Global Cloud Index?: Forecast and Methodology, 2015–2020*. Accessed: Mar. 2, 2019. [Online]. Available: https://www.cisco.com/c/dam/m/en_us/service-provider/ciscoknowledgenetwork/files/622_11_15-16-Cisco_GCI_CKN_2015-2020_AMER_EMEAR_NOV2016.pdf
- [2] N. Fonseca and R. Boutaba, *Cloud Services, Networking, and Management*. Hoboken, NJ, USA: Wiley, 2015.
- [3] S. Antonio. (2014). Cisco delivers vision of fog computing to accelerate value from billions of connected devices. Cisco. Accessed: Sep. 27, 2017. [Online]. Available: <https://newsroom.cisco.com/press-release-content?articleId=1334100>
- [4] A. Yousefpour, G. Ishigaki, and J. P. Jue, "Fog computing: Towards minimizing delay in the Internet of Things," in *Proc. IEEE 1st Int. Conf. Edge Comput.*, Jun. 2017, pp. 17–24.
- [5] *Fog Computing and the Internet of Things: Extend the Cloud to Where the Things Are*, Cisco Syst., San Jose, CA, USA, 2016, p. 6. [Online]. Available: <https://www.cisco.com>
- [6] O. Skarlat, M. Nardelli, S. Schulte, and S. Dustdar, "Towards QoS-aware fog service placement," in *Proc. IEEE 1st Int. Conf. Fog Edge Comput. (ICFEC)*, May 2017, pp. 89–96.
- [7] N. Iotti, M. Picone, S. Cirani, and G. Ferrari, "Improving quality of experience in future wireless access networks through fog computing," *IEEE Internet Comput.*, vol. 21, no. 2, pp. 26–33, Mar. 2017.
- [8] J. Wu, M. Dong, K. Ota, J. Li, W. Yang, and M. Wang, "Fog-computing-enabled cognitive network function virtualization for an information-centric future Internet," *IEEE Commun. Mag.*, vol. 57, no. 7, pp. 48–54, Jul. 2019.
- [9] 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, May 2016.
- [10] R. Deng, R. Lu, C. Lai, and T. H. Luan, "Towards power consumption-delay tradeoff by workload allocation in cloud-fog computing," in *Proc. IEEE Int. Conf. Commun. (ICC)*, Jun. 2015, pp. 3909–3914.
- [11] S. Sarkar and S. Misra, "Theoretical modelling of fog computing: A green computing paradigm to support IoT applications," *IET Netw.*, vol. 5, no. 2, pp. 23–29, Mar. 2016.
- [12] M. O. I. Musa, T. E. H. El-Gorashi, and J. M. H. Elmghani, "Bounds on GreenTouch GreenMeter network energy efficiency," *J. Lightw. Technol.*, vol. 36, no. 23, pp. 5395–5405, Dec. 1, 2018.

- [13] X. Dong, A. Lawey, T. E. H. El-Gorashi, and J. M. H. Elmirghani, "Energy-efficient core networks," in *Proc. 16th Int. Conf. Opt. Netw. Des. Modelling (ONDM)*, 2012, pp. 1–9.
- [14] X. Dong, T. E. H. El-Gorashi, and J. M. H. Elmirghani, "On the energy efficiency of physical topology design for IP over WDM networks," *J. Lightw. Technol.*, vol. 30, no. 12, pp. 1931–1942, Jun. 15, 2012.
- [15] X. Dong, T. El-Gorashi, and J. M. H. Elmirghani, "IP over WDM networks employing renewable energy sources," *J. Lightw. Technol.*, vol. 29, no. 1, pp. 3–14, Jan. 1, 2011.
- [16] M. Musa, T. Elgorashi, and J. Elmirghani, "Energy efficient survivable IP-over-WDM networks with network coding," *J. Opt. Commun. Netw.*, vol. 9, no. 3, pp. 207–217, Mar. 2017.
- [17] M. Musa, T. Elgorashi, and J. Elmirghani, "Bounds for energy-efficient survivable IP over WDM networks with network coding," *J. Opt. Commun. Netw.*, vol. 10, no. 5, p. 471, May 2018.
