Economic and environmental sustainability for cloud computing systems : strategies and algorithms

Paul, Debdeep

2016

Paul, D. (2016). Economic and environmental sustainability for cloud computing systems : strategies and algorithms. Doctoral thesis, Nanyang Technological University, Singapore.

<https://hdl.handle.net/10356/69346>

<https://doi.org/10.32657/10356/69346>

Thesis

PhD

**Economic and environmental sustainability for cloud
computing systems: Strategies and Algorithms**

**DEBDEEP PAUL
G1101649F**

SCHOOL OF ELECTRICAL AND ELECTRONIC ENGINEERING

Jan, 2016

Acknowledgements

This research work received both financial and technical supports from the School of Electrical & Electronic Engineering, Nanyang Technological University. First and foremost, I would like to express my deepest gratitude to my supervisor, Professor Zhong Wende for his professional guidance, patience and encouragement throughout my PhD life. Next, I would like to take this opportunity to thank Professor Sanjay Bose from Indian Institute of Technology, for his precious inputs and valuable suggestions. They are the ones who taught me the art of writing scientific articles and helped me navigating through the process of publication. I would also like to thank the laboratory technicians at Communication Lab-III.

My friends, peers, juniors and seniors owe me big thanks for supporting me in various ways.

Finally, I would like to thank my parents and sister whose love and support have always been the greatest inspiration for me. I would like to dedicate this thesis to my ‘Dadu’ who taught me the beauty of Mathematics, during my childhood days.

Table of Contents

Acknowledgements.....	2
Summary.....	6
List of Figures.....	8
List of Tables.....	11
List of Acronyms.....	12
1. Introduction.....	14
1.1 Background and Motivation.....	14
1.2 Objectives	24
1.3 Major Challenges.....	26
1.3.1 Economic Sustainability.....	26
1.3.2 Environmental Sustainability.....	27
1.4 Novelty and Contributions.....	28
1.5 Organization of this Thesis.....	31
 2. Literature Review.....	 34
2.1 Introduction.....	34
2.2 Pricing for Cloud Services.....	35
2.3 Energy efficiency in Cloud Computing.....	39
2.3.1 Energy efficiency in servers.....	40
2.3.2 Energy efficiency in a single data center.....	40
2.3.2.1 Scheduling.....	44
2.3.3 Energy efficiency in the case of multiple data centers owned by a single organization	46
2.4 Demand Response (DR) in data centers.....	47

2.5 Cost Volatility Control.....	49
3. Energy Efficient Cloud Service Pricing: A Two Timescale Optimization Approach	52
3.1 Introduction.....	52
3.2 Cloud Service Market Model with Multiple CSPs.....	57
3.3 Distribution of Service Demand between Multiple CSPs: Slow Timescale Optimization Procedure (<i>STOP</i>).....	60
3.4 Energy Aware Scheduling for Maximizing Renewable Energy Integration Fast Timescale Optimization Procedure (<i>FTOP</i>).....	71
3.5 Performance Evaluation.....	76
3.5.1 Benchmark strategy for comparing fairness of revenue and profit.....	78
3.5.2 Results and Discussions	80
3.6 Summary.....	95
4. Demand Response in Data Centers Through Energy Efficient Scheduling and Simple Incentivization.....	97
4.1 Introduction.....	97
4.2 System Model.....	100
4.2.1 Model of the tiered architecture of a data center.....	100
4.3 Problem formulation.....	104
4.3.1 Minimization of operational cost (<i>Min-OpC</i> Procedure).....	106
4.3.2 Joint minimization of capital and operational cost (<i>Min-CapOpC</i> Procedure).....	115
4.4 Heuristic algorithm to improve on cooling system energy efficiency.....	117
4.5 Incentivazition for Deadline Deferral.....	118
4.6 Simulation Setup and Results.....	122
4.6.1 Model parameters.....	122

4.6.2 Comparison between <i>Min-OpC</i> , <i>Min-CapOpC</i> and their respective offline counterparts.....	127
4.6.3 Improvements after applying <i>IDD</i>	132
4.7 Summary.....	135
5. Energy Efficiency Aware Load Distribution and Electricity Cost Volatility Control for Cloud Service Providers.....	137
5.1 Introduction.....	137
5.2 System Model of a CSP with Multiple Distributed Data Centers.....	140
5.2.1 Optimization with Linear Switching Cost (<i>LB-L</i> Procedure).....	144
5.2.2 Optimization Based on Second Moments of the Switching Cost (<i>LB-SM</i> Procedure).....	151
5.3 Electricity Cost Variability Control (<i>ECVC</i>).....	153
5.4 Performance Evaluation.....	160
5.4.1 Locations and Interconnection Topology.....	160
5.4.2 Electricity Price, Renewable Energy Generation and Job Request Distribution over Time.....	162
5.4.3 Implementation of <i>LB-L</i> and <i>LB-SM</i>	164
5.5 Results.....	165
5.6 Summary.....	176
6. Conclusions and Future Research	178
6.1 Conclusions.....	178
6.2 Future Research.....	181
Author's Publications	184
Bibliography.....	185

Summary

With the enormous growth of cloud computing technology the data centers, which serve as the primary physical resources to provide various services, are also growing in size and complexity at a tremendous pace. This thesis aims to contribute towards the two important issues faced by cloud service operators; namely, i) economic sustainability and ii) environmental sustainability. Firstly, the scenario of computing the optimal service prices for a cloud service provider operating in a competitive cloud service market, with multiple cloud service providers, is considered. In this paradigm, all the cloud service providers operating in such markets compete among themselves to attract customers. The market consists of price sensitive, rational customers. As a result, eventually all the customers will move to that cloud service provider, which offers better service quality at a lower cost. Hence, each cloud service provider needs to strategically determine the optimal set of actions (including the service quality offered and the price charged for the same). These actions have to consider the set of actions chosen by the other competitors, operating in the same market. This will ensure a healthy competition among the cloud service providers and will eliminate the possibility of formation of a monopolistic market. The contribution of this thesis is that, through a bi-level optimization strategy, the environmental sustainability issue (through renewable energy integration) is connected to the long term economic sustainability issue, which deals with the revenue earned and the market share in terms of number of customer requests served. Pricing decisions are taken in the slower timescale and energy efficient job scheduling is performed in the faster timescale. To enable the above mentioned bi-level optimization strategy a modeling framework is provided to determine the appropriate pricing strategy based on the preferences strategically chosen by each of the cloud

service providers. Through this framework the environmental sustainability issue is addressed by appropriately pricing the renewable energy. All the proposed algorithms are evaluated with real life data traces for electricity price, workload, renewable generation and all the other model parameters. Then, the case of a single data center is considered where different applications are deployed in a tiered architecture, which is usually the case of real systems. Two job scheduling algorithms are developed and a rigorous investigation is carried out to understand their performance. Under this framework, thereafter a novel incentivization strategy to achieve more aggressive renewable energy integration and electricity cost reduction is designed. The proposed algorithms are evaluated with real life data traces for electricity price, workload, renewable generation and all the other model parameters. Then the case of a cloud computing system with multiple data centers is considered. A rigorous investigation is carried out on the load distribution strategy for a cloud service provider owning multiple geographically distributed data centers, which is the case for a typical mid to large cloud service provider. In this case two algorithms for optimal load distribution are presented and compared. In order to hedge the risk associated with the volatility of average cost of electricity a simple derivate, namely forward contract, is designed and implemented. In order to address the economic sustainability issue a few relevant topics in economics (such as price sensitivity, market equilibria, competition etc.) are leveraged. To address the environmental sustainability issue, a promising technique prevalent in the smart grid arena namely, demand response, is leveraged. Since the cloud computing technology is still in a nascent stage, the strategies and algorithms proposed in this thesis will indeed make some meaningful contributions in shaping the future of cloud computing.

List of Figures

Figure 1.1: Energy Spaghetti chart 2014 (published by Lawrence Livermore National Laboratory).....	20
Figure 1.2: The objectives of this thesis.....	25
Figure 3.1: Interaction between the three major entities in the cloud service markets.....	58
Figure 3.2: Profit earned by the CSPs before and after implementation of CR: (a) before, (b) after.....	81
Figure 3.3: Revenue earned by the CSPs before and after implementation of CR: (a) before, (b) after.....	82
Figure 3.4: Comparison of total profit and revenue earned by individual CSPs before and after implementation of CR: (a) Profit, (b) Revenue.....	84
Figure 3.5: Average profit per request earned by the CSPs and amount of demand served by the CSPs before and after implementation of CR: (a) Before, (b) After.....	86
Figure 3.6: Profit earned by the CSPs and the amount of demand (no. of requests serviced) before and after implementation of CR: (a) before, (b) after.....	89
Figure 3.7: Revenue earned by the CSPs and the amount of demand (no. of requests serviced) before and after implementation of CR: (a) before, (b) after.....	91

Figure 3.8: Cumulative electricity cost incurred and cumulative renewable energy used by each CSP under our proposed technique for job scheduling <i>FTOP</i> and a benchmark scheme: a) Cumulative electricity cost incurred, b) Cumulative renewable energy used.....	93
Figure 3.9: Cumulative number of server switchings and comparison with a baseline strategy: a) Cumulative switching from OFF to ON, b) Cumulative switching from ON to OFF.....	94
Figure 4.1: Tiered architecture of data center.....	99
Figure 4.2: The effect of scaling factor of electricity cost in terms of: (a) Cumulative electricity cost, (b) Cumulative green energy used and (c) Cumulative switchings.....	124
Figure 4.3: Sensitivity of the proposed optimization procedures with respect to the window size of the MPC. (a) Cumulative electricity cost, (b) Cumulative renewable energy used, (c) Cumulative no. of switchings.....	126
Figure 4.4: Workload trace: (a) Non-deferrable workload, (b) Deferrable workload, (c) Distribution of deadlines of deferrable loads.....	127
Figure 4.5: Comparison of no. of servers provisioned under online and offline cases under different schemes.....	129
Figure 4.6: Comparison of cumulative electricity cost under different schemes.....	130
Figure 4.7: Comparison of cumulative green energy consumption under different schemes....	131
Figure 4.8: Comparison of cumulative no. of switchings under different schemes.....	131
Figure 4.9: Performance of the proposed <i>IDD</i> (a) Cumulative electricity cost, (b) Cumulative renewable energy used, (c) Cumulative switchings.....	133
Figure 4.10: Effect of electricity price forecast error in the cumulative electricity cost.....	135
Figure 5.1: Operation of a CSP with multiple geographically distributed data centers.....	142
Figure 5.2: Topology of the distributed data center network used in the simulation.....	161
Figure 5.3: Electricity price in three different locations.....	163
Figure 5.4: Renewable energy generation at three different locations (in percentage of the total capacity).....	163
Figure 5.5: Carbon emission rate at different locations used in the simulation.....	165
Figure 5.6: Comparison of cumulative electricity cost.....	166

