Master of Science in Telecommunication Systems
February 2018

Virtualization of Data Centers

study on Server Energy Consumption Performance

Praneel Reddy Padala

Department of computing
Blekinge Institute of Technology
SE–371 79 Karlskrona, Sweden

This thesis is submitted to the Faculty of Computing at Blekinge Institute of Technology in partial fulfillment of the requirements for the degree of Master of Science in Telecommunication Systems. The thesis is equivalent to 20 weeks of full time studies.

Contact Information:
Author(s):
Praneel Reddy Padala
E-mail: prpa16@student.bth.se

University advisor:
Prof. Adrian Popescu
Department of TeleCommunications

Faculty of computing	Internet :	www.bth.se
Blekinge Institute of Technology	Phone :	+46 455 38 50 00
SE–371 79 Karlskrona, Sweden	Fax :	+46 455 38 50 57

Abstract

Due to various reasons data centers have become ubiquitous in our society. Energy [19] costs are significant portion of data centers total lifetime costs which also makes financial sense to operators. This increases huge concern towards the energy costs and environmental impacts of data center. Power costs and energy efficiency are the major challenges front of us. From overall cyber [14] energy used, 15% is used by networking portion of a data center. Its estimated that the energy used by network infrastructure in a data center world wide is 15.6 billion kWh and is expected to increase to around 50%.

Power costs and Energy Consumption plays a major role throughout the life time of a data center, which also leads to the increase in financial costs for data center operators and increased usage of power resources. So, resource utilization has become a major issue in the data centers.

The main aim of this thesis study is to find the efficient way for utilization of resources and decrease the energy costs to the operators in the data centers using virtualization. Virtualization technology is used to deploy virtual servers on physical servers which uses the same resources and helps to decrease the energy consumption of a data center.

Keywords: Energy consumption, Data Centers, Virtualization, Servers.

Acknowledgements

Firstly, I would like to thank supremacy, God for blessing me with knowledge, strength and vision in my life. I would also thank my parents and other family members for unconditional love and support in my life.

I express my deepest gratitude to my supervisor Dr. Adrian Popescu for help, encouragement and support offered throughout the thesis with his valuable suggestions. I feel grateful working under his guidance. From the discussions and conversations with my professor I have increased my thinking capabilities and problem solving in different ways.

I would like to thank my examiner Prof. Kurt Tutschku for his valuable feedback and suggestions, which helped me a lot in improving this thesis work.

I also thank my friends for their help throughout my studies. I would also like to thank all my professors from JNTUK and BTH institutions for guiding and shaping my career.

Contents

Abstract **i**

Acknowledgements **ii**

1 Introduction **1**
1.1 Data Center 1
1.2 Need for Data Center 1
1.3 Major Components of a Data Center 1
1.3.1 Electrical Infrastructure 1
1.3.2 IT Infrastructure 2
1.3.3 Building Shells 3
1.4 Energy Consumption 3
1.5 Virtualization 3
1.6 Server Consolidation Technique 4
1.7 Types of Virtualization 4
1.7.1 Server Virtualization 4
1.7.2 Operating System Virtualization 5
1.7.3 Network Function Virtualization 6
1.7.4 Desktop Virtualization 6
1.8 Virtualization Vs Cloud Computing 7
1.9 Hypervisors 8
1.10 Types of Hypervisors 8
1.10.1 Type1 Hypervisor 8
1.10.2 Type2 Hypervisor 9
1.11 Problem Statement 9
1.12 Motivation 9
1.13 Outline 10
1.14 Research Questions 10

2 Conceptual Background **11**
2.1 Data Center Efficiency Metrics 11
2.1.1 Power Usage Efficiency (PUE) 12
2.1.2 Data Center Infrastructure Efficiency (DCiE) 12
2.1.3 Carbon Footprint (C) 13

2.2 Functioning of a Data Center 13
2.2.1 Power Supply Unit 13
2.2.2 Cooling Unit 14
2.2.3 Controlled Access 15
2.2.4 Control Center 15
2.2.5 Safety Measurements 16
2.3 Simulation Tools 16
2.3.1 Cloud Sim 16
2.3.2 Green Cloud 16

3 Related Work **18**

4 Method **20**
4.1 CloudSim 20
4.2 Implementation 21

5 Results **24**

6 Analysis and Discussion **29**
6.1 What is the impact on energy consumption of virtual data centers when compared with physical data centers? 29
6.2 Will KVM perform better than Xen hypervisor in server virtualization in terms of energy consumption while varying number of virtual machines? 30

7 Conclusions and Future Work **31**

References **32**

List of Figures

1.1 Electrical Infrastructure 1
1.2 Virtualization 3
1.3 Server Consolidation 4
1.4 Server Virtualization 5
1.5 Operating System Virtualization 5
1.6 Network Function Virtualization 6
1.7 Desktop Virtualization 7
1.8 Type1 and Type2 Hypervisors 8

2.1 Demands and Constraints 11
2.2 For calculating Energy Metrics 12
2.3 Power Supply Unit 13
2.4 Cooling Unit 14
2.5 Controlled Access 15
2.6 Cloud Sim Architecture 16
2.7 Green Cloud Architecture 17

4.1 layered architecture of cloudsim 21
4.2 Simulating with one vm 22
4.3 Simulating with three vms 22
4.4 Simulating with five vms 23
4.5 Simulating with ten vms 23
4.6 Simulating with fifteen vms 23

5.1 KVM with varying number of virtual machines 24
5.2 Xen with varying number of virtual machines 25
5.3 Comparing KVM and Xen when 3 virtual machines are deployed . 26
5.4 Comparing KVM and Xen when 5 virtual machines are deployed . 27
5.5 Comparing KVM and Xen when 10 virtual machines are deployed 27
5.6 Comparing KVM and Xen when 15 virtual machines are deployed 28

Chapter 1

Introduction

1.1 Data Center

[6] A Data Center is a facility used to house computer systems and associated components, such as telecommunications and storage systems and uses redundant backup power supplies and redundant data communication connections.

1.2 Need for Data Center

Every organization,whether big or small, has a large amount of data that are needed to be stored somewhere. A Data Center is essential to store, manage, process, and distribute large amounts of data of an organization. In today's modern era, every organization needs to have a Data Center for its smooth running. In the fast growing world where everything is going online, You need to make your presence over the internet for more connectivity and convenience. So, we are growing with huge requirement for data centers.

1.3 Major Components of a Data Center

1.3.1 Electrical Infrastructure

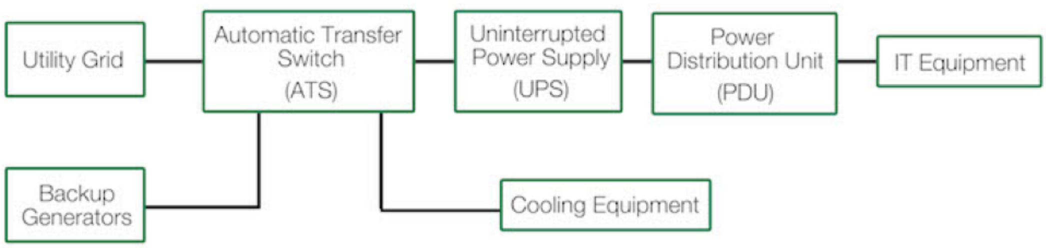

Figure 1.1: Electrical Infrastructure

Utility Grid

[18] utility grid is usually a commercial electric power distribution system that takes electricity from a generator and transmits it over a certain distance and then takes the electricity down to the consumer through a distribution system. The entire system is referred as Grid.

Automatic Transfer Switch (ATS)

ATS is an electrical switch that switches load between two sources. This ATS is installed where a backup generator is located. When there was a failure with electrical power, it switches the load and electrical input is taken from the generator.

Uninterrupted Power Supply (UPS)

[17] Power is very essential for a data center and our IT infrastructure is designed to be supplied with uniform flow. But, we can't expect a steady flow from the utility. They are used as backup during brief power outages. Although backup generators helps during extended power outages, they can't help you with power fluctuations. UPS helps us during power fluctuations and brief power outages providing clean and high quality power.

Power Distribution Unit (PDU)

Demand for more computing power in less space means rack cabinets are more densely packed than ever. A network PDU unit distributes power supplied to the rack via multiple outlets to the rack cabinet's servers and networking equipment.

Backup Generators

Backup Generators are used to supply power for many hours or many days during extended power outages.

Cooling Equipment

Cooling equipment is used to reduce the heat generated by the servers and networking unit.

1.3.2 IT Infrastructure

IT equipment is the stuff that fills up the data center and is the reason the facility exists in the first place. The IT equipment is where all the magic happens that makes the internet, mobile communication, and apps what they are today.

At a high level, IT equipment can be categorized into three buckets: servers, communication gear, and storage equipment.

1.3.3 Building Shells

The architectural and structural components of a data center are pretty basic, consisting of four walls and a roof. A data centers is very similar to a warehouse where the middle of the building is empty (but will be filled up with IT equipment instead of boxes). Data centers are built using the same techniques and materials as typical office buildings, however the two structures differ in robustness. Their structural components will be bigger and stronger in order to sustain natural disasters or explosives. The most critical data centers, such as those used by defense departments for national security, might even be located within a mountain or deep underground for protection.

1.4 Energy Consumption

Energy consumption is the amount of energy or power used. Energy Consumption of a data center is defined as energy consumed by the IT infrastructure like servers, switches, computer components and cooling infrastructure like air conditioners, cooling turbines and electrical infrastructure.

1.5 Virtualization

[7] Virtualization is a technology that helps us to use the available hardware resources efficiently. This technology allows us to use the full capacity of physical machines by distributing its capabilities among many users or environments.