- [18] X. Dong, T. El-Gorashi, and J. M. H. Elmirghani, "Green IP over WDM networks with data centers," *J. Lightw. Technol.*, vol. 29, no. 12, pp. 1861–1880, Jun. 15, 2011.
- [19] A. Q. Lawey, T. E. H. El-Gorashi, and J. M. H. Elmirghani, "BitTorrent content distribution in optical networks," *J. Lightw. Technol.*, vol. 32, no. 21, pp. 4209–4225, Nov. 1, 2014.
- [20] H. M. M. Ali, T. E. H. El-Gorashi, A. Q. Lawey, and J. M. H. Elmirghani, "Future energy efficient data centers with disaggregated servers," *J. Lightw. Technol.*, vol. 35, no. 24, pp. 5361–5380, Dec. 15, 2017.
- [21] N. I. Osman, T. El-Gorashi, L. Krug, and J. M. H. Elmirghani, "Energy-efficient future high-definition TV," *J. Lightw. Technol.*, vol. 32, no. 13, pp. 2364–2381, Jul. 1, 2014.
- [22] A. M. Al-Salim, T. E. H. El-Gorashi, A. Q. Lawey, and J. M. H. Elmirghani, "Greening big data networks: Velocity impact," *IET Optoelectron.*, vol. 12, no. 3, pp. 126–135, Jun. 2018.
- [23] A. M. Al-Salim, A. Q. Lawey, T. E. H. El-Gorashi, and J. M. H. Elmirghani, "Energy efficient big data networks: Impact of volume and variety," *IEEE Trans. Netw. Service Manage.*, vol. 15, no. 1, pp. 458–474, Mar. 2018.
- [24] M. S. Hadi, A. Q. Lawey, T. E. H. El-Gorashi, and J. M. H. Elmirghani, "Patient-centric cellular networks optimization using big data analytics," *IEEE Access*, vol. 7, pp. 49279–49296, 2019.
- [25] M. S. Hadi, A. Q. Lawey, T. E. H. El-Gorashi, and J. M. H. Elmirghani, "Big data analytics for wireless and wired network design: A survey," *Comput. Netw.*, vol. 132, pp. 180–199, Feb. 2018.
- [26] 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. 347–357, Sep. 2017.
- [27] H. A. Alharbi, T. E. H. Elgorashi, A. Q. Lawey, and J. M. H. Elmirghani, "The impact of inter-virtual machine traffic on energy efficient virtual machines placement," in *Proc. IEEE Sustainability Through ICT Summit (StICT)*, Jun. 2019, pp. 1–7.
- [28] A. Q. Lawey, T. E. H. El-Gorashi, and J. M. H. Elmirghani, "Distributed energy efficient clouds over core networks," *J. Lightw. Technol.*, vol. 32, no. 7, pp. 1261–1281, Apr. 1, 2014.
- [29] A. N. Al-Quzweeni, A. Q. Lawey, T. E. H. Elgorashi, and J. M. H. Elmirghani, "Optimized energy aware 5G network function virtualization," *IEEE Access*, vol. 7, pp. 44939–44958, 2019.
- [30] L. Nonde, T. E. H. El-Gorashi, and J. M. H. Elmirghani, "Energy efficient virtual network embedding for cloud networks," *J. Lightw. Technol.*, vol. 33, no. 9, pp. 1828–1849, May 1, 2015.
- [31] Z. T. Al-Azez, A. Q. Lawey, T. E. H. El-Gorashi, and J. M. H. Elmirghani, "Energy efficient IoT virtualization framework with peer to peer networking and processing," *IEEE Access*, vol. 7, pp. 50697–50709, 2019.
- [32] VMware. (2014). *Virtualization Essentials*. Accessed: Mar. 22, 2017. [Online]. Available: <http://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/ebook/gated-vmw-ebook-virtualization-essentials.pdf>
- [33] Microsoft. (2018). *Virtual Machine Load Balancing Overview*. Accessed: Feb. 24, 2018. [Online]. Available: <https://docs.microsoft.com/en-us/windows-server/failover-clustering/vm-load-balancing-overview>
- [34] I. Pietri and R. Sakellariou, "Mapping virtual machines onto physical machines in cloud computing: A survey," *ACM Comput. Surv.*, vol. 49, no. 3, pp. 1–30, Dec. 2016.