Figure 5.7: Comparison of cumulative renewable energy used.....	167
Figure 5.8: Comparison of cumulative number of switchings.....	168
Figure 5.9: Comparison of cumulative electricity cost under <i>LB-L</i> and <i>LB-SM</i>	169
Figure 5.10: Comparison of cumulative renewable energy used under <i>LB-L</i> and <i>LB-SM</i>	170
Figure 5.11: Comparison of cumulative carbon emission under <i>LB-L</i> and <i>LB-SM</i>	170
Figure 5.12: Comparison of cumulative no. of switchings.....	171
Figure 5.13: Percentage of jobs queued.....	172
Figure 5.14: Comparison of average cost per request before and after <i>ECVC</i> for <i>LB-L</i>	174
Figure 5.15: Comparison of average cost per request before and after <i>ECVC</i> in <i>LB-SM</i>	174
Figure 5.16: Comparison of cumulative electricity cost before and after <i>ECVC</i>	175

List of Tables

Table 3.1: Deviation of profit and revenue earned by the CSPs from the benchmark pricing scheme with and without CR (in percentage).....	83
Table 3.2: Cumulative profit, revenue earned by the CSPs and the amount of demand serviced by them.....	87
Table-3.3: Comparison of <i>FTOP</i> with a baseline scheme.....	94
Table: 4.1: Classification of different types of workload in terms of deferability.....	102
Table 4.2: Notations for incentivized deadline deferral model.....	119
Table 4.3: Summary of results in terms of percentage.....	134
Table 5.1: Description of notations used in <i>ECVC</i>	156
Table 5.2: The service rate of servers at different locations.....	164

List of Acronyms

ERP	Enterprise Resource Planning
CSP	Cloud Service Provider
CSM	Cloud Service Market
IoT	Internet of Things
IT	Information Technology
GLB	Geographical Load Balancing
QoS	Quality of Service
SLA	Service Level Agreement
IaaS	Infrastructure-as-a-Service
PaaS	Platform-as-a-Service
SaaS	Software-as-a-Service
MPC	Model Predictive Control
CPU	Central Processing Unit
DVS	Dynamic Voltage Scaling
DVFS	Dynamic Voltage Frequency Scaling
MVA	Mean Value Analysis
PUE	Power Usage Effectiveness

OR	Operations Research
FCFS	First Come First Serve
PS	Processor Sharing
SRPT	Shortest Remaining Processing Time
FB	Foreground Background
EDF	Earliest Deadline First
MINLP	Mixed Integer Nonlinear Programming
GBD	Generalized Benders Decomposition
DR	Demand Response
KKT	Karush–Kuhn–Tucker conditions
IEEE	Institute of Electrical and Electronics Engineers
ACM	Association for Computing Machinery

Chapter 1

Introduction

1.1 Background and motivations

The use of Internet technology has become pervasive in almost all aspects of our life-styles and in our economy, starting from web browsers catering to the needs of retail customers to Enterprise Resource Planning (ERP) systems [1] which are extensively used by medium to large sized corporations. To support various complex processes crucial to their business, virtually all participants across the entire business ecosystem use the Internet fairly extensively in some way or the other. Given this deep penetration of the Internet in our personal lives and in commercial activities, the data transmitted over the Internet is also growing at an extremely rapid rate. The

various different services and applications developed and implemented by IT companies for e-commerce, social networking and intelligent web analytics need to tackle exchange of massive amounts of data between the servers and between servers and clients. This unprecedented growth of the volume of data carried over the Internet has given rise to the new paradigm of Big Data which is receiving an increasing amount of attention from both academia and industry. It is widely believed that the big data paradigm is still in its nascent phase and many different techniques are evolving to support the enormous computational requirements [2] that this entails. The business growth of almost every company (across various industry segments such as automotive, IT, finance & banking, manufacturing, consumer goods, electronics & semiconductors, telecommunications, healthcare, hospitality to name a few) largely depends on the effectiveness with which it can leverage data analytics to make sensitive and crucial business decisions [3], [4]. Motivated by the enormous value to the business processes, big data analytics is emerging as a fairly new interdisciplinary field of study [5]. This requires expertise from a variety of disciplines such as communications, data storage, information theory, software engineering, and statistical analysis to handle these heterogeneous data streams of high volume from multiple autonomous sources. The review paper in [5] reported that 2.5 quintillion bytes (i.e., 2.5×10^{18} bytes) of data are generated every day. Very interestingly, 90% of the total data in the world today has only been generated over the last couple of years.

A variety of different data communication and processing technologies are needed to handle this explosive growth in data volume. The two emerging technologies that will continue this growth in the coming years are: i) Mobile computing [6] and ii) Internet of Things (IoT) [7], which are probably still at a nascent stage at this time. As a result of this incredible growth in the overall cloud computing technology, the volume of data transmitted over the Internet is growing

at the rate of more than 20% annually [8]. It is argued in [6] and [8] that a significant proportion of this huge explosion of data over the Internet can be attributed to handheld mobile devices. The data traffic created by wireless devices is also increasing rapidly at a high rate and this growth is expected to continue till 2019 [8].

Reference [9] presented extensive discussions on the business aspects of cloud computing and the Cloud Service Market (CSM). Cloud based hosting for a diverse range of services allows a company to focus more on its core competencies since the storage and processing requirements of its applications are serviced and managed by its partner Cloud Service Provider (CSP). Moreover, in this case the entire system can take advantage of the economies of scale, since the capital cost is depreciated across all service tenants. However, before migrating to cloud, every company should judiciously calculate the cost implications of the various options available to make a prudent decision [10]. A recent analysis from a renowned market research firm reported that the net worth of cloud computing market is expected to grow to \$121 billion (USD) by 2015 with a 26% compounded annual growth rate from its \$37 billion (USD) value back in 2010 [11]. This growth has been mainly driven by rapid adoption of Software as a Service (SaaS) platforms, resulting in saving businesses on the cost of deployment and management of software and hardware for collaborative knowledge sharing and productivity enhancement. In the context of Singapore, there are many small to medium scale start-ups, which are participating as SaaS providers in the cloud service market. An example of a SaaS provider operating in Singapore is Graymatics [12], which is a scalable cloud platform providing customers real-time indexing, analysis and classification of videos and images with a suite of search, advertising and recommendation tools. It provides a wide spectrum of services to its corporate clients, ranging from e-commerce, consumer data analytics, image and video content mining, etc. SaaS providers

do not typically own cloud computing resources and data centers, since that would require an enormous capital investment, which is typically not affordable for small to medium size organizations whose major expertise lie in software based application development and maintenance. To support this massive growth, the market architecture is constantly evolving with a number of different stakeholders operating at various levels. In these types of massive CSMs, a very crucial issue faced by the CSPs is the issue of pricing for the services provided to the customers. With the rapid growth of the cloud service market, a thumb rule based pricing strategy is becoming increasingly inefficient. In this scenario, there is an obvious need for structuring and implementing appropriate pricing mechanisms for the prices charged by service providers to consumers, which would result in an economically sustainable CSM. This thesis aims to develop a novel theoretical framework to implement sustainable optimal pricing policies and the corresponding optimal resource provisioning policies in the CSM, with multiple CSPs and price sensitive rational consumers. The CSPs compete to attract customer demands where the customers are price sensitive, rational and strategic, which is indeed the case in reality. As a result, eventually all the customers will move to the cloud service provider which offers better service quality at lower cost. Hence, each CSP needs to strategically determine the optimal set of actions (i.e. the service quality offered and the price charged for the same). These actions taken by an individual CSP also has to consider the set of actions chosen by other competitors operating in the same market. In the proposed model this is being monitored and implemented with the intermediation of the regulators, who are responsible to ensure a healthy competition in the market.

As is the case with the electricity market or the stock market, there would typically be a regulatory body which regulates the business operations of the CSPs and the prices they are

charging from their customers. The market regulator would be responsible for auditing the service quality provided to the customers and ensuring a healthy competition between the CSPs operating in the market. They would also be responsible for maintaining a stable and efficient market where no CSP can become a monopoly by acquiring the power to manipulate the market price by behaving strategically [13]. On the other hand, if the competition is too fierce then some of the CSPs would not be able to sustain themselves in the market and that would also be detrimental to the ecology of the CSM. This thesis provides some insightful information to the practitioners about the optimal way of determining the service prices to strike a preferred balance between earning profits and attracting more customers. The inherent trade-off is that a CSP should earn sufficient profit (revenue – operational cost incurred) to sustain its business while at the same time it should not charge too much from the customers as that would then result in subsequent loss of market share (demand for services). It may be noted that the market share of CSP i may be represented by the proportion of the total demand served by CSP i . Most importantly, the CSPs should provide an acceptable level of Quality of Service (QoS) to their customers. Even though the ultimate objective of a CSP is to make more profit, pricing services too aggressively will be detrimental to its long-term business sustainability and survival in the market. Since we assume the customers to be price sensitive, they would in that case choose a CSP which can provide good QoS at lower cost than the others.

Here, economic sustainability would imply for both long term and short term economic sustainability for all the CSPs in the market. This would indeed be the first step to establishing a sustainable cloud service ecosystem. In this sustainable cloud service ecosystem, the various players at one level should charge their next customers reasonably enough, so as to (i) make business profitable (which we call *short term economic sustainability*) and (ii) attract more

customers and increase its market share (which we refer to as *long term economic sustainability*). It may be noted that objectives (i) and (ii) are naturally contradictory with each other. Therefore there has to be a well-chosen strategy to determine the optimal price at various levels in the market. This thesis proposes a well-engineered modeling paradigm where each individual CSP may have different preferences on the two objectives (i) and (ii) mentioned above and can determine the service price to be charged to the customers in this competitive CSM.

The data centers are the primary physical resources that support the operations of cloud computing. With the tremendous growth of the Internet, these data centers are also growing at a similar pace. In its 2014 annual report, Google declared there will be even more investment in the data center assets owned by the organization [14]. This naturally implies that the electricity required to operate these data center facilities is also growing enormously. It was predicted that, the aggregate electricity demand of the data center infrastructure would have ranked in sixth position, if it would have been considered as a country [15, 16]. This prediction has proved to be true by now. However, the more alarming fact about this energy consumption is that it is also growing at an incredible pace. It is reported that the electricity demand for cloud computing infrastructure is growing ten times faster than the overall rate of growth of the electricity demand in the US [15], [17]. As a result, the electricity cost incurred by a medium to large sized CSP constitutes a significant proportion of its overall cost of operations. As per an official report, Google alone consumed \$138M worth of electricity annually [18, 19] in 2011, which would have increased even further by now. Reducing the electricity cost would result in significant reduction in the operational cost incurred by a CSP. By saving on electricity cost, a CSP may i) provide same level of quality of service (QoS) at a lower cost, or ii) make even more profit by charging the same cost, or a combination of the two. Providing the same level of QoS at a lower cost

would help to increase its market share and thereby improve its long term economic viability. On the other hand, making more profit would improve its short term economic sustainability without affecting the long term economic sustainability of the CSP.