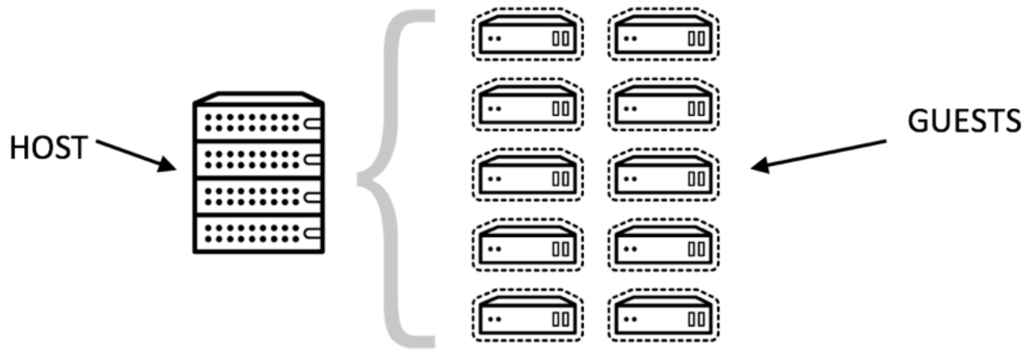

Figure 1.2: Virtualization

The physical machine holding the hypervisors is called host, while the virtual machines sharing the resources of the cost are called guests. Guests generally treat computing resources like CPU, memory and storage as hanger of resources. Operators treat these guests by controlling virtual instances of computing resources and providing them the required resources when they need them.

1.6 Server Consolidation Technique

Server consolidation stands for the efficient way to use the computer resources in order to reduce the total number of servers or server locations which an organization require. Here we can observe that each server performs its own tasks relating to that particular operating system. To the maximum extent each server can use 30% of its total computer resources allocated to it. The remaining 70% of the resources are left without use. Virtualization is a technique which helps us to consolidate more number of servers on the same underlying physical hardware. This helps in resource utilization and decrease the number of servers in an organization and there by reduces the operation and maintenance costs.

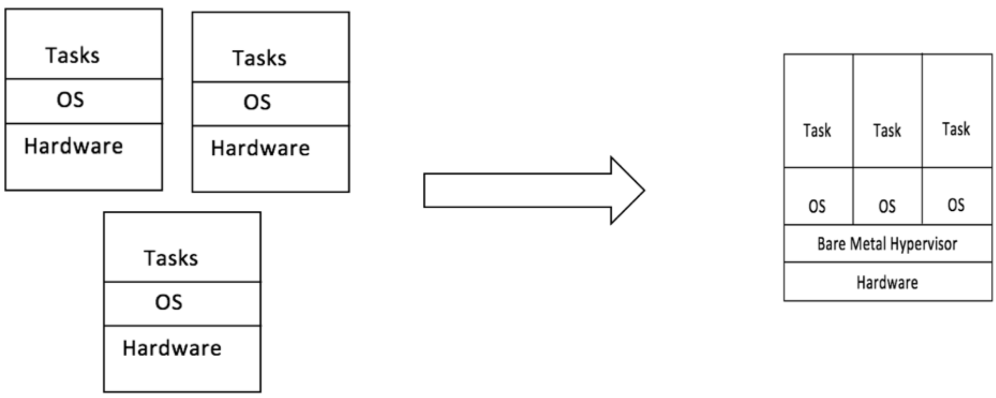

Figure 1.3: Server Consolidation

1.7 Types of Virtualization

1.7.1 Server Virtualization

[2] Servers are computers designed to process a high volume of specific tasks really well so other computers—like laptops and desktops—can do a variety of other tasks. Virtualizing a server lets it to do more of those specific functions and involves partitioning it so that the components can be used to serve multiple functions.

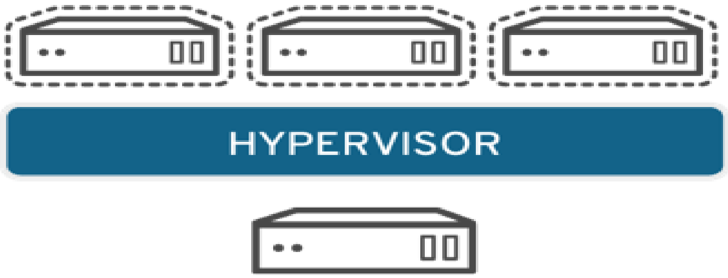

Figure 1.4: Server Virtualization

1.7.2 Operating System Virtualization

Operating system virtualization happens at the kernel and the central task managers of operating systems. It's a useful way to run Linux, Windows or any other operating system environments side-by-side. Enterprises can also push virtual operating systems to computers, which

- Reduces bulk hardware costs, since the computers don't require such high out-of-the-box capabilities.
- Increases security, since all virtual instances can be monitored and isolated.
- Limits time spent on IT services like software updates.

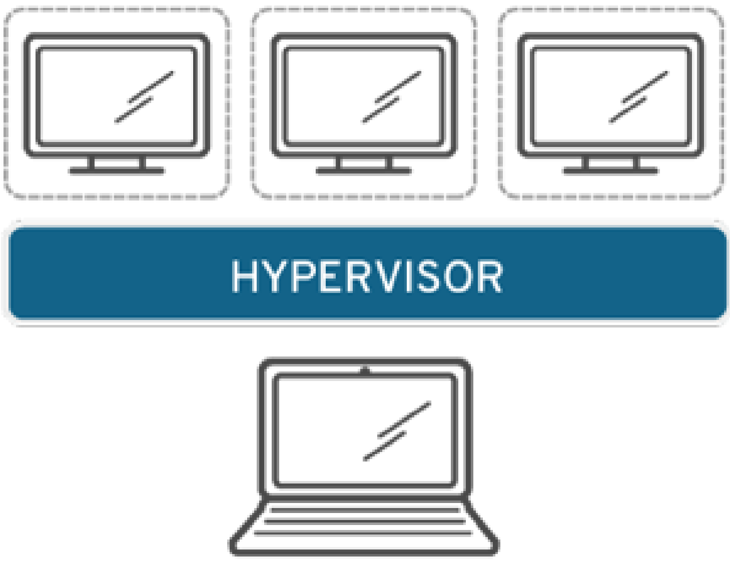

Figure 1.5: Operating System Virtualization

1.7.3 Network Function Virtualization

Network functions virtualization separates a network's key functions (like directory services, file sharing, and IP configuration) so they can be distributed among environments. Once software functions are independent of the physical machines they once lived on, specific functions can be packaged together into a new network and assigned to an environment. Virtualizing networks reduces the number of physical components like switches, routers, servers, cables, and hubs that are needed to create multiple, independent networks, and it's particularly popular in the telecommunications industry.

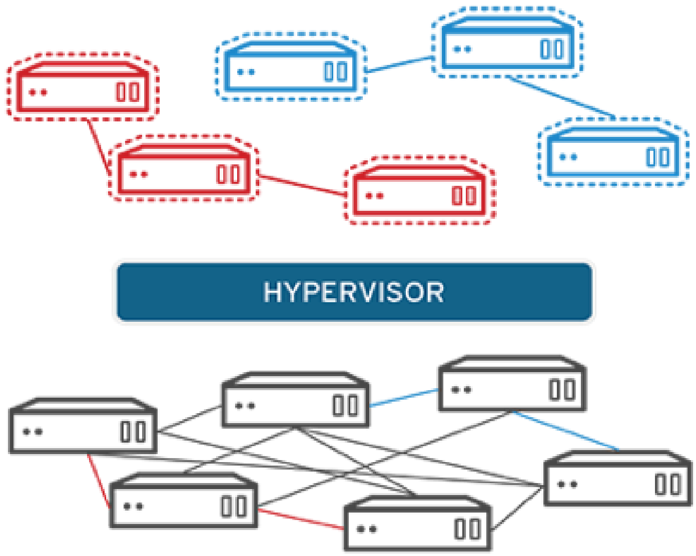

Figure 1.6: Network Function Virtualization

1.7.4 Desktop Virtualization

Easily confused with operating system virtualization—which allows you to deploy multiple operating systems on a single machine—desktop virtualization allows a central administrator (or automated administration tool) to deploy simulated desktop environments to hundreds of physical machines at once.[20] Unlike traditional desktop environments that are physically installed, configured, and updated on each machine, desktop virtualization allows admins to perform mass configurations, updates, and security checks on all virtual desktops.

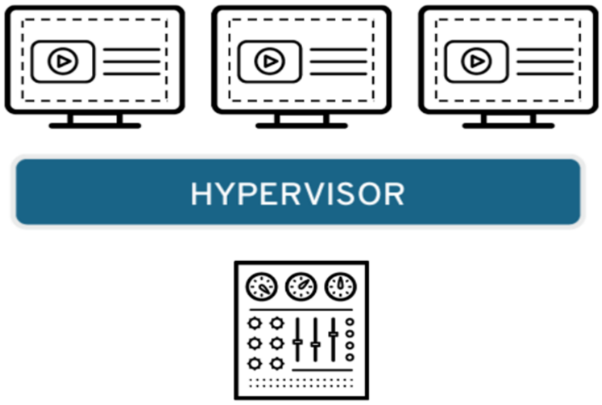

Figure 1.7: Desktop Virtualization

1.8 Virtualization Vs Cloud Computing

Since both virtualization and cloud computing revolve around separating resources from hardware, it creates us a confused state. Virtualization is a technology that separates functions from hardware.

- **Cloud computing** is more of a solution that relies on that split
- **Virtualization** helps create clouds, but that doesn't make it cloud computing.