- [35] J. M. H. Elmirghani, T. Klein, K. Hinton, L. Nonde, A. Q. Lawey, T. E. H. El-Gorashi, M. O. I. Musa, and X. Dong, "GreenTouch Green-Meter core network energy-efficiency improvement measures and optimization," *J. Opt. Commun. Netw.*, vol. 10, no. 2, p. A250, Feb. 2018.
- [36] 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," *Econ. Benefits Deploy. Challenges. HotCloud*, pp. 8–15, 2010.
- [37] A. Bianco, L. Girardo, and D. Hay, "Optimal resource allocation for disaster recovery," in *Proc. IEEE Global Telecommun. Conf. (GLOBECOM)*, Dec. 2010, pp. 1–5.
- [38] A. Satpathy, S. K. Addya, A. K. Turuk, B. Majhi, and G. Sahoo, "Crow search based virtual machine placement strategy in cloud data centers with live migration," *Comput. Electr. Eng.*, vol. 69, pp. 334–350, Jul. 2018.
- [39] 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, nos. 5–6, pp. 1222–1235, Sep. 2013.
- [40] F. Alharbi, Y.-C. Tian, M. Tang, W.-Z. Zhang, C. Peng, and M. Fei, "An ant colony system for energy-efficient dynamic virtual machine placement in data centers," *Expert Syst. Appl.*, vol. 120, pp. 228–238, Apr. 2019.
- [41] S. K. Addya, A. K. Turuk, B. Sahoo, A. Satpathy, and M. Sarkar, "A game theoretic approach to estimate fair cost of VM placement in cloud data center," *IEEE Syst. J.*, vol. 12, no. 4, pp. 3509–3518, Dec. 2018.
- [42] N. Tziritis, M. Koziri, A. Bachtsevani, T. Loukopoulos, G. Stamoulis, S. U. Khan, and C.-Z. Xu, "Data replication and virtual machine migrations to mitigate network overhead in edge computing systems," *IEEE Trans. Sustain. Comput.*, vol. 2, no. 4, pp. 320–332, Oct. 2017.
- [43] A. Manzalini, R. Minerva, F. Callegati, W. Cerroni, and A. Campi, "Clouds of virtual machines in edge networks," *IEEE Commun. Mag.*, vol. 51, no. 7, pp. 63–70, Jul. 2013.
- [44] L. F. Bittencourt, J. Diaz-Montes, R. Buyya, O. F. Rana, and M. Parashar, "Mobility-aware application scheduling in fog computing," *IEEE Cloud Comput.*, vol. 4, no. 2, pp. 26–35, Mar. 2017.
- [45] AT&T. *AT&T's 38 Global Internet Data Centers*. Accessed: Aug. 15, 2018. [Online]. Available: https://www.beatriceco.com/pdf/AB-0119-09Eb_idcmap.pdf
- [46] W. Dargie, "A stochastic model for estimating the power consumption of a processor," *IEEE Trans. Comput.*, vol. 64, no. 5, pp. 1311–1322, May 2015.
- [47] M. Pedram and I. Hwang, "Power and performance modeling in a virtualized server system," in *Proc. 39th Int. Conf. Parallel Process. Workshops*, Sep. 2010, pp. 520–526.
- [48] Jarrod. *Linux Web Server Performance Benchmark—2016 Results*. Accessed: Nov. 24, 2017. [Online]. Available: <https://www.rootusers.com/linux-web-server-performance-benchmark-2016-results/>
- [49] P. Padala, X. Zhu, Z. Wang, S. Singhal, and K. G. Shin, "Performance evaluation of virtualization technologies for server consolidation," Tech. Rep., 2007.
- [50] P. R. M. Vasconcelos, G. A. de Araujo Freitas, and T. G. Marques, "Virtualization technologies in Web conferencing systems: A performance overview," in *Proc. 11th Int. Conf. Internet Technol. Secured Trans. (ICITST)*, Dec. 2016, pp. 376–383.