The immediately obvious detrimental effect of the enormous amount of energy consumption is the emission of greenhouse gases, which is a major concern of many private and government agencies because of the harmful effects it has on our environment. A very inconvenient truth about human civilization is that our technologies are too dependent on fossil fuels which results in the release of greenhouse gases such as carbon dioxide, methane and nitrous oxide. These issues have raised concerns of climate change across the globe. Figure 1.1 depicts the energy spaghetti chart released by Lawrence Livermore National Laboratory, California [20] which reflects the sector wise details of carbon emission.

Figure 1.1: Energy Spaghetti chart 2014 (published by Lawrence Livermore National Laboratory)

It is commonly accepted that the alarming issue of global warming should be considered to be a major crisis that mankind would have to face [21]. Several governmental and legislative bodies in a number of countries have launched aggressive drives to tackle this. For example, the state of California, USA has set itself the very ambitious target of achieving an energy portfolio where 33% of the total energy consumption will be supplied by renewable energy sources [22]. These have had a big influence in driving renewable power in several key markets and a growing number of utilities have begun to shift their investments to renewable energy to meet this new demand. In order to contribute towards handling the important issue of environmental sustainability, a growing number of companies in the cloud computing business ecology have begun to integrate different renewable energy sources so as to eliminate (or at least reduce) these detrimental emissions. By leveraging appropriate technologies and strategies, over a half-a-dozen large companies in the cloud service ecosystem, including Google, Apple and Facebook, are now committed to eventually becoming 100% renewably powered [16]. As an example, in 2013 Google invested USD 213 million towards renewable energy integration and increased this to USD 430 million in 2014 [14]. This is indeed quite dramatic and demonstrates the importance of this issue. As per the report in [8], Apple has also demonstrated very aggressive steps towards environment sustainability. All three of its data center expansions announced in 2014, are supposed to be powered entirely by renewable energy. Additionally, Apple has made a significant effort on pushing its major colocation providers to help it maintain progress toward its greenhouse gas emissions goal. However, companies in the cloud service market committed to such renewable energy goals will need to put in substantial efforts for this and need to collaborate closely with policy makers and electricity providers (utility providers) to make further progress in energy efficiency [8]. A detailed discussion with associated statistics on

energy emission and environmental sustainability practices prevalent in different large companies may be found in [8].

To consume energy in a more efficient way and to address the issues of increased renewable energy integration, *smart grid technology* [23] is another area that has emerged recently. The smartness in this context is associated with the various practices associated with energy generation, transmission, distribution and consumption and is of particular interest when there is a need for integrating renewable energy into the grid. It is reported that the main challenge associated with renewable energy is the randomness associated with the generation of the commonly available renewable energy sources such as solar, wind etc. This randomness is often attributed to the inherent intermittency, uncontrollability and uncertainty in the generation of renewable energy from these sources [24], [25]. A number of strategies have been successfully designed, developed and implemented over the past decade to tackle these problems. In the smart grid arena, a very effective and well-studied strategy is *demand response* (DR). This strategy can be of immense value in the context of data centers and cloud computing services [26]. A recent survey on the current state of the art practices in implementing demand response in the context of different applications and users may be found in [27].

Given the aggressive commitments made by different nations, it is imperative that an adequate amount of economic incentivization is provided in the future to attract different participants. It may be noted that to attract different organizations to invest in renewable energy infrastructure some economic and financial plans are already in place. The report released by the market research team from Deutsche Bank, provides a set of futuristic insights related to the investments in different areas of the energy infrastructure [28]. They carried out a critical analysis on energy market dynamics and concluded that in the near future some changes in

policies would be needed to attract more investments in renewable infrastructure so as to adequately address the pressing issue of environmental sustainability.

This thesis aims to combine these two types of sustainability (economic and environmental) through appropriate pricing mechanisms and implement this through appropriate modeling paradigms while considering the major aspects prevalent in the market. The practices presently followed on the pricing of cloud services may not be optimal in terms of many crucial parameters. For example, after the deregulation of electricity markets in US and some other countries, the price of electricity in the wholesale market varies dynamically over time. In this scenario, the key ideas will be to capture crucial attributes like the fluctuations in demand and resource congestion in order to evolve an appropriate pricing mechanism. This pricing scheme would allow the service providers to reflect the marginal cost of providing resources in a way so as to have some control on the consumption pattern. This thesis relates the two aspects of pricing and the resource provisioning cloud services to each other and analyzes them to evolve an appropriate pricing strategy. The area of environmental sustainability has been relatively well explored with multiple approaches proposed from different fields to achieve better energy efficiency, e.g. hardware technology, control engineering, machine learning, mechanical engineering, and physics. When pricing and resource consumption (or energy consumption) are considered together, the resource utilization can be improved substantially and the service provider is also able to better allocate resources with increased economic efficiency.

1.2 Objectives

As mentioned earlier the two main objectives of this thesis are to contribute towards, i) economic sustainability, and, ii) environmental sustainability of the cloud computing system. This thesis aims to achieve the first objective by adopting two lines of actions. We first develop strategies to compute optimal prices in the competitive CSM. This aims to address the CSM with multiple CSPs in a holistic way. The other approach is to figure out strategies which would reduce the operational and capital cost incurred by the individual CSPs to continue their operations. This should improve profitability and thereby economic sustainability.

The second objective is to improve the environmental friendliness of the cloud computing system. This thesis aims to achieve this objective by developing strategies to increase the usage of renewable energy to power the servers in the data centers operated by the CSPs. This is done by proposing a coherent set of strategies, similar to those prevalent in the smart grid arena, to design and implement effective and efficient algorithms to achieve the objective of greening the cloud computing system.

It is important to note that environmental sustainability is very closely related with the issue of economic sustainability. Environmental sustainability through better engineering approaches has become reasonably feasible today with the tremendous progress achieved in this area over the past several decades. However, the issue of environmental sustainability through appropriate pricing and economic motivation has not been adequately explored or even properly understood. Even though the economics related to energy systems is quite a rich topic, when it

comes to cloud service pricing there has been surprisingly little synergy between the economic foundations of cloud service pricing and energy (more particularly electricity).

The objectives of the thesis are summarized in Figure 1.2 given below.

Figure 1.2: The objectives of this thesis

1.3 Major Challenges

In this section the major challenges associated with the two issues addressed in this thesis are discussed. These two issues are, i) economic sustainability and, ii) environmental sustainability.

1.3.1 Economic Sustainability

The primary challenge associated with service pricing in a competitive CSM is that there is an inherent tradeoff between the price charged (which is directly related to the profit earned) and the QoS provided. Here, the relationship between the prices charged (for a service) and the corresponding profits that are earned are functions of the cost incurred by the CSP to provide that level of QoS. If one CSP is charging more to make higher profits then the customers, being rational and price conscious, will eventually migrate to some other CSP and the CSP targeting profit making alone will eventually lose its market share. The challenge in this is to develop an appropriate strategy for pricing in these types of markets which would provide economic sustainability to all the players in the market.

Typically the QoS offered and the price charged are very closely related and will largely be driven by the internal strategies of the customers, i.e. to determine how much the customer is willing to pay for what level of QoS. To address this issue, the pricing strategy must consider the resource requirement of the demand. The current practices followed are mostly that different CSPs implement pricing plans on an ad hoc basis and there is no established modeling paradigm based on which the CSPs can compute their pricing strategies. At present the most common

pricing scheme charges customers on the basis of usage of resources in a pay-as-you-go manner. Additionally, it may be noted that pricing today is becoming more sensitive to long term economic sustainability than ever before as a lot of information is available to the users (customers) very easily through social networking platforms and other electronic means. Given the ease with which one can be informed about the pricing strategies deployed by different CSPs, a customer has easy options of choosing a CSP which is charging less or providing a better level of QoS, depending on the individual preferences of the customer.

1.3.2 Environmental Sustainability

The greatest challenge in implementing a strategy to dynamically allocate resources as per the service demand is that the demand varies quite significantly, and even randomly, over time. It is reported that the average utilization of servers in a typical commercial data center is only around 10% - 30% [9, 29]. The commercial servers typically deployed in the data centers are far from being energy proportional. Even during idling conditions, when the server is not performing any work, the power consumption of a typical server is close to 60% of its peak power consumption [29]. The long term economic sustainability of a CSP is also dramatically sensitive to the QoS provided to its customers. For example, Amazon reported that it tends to incur a loss of 1% of sales revenue for every increase of 100 ms in response time, whereas Google reported that a 500 ms increase in response time translates into a 20% reduction in its revenue [30]. It is evident that a CSP has to be very careful (and possibly even conservative) while determining the amount of resources to be provisioned to the customers (as per its requirement) since resource under provisioning would be highly detrimental to its business.

On the other hand, server switching (i.e. transition from OFF to ON and vice versa) also has some undesirable effects. It takes some time to switch a server from OFF to ON state and during that time the server does no work, but consumes energy which is very close to its maximum energy consumption [31]. More importantly, there are issues like data localization, data migration (most often in the form of virtual machine migration in case of data centers) which also carries extra overhead along with the server switching. Last but not the least, IT operators are often skeptical about server switching since there can be adverse effects on server lifetime due to frequent switching from OFF to ON and vice versa [32]. As a result, there is an inherent tradeoff between renewable energy integration and cumulative number of switchings. If there would have been no adverse effects of server switchings then the CSP can simply turn on more servers when renewable energy is high and use the bare minimum number of servers when the renewable energy is low. On the other hand, if server switching cost is too high then the CSP may choose to keep active a constant number of servers and in this case the performance in terms of renewable energy integration will be poor.

1.4 Novelty and Contributions

In this section the contributions of this thesis are briefly discussed. The details of the original contributions of this thesis are presented in chapters 3, 4 and 5. As mentioned above, the overall high level objectives of this thesis are to improve two critical attributes of cloud computing sustainabilities, namely, i) economic sustainability and, ii) environmental sustainability. In different chapters we consider different scenarios, design appropriate models, develop and implement effective algorithms to achieve these two high level objectives.

Concerted efforts are made to connect these two issues together. The original contributions of this thesis are given below:

- In the context of a competitive CSM with multiple CSPs, an analytical framework is developed to compute the equilibrium load balancing condition under reasonable assumptions as would be applicable in a competitive market where the customers are both rational and price sensitive. Under this framework an energy aware resource pricing scheme is designed to compute the price which is a good trade-off between charging too little (thereby operating on low profit margins) and charging too much (which will lead to loss of customers and eventually to loss in the CSP's market share). This is achieved by implementing a bi-level optimization strategy. On a slower timescale (first level) the service demands are distributed optimally to the multiple CSPs competing with each other. Once the service demands are distributed, an efficient algorithm is used to perform a fast timescale optimization (second level) to schedule deferrable jobs so that renewable energy integration can be maximized. The proposed methodologies and algorithms are evaluated with actual traces of workloads collected from real server clusters, real life renewable traces, and real electricity prices in order to emulate the practical operations of a CSP.
- In the context of a single data center, the resource provisioning for both deferrable jobs and non-deferrable jobs are tackled. To focus on the practical scenario of a real data center, a realistic tiered architecture of a data center is considered, which enables us to address the requirement of provisioning different types of resources to different types of applications. This also allows capturing the constraints of

executing some portion of a computation-intensive job only on a selected set of servers due to some additional security requirements or data localization issues. Under this framework, investigation and analysis are carried out on the opportunity of peak shaving. To achieve peak shaving, the proposed scheme minimizes a weighted combination of capital cost (peak server provisioning), operational cost and switching cost, subject to operational and deadline constraints. Additionally, a simple mechanism to incentivize deadline deferral for more aggressive cost saving and renewable incorporation is proposed. This incentivization mechanism is novel, scalable and can indeed be implemented in a real life large scale system. Moreover, this incentivization mechanism lies in the intersection between economic sustainability and environmental sustainability which is the central theme of this thesis.