The National Institute of Standards and Technology cites 5 features of Cloud computing:

- A network
- Pooled resources
- A user interface
- Provisioning capabilities and
- Automatic resource control (or allocation).

While virtualization creates the network and pooled resources. Further to create a user interface, provision VMs, and control/allocate resources, additional management and operating system software are needed.

1.9 Hypervisors

A hypervisor or virtual machine monitor(VMM) is the one that creates and runs virtual machines.[10] A system on which a hypervisor runs one or more virtual machines is called a host system, and each virtual machine which runs on the hypervisor is called as guest system. The hypervisor provides the guest operating systems with a virtual platform which manages the execution of the guest operating systems. Multiple instances of a variety of operating systems like linux, windows and mac may share this virtualized hardware resources.

1.10 Types of Hypervisors

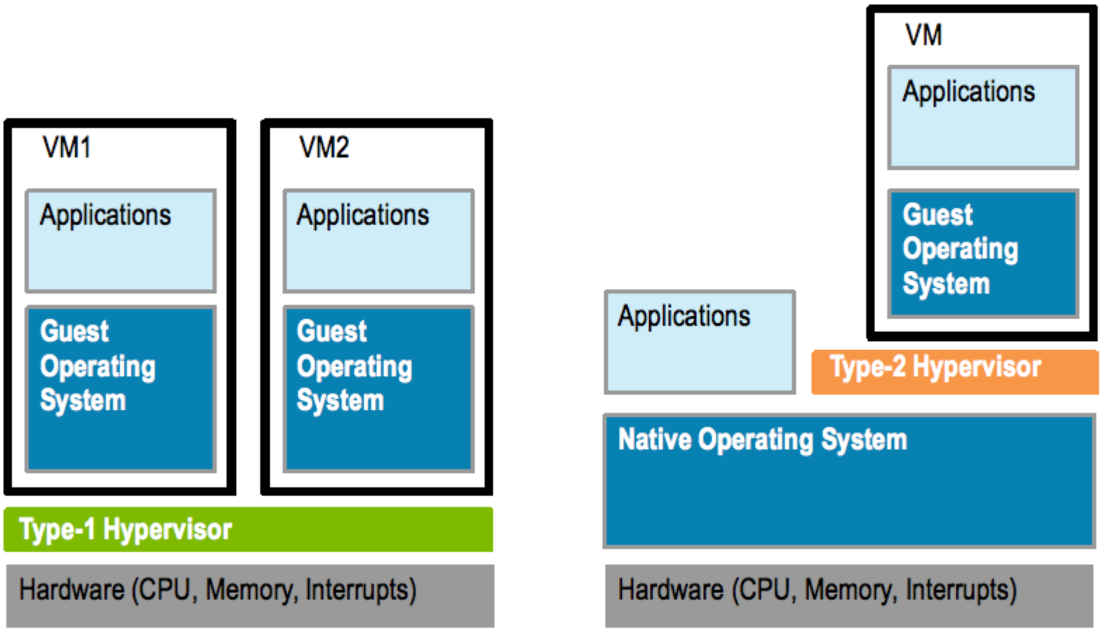

Figure 1.8: Type1 and Type2 Hypervisors

1.10.1 Type1 Hypervisor

These Hypervisors run directly on host hardware for controlling the hardware and guest operating systems. So they are called bare metal hypervisors.[4] They are also referred as native or embedded hypervisors. Xen, Microsoft Hyper-V, VMware ESXI are some of the type1 hypervisors.
Day by day type1 hypervisors are gaining popularity as they are proving more efficient than type2. According to IBM, type1 hypervisors provide higher performance, availability and security compared to type2 hypervisors.

1.10.2 Type2 Hypervisor

These hypervisors run on the host operating system. Guest operating systems runs on the host as a process. VMware workstation, Virtual Box, VMware Player are some of the type2 hypervisors. IBM recommends that type2 hypervisors can mainly be used on client systems where efficiency is less critical.

1.11 Problem Statement

As discussed above, need for data centers is increasing exponentially. With increase in data centers there will be proportional increase in energy consumption of the data centers. In 2014, [9] [23] energy consumption of data centers in the U.S is estimated 70 billion kWh, which represents about 1.8% of total U.S. electricity consumption. By 2020, its expected that energy use is increased by 4% which is approximately 73 billion KWh. So, energy consumption has become a serious concern in data centres. This also leads to huge environmental impacts and huge operational costs to the data center operators.

1.12 Motivation

Using server consolidation technique, we can reduce the energy consumption of servers as servers constitute more than 40% of the total data center energy. We can deploy either type1 or type 2 hypervisors on the servers and add virtual machines which are equivalent to the servers. These virtual machines help in best use of physical resources.

In this paper we consider two type 1 hypervisors, Xen and KVM on a physical server and gradually increased the number of virtual machines on the server and calculated the power consumption of the server and compared the performance against the physical server.

1.13 Outline

This thesis paper is structured as follows.

Chapter 1: This chapter contains the introduction for data centers and its components. We also discuss the problem statement, motivation to solve the problem along with research questions.
Chapter 2: In this chapter, we discuss about the conceptual aspects related to our thesis and the study work done regarding this thesis work.
Chapter 3: Study of various related work helpful for our thesis are joined in this chapter.
Chapter 4: The method and implementation of the thesis are discussed.
Chapter 5: The calculated results and graphs are added in this chapter and the graphs are analyzed.
Chapter 6: The analysis of results and the answers relating the research questions are discussed in this chapter.
Chapter 7: The conclusions and future work are discussed here.

1.14 Research Questions

1. What is the impact on energy consumption of virtual data centers when compared with physical data centers?

2. Will KVM perform better than Xen hypervisor in server virtualization in terms of energy consumption while varying number of virtual machines?

Chapter 2

Conceptual Background

2.1 Data Center Efficiency Metrics

As this report is a study on the energy consumption performance of the data centers, here we will discuss about some of the efficiency measurements of a data center. fig: 2.1 shows us the demands and constraints upon a data center operator.

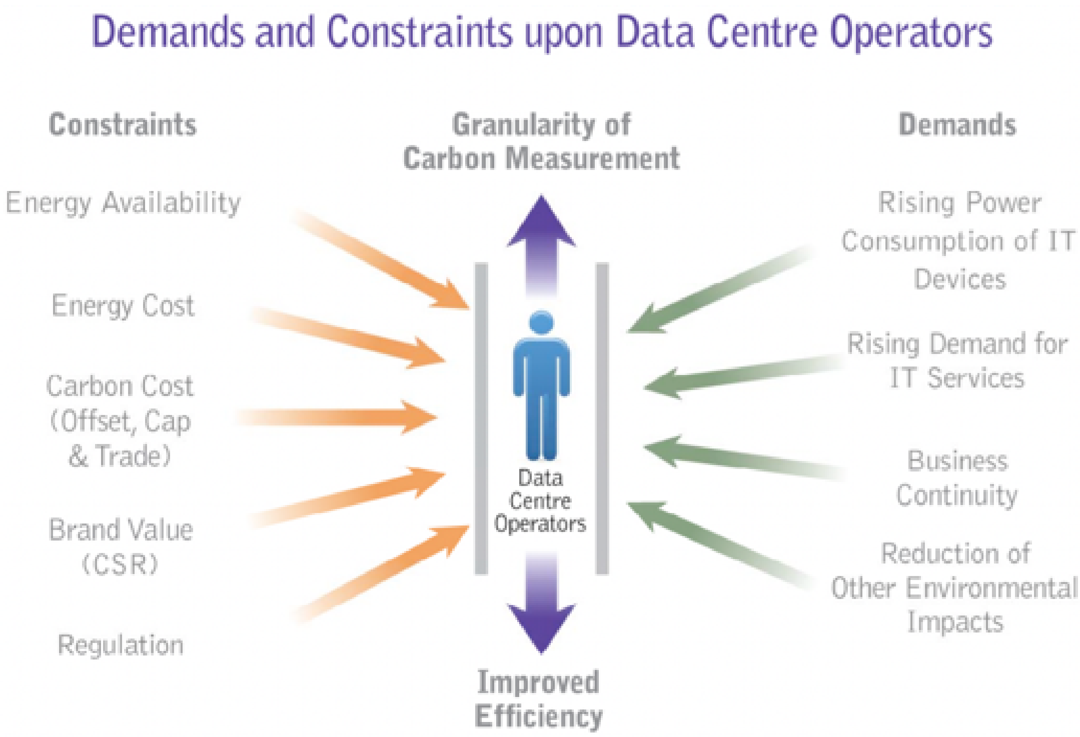

Figure 2.1: Demands and Constraints

2.1.1 Power Usage Efficiency (PUE)

Since the Green Grid first published its white paper on Power Usage Effectiveness (PUE) and Data Center Infrastructure Efficiency (DCIE), the term PUE and DCIE has become the predominant metric for discussing about data center efficiency.

$$PUE = \frac{TotalFacilityEnergy}{ITEquipmentEnergy}$$

Total facility power is the power measured at the utility meter. The IT equipment power includes all the actual load of IT equipment such as work stations, servers, storage, switches, printers and other service delivery equipment. PUE is determined on the scale from 1 to 4 where 1 is most efficient and 4 being most inefficient.