- [51] A. Bharambe, J. Pang, and S. Seshan, "Colyseus: A distributed architecture for online multiplayer games," in *Proc. 3rd Symp. Netw. Syst. Design Implement. (NSDI)*, 2006, pp. 155–168.
- [52] IBM System Energy Estimator. Accessed: Mar. 15, 2017. [Online]. Available: <http://see.au-syd.mybluemix.net/see/EnergyEstimator>
- [53] M. Forzati, A. Bianchi, J. Chen, K. Grobe, B. Lannoo, C. M. Machuca, J.-C. Point, B. Skubic, S. Verbrugge, E. Weis, L. Wosinska, and D. Breuer, "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, Feb. 2015.
- [54] Huawei. (2010). *Next-Generation PON Evolution*. Accessed: Mar. 7, 2017. [Online]. Available: http://www.huawei.com/ilink/en/download/HW_077443
- [55] S. B. Weinstein, Y. Luo, and T. 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.
- [56] Y. Zhang, P. Chowdhury, M. Tornatore, and B. Mukherjee, "Energy efficiency in telecom optical networks," *IEEE Commun. Surveys Tuts.*, vol. 12, no. 4, pp. 441–458, 4th Quart., 2010.
- [57] J. Baliga, R. Ayre, K. Hinton, W. V. Sorin, and R. S. Tucker, "Energy consumption in optical IP networks," *J. Lightw. Technol.*, vol. 27, no. 13, pp. 2391–2403, Jul. 2009.
- [58] G. Shen and R. S. Tucker, "Energy-minimized design for IP over WDM networks," *J. Opt. Commun. Netw.*, vol. 1, no. 1, pp. 176–186, Jun. 2009.

- [59] F. Parmigiani, L. Provost, P. Petropoulos, D. J. Richardson, W. Freude, J. Leuthold, A. D. Ellis, and I. Tomkos, "Progress in multichannel all-optical regeneration based on fiber technology," *IEEE J. Sel. Topics Quantum Electron.*, vol. 18, no. 2, pp. 689–700, Mar. 2012.
- [60] K. T. Malladi, B. C. Lee, F. A. Nothaft, C. Kozyrakos, K. Periyathambi, and M. Horowitz, "Towards energy-proportional datacenter memory with mobile DRAM," *ACM SIGARCH Comput. Archit. News*, vol. 40, no. 3, pp. 37–48, Sep. 2012.
- [61] ZTE Corporation. (2013). *ZXA10 C300 Hardware Description*. Accessed: Mar. 15, 2017. [Online]. Available: http://enterprise.zte.com.cn/en/products/network_Infrastructure/broadband_access/xpon_onu/201312/t20131209_414454.html
- [62] A. Shehabi, "United States data center energy usage report," Tech. Rep., 2016.
- [63] *Cisco CRS-1 4-Slot Single-Shelf System*, Cisco Syst., San Jose, CA, USA, 2014.
- [64] *Cisco Network Convergence System 5500 Series Data Sheet*, Cisco, San Jose, CA, USA, 2018, pp. 1–11.
- [65] *Cisco Nexus 9300-EX and 9300-FX Platform Switches*, Cisco Syst., San Jose, CA, USA, 2018.
- [66] (2016). *Facilities | N8HPC*. Accessed: Feb. 8, 2016. [Online]. Available: <http://n8hpc.org.uk/facilities/>
- [67] *Cisco ONS 15454 40 Gbps CP-DQPSK Full C-Band Tuneable Transponder Card*, Cisco Syst., San Jose, CA, USA.
- [68] *OTS-4400 Regenerator*, Oclaro, San Jose, CA, USA, 2014.
- [69] *EDFA Optical Amplifiers*, MRV, New York, NY, USA, 2012.
- [70] Glimmerglass. (2013). *Intelligent Optical System 600*. Accessed: Mar. 16, 2017. [Online]. Available: <http://www.glimmerglass.com/products/intelligent-optical-systems/>
- [71] "GreenTouch green meter research study?: Reducing the net energy consumption in communications networks by up to 90% by 2020," White Paper, 2015, pp. 1–25.