- In the context of a CSP having multiple geographically distributed data centers, located across the globe, a holistic framework for modeling and distributing the jobs among geographically distributed data centers is presented. This model considered the important attributes including the electricity price in the wholesale market, local renewable availability, cooling related conditions (at a macroscopic level), carbon emission effects and, most crucially, the adverse effects of server switching. The proposed model is leveraged to investigate the implications of different formulations of the optimization problem to perform Geographical Load Balancing (GLB). This provides more insights to help practitioners making the decision based on priority towards electricity cost reduction, renewable energy incorporation and server switchings. The relative performance between two

algorithms for GLB based on Model Predictive Control (MPC), are analysed and compared in terms of the cost of electricity, renewable integration and switching. While one of them is quite close to the existing works [32, 33], the other approach is a novel one which has not been explored as yet. Furthermore a strategy to determine the optimal trading in the forward electricity market is developed to deal with the issue of risk associated with the time variability in electricity price, renewables and amount of job requests. Extensive simulations are carried out with real data traces on electricity price, renewable generation, cooling related conditions, workload, and carbon emission effects at ten locations across the globe, to emulate the complexity of a real system.

1.5 Organization of the Thesis

This thesis focuses on two issues, namely, i) economic sustainability and ii) environmental sustainability and also finds the overlaps between the two. This section provides an overview of the organization of the contents of this thesis. Throughout the thesis, there is a fair amount of synergy between the two attributes mentioned above.

In Chapter 2, an extensive literature review in the context of the holistic view of this thesis is provided. In this chapter the previous works relevant to the work in this thesis are presented and discussed. However, subsequently in chapters 3, 4 and 5, only the relevant references are mentioned to differentiate our work and highlight the novelties and contributions of our research.

Chapter 3 focuses on economic sustainability in the competitive CSM with price sensitive rational customers. This chapter presents a bi-level optimization problem to connect the

economic and environmental issues together, which is solved at two different timescales. In the slower timescale the prices charged by the CSPs to the customers are computed. In the faster timescale, the energy efficiency aware job scheduling is performed. In the slower timescale, the optimization procedure aims to enhance economic sustainability while in the faster timescale the optimization procedure is designed to improve both economic sustainability and environmental sustainability. In the faster timescale, economic sustainability is associated with electricity cost reduction while environmental sustainability is associated with renewable energy integration.

In chapter 4, the case of a single data center is considered more rigorously and in greater detail. In this chapter, we model a data center where several different applications are implemented in different tiers as this is typically the way modern data centers are operated [34], [35]. Under this framework, we consider both non-deferrable and deferrable types of workloads and propose a renewable energy aware job scheduling strategy. In this chapter, we also propose a novel incentivization strategy to encourage the customers to further shift the stipulated deadline of the deferrable jobs. In return, the CSP provides monetary rewards based on the amount of delay the customer is willing to tolerate and the size of the job. We demonstrate that, through this strategy the CSP can actually incorporate more renewable energy with no additional net cost (operational cost incurred by the CSP) and no extra server switchings. It may be noted that this incentivization strategy presented in chapter 4 lies in the intersection of economic sustainability and environmental sustainability.

In Chapter 5, we consider the case of a CSP with multiple geographically distributed data centers. This issue has already been addressed in the literature to some extent and is referred to as Global Load Balancing (GLB) by various authors. In this chapter, some insights are provided about the performance of different variants of GLB algorithms with respect to cumulative

electricity cost, cumulative renewable energy integration and cumulative number of switchings. Additionally, the issue of emission is also considered which depends on the electricity generation portfolio of the place where the data center is located. It may be noted that the average electricity cost to service a job request varies dramatically over time. To deal with this, an effective and efficient strategy to mitigate the risk associated with average electricity cost to service a job request is proposed. Being a price risk mitigation strategy, this would help in improving economic sustainability. The effectiveness of the proposed algorithms is demonstrated through simulations with real data traces for electricity price in the wholesale electricity market, renewable energy generation, user request etc.

Finally, in chapter 6, the conclusions of this thesis are presented. Furthermore some meaningful research directions are identified to address similar issues when the architecture of the CSM is more complex.

Chapter 2

Literature Review

2.1 Introduction

In this chapter, important research contributions from existing literature relevant to this thesis are reviewed. In Section 2.2, we discuss existing works related to pricing for cloud services are discussed. In Section 2.3, existing works related to environmentally sustainable cloud computing and energy efficiency related contributions are summarized. In Section 2.4, the existing literature on demand response in data centers are discussed. This also relates to overall energy efficiency and is an area where ideas being developed under *smart grid technology* may be combined with ideas from optimal control and economical pricing in cloud computing. This is an important direction where significant developments and progress are expected in the near future. Therefore, we discuss the recent developments in this particular topic in a separate subsection. In Section 2.5 the existing works related to cost volatility control and financial hedging are discussed.

Various aspects of cloud computing have been presented in [9, 36, 37], while [38] provides details on engineering, operational and business issues of the data centers which are used to implement the cloud.

2.2 Pricing for Cloud Services

The topic of economic sustainability is very rich and has been studied in the context of different system settings. Insightful discussions about economic sustainability from the perspective of sustainable development and optimal growth as viewed by an economist can be found in [39] and [40]. In the context of CSM, the economics and pricing issues driven by different end objectives have been studied from economic and engineering perspectives and different pricing strategies, are proposed and deployed. However, the current pricing approach for cloud services is not very well aligned with economic theories of pricing and market equilibria. Our strong conviction is that with the growth of CSM, more economically sustainable pricing strategies will eventually be developed and adopted by the practitioners.

During the last decade, most of the CSPs charged their customers based on fixed prices [41] which fails to address issues like resource congestion, marginal resource cost incurred by the CSP, energy usage cost etc. This pricing scheme is inefficient in terms of resource congestion, marginal resource cost incurred by the CSP and energy cost. Nevertheless, due to simplicity in implementation, the practice of charging fixed price based on usage is still existent in the CSM [41]. The motivation for implementing intelligent pricing strategy is to differentiate Quality of Service (QoS), across different customers or applications. This can capture the difference in *value* of different services to different customers. Another crucial attribute with respect to economic sustainability from the perspective of CSPs is *market competition*. In reality, there would be multiple CSPs operating in the market that would compete with each other to service the demands generated from the customers, who are themselves price sensitive, rational and informed.

There are various types of pricing strategies reported in the literature and some of them have been adopted by the industry. However, static pricing strategies are the ones mostly implemented by the CSPs currently operating in the market. A pioneer work suggesting an efficient autonomic pricing mechanism which dynamically adjusts the pricing parameters to discriminate between the different service and application requirements of different customers can be found in [41]. It proposed a dynamic pricing strategy to differentiate among customers based on their requirements. The most important advantage of customized, dynamic pricing is that it can address resource contention in a better way, and can therefore improve the resource usage efficiency to generate higher business revenue for the CSP. Being a pioneer in the CSM, Amazon EC2 has introduced a spot pricing strategy where excess capacity (servers, storage capacity) are sold at a dynamic price to better leverage supply and demand using an auction based method [42]. Though this pricing strategy is quite straightforward and easy to implement, the authors of [43] performed an exhaustive empirical analysis and reported that spot pricing may be inefficient in reflecting the actual market demand. Instead, they tackled the revenue management issue through a framework based on economics and formulated a stochastic optimization problem for revenue maximization with dynamic pricing. Substantial past research works do exist which formulates a theoretical framework for service pricing. The work in [44] presented an analytical framework focusing on the fairness and revenue trade-off that arises for a heterogeneous set of physical and virtual resources of a data center and the influence various pricing strategies can have on the trade-off between fairness and revenue maximization. It characterized the implications of different pricing strategies on different fairness metrics and derived the analytical bounds on the CSP's strategies to balance between fairness and revenue. The work presented in [45] focussed on maximizing the long term social surplus (which is

computed by the users' utility for getting the job done minus the operational expenditure incurred to execute the job such as resource usage cost, electricity cost etc. The authors carried out mathematically rigorous investigation to characterize the existence and uniqueness of the social optimum under some reasonable technical assumptions.

The work which is closest to the one presented in this thesis is [46]. In that work, the authors considered the competitive CSM with more than one service provider and price sensitive customers. They provided an analytical framework which is adopted in our work. However, they primarily focus on the economics of the CSM whereas this thesis focuses on formulating pricing strategies which make the system greener. A major aim of this thesis is to bridge the issue of economically sustainable pricing with environmental sustainability by appropriately pricing the renewable energy.

It may be noted that the issue of resource provisioning is very closely related with pricing, which is also leveraged in this thesis. In general, the literature on cloud resource provisioning is quite rich. However, as far as resource provisioning is concerned, very often pricing does not reflect the true marginal cost incurred by the provider. It is evident that synergy between pricing and resource provisioning would indeed make the system more efficient from economic perspectives. The recent developments in different domains of cloud services and resource management practices can be found in [36, 47]. Some promising and meaningful future research directions have also been presented in these papers. The work in [48] presented a sophisticated method for cloud computing resource evaluation based on entropy optimization considering the inherent uncertainty. The work in [49] proposed an analytical model for resource provisioning for multi-tiered applications. The work in [50] presented a fairly exhaustive survey on various

resource management strategies proposed in the literature and the practices prevalent in the industry by various CSPs.

There have also been some efforts to consider the energy efficiency issue through an appropriate pricing mechanism. The work conducted by researchers from IBM [51] considered a combined framework where scheduler, charge back model and customer's preferences work in harmony for a green and energy efficient cloud computing system. However, they considered the case of a single CSP, whereas we consider the entire CSM with many CSPs and also address the issue of competition which will inherently arise between them. The work reported in [52] presented a similar cost optimization framework, with collaboration between the CSP and the utility provider to increase renewable energy penetration. They proposed a distributed collaborative cost optimization framework via dynamic pricing. However, this work did not deal with the issue of pricing and market competition which is the central focus of this thesis. They also did not consider practical issues in data center operations like the adverse effects of server switching, the multi-tier architecture of a large scale data center etc., which are appropriately addressed in this thesis.

Intuitively, the economics of emerging technologies like the Internet of Things (IoT) is also experiencing similar challenges, and would require multidisciplinary approaches from computer science, information theory, communication theory and economics, to support and realize their potential growth as a new technological paradigm. The work presented in [53] reviews literature in the economics of IoT and presents some meaningful research directions. The work in [54] presents a comprehensive literature review and discussion in the context of pricing the data over broadband and wireless networks. We believe that the insights obtained from new areas like IoT and data pricing (over broadband and wireless networks) will be valuable in

structuring sustainable pricing strategies, appropriate in the context of CSM, to ensure the overall growth of the entire cloud service ecosystem.