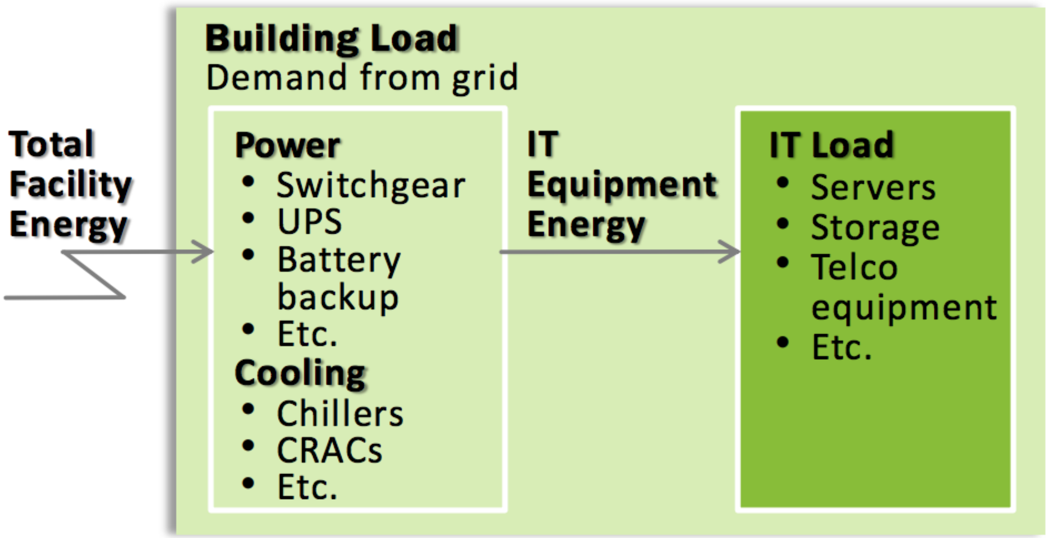

Figure 2.2: For calculating Energy Metrics

2.1.2 Data Center Infrastructure Efficiency (DCiE)

Data Center Infrastructure Efficiency most commonly known as DCIE, is another popular metric used to benchmark energy efficiency of a data center and is most commonly used by many data centers in the world.

$$DCiE = \frac{ITEquipmentEnergy}{TotalFacilityEnergy} * 100\%$$

Total facility power is the power measured at the utility meter. The IT equipment power includes all the actual load of IT equipment such as work stations, servers, storage, switches, printers and other service delivery equipment.

2.1.3 Carbon Footprint (C)

The total amount of co2 emitted from a data center is known as the carbon footprint of that data center.

$$C = \frac{E(N)}{T}$$

In this equation C stands for carbon footprint, E is the electrical usage in kWh per year, N is the annual co 2 emissions and T stands for metric tons. From this equation, the amount of co2 emitted annually can be calculated and necessary measures to control these emissions can be implemented.

2.2 Functioning of a Data Center

2.2.1 Power Supply Unit

[3] The data center is connected to two separate grid sectors operated by the local utility company. If one sector were to fail, then the second one will ensure that power is still supplied.

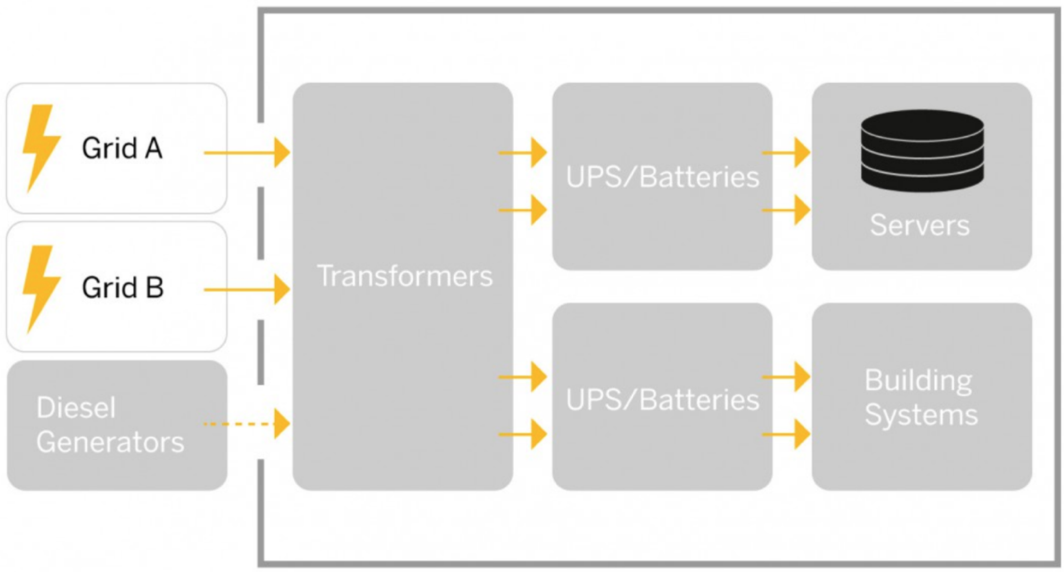

Figure 2.3: Power Supply Unit

In addition to this, Data Centers are equipped with diesel generators to support electricity demand in an emergency. These generators take an outage in just one of the external grid to automatically actuate the generators. Both local utility company and the diesel generators deliver electricity with a voltage of 20kV, which is then step downed to 220 or 380Volts in the data center.
Within the data center, block batteries ensure that all operating applications can run for15 minutes. This backup system makes it possible to provide power from the time a utility company experiences a total blackout to the time that the diesel generators start up.
The uninterruptible power supply (UPS) also ensures that the quality remains constant. It compensates for voltage and frequency fluctuations and thereby effectively protects sensitive computer electronic components and systems.
A redundantly designed power supply system is another feature of the data center. This enables one to perform repairs on one network, for example, without having to turn off servers, databases, or electrical equipment.

2.2.2 Cooling Unit

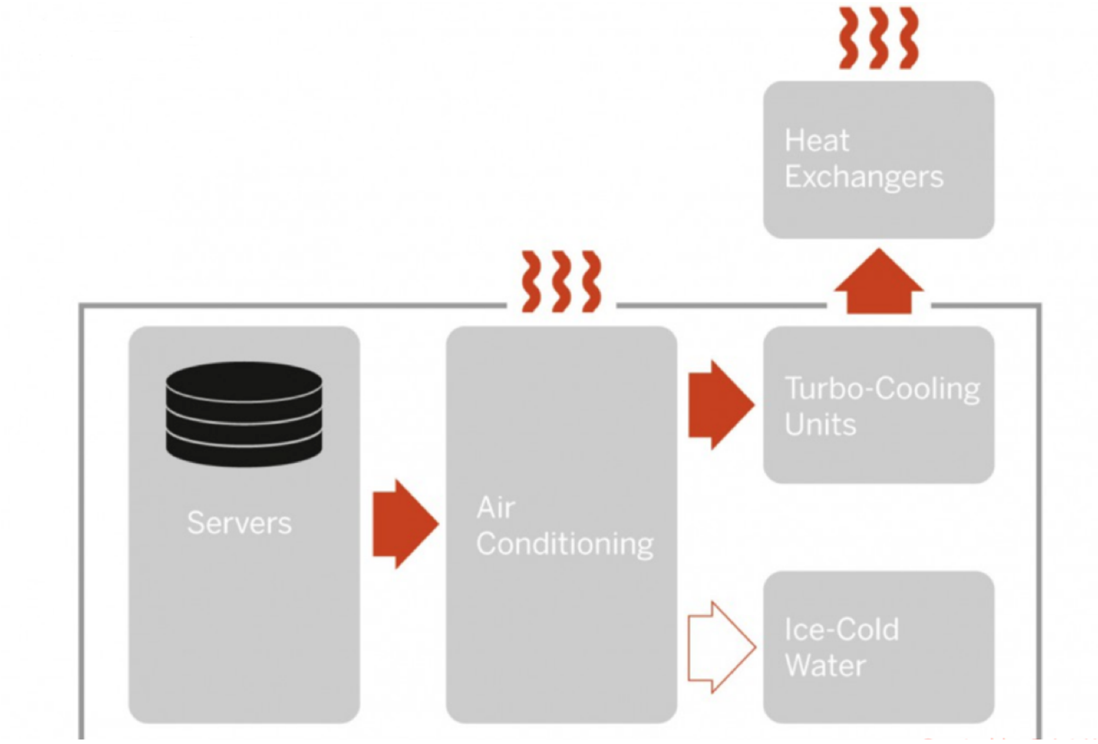

Figure 2.4: Cooling Unit

All electronic components especially the processors generate heat when they are in operation. There is a chance that the efficiency of the component may

decrease or component may fail if this heat is not dissipated. So, cooling a data center is very much essential, and because of the concentrated computing power the costs to do so are also considerable.
For this reason, servers are installed in racks, which basically resemble specially standardized shelves. They are laid out so that two rows of racks face each other, thereby creating an aisle from which the front side of the server is accessible. The aisles are covered above and closed off at the ends by doors. Cool air set to a temperature of 24 to 26°C is blown in through holes in the floor, flows through the racks, and dissipates the heat emitted by the servers.
Generally, a server room will contain several such “enclosed” server rows. The warm air from the server room is removed by the air-conditioning system. Yet, even the air-conditioning system has to dissipate the heat. When the outside temperature is below 12 to 13°C, outside air can be used to effectively cool the heat absorbed by the air-conditioning systems.
If outside temperatures are above 26°C, the heat exchangers are sprinkled with water in order to make heat dissipation more effective through evaporative cooling. The municipal water supply system provides a reserve supply of water in this case and acts as a failsafe.

2.2.3 Controlled Access

1. Single-person access and mantrap systems allow authorized individuals to enter the data center.
2. Technicians can access rooms for facilities maintenance via separate entrances.
3. Maintenance staff has further authorization to enter the server rooms.