- [72] C. Gray, R. Ayre, K. Hinton, and R. S. Tucker, "Power consumption of IoT access network technologies," in *Proc. IEEE Int. Conf. Commun. Workshop (ICCW)*, Jun. 2015, pp. 2818–2823.
- [73] *IBM Systems Energy Estimator*, 2018.
- [74] *United States—2021 Forecast Highlights*, Cisco, San Jose, CA, USA, 2018.
- [75] (2018). *Similar Web*. Accessed: May 10, 2019. [Online]. Available: <https://www.similarweb.com/>
- [76] L. A. Adamic and B. A. Huberman, "Zipf's law and the Internet," *Glottometrics*, vol. 3, no. 1, pp. 143–150, 2002.
- [77] A. Venkatraman, V. Pandey, B. Plale, and S. Shei, "Benchmarking effort of virtual machines on multicore machines," Tech. Rep., 2007.
- [78] N. Mi, G. Casale, L. Cherkasova, and E. Smirni, "Injecting realistic burstiness to a traditional client-server benchmark," in *Proc. 6th Int. Conf. Auton. Comput. (ICAC)*, 2009, pp. 15–19.
- [79] *Cisco Global Cloud Index?: Forecast and Methodology, 2016–2021*, Cisco, San Jose, CA, USA, 2018.
- [80] Google. (2018). *Google Help—Network Check*. Accessed: Feb. 2, 2018. [Online]. Available: <https://support.google.com/wifi/answer/6246634?hl=en-GB>
- [81] YouTube. (2018). *Video Quality Report*. [Online]. Available: <https://www.google.com/get/videoqualityreport/#methodology>
- [82] YouTube. (2017). *Live Encoder Settings, Bitrates, and Resolutions*. [Online]. Available: <https://support.google.com/youtube/answer/2853702>
- [83] A. Bharambe, "Donnybrook: Enabling large-scale, high-speed, peer-to-peer games," in *Proc. SIGCOMM*, 2008, pp. 389–400.
- [84] Netflix. (2014). *Internet Connection Speed Recommendations*. Accessed: Jan. 27, 2018. [Online]. Available: <https://help.netflix.com/en/node/306>



HATEM A. ALHARBI received the B.Sc. degree (Hons.) in computer engineering from Umm Alqura University, Mecca, Saudi Arabia, in 2012, and the M.Sc. degree (Hons.) in digital communication networks from the University of Leeds, U.K., in 2015, where he is currently pursuing the Ph.D. degree with the School of Electronic and Electrical Engineering. He is currently a Lecturer with the Computer Engineering Department, School of Computer Science and Engineering, University of Taibah, Saudi Arabia.



TAISIR E. H. ELGORASHI received the B.S. degree (Hons.) in electrical and electronic engineering from the University of Khartoum, Khartoum, Sudan, in 2004, the M.Sc. degree (Hons.) in photonic and communication systems from the University of Wales, Swansea, U.K., in 2005, and the Ph.D. degree in optical networking from the University of Leeds, Leeds, U.K., in 2010. She was a Postdoctoral Researcher with the University of Leeds, from 2010 to 2014, where she focused on the energy efficiency of optical networks investigating the use of renewable energy in core networks, green IP over WDM networks with datacenters, energy efficient physical topology design, energy efficiency of content distribution networks, distributed cloud computing, network virtualization, and big data. In 2012, she was a BT Research Fellow, where she developed energy efficient hybrid wireless-optical broadband access networks and explored the dynamics of TV viewing behavior and program popularity. She is currently a Lecturer in optical networks with the School of Electrical and Electronic Engineering, University of Leeds. The energy efficiency techniques developed during her Postdoctoral Research contributed three out of the eight carefully chosen core network energy efficiency improvement measures recommended by the GreenTouch consortium for every operator network worldwide. Her work led to several invited talks at GreenTouch, Bell Labs, Optical Network Design and Modelling Conference, Optical Fiber Communications Conference, International Conference on Computer Communications, EU Future Internet Assembly, IEEE Sustainable ICT Summit, and IEEE 5G World Forum and collaboration with Nokia and Huawei.