2.3 Energy Efficiency in Cloud Computing

Financial and social pressures have brought about significant operational improvements which have reduced the energy consumption of data centers and increased their energy efficiency. As discussed in the previous chapter, the issue of energy consumption in cloud computing system has received considerable attention from both industry and academia over the past several decades. Given the many research contributions reported in this particular area, this section discusses the existing literature is discussed in three categories, i.e., i) energy efficiency in servers, ii) energy efficiency in a single data centers, and iii) energy efficiency in multiple data centers.

This thesis develops strategies based on the idea of turning off extra servers when the service requirement is low. We aim to design algorithms to dynamically determine the number of servers to be used, as per the service requirement. Driven by this idea, strategies are developed for both single data center and multiple data center scenarios. There are two benefits of this. The first benefit is that this would reduce the electricity cost incurred by the CSP. In fact, given the huge electricity cost paid by the CSP (to operate big data centers each having a large number of servers), even small savings here will translate into significant monetary benefits [18], [19].

2.3.1 Energy Efficiency in Servers

Though our work does not target this area, any literature survey would not be complete without discussing the recent research developments in energy efficiency in the server level. A survey on the developments in architectural techniques to reduce power consumption at the hardware level can be found in [55]. The notion of adjusting the CPU speed as per the requirement remains at the core of most commonly used strategies for improving energy efficiency within a server. It is evident that the power consumption would get lower as the speed is reduced. Some of the techniques to achieve power efficiency at the server level include Dynamic Voltage Scaling (DVS), Dynamic Voltage Frequency Scaling (DVFS) [56] etc. The recent survey paper [57] and the references therein discusses the developments and research contributions in this area.

2.3.2 Energy Efficiency in a Single Data Center

Over the past decade, there have been remarkable improvements in the data center power efficiency by dynamically controlling the number of servers provided. When the demand is low, some servers can be put into low power state (sleep state), which is the notion of power-performance proportionality at the whole data center level [32]. Fairly comprehensive surveys on some algorithmic approaches for data center and network power reduction, efficiency and control are found in [58-61].

The pioneer work, proposing an algorithm based on the idea of energy saving by turning off extra servers as the resource demand decreases is [62]. In that work, the authors characterized the properties based on a real data trace collected from a live messenger service owned by

Microsoft. Using the model, they implemented and evaluated server provisioning and load dispatching algorithms in a large cluster of servers owned by Microsoft. This real large scale system implementation indeed provided extra confidence to make further progress along this idea of turning off servers when the demand is low. The seminal work in [32] aimed to achieve energy proportionality in a single data center. This paper proposed an elegant online algorithm to dynamically control the number of active servers in a data center, when the demand is stochastic. The work in [63] considered multiple sleep states (with different power consumptions and time required to turn on) and presented some insightful theoretical results on the characteristics of optimal policies when the service demand is stochastic and there is a cost associated with server switching, in terms of wake-up time and the energy consumed during the wake-up time while doing no actual work. This naturally demands understanding of the theories behind online algorithms and their implementation for energy saving purposes in computer systems. An online algorithm would be one without perfect future knowledge. A good overview and summary of prominent contributions in the areas of algorithms focusing on energy efficient aspects is presented in [64]. This paper discusses the performance of various types of online algorithms (with knowledge up to the present) with respect to the offline (with perfect future knowledge) in terms of competitive ratio. However, none of these works considers issues like energy cost, renewable energy generation etc., which are addressed by us in this thesis.

In data centers, a very widely used and well developed technique is virtualization [65]. In simple terms, virtualization is the method by which a single server can be divided into multiple isolated virtual environments and the tasks allocated to different virtual portions, are carried out independently. As a result, virtualization enables more efficient resource provisioning for heterogeneous jobs which require different amounts of various resources such as memory,

storage, processor share etc. It may be noted that, virtualization is one of the technologies which facilitate enhance energy efficiency at the data center level. As an example, to achieve power proportionality at the data center level the jobs required to be serviced can be consolidated into fewer servers through the virtualization technique and the rest can be powered off. The recent work in [66] presented an exhaustive survey on the recent developments in the context of achieving power efficiency in data centers by leveraging virtualization techniques.

As mentioned earlier, the jobs in cloud computing can be divided into two categories, namely, i) non-deferrable jobs, and, ii) deferrable jobs. The non-deferrable jobs are the jobs which have to be processed immediately on arrival. In contrast, deferrable jobs are those jobs which do not have to be done immediately but do need to be executed within a timeframe specified by the respective customer. A few examples of non-deferrable jobs are: search query, online shopping, online video streaming etc. On the other hand, the deferrable jobs are ones which can be delayed, as long as their deadlines are met. The examples of deferrable jobs are scientific simulations, large scale data analytics etc.

Some notable references which deal with non-deferrable jobs are [67], [32], [35]. For non-deferrable jobs, the SLA is commonly based on the upper bound on the expected response time, or the upper bound on some percentile of the response time [32]. It may be noted that, as compared to non-deferrable jobs, the literature addressing deferrable jobs is less developed. Some works considering deferrable jobs are available in [68] and [69]. The former is based on advanced queueing theoretic analysis and the latter utilizes a Lyapunov optimization technique based algorithm.

This thesis makes a novel contribution in designing energy efficient scheduling algorithms for multi-tier data centers. The notable previous work which considered the resource provisioning issue in a multi-tier server cluster is [35]. The work in [35] proposed an analytical framework for resource provisioning in a multi-tier server cluster for session based non-deferrable jobs, which is a pioneer work in this area. They used Mean Value Analysis (MVA) to implement an efficient iterative algorithm to address session based traffic which triggers multiple jobs at various tiers, by modeling the system as a closed network of queues. It may be noted that, modern Internet workload is session based where a session typically comprises a series of requests from the customer. Moreover, each of these requests may trigger multiple requests at different tiers and there can be think times in between which correspond to phenomena such as network delay etc. This characteristic of Internet traffic was appropriately modeled in [35], and has also been used in this thesis for modeling resource provisioning for non-deferrable workload. Additionally, the proposed strategy was rigorously evaluated through experiments in a small scale test system. However, deferrable workloads were not considered which is appropriately addressed in this thesis.

The other area which has received a lot of attention over the past few decades is the efficiency of the cooling systems. The research contributions in this area were driven by researchers and practitioners from various disciplines including mechanical engineering, physics, electrical engineering and computer science. This thesis does not make any specific novel contribution towards improving the cooling system efficiency. However, almost all the optimization procedures proposed in this thesis include the macro level statistics of the cooling system efficiency. It may be mentioned that this thesis considers the most commonly used parameter Power Usage Effectiveness (PUE) to capture the dynamics of the cooling efficiency.

PUE is defined as the ratio between the total power consumed in the data center and the total power consumed by the IT equipment [38]. This is an area which has seen a significant number of improvements over the past few decades. A fairly accurate modeling framework is available in [70] to understand the interaction between the computational domain (such as workload placement) and the physical domain (such as the resulting temperature profile), in the context of a single data center. The strategies of allocating the incoming jobs in an intelligent way so as to reduce the overall cooling energy are found in [71], [72]. The various strategies proposed in the literature to improve the cooling system efficiency can be integrated with the strategies proposed in this thesis to make the system more efficient overall.

2.3.2.1 Scheduling

Since this thesis uses the theory of scheduling, a separate discussion is provided to elaborate the literature in the topic of scheduling and their connections with this thesis. The general topic of scheduling is quite well studied in Operations Research (OR) [73], [74] and in computer science [75]. In the context of computer systems, scheduling is very critical and is well-studied with respect to different applications such as scheduling in distributed systems [76], parallel program scheduling [77] etc. The application which is most relevant to this thesis is queueing systems. Scheduling plays a very crucial role in performance modeling in queueing systems. In fact, scheduling determines important performance parameters like average waiting time and average response time in queueing systems. The most commonly used queueing theoretic models include First Come First Serve (FCFS) and Processor Sharing (PS) scheduling policies [78]. In computer systems, other sophisticated scheduling policies like Shortest Remaining Processing Time (SRPT) and Foreground Background (FB) scheduling have also

been well studied. An extensive discussion on different scheduling policies in computer systems can be found in [75]. Chapter 4 of this thesis designs an energy aware scheduling strategy in data centers. Based on a Model Predictive Control (MPC) framework, the proposed algorithm leverages Earliest Deadline First (EDF) scheduling policy. The previous works which investigated the effectiveness of EDF scheduling policy in the context of scheduling of deferrable jobs are [79], [80]. It may be noted that under the EDF policy, there is a theoretical guarantee of completion of all the jobs before the stipulated deadline as long as there is no resource bottleneck [81], [82].

A few references on the theoretical background and various algorithms related to MPC and their implementations can be found in [83]. MPC is a classical topic originally coined by the control engineering community and has been widely applied in many application domains. In the context of resource provisioning and scheduling, the work in [84] proposed a strategy to optimally schedule the available hydro power from cascaded river power plants, subjected to some predefined operational constraints.

In various different fields, scheduling plays a crucial role to achieve energy efficiency. For example, in [85] the authors proposed an adaptive scheduling approach by leveraging extended Kalman filter based estimation techniques to save the energy cost in wireless sensor networks.

2.3.3 Energy Efficiency in the Case of Multiple Data Centers Owned by a Single Organization

Today's large to medium sized CSPs typically have multiple data centers that are located at different geographic locations across the globe. These CSPs, normally replicate data, processes and systems in different data centers to offer short response times, robust performance, enhanced security and operational redundancy [86]. They typically have front end servers to collect the job requests from their customers. All the data centers (owned or leased by a single organization) are connected through a high bandwidth communication network (often optical) that offers high availability and reliability.

To achieve power cost reduction, a well-recognized approach is Geographical Load Balancing (GLB) based on the spatial variation of electricity price [19, 67]. However, this approach is applicable only to a very large CSP owning data centers in multiple locations across the globe. As per current industry practices, either some appropriate load distribution algorithm is employed to minimize the response time [86] or it is simply done based on local geographical proximity. In the past few years, several load distribution strategies have been proposed based on electricity cost reduction [19], carbon emission reduction [87], [88], [89], or a combination of the two [67, 90]. A pioneer work proposing the idea of exploiting spatial variation of electricity price to reduce the cost of electricity incurred by a CSP is [19] where the authors studied temporal and geographical variation in electricity price and proposed heuristic based algorithms for distance constrained electricity price aware routing. Another work in [91] with the same goal, modeled the system as a constrained mixed integer nonlinear programming problem (MINLP) and solved it using the Generalized Benders Decomposition (GBD) technique to obtain optimal load

balancing and power control. It may be noted that MINLP is computationally expansive, which we address in this thesis appropriately. Subsequently, several other works have explored the diversity in green energy availability by considering renewable energy aware load distribution along with overall cost reduction. In [67, 90] the authors merged the two objectives namely, electricity cost reduction and renewable energy incorporation. The seminal work in [67] presented two distributed algorithms to achieve optimal load distribution by minimizing a linear combination of the energy cost and the revenue lost due to delay. They used Gauss-Seidel iteration and Distributed Gradient projection methods [92] to design two proposed scalable, distributed algorithms. A rigorous investigation on the feasibility of powering data centers entirely by renewable energy is carried out in [90] where the authors demonstrated that GLB can optimally be integrated with a reasonable amount of storage to achieve almost zero use of conventional brown energy from the grid. They also provided insights on the optimum portfolio of solar and wind energy to achieve this ambitious target. The work in [88] tackled the issue of carbon emission and proposed algorithms for carbon footprint reduction, which is also addressed in this thesis. An extensive survey on the various strategies and algorithms proposed to perform GLB with different objectives can be found in [93].