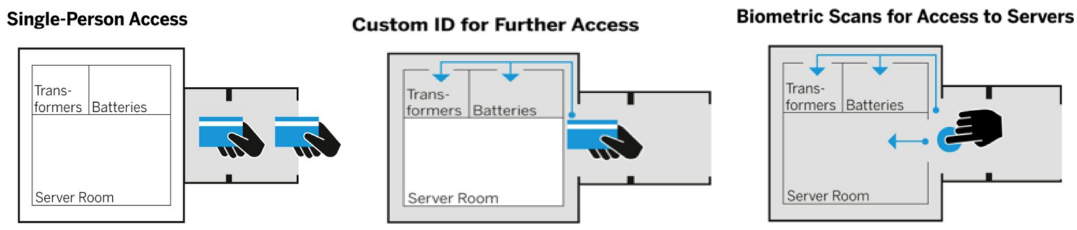

Figure 2.5: Controlled Access

2.2.4 Control Center

Control stations in the data centers serve as command centers for both IT maintenance and building security. All the information about IT infrastructure are monitored here and any difference from standard operation is promptly reported.

2.2.5 Safety Measurements

During fire mishaps water, extinguishing foam or powder fire suppression systems can cause more damage in a data center than a charred cable. So, special extinguishing gas called INERGEN which displaces oxygen in the air and smothers fire is used. This gas is harmless to people and the equipment.
Closed Circuit Cameras are installed around the building and can be monitored from the control center.

2.3 Simulation Tools

2.3.1 Cloud Sim

CloudSim is a library for the simulation of cloudscenarios. It provides essential classes for describing data centres, computational resources, virtual machines, applications, users, and policies for the management of various parts of the system such as scheduling and provisioning.

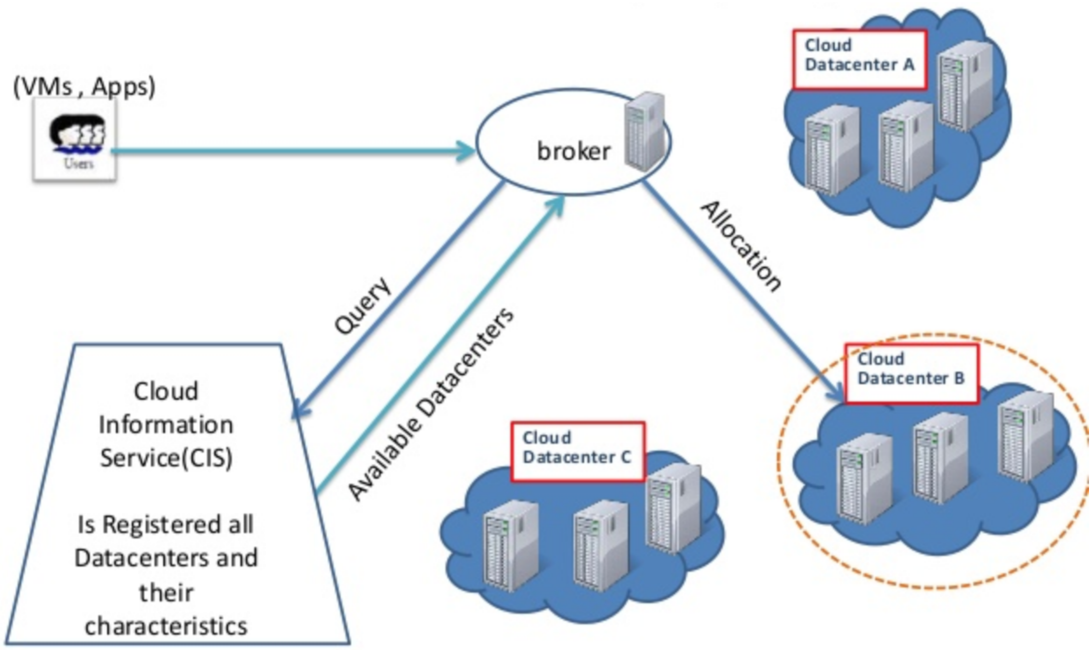

Figure 2.6: Cloud Sim Architecture

2.3.2 Green Cloud

Green cloud is a buzzword that refers to the potential environmental benefits that information technology services delivered over the Internet can offer society.

The term combines the words green – meaning environmentally friendly – and cloud, the traditional symbol for the Internet and the shortened name for a type of service delivery model known as Cloud Computing.

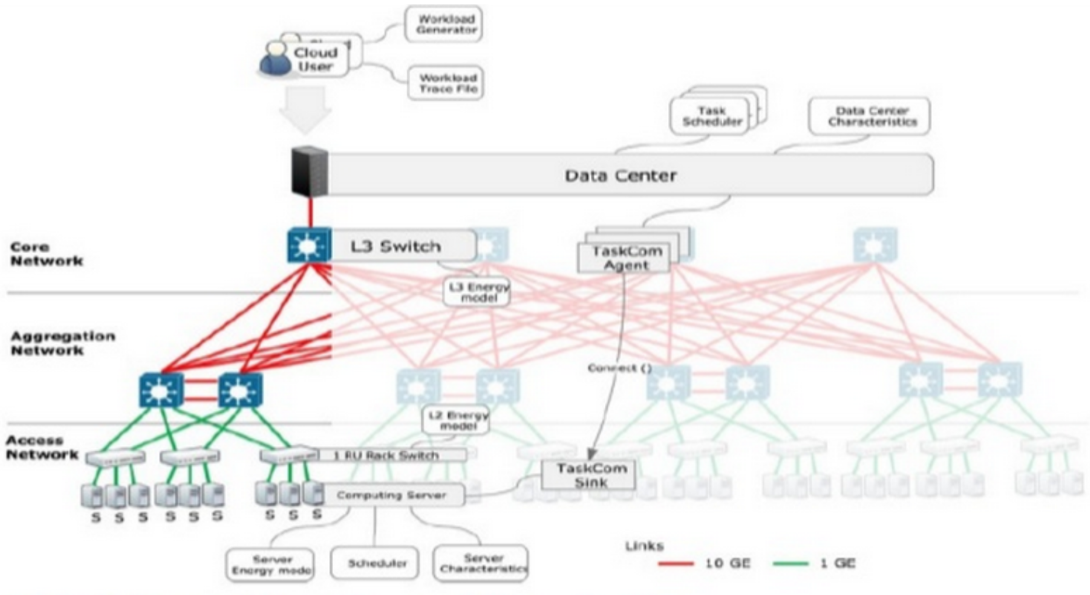

Figure 2.7: Green Cloud Architecture

Chapter 3

Related Work

Boccaletti et al. [15] discussed about how Greenhouse gas emissions are associated with making and powering the world's computers and telecom networks which are growing fast creating big carbon footprint especially in developing nations. They also suggest in their research that there are opportunities to use certain technologies to make the world economy more energy and carbon efficient.

Koomey at al. [21] have estimated total electricity used by servers in the U.S. and world wide by combining measured data and estimates of power used by the most popular servers with data on the server installed base. These estimates are based on more detailed data than are previous assessments, and they will be of use to policy makers and businesses attempting to make sense of recent trends in this industry. It does not attempt to estimate the effect of structural changes in the economy enabled by increased use of IT.

Rao et al. [22] discussed about the challenges and opportunities for researchers to explore the issues to quantify the electricity savings and methods for improving the energy efficiency. In this paper, they systematically explored about the opportunities and adoptable methods for achieving the energy efficiency in data centers. They explained about the study on data center virtualization and scaled it to realize significant savings for an enterprise.

Brown et al. [16] The United States Environmental Protection Agency(EPA) developed this report, which assesses current trends in energy use and energy costs of data centers and servers in the U.S. and emerging opportunities for improved energy efficiency. It provides the opportunities for reducing those costs through improved efficiency. It also makes recommendations for pursuing these energy-efficiency opportunities broadly across the country through the use of information and incentive-based programs.

Umeno et al. [26] evaluated the virtualization overhead by measuring the performance of the transaction systems running database servers on VMs. The objectives are to evaluate the VM overhead and the virtualization overhead for heavy database servers running in VMs. They have found that the total CPU utilization in the two VMs running database servers increases to 2.5 times than that in the native system with the same hardware resources.

Uddin et al. [25] suggests about implementing server virtualization to achieve energy efficient data centers. The proposed technique increases the utilization ratio of underutilized servers up to 50%, saving a huge amount of power and at the same time reducing the emission of greenhouse gases.

Avelar et al. [12] Green Grid proposed the use of Power Usage Effectiveness (PUE) and its reciprocal, Data Center Infrastructure Efficiency (DCiE) metrics, which enable data center operators to quickly estimate the energy efficiency of their data centers, compare the results against other data centers, and determine if any energy efficiency improvements need to be made.

Avgerinou et al. [13] The aim of this paper is to evaluate, analyse and present the current trends in energy consumption and efficiency in data centres in the European Union based on the data submitted by companies participating in the European Code of Conduct for Data Centre Energy Efficiency programme, a voluntary initiative created in 2008. The analysis shows that the average PUE of the facilities participating in the programme is declining year after year. This confirms that voluntary approaches could be effective in addressing climate and energy issue.

Meeun et al. [24] an average server environment, 30% of the servers are only consuming energy, without being properly utilized. This paper proposes a five step model using an emerging technology called virtualization to achieve energy efficient data centers. The proposed model helps Data Center managers to properly implement virtualization technology in their data centers to make them green and energy efficient so as to ensure that IT infrastructure contributes as little as possible to the emission of greenhouse gases.