JAAFAR M. H. ELMIRGHANI (Senior Member, IEEE) received the Ph.D. degree in the synchronization of optical systems and optical receiver design from the University of Huddersfield, U.K., in 1994, and the D.Sc. degree in communication systems and networks from the University of Leeds, U.K., in 2014. From 2000 to 2007, he was the Chair in optical communications with the University of Wales Swansea. He is currently the Director of the School of Electronic and Electrical Engineering, Institute of Communication and Power Networks, University of Leeds. He joined Leeds, in 2007. He founded, developed, and directed the Institute of Advanced Telecommunications and the Technium Digital (TD), a technology incubator/spin-off hub. He has provided outstanding leadership in a number of large research projects at the IAT and TD. He has coauthored *Photonic switching Technology: Systems and Networks* (Wiley). He has published over 500 articles. His research interests include optical systems and networks. He is a Fellow of the IET and the Institute of Physics. He was a Principal Investigator (PI) of the £6m EPSRC INTElligent Energy awaRe NETworks (INTERNET) Programme Grant, from 2010 to 2016. He is currently the PI of the £6.6m EPSRC Terabit Bidirectional Multi-user Optical Wireless System (TOWS) for 6G LiFi Programme Grant, from 2019 to 2024. He has been awarded in excess of £30 million in grants to date from EPSRC, the EU and industry, and has held prestigious fellowships funded by the Royal Society and by BT. He was an IEEE Comsoc Distinguished Lecturer, from 2013 to 2016. He received the IEEE Communications Society Hal Sobol Award, the IEEE Comsoc Chapter Achievement Award for excellence in chapter activities, in 2005, the University of Wales Swansea Outstanding Research Achievement Award, in 2006, the IEEE Communications Society Signal Processing and Communication Electronics Outstanding Service

Award, in 2009, the Best Paper Award at IEEE ICC'2013, the IEEE Comsoc Transmission Access and Optical Systems Outstanding Service Award, in 2015, in recognition of Leadership and Contributions to the Area of Green Communications, received the GreenTouch 1000x Award, in 2015, for pioneering research contributions to the field of energy efficiency in telecommunications, the 2016 IET Optoelectronics Premium Award and shared with six GreenTouch innovators the 2016 Edison Award in the Collective Disruption Category for their work on the GreenMeter, the international competition, and clear evidence of his seminal contributions to Green Communications, which have a lasting impact on the environment (green) and society. He was the Co-Chair of the GreenTouch Wired, Core and Access Networks Working Group, an Adviser to the Commonwealth Scholarship Commission, a member of the Royal Society International Joint Projects Panel, and a member of the Engineering and Physical Sciences Research Council (EPSRC) College. He was the Chairman of the IEEE Comsoc Transmission Access and Optical Systems Technical Committee and the IEEE Comsoc Signal Processing and Communications Electronics Technical Committee. He was the founding Chair of the Advanced Signal Processing for Communication Symposium

which started at IEEE GLOBECOM'99 and has continued since at every ICC and GLOBECOM. He was also the founding Chair of the first IEEE ICC/GLOBECOM Optical Symposium at GLOBECOM'00, the Future Photonic Network Technologies, Architectures, and Protocols Symposium. He chaired this Symposium, which continues to date under different names. He was the Founding Chair of the first Green Track at ICC/GLOBECOM at GLOBECOM 2011. He is the Chair of the IEEE Sustainable ICT Initiative within the IEEE Technical Activities Board (TAB) Future Directions Committee (FDC) and within the IEEE Communications Society, a pan IEEE Societies Initiative responsible for Green and Sustainable ICT activities across IEEE, since 2012. He was on the technical program committee of 38 IEEE ICC/GLOBECOM conferences, from 1995 to 2019, including 18 times as the Symposium Chair. He was an Editor of the *IEEE Communications Magazine*. He is currently an Editor of the *IET Optoelectronics*, the *Journal of Optical Communications*, IEEE COMMUNICATIONS SURVEYS AND TUTORIALS, and the IEEE Journal on Selected Areas in Communications series on Green Communications and Networking.

• • •