2.4 Demand Response (DR) in Data Centers

In the smart grid arena, at macro level, the techniques for controlling the load according to the supply fall into one of two categories namely, *direct load control* (DLC) and *indirect load control* (ILC) [94]. With DLC, the electricity service provider possesses the right to control the loads in a pre-specified manner. The work presented in [95] proposed a DLC based adaptive strategy to improve the operations of ancillary services. In contrast, when ILC is used, the utility

provider does not have a direct control over the loads. In this case the utility provider controls the loads through various pricing signals. The various pricing schemes include time-of-use pricing [96], coincident peak pricing [97], critical peak pricing [98], real time pricing [99], [100]. The general idea behind all these pricing strategies is to ensure that the utility providers can reflect and recover, as a function of time, the true marginal cost for producing one extra unit of electricity. The work in [101] discussed various challenges and objectives to create reference architectures for smart energy systems, especially in the context of Europe. The mechanism proposed in this thesis falls into the ILC category where the consumer reacts by deferring the deadline of a deferrable task further and receives some monetary incentives in return. Functionally, the time varying price acts like a signal to have a decentralized control on the individual customers. The authors in [102] rigorously investigated the economic value of deferring loads to more favorable timeslots. They developed a stochastic process to replicate the electricity spot price instances, using real option theory.

Apple, Facebook, Google and some other IT giants have made significant progress in investing on renewable energy sources over the past few years. A very recent report in [16] presented a comprehensive and exhaustive survey on this. Since the data centers are large consumers of electricity and a significant portion of the workload is deferrable, the data centers are potentially eminent candidates for implementing DR.

The early papers in this area aimed to increase renewable energy usage in powering data centers, which is an issue addressed in this thesis. The central idea is to use the deferability of the deferrable workload in order to increase the usage of renewable energy. The work in [103] presented an implementation oriented work based on a heuristic algorithm to integrate more renewables to power a data center. A good summary of recent developments related to the

integration of renewable energy in data centers is presented in [104]. The reference in [105] considered the case of distributed systems and proposed a novel strategy for renewable energy integration.

It may be noted that, data centers are not only good candidates for increasing the renewable energy usage, but they can potentially contribute much more in realizing the full potential of smart grid. Perhaps more importantly, a medium to big sized data center can serve as a large scale storage system in the grid, which can be even more valuable in stabilizing the grid with a lot of Distributed Energy Resources (DER), which is growing very rapidly [106]. The recent work in [107] presented the potential research directions to employ effective DR mechanisms in data centers.

2.5 Cost Volatility Control

Chapter 5 of this thesis, designs a method to reduce the price risk a CSP is exposed to because of large variations in electricity price, renewable energy generation and service demand. To realize this aim, *hedging* is implemented through a simple derivative, namely forward contract. Forward contracts are contracts which specify future delivery of a given quantity (amount) of a particular commodity at a given fixed price (known at the beginning of the contract), before the expiry.

Following the deregulation of electricity market, various new market participants came up to replace the conventional vertically integrated structure [108, 109]. The major economic issue with this type of market is that the price varies over time. This is quite similar to real time price of commodities such as coal, gold etc. However, a major difference between inherent natures of

electricity market with other commodity markets is that electricity is not storable, so the consumers do not have the option of operating with an appropriately chosen level of inventory. The various models for constructing different derivatives in the commodity market with storable commodities may be found in [110]. It may be noted that the market for electricity trading is still in its nascent phase. However, as renewable penetration increases, the need for applying tools to mitigate the operational risk will become more important.

In finance, the strategies to mitigate risk associated with price volatility are referred to as hedging. Hedging is an efficient tool to deal with the uncertainty associated with a wide array of parameters whose future values are not known in advance. This can be realized through various kinds of tools such as forward contracts, options, futures and other derivatives [110]. A fairly rigorous mathematical background of these derivative products may be found in [111]. There is an extremely rich set of literature for dealing with price volatility through hedging in financial markets. This area is theoretically and practically quite mature in quantitative finance [110] and some strategies have also been developed and implemented on electricity trading in real time markets. A good summary of various products to mitigate price risk in the deregulated electricity market may be found in [112, 113].

Nevertheless, very little work has been done in cloud computing to mitigate the price risk the CSPs are exposed to. This is a topic which falls under the umbrella of *economic sustainability* of an individual CSP. Recently, the authors of [114] proposed a mechanism to deal with this uncertainty. Under a stochastic optimization framework they proposed a strategy leveraging the electricity price diversity in the day ahead and in real time electricity markets.

The work which is most relevant to the strategy presented in this thesis is [115], which deals with the electricity price risk the CSP is exposed to. It may be noted that the model of [115] did not consider switching costs, which would be of major concern to a CSP for implementing any power saving strategy. Additionally, that work did not consider the issue of integrating variable renewable energy sources. This thesis augments the work of [115] by additionally considering switching costs, renewable energy and carbon emission aspects, which make our work more comprehensive. Moreover, the optimization problem formulated in [115] is a Mixed Integer Linear Program (MILP), which is NP complete and computationally expensive. In contrast, the optimization problems we formulate for load distribution are convex problems, which can be solved efficiently and reliably by well-known polynomial time algorithms [116].

Our strong conviction is that this endeavor will act as a bridge between the community working on stochastic electricity pricing models to apply various tools [112, 113, 117, 118] prevalent in quantitative finance and the engineering community aiming at power cost reduction for the CSP.

Chapter 3

Energy Efficient Cloud Service Pricing: A Two-Timescale Optimization Approach

3.1 Introduction

This chapter considers the scenario of a Cloud Service Market (CSM) where several CSPs in operation are competing against each other to attract and serve the demands generated by the customers. The competition arises from the fact that the customers (the individuals or organizations) who generate demands for services in the market rationally choose the CSP which offers good Quality of Services (QoS) at a lower price. For pricing the services offered by the CSPs operating in the competitive CSM, an analytical framework is designed by considering the operational cost incurred by the CSP to service the given demand, the QoS offered and the prices other CSPs are charging. The proposed pricing strategy strikes a reasonable balance between charging too little which would result in irrationally low business profit and charging too much which would result in customer loss and thereby eventually loss of market share. In addition, the pricing strategy proposed promotes energy efficiency and renewable energy integration using a bi-level optimization strategy. The first level consisting of a *Slow Timescale Optimization Procedure (STOP)* addresses the economic efficiency related issues for cloud service pricing.

The second level involving a *Fast Timescale Optimization Procedure (FTOP)* performs energy efficient job scheduling. Numerical simulations are carried out to validate the proposed pricing strategy which is then compared with an oracle benchmark policy which enables a fair profit sharing among the CSPs. It may be noted that the oracle policy is a hypothetical one and it is not possible to implement that in practice. Nevertheless to carry out a conservative performance evaluation the comparison is carried out with this ideal policy. Furthermore, the proposed energy-aware scheduling policy is compared with a baseline scheduling policy, which is to deploy a constant number of servers throughout one timeslot of *STOP*. Since the process of turning off extra servers during low load period are far from being matured, at present many CSPs tend to always deploy constant number of servers [32]. The efficiency and effectiveness of the proposed bi-level strategy are demonstrated and the results are discussed.

In this chapter, economic sustainability is connected together with environmental sustainability, through a bi-level optimization approach. In the proposed bi-level optimization approach, the first problem which is solved in a slow timescale, addresses the distribution of service demand among the competing CSPs (based on a discrete time model). Since, in the wholesale electricity market, the price of electricity typically varies in hourly intervals, this slow timescale optimization problem is formulated and solved once every hour. The outputs of the slow timescale optimization problem are the prices and distribution of demands among the CSPs. To address the competition between the different CSPs, the demand serviced by a CSP is modeled as a function of the service quality offered by it, the price charged to the customer and the price charged by other competing CSPs. This approach provides a fairly exhaustive, practical model of a CSM where multiple CSPs compete among themselves to serve price-sensitive, rational customers. Inside this slow timescale problem, the fast timescale problem is solved

multiple times. To tackle this fast timescale optimization a MPC based strategy is implemented to perform the scheduling of deferrable jobs so as to maximize renewable energy integration into the grid. The approach of bi-level optimization over two time scales is a novel one which can conveniently take into account the variability, intermittency and randomness in the supply of renewable energy in the system.

In this chapter (more generally in this thesis) some ideas from smart grid operations are applied to increase the renewable energy integration. It is well recognized that big data centers are excellent candidates to contribute towards environmental sustainability. The first reason is that, they are big customers of electricity. These days a large data center may have hundreds of thousands of servers and other computing and networking infrastructures. The other reason is that, a significant proportion of the workload experienced by today's CSPs is deferrable [119], i.e., there is a time flexibility in execution of those jobs. This would provide an excellent opportunity to implement *demand response* (DR) to shape the demand to fall in line with the supply. This can be a very effective tool to tackle the joint effects of the intermittency, uncontrollability and uncertainty associated with the generation of renewable energy. The work in [120] presented a summary of various DR strategies proposed in the literature and the way these have been applied in the operations of a smart electric grid.

The pricing mechanism is formulated to incentivize renewable energy integration into the grid. A CSP with greater access to renewable energy can offer a lower price to its customers, thereby increasing the demand for that CSP. Here, the main challenge would be to determine appropriate pricing policies to effectively implement this. While a small change in price may not be effective in incentivizing renewable energy integration, too large a change in price may result in undesirable system oscillations which would make the CSM unstable and inefficient. It may

be noted that most of the previous strategies presented in the literature [90, 121, 122] are for cases where the data center and the renewable energy resource are co-located. This is typically the case when the same organization owns both the data center and its co-located renewable energy resource. The model proposed in this chapter is more general and can incentivize both onsite and offsite sources of renewable energy. This would indeed be the case in current scenarios where a large organization owning one or more data centers would rely on utility service providers for supplying the renewable energy. Typically they have a long term contract with the utility company to contribute towards environmental sustainability and meet the (sometimes quite aggressive) carbon emission goals set by the various regulating bodies [16], [8].