Chapter 4

Method

4.1 CloudSim

Now-a-days it is important to measure the cost of large-scale cloud infrastructure as it is hard to manage large-scale experiments on a real infrastructure, especially when it is essential to repeat the experiments with the same conditions. So, simulations (which is the imitation of the real model) have been selected as a way to evaluate the proposed model.

CloudSim [11] [1], a platform in which cloud strategies can be tested in a controlled and reproducible manner. Therefore, simulation frameworks such as CloudSim are important, as they allow for the evaluation of the performance of resource provisioning and application scheduling techniques under different infrastructure configurations.

Fig shows the layered architecture of cloudsim. The top-most layer in the simulation stack is the user code that exposes configuration related functionalities for hosts (number of machines, their specifications and so on), applications (number of tasks and their requirements), vms, number of users and their application types, and broker scheduling policies.

A cloud application developer can generate:

- Mix of user request distributions and application configurations.
- Cloud availability scenarios at this layer and perform robust tests based on the custom configurations already supported within the cloudsim.

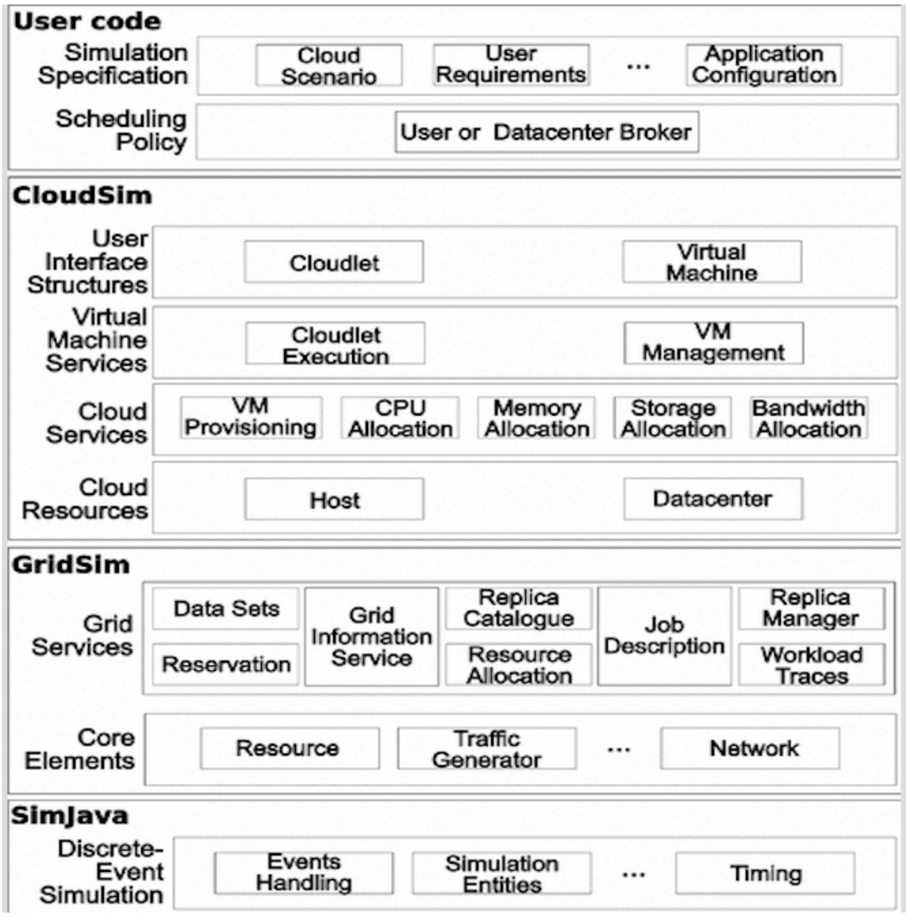

Figure 4.1: layered architecture of cloudsim

4.2 Implementation

We first installed the cloudsim for the purpose of collecting the energy consumption of server in a data center. We consider a host in a data center and a customer with one virtual machine. From fig we can see that one vm with round robin policy is considered. We first calculate the power consumption for this physical server. Then we start increasing the number of virtual machines and power consumption is calculated for both KVM [5] installed vms and Xen [8] installed vms.

CloudSim configuration:

Server Architecture: x86; Operating System: Linux

Hypervisors: Xen, KVM; Broker Policy: Round Robin

Motivation for Round Robin Policy:

We have many broker policies like round robin policy, service proximity policy, performance optimization policy, dynamic configuration policy and more. We chose round robin policy because from various studies it shows better performance when compared with other policies in terms of overall server response time and data center processing time.

Figure 4.2: Simulating with one vm

Figure 4.3: Simulating with three vms

Figure 4.4: Simulating with five vms

Figure 4.5: Simulating with ten vms

Figure 4.6: Simulating with fifteen vms

Chapter 5

Results

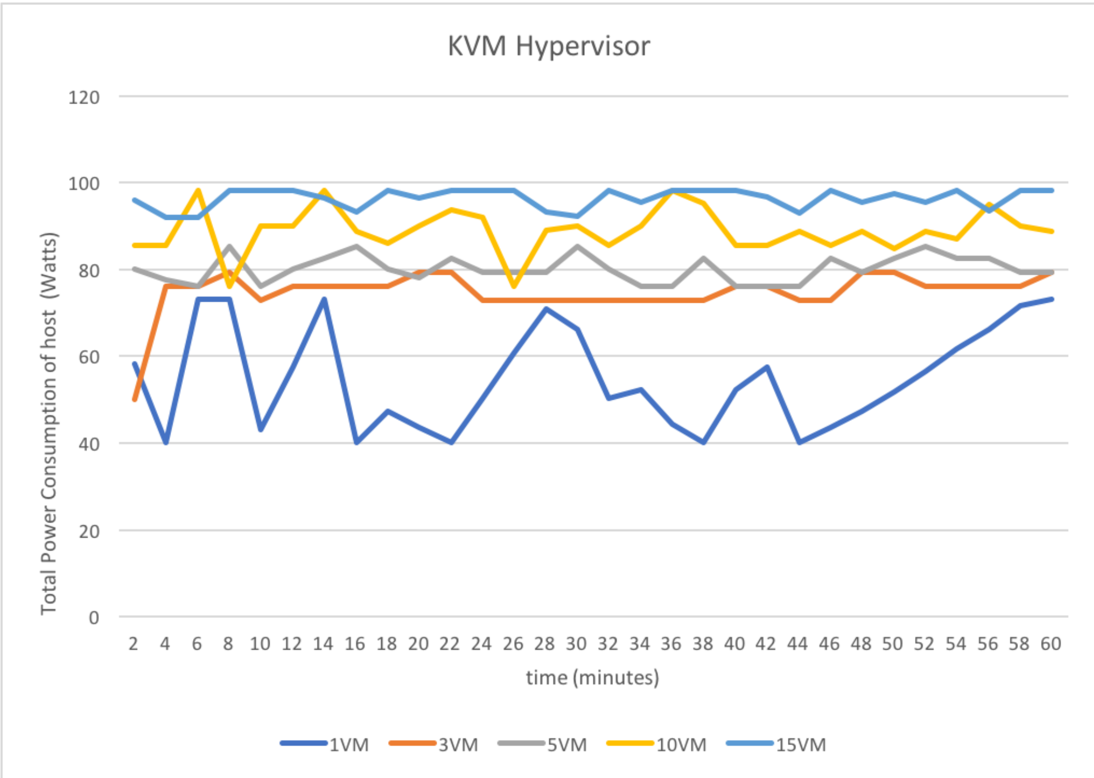

Figure 5.1: KVM with varying number of virtual machines

Figure 5.1 shows us the power consumption of a server with KVM hypervisor and varying number of virtual machines. With increase in number of vms we can observe a small increase in power consumption of server. The different lines represent various number of vm's. The graph having 1vm shows heavy fluctuations due to the change in cloudlets. During the hand off process of cloudlets, resource utilization will be minimum resulting in the less power consumption at that instance. So, we can observe heavy fluctuations in graph holding 1vm.

X axis represents the time in minutes and Y axis represents the total power consumption of the host with Xen hypervisors. The average readings for 60min of time interval are calculated and the values are tabulated and graphs are plotted.

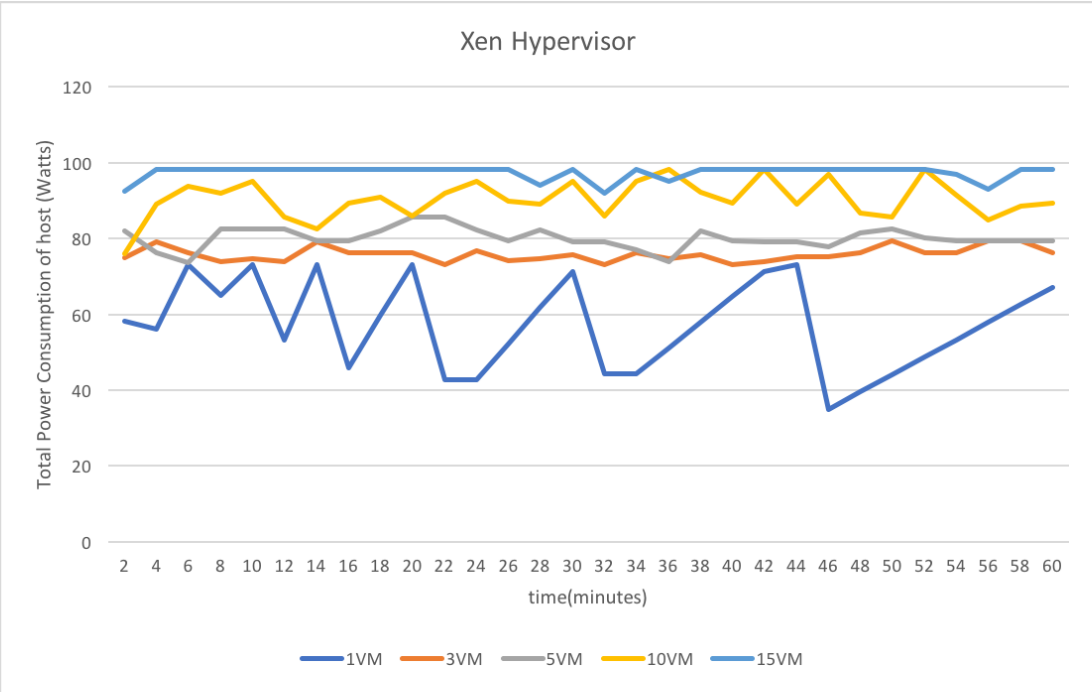

Figure 5.2: Xen with varying number of virtual machines

Figure 5.2 shows us the power consumption of a server with Xen hypervisor and varying number of virtual machines. With increase in number of vms we can observe a small increase in power consumption of server. The graph having 1vm shows heavy fluctuations due to the change in cloudlets. During the hand off process of cloudlets, resource utilization will be minimum resulting in the less power consumption at that instance. So, we can observe heavy fluctuations in graph holding 1vm.

X axis represents the time in minutes and Y axis represents the total power consumption of the host with Xen hypervisors. The average readings for 60min of time interval interval are calculated and the values are tabulated and graphs are plotted.

The tabulated values for total power consumption of host for Xen and KVM hypervisors with varying number of vms are shown here "Results table".

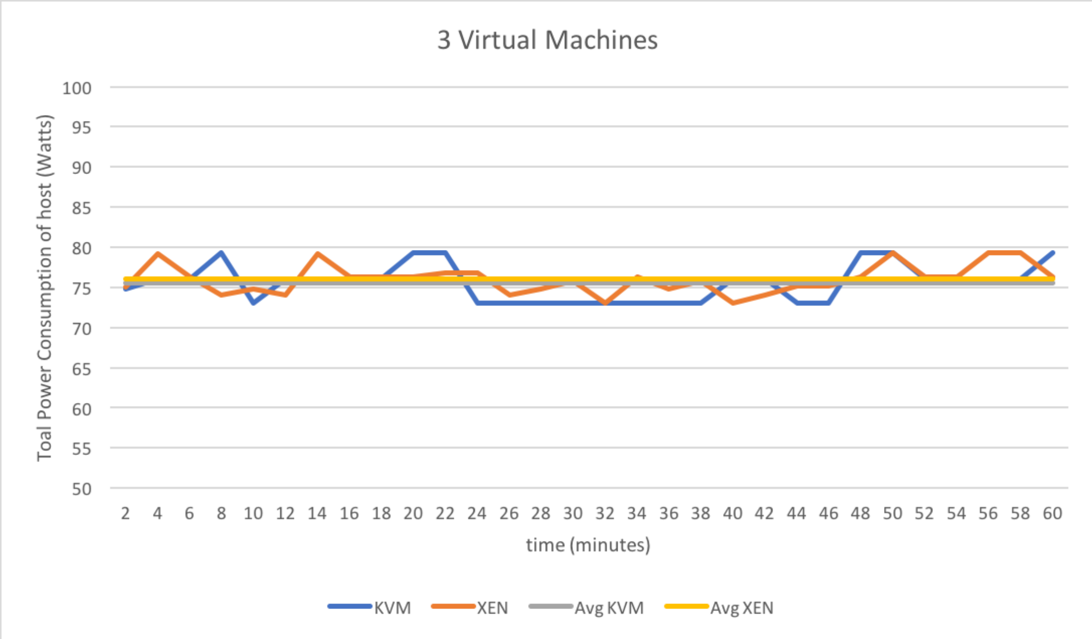

Figure 5.3: Comparing KVM and Xen when 3 virtual machines are deployed

Figure 5.3 shows the comparison in total power consumption of the host between KVM and Xen when 3 virtual machines are deployed on a server. The average power consumption of host with KVM is 75.5, compared to 76.1 when we consider Xen hypervisor. The average power consumption value is also considered and the graph is plotted for both KVM and Xen for better comparision.

Figure 5.4 shows the comparison in total power consumption of the host between KVM and Xen when 5 virtual machines are deployed on a server. The average power consumption of host with KVM is 79.8, compared to 80.8 when we consider Xen hypervisor. The average power consumption value is also considered and the graph is plotted for both KVM and Xen for better comparision.

Figure 5.5 shows the comparison in total power consumption of the host between KVM and Xen when 10 virtual machines are deployed on a server. The average power consumption of host with KVM is 88.8, compared to 90.3 when we consider Xen hypervisor. The average power consumption value is also considered and the graph is plotted for both KVM and Xen for better comparision.

Figure 5.6 shows the comparison in total power consumption of the host between KVM and Xen when 10 virtual machines are deployed on a server. The average power consumption of host with KVM is 88.8, compared to 90.3 when we consider Xen hypervisor. The average power consumption value is also considered and the graph is plotted for both KVM and Xen for better comparision.