The work which is closest the one presented in this chapter is [46]. In that work the authors considered the competitive cloud service markets with more than one service provider and price sensitive customers. They provided an analytical framework which is partially adopted and appropriately improved upon in our work. They primarily focused on the economics of the CSM whereas the focus of this chapter is on formulating pricing strategies which increases renewable energy usage. The issue of economically sustainable pricing is bridged with the issue of environmental sustainability by appropriately incentivizing the usage of renewable energy. The novelty of this chapter lies in the fact that, a two timescale optimization approach is designed where the service pricing issue is addressed in a slower timescale and the energy aware job scheduling is performed in a faster timescale. The work in [46] did not consider the issue of energy efficiency. To the best of our knowledge, the work presented in [123] was the first to provide a rigorous method to consider the energy efficiency issue through service pricing, in a competitive CSM with price sensitive customers, which is presented in this chapter.

Although there have been some efforts to consider the energy efficiency issue through an appropriate pricing mechanism, none of these considered the holistic CSM. The work in [51] considered a combined framework where scheduler, charge back model and customer's preferences work in harmony for a green and energy efficient cloud computing systems. However, they have considered only the case of a single CSP, whereas we consider the entire CSM with many CSPs and inherently address the issue of competition between them. The work reported in [52] presented a similar cost optimization framework, with collaboration between the CSP and the utility provider to increase renewable energy penetration. They considered the situation where multiple geographically distributed data centers are owned by a single CSP. Both the works in [52] and [51] did not consider the cloud service market where there are many CSPs in a competitive cloud service market with price sensitive rational customers. Moreover none of the above mentioned works considered the adverse effects of switching a server from OFF to ON or vice versa, which is addressed in this chapter by considering the cost for server switching. Additionally, they have captured the renewable generation only indirectly through price, demand and supply difference, which in practice, may be less effective in terms of renewable energy integration. The previous works which aimed for green cloud computing through appropriate pricing did not consider the competitive CSM [67], [90], with competition between different CSPs. On the other hand the works focusing on cloud service pricing did not address the energy efficiency issue [41], [44], [43]. Additionally, in this chapter the pricing (performed in slow timescale) is merged with job scheduling (performed in fast timescale) to achieve higher renewable energy integration. This chapter augments the work in [46] by considering a more realistic model which reflects the adverse effects of server switching. This chapter formulates

and solves convex optimization problems at different timescales in a computationally efficient manner.

The rest of this chapter is organized as follows. In Section 3.2, the actual problem is introduced and the model of a CSM is described with multiple CSPs and price sensitive customers. In Section 3.3 the approach for load distribution is presented over a slow timescale. In Section 3.4, energy efficient scheduling problem is described which is applied in the fast timescale. In Section 3.5, the setup for evaluation of the proposed approaches is described and the results are discussed. Section 3.6 concludes the chapter.

3.2 Cloud Service Market Model with Multiple CSPs

In this section, the characteristics of a typical CSM are discussed and the model to capture the inherent dynamics of this CSM is presented. Presently, the market for cloud services is indeed in a nascent stage. As this market grows in size, different trading practices and pricing mechanisms are likely to be developed [124] so that an efficient market will be established with healthy competition between the participating CSPs and the price-sensitive, rational customers. The CSM consists of three fundamentally different entities, i) a set of regulatory bodies, ii) multiple CSPs competing with each other to generate business profit and, iii) price-sensitive, rational customers who generate service demands in the market. These three entities have different objectives and accountabilities as shown in Figure 3.1 given below.

Figure 3.1: Interaction between the three major entities in the cloud service market

The market regulator is responsible for auditing the service quality provided to the customers and ensuring a healthy competition between the CSPs operating in the CSM. They are also responsible for maintaining a stable and efficient market where no CSP can become a monopoly by acquiring the power to manipulate the market price by behaving strategically. On the other hand, if the competition is too fierce, then some of the CSPs would not be able to sustain themselves in the market and that would be detrimental to the cloud service ecosystem. The objective of this chapter is to provide some insightful information to the CSPs about the optimal way of determining the service prices to strike a preferred balance between earning

profits and attracting more customers. The inherent trade-off is that a CSP should earn sufficient profit (revenue – operational cost incurred) to sustain the business and at the same time it should not charge too much to the customers as that would result in subsequent loss of market share (demand for services). Additionally, and very importantly, the CSPs should provide an acceptable level of QoS to the customers. Even though the ultimate objective of a CSP is to make more profit, pricing its services too aggressively will be detrimental to its long-term business sustainability and survival in the CSM. Since customers are price sensitive, they would choose a CSP which can provide good QoS at lower cost than the others. For simplicity, it is assumed that the price information of the different CSPs is available to the regulatory bodies, which is a reasonably realistic assumption, since almost all the CSPs advertises the price charged for different services (resources) in their official website [125]. It may be noted that the various services provided by a CSP are further categorized into Infrastructure as a Service (IaaS) [126], Platform as a Service (PaaS) and Software as a Service (SaaS) as given in [9, 127]. However, in this chapter we do not consider these service specific details and our model is generally applicable to all these frameworks as long as the resource requirement for the given amount of demand is known.

3.3 Distribution of Service Demand between Multiple CSPs: *Slow Timescale Optimization Procedure (STOP)*

As mentioned earlier, the first level procedure of our proposed bi-level optimization strategy deals with the distribution of service demands among the multiple CSPs in operation. The output of this algorithm is the distribution of service demands among the multiple CSPs operating in the CSM. The objective is to address the economic aspect of the market and to investigate the optimal strategy to be adopted by the CSPs to determine the price for their services in this market. We refer to this as the *Slow Timescale Optimization Procedure (STOP)*, since this optimization is carried out at uniform intervals of a longer time period. At this level, the CSP has information about the price of electricity and can predict the renewable energy generation. The service price can either be determined by the regulator in collaboration with the CSPs or can be determined by exchange of information among the CSPs in a distributed manner as in [67]. We call this as pricing through coordinated regulation (CR). For the sake of simplicity, we use a centralized model here for implementing the two-stage optimization procedure. However, this can be extended to operate with periodic information exchange among the CSPs in a distributed and iterative manner, as done for the strategy proposed in [67]. This chapter considers both deferrable and non-deferrable jobs.

In this model, the time index for timeslots in *STOP* is denoted by T . The optimization problem is solved for Ω number of timeslots, in other words T goes from 1 to Ω . The deferrable jobs need to be completed before the stipulated deadline. For the non-deferrable jobs, the QoS (denoted by $s_i(T)$) offered by CSP i , is measured in terms of the average response time for the

execution of the jobs during timeslot T . Usually, there would be a maximum delay specified as per the Service Level Agreement (SLA) and the QoS is assumed to be equal to the difference between the upper bound on the average response time and the actual average response time [32].

$$s_i(T) = d_{\max}(T) - E[d_i(T)] \quad (3.1)$$

However, to address a more general case, this model can be extended to consider some percentile delay, like 99th percentile delay instead of mean, which would impose a tighter control on QoS. In that case, there may be an upper bound mentioned on the 99th percentile delay, say $d_{\max,99}(T)$ and then $s_i(T)$ may be defined as

$$s_i(T) = d_{\max,99}(T) - \phi_{99}\{d_i(T)\} \quad (3.2)$$

where $\phi_{99}\{.\}$ denotes the 99th percentile delay during timeslot T .

As $s_i(T)$ increases, the QoS improves and the demand for services to CSP i increases. This implies that demand function for CSP i is an increasing function in s_i and is described later in this section. For the service, we consider the average service demand requested by the customers in aggregate and the amount of demand serviced by CSP i , respectively, over timeslot T . Mathematically, these are obtained by:

$$\frac{1}{T} \int_0^T \lambda(\tau) d\tau = \Lambda(T) \quad (3.3)$$

$$\frac{1}{T} \int_0^T \lambda_i(\tau) d\tau = \Lambda_i(T) \quad (3.4)$$

where $\lambda(\tau)$ is the aggregated service demand generated from the customers at time instant τ and $\Lambda(T)$ is the average aggregated service demand generated from the customers during timeslot T . To model the demand function of a particular CSP, three crucial attributes are considered. These are i) the QoS (delay or response time) provided by the CSP, ii) the price it is charging to the customers and iii) the price charged by the other competing CSPs in the market. This modifies the model presented in [46] to consider the *profit* made by the other competing CSPs, i.e. how much are they charging their customers above the operational cost incurred by them to provide the service as per their SLAs. We are convinced that this would be a more effective parameter which would enable us to impose a tighter control in the market. The computation of the service price involves an inherent trade-off between earning profit and attracting demand in a competitive market with price sensitive rational customers. If one CSP chooses to go for a higher profit then it would subsequently face a reduction in demand. The sensitivity of the demand with respect to other CSPs' pricing is controlled through a parameter β_{ij} .

A high value of β_{ij} implies that, the demand for CSP i increases sharply when CSP j charges more to the customers. The demand function is formulated from the point of view of the CSPs. The inherent assumption is that the customers are generating service demands in the market and a set of CSPs are servicing these demands. The customers are actually insensitive to the profit the individual CSPs are making. They only care about the price they are paying. Once the customers submit their demands to the market, it is to be distributed among the CSPs following certain rules and the service price has to be determined. That rule is to minimize the total price paid by the customers, which is discussed later in this section. The demand function is modeled as per the following:

$$\Lambda_i(T) = x_i\{s_i(T)\} - y_i(pr_i\{\Lambda_i(T)\} - pr_i^{\min}\{\Lambda_i(T)\}) + \sum_{j \neq i} \beta_{ij}(pr_j\{\Lambda_j(T)\} - pr_j^{\min}\{\Lambda_j(T)\}) \quad (3.5)$$

where $\Lambda_i(T)$ is the average service demand (amount of jobs serviced) experienced by CSP i during timeslot T and $x_i(\cdot)$ is a function determined by CSP i to reflect the sensitivity of demand with respect to service quality. The variability in demand with the price charged by CSP i is captured by the second term of the RHS of equation (3.5) where y_i is a constant which may be different for different CSPs. This type of *scaled* decrease in demand with increase in price was also suggested in [46]. The price charged by CSP i to service the demand $\Lambda_i(T)$, during timeslot T is denoted by $pr_i\{\Lambda_i(T)\}$. The operational cost incurred by CSP i , for serving the demand, during timeslot T is denoted by $pr_i^{\min}\{\Lambda_i(T)\}$, which is the minimum price the CSP i must charge to the customers.

The *utility* function $U_i\{.\}$ is framed in a novel way to capture the two most important factors affecting the business performance of a CSP, namely,

- i) the *profit* earned, which is a good indicator of financial performance in the short term.
- ii) the amount of *demand serviced* by the CSP, which is a good indicator of the market share of the CSP in the competitive market and also reflects other comparatively less tangible parameters like, good will, market penetration etc., all of which translates into business sustainability over a longer period of time.