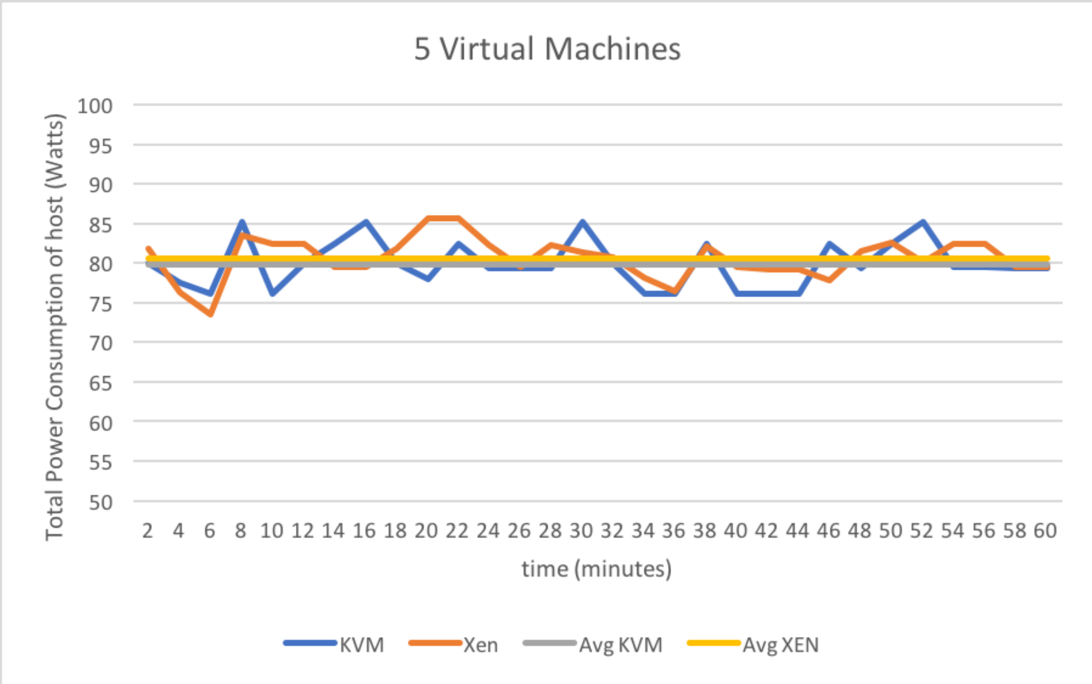

Figure 5.4: Comparing KVM and Xen when 5 virtual machines are deployed

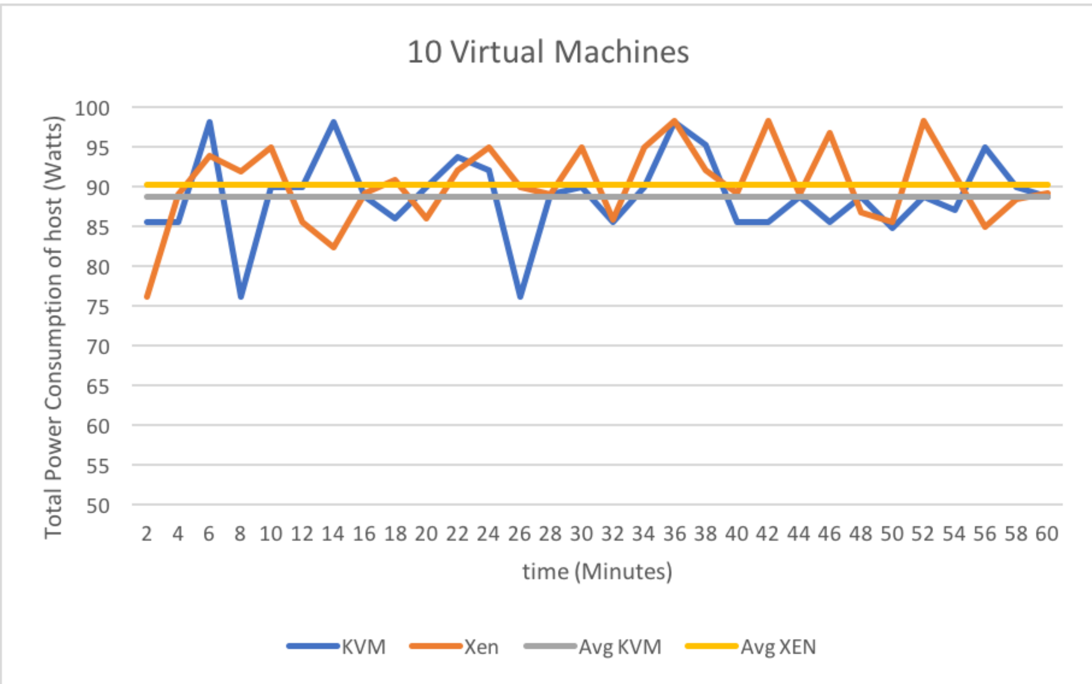

Figure 5.5: Comparing KVM and Xen when 10 virtual machines are deployed

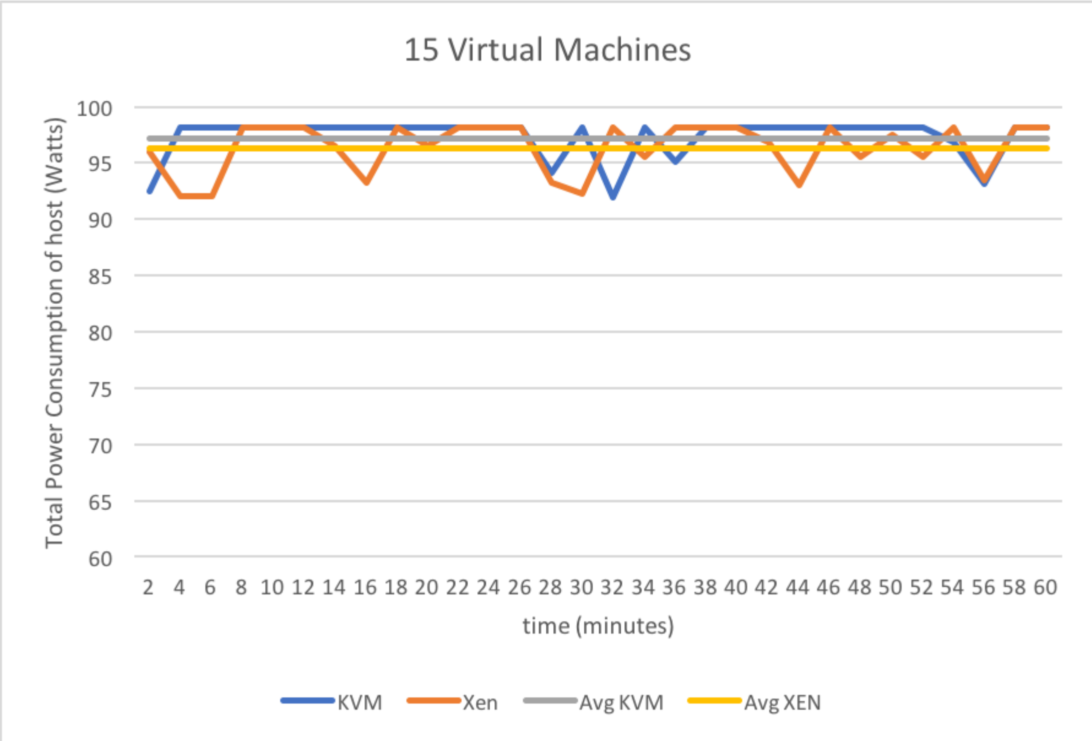

Figure 5.6: Comparing KVM and Xen when 15 virtual machines are deployed

Chapter 6

Analysis and Discussion

6.1 What is the impact on energy consumption of virtual data centers when compared with physical data centers?

As discussed earlier, energy consumption plays a key role in a data center. So our aim is to efficient use of resources by deploying more number of virtual machines on the same hardware with a minor changes in energy consumption.
From the above graphs 5.1 and 5.2 we can observe the power consumption of a server when KVM and Xen hypervisors are installed and varying the number of virtual machines.
First we considered single virtual machine in a server and calculated the power consumption using cloudsim. From graph we can observe uneven increase and decrease in power consumption which is due to drop in cloudlets. There on we increased the number of vms and corresponding power consumption values are plotted. We can observe that there is small increase in power consumption with increase in number of vms.
If a server is being used instead of the 15 virtual machines more power is consumed which is equivalent to power consumed by 15 servers. Here we can see that deploying 15vms on a single server has consumed power which not even greater than the power consumed by two servers. Hence server consolidation method is very much useful to reduce the energy consumption of a server.

6.2 Will KVM perform better than Xen hypervisor in server virtualization in terms of energy consumption while varying number of virtual machines?

Here we considered comparison between KVM and Xen hypervisors. First we deployed 3 virtual machines on both KVM and Xen and their corresponding power consumption is calculate using cloudsim. There after we repeat the same experiment with increasing the number of vms from 5 to 10 and to 15 and their values are plotted.
From graphs 5.3 5.4 5.5 5.6 we observed that KVM has less power consumption when compared with Xen hypervisors. From 5.6 we can observe that the power consumption of a server with KVM hypervisor is more stable than the server with Xen hypervisor.

Chapter 7

Conclusions and Future Work

From the above results we can conclude that power consumption of the server does not vary much even after deploying VMs. So, we can decrease the use of additional physical servers which decreases the overall power consumption.

From graph we can observe that the average power consumed by the server when Xen is deployed is greater than when KVM is deployed. Both of them being bare metal hypervisors KVM consumes less power than Xen. So, we can conclude that KVM shows better performance than Xen in terms of power consumption.

We have done in this paper on server virtualization. It can further be extended to storage virtualization, network function virtualization which helps in data center virtualization. With this we can furthermore decrease the energy consumption of a data center.

We can also calculate the energy metrics of data center before and after virtualization and find energy efficiency of data centers. We can also find the excess energy before virtualization and suggest ways to effective use of this excess energy.

References

[1] The CLOUDS Lab: Flagship Projects - Gridbus and Cloudbus. http://cloudbus.org/cloudsim/.

[2] Electric Grid. https://en.wikipedia.org/wiki/Electrical_grid.

[3] How a Data Center Works. http://www.sapdatacenter.com/article/data_center_functionality/.

[4] Hypervisor types. https://www.vgyan.in/type-1-and-type-2-hypervisor/.

[5] KVM. https://www.linux-kvm.org/page/Main_Page.

[6] Virtual Data Center. http://study.com/academy/lesson/what-is-a-virtual-data-center.html.

[7] Virtualization. https://www.redhat.com/en/topics/virtualization/what-is-virtualization.

[8] VS16: Video Spotlight with Xen Project's Lars Kurth. https://www.xenproject.org/.

[9] Here's How Much Energy All US Data Centers Consume, June 2016.

[10] Hypervisor, November 2017. https://en.wikipedia.org/w/index.php?title=Hypervisor.

[11] cloudsim: CloudSim: A Framework For Modeling And Simulation Of Cloud Computing Infrastructures And Services, January 2018. https://github.com/Cloudslab/cloudsim.

[12] Victor Avelar, Dan Azevedo, Alan French, and Emerson Network Power. PUE: a comprehensive examination of the metric. *White paper*, 49, 2012.

[13] Maria Avgerinou, Paolo Bertoldi, and Luca Castellazzi. Trends in Data Centre Energy Consumption under the European Code of Conduct for Data Centre Energy Efficiency. *Energies*, 10(10):1470, September 2017.

[14] Kashif Bilal. Green Data Center Networks: Challenges and Opportunities. December 2013.

[15] Giulio Boccaletti, Markus Löffler, and Jeremy M. Oppenheim. How IT can cut carbon emissions. *McKinsey Quarterly*, 37:37–41, 2008.

[16] Richard E. Brown, Richard Brown, Eric Masanet, Bruce Nordman, Bill Tschudi, Arman Shehabi, John Stanley, Jonathan Koomey, Dale Sartor, and Peter Chan. Report to congress on server and data center energy efficiency: Public law 109-431. Technical report, Ernest Orlando Lawrence Berkeley National Laboratory, Berkeley, CA (US), 2007.

[17] Jeff Clark. Basics of a UPS System.

[18] DatacenterDynamics. Utility data center architectures part of smart grid project.

[19] Juho Kaivosoja. Green Data Centers: Evaluating Efficiency. October 2016.

[20] Kelser. Types of Virtualization. https://inbound.kelsercorp.com/blog/the-7-types-of-virtualization.

[21] Jonathan G. Koomey. *Estimating total power consumption by servers in the US and the world*. February, 2007.

[22] K. Thirupathi Rao, P. Sai Kiran, and L. Siva Shanker Reddy. Energy efficiency in datacenters through virtualization: A case study. *Global Journal of Computer Science and Technology*, 2010.

[23] Arman Shehabi, Sarah Smith, Dale Sartor, Richard Brown, Magnus Herrlin, Jonathan Koomey, Eric Masanet, Nathaniel Horner, Inês Azevedo, and William Lintner. United States Data Center Energy Usage Report, 2016.

[24] Mueen Uddin and Azizah Abdul Rahman. Virtualization implementation model for cost effective & efficient data centers. *arXiv preprint arXiv:1206.0988*, 2012.

[25] Mueen Uddin, Asadullah Shah, Adamu Abubakar, and Imran Adeleke. Implementation of Server virtualization to Build Energy Efficient Data Centers. *Journal of Power Technologies*, 94(2):85, 2014.

[26] Hidenori Umeno, Maria LC Parayno, Koudai Teramoto, Masato Kawano, Hiraku Inamasu, Shuuhei Enoki, Masato Kiyama, Tomoo Aoyama, and Takafumi Fukunaga. Performance evaluation on server consolidation using virtual machines. In *SICE-ICASE, 2006. International Joint Conference*, pages 2730–2734. IEEE, 2006.