It is to be noted that, in the proposed model the price $pr_i\{\Lambda_i(T)\}$ already takes the demand served into account. Nevertheless, to explicitly consider the market share, the utility function is formulated as the product of the *profit* earned and the *demand serviced* by CSP i . Accordingly the utility function of CSP i , is given by:

$$U_i\{\Lambda_i(T)\} = \Lambda_i(T)(pr_i\{\Lambda_i(T)\} - pr_i^{\min}\{\Lambda_i(T)\}) \quad (3.6)$$

It may be noted that like $pr_i\{\Lambda_i(T)\}$, this is also a function of the demand $\Lambda_i(T)$. Since, this function is separable in price; we take the first order partial derivative with respect to $pr_i\{\Lambda_i(T)\}$,

$$\frac{\partial U_i\{\Lambda_i(T)\}}{\partial pr_i\{\Lambda_i(T)\}} = -y_i(pr_i\{\Lambda_i(T)\} - pr_i^{\min}\{\Lambda_i(T)\}) + \Lambda_i(T) \quad (3.7)$$

Now, equating (3.7) to zero, the relation between the optimum demand and the corresponding profit for CSP i is obtained. The second order derivative of the utility function, with respect to $pr_i\{\Lambda_i(T)\}$ is negative, indicating that to be the point of maxima.

The optimal price and optimal load are denoted by (*) with the respective symbol.

$$\Lambda_i^*(T) = y_i(pr_i^*\{\Lambda_i^*(T)\} - pr_i^{\min}\{\Lambda_i^*(T)\}) \quad (3.8)$$

Equating (3.8) to (3.5) we obtain the expression for the optimum profit to be earned by CSP i as

$$y_i(pr_i^*\{\Lambda_i^*(T)\} - pr_i^{\min}\{\Lambda_i^*(T)\}) = x_i\{s_i^*(T)\} - y_i pr_i^*\{\Lambda_i^*(T)\} + \sum_{j \neq i} \beta_{ij} pr_j^*\{\Lambda_j^*(T)\} \quad (3.9)$$

Now simple queueing theoretic formulae are used to get more insight on the pricing strategy as this would allow us to determine the number of servers to be deployed, given the average workload (here service requests) and the average response time per request. For this we use the M/M/k model as in [78]. This allows us to compute both the average response time $E[d_i(T)]$ and the number of servers required (denoted by $m_i(T)$) as given (3.10) and (3.11) ;

$$E[d_i(T)] = \frac{1}{\mu_i m_i(T) - \Lambda_i(T)} \quad (3.10)$$

where the service rate of a server in a data center owned by CSP i is denoted by μ_i . It is to be noted that the number of servers should be an integer. However, if the algorithm returns a non-integer value we round it up to the nearest integer greater than that. In data centers, typically $m_i(T)$ is quite large, often of the order of tens of thousands. This assumption is quite meaningful and would not impose any significant performance penalty in practice [32, 67].

$$m_i(T) = \left\lceil \frac{1}{\mu_i} (\Lambda_i(T) + \frac{1}{E[d_i(T)]}) \right\rceil \quad (3.11)$$

It is assumed that $x_i\{s_i(T)\}$ is a simple linear function in $s_i(T)$ and $\Lambda_i(T)$, with the multiplication constant α_i used to control the weightage of $x_i\{s_i(T)\}$ in the demand function of (3.5). In general, $x_i\{.\}$ has to be strictly increasing in $s_i(T)$, as per the form of $s_i(T)$ defined earlier.

Substituting the respective values we get:

$$x_i\{s_i(T)\} = \alpha_i \Lambda_i(T) s_i(T) = \alpha_i \Lambda_i(T) (d_{\max}(T) - E[d_i(T)]) = \alpha_i \Lambda_i(T) (d_{\max}(T) - \frac{1}{\mu_i m_i(T) - \Lambda_i(T)}) \quad (3.12)$$

Combining (3.8) to (3.12) the analytical expression for the optimum price to be charged by CSP i is given by (3.13).

$$\left. \begin{aligned} pr_i^* \{\Lambda_i^*(T)\} &= pr_i^{\min} \{\Lambda_i^*(T)\} + \frac{1}{2y_i} x_i\{s_i^*(T)\} + \frac{1}{2y_i} \sum_{j \neq i} \beta_{ij} (pr_j \{\Lambda_j^*(T)\} - pr_j^{\min} \{\Lambda_j^*(T)\}) \\ &= pr_i^{\min} \{\Lambda_i^*(T)\} + \frac{\alpha_i \Lambda_i^*(T)}{2y_i} (d_{\max}(T) - \frac{1}{\mu_i m_i(T) - \Lambda_i^*(T)}) + \frac{1}{2y_i} \sum_{j \neq i} \beta_{ij} (pr_j \{\Lambda_j^*(T)\} - pr_j^{\min} \{\Lambda_j^*(T)\}) \end{aligned} \right\}$$

(3.13)

The overall revenue generated by CSP i , from $T=1$ to Ω and its overall profit over this time interval are denoted by $\Theta_i(\Omega)$ and $\Pi_i(\Omega)$, respectively. These are given by (3.14) and (3.15), respectively.

$$\Theta_i(\Omega) = \sum_{T=1}^{\Omega} pr_i\{\Lambda_i(T)\} \quad (3.14)$$

$$\Pi_i(\Omega) = \sum_{T=1}^{\Omega} (pr_i\{\Lambda_i(T)\} - pr_i^{\min}\{\Lambda_i(T)\}) \quad (3.15)$$

In (3.13) presented above, y_i and β_{ij} control the sensitivity of price charged by CSP i , with respect to the other CSPs operating in the market. This attribute is not considered by the previous works [51] and [52]. The recent work presented in [128] connects the resource provisioning with profit earned by the CSP. They have leveraged economics and finance to capture the time value of money and proposed a risk adjusted resource pricing model. However, the work presented in [128] does not consider the competition among the CSPs, which we believe is a crucial issue in the competitive cloud service market with price sensitive, rational customers.

These results may now be used to compute the operational cost incurred by the CSP to service the demands. In order to focus on the operational energy efficiency of the CSPs, we consider both the electricity cost and renewable energy integration keeping environmental sustainability in mind. It may be noted that the cost implication for integrating renewable energy may vary significantly from one CSP to the other and is difficult to model under a common, unified framework. Driven by this fact the model is intentionally kept quite general so that it can incorporate any kind of operational cost that may be relevant for a CSP and for the system. For the renewable energy, we extend this model to incorporate long term contracts from the utility provider of renewable energy generator. This is a typical approach followed by large service providers which want to incorporate as much renewable energy as possible to demonstrate their

commitments towards environmental sustainability [16]. To realize this, we assume a simple incentive based model for integrating renewable energy, where the incentive for integrating a unit of renewable energy is denoted by δ_i , which can be adjusted as per the value of integration of one unit of renewable energy to CSP i . The special case of collocated renewable energy sources owned by the CSP can also be accommodated by choosing δ_i equal to the spot market electricity price $p_i(T)$.

$$pr_i^{\min}\{\Lambda_i(T)\} = (p_i(T)m_i(T)\gamma_i(T) - \delta_i(\min(m_i(T)\gamma_i(T), \int_0^T w_i(\tau)d\tau)))^+, \text{ where } (x)^+ = \max(x, 0) \quad (3.16)$$

To capture some operational details of the data center the parameter Power Usage Effectiveness (*PUE*) is chosen. This is calculated as the ratio between the total power consumed by the data center and the amount of power consumed by the IT systems. Naturally, it is always greater than 1. The PUE of the data center operated by CSP i , for the timeslot T , is denoted by $\gamma_i(T)$.

We now present the optimization procedure to be implemented, which we call as *STOP*. The aim is to minimize the total service cost paid by the customers to service their aggregate demand. This is equivalent to maximizing the *value for money* from the perspective of the customers, which can be perceived as the customer's utility function. We feel this would be the best way to design a stable, transparent market with healthy competition among the CSPs, where customers receive fair services, maximizing their utility.

We construct the *STOP* as,

$$\text{Minimize } \sum_{i=1}^N pr_i\{\Lambda_i(T)\}, \forall T \quad (3.17)$$

Subject to

$$\frac{1}{T} \int_0^T \lambda_i(\tau) d\tau = \Lambda_i(T), \forall T \quad (3.18)$$

$$\sum_{i=1}^N \Lambda_i(T) = \Lambda(T), \forall T \quad (3.19)$$

$$\Lambda_i(T) \geq 0, \forall i, T \quad (3.20)$$

Here (3.17) represents the objective function and (3.18) to (3.20) represent the constraints of the optimization problem. The constraint in (3.18) accounts for the average demand rate experienced by CSP i , during timeslot T . The constraint in (3.19) ensures that the total demand generated by the customers is serviced by some CSP. The constraint (3.20) represents an operational feasibility constraint that for every timeslot demand experience by all the CSPs is non-negative. The computational complexity of the above problem depends on the structure of the functions involved. However, in the proposed model, this is easy to tackle since the objective function and the feasible set (defined by the intersection of the constraints) are convex in the decision variables. Therefore, there are reliable, computationally efficient and low footprint algorithms to solve the above problem numerically with very high degree of accuracy [129]. Moreover, since this is a convex optimization problem the optimal solutions can be constructed by expanding as per Karush–Kuhn–Tucker (KKT) conditions [129]. Since the price charged by each CSP is affected by the price charged by other CSPs, we call this optimization problem as pricing through coordinated regulation (CR) with strictly positive β_{ij} . In fact under a distributed setting this optimization problem can even be solved by iterative exchange of information among all the CSPs.

To obtain additional insights about the structure of the optimal pricing policy, a simple market is considered with two CSPs and a finite number of customers generating service demands. Following the steps described here, the closed form expression for the optimum price to be charged by the two CSPs, are derived in (3.21) to (3.24). It is assumed that the demand does not exceed the sum of the individual capacities of the two CSPs.

$$\begin{aligned}
 pr_1\{\Lambda_1(T)\} &= pr_1^{\min}\{\Lambda_1(T)\} + \frac{\alpha_1\Lambda_1(T)}{2y_1} \left(d_{\max}(T) - \frac{1}{\mu_1 m_1(T) - \Lambda_1(T)}\right) \\
 &+ \frac{\beta_{12}}{2y_1} (pr_2\{\Lambda_2(T)\} - pr_2^{\min}\{\Lambda_2(T)\})
 \end{aligned} \tag{3.21}$$

$$\begin{aligned}
 pr_2\{\Lambda_2(T)\} &= pr_2^{\min}\{\Lambda_2(T)\} + \frac{\alpha_2\Lambda_2(T)}{2y_2} \left(d_{\max}(T) - \frac{1}{\mu_2 m_2(T) - \Lambda_2(T)}\right) \\
 &+ \frac{\beta_{21}}{2y_2} (pr_1\{\Lambda_1(T)\} - pr_1^{\min}\{\Lambda_1(T)\})
 \end{aligned} \tag{3.22}$$

More compactly,

$$pr_1\{\Lambda_1(T)\} = \begin{bmatrix} 1 \\ \frac{\alpha_1}{2y_1} \\ \frac{\beta_{12}}{2y_1} \end{bmatrix}^T \begin{bmatrix} pr_1^{\min}\{\Lambda_1(T)\} \\ \Lambda_1(T) \left(d_{\max}(T) - \frac{1}{\mu_1 m_1(T) - \Lambda_1(T)}\right) \\ pr_2\{\Lambda_2(T)\} - pr_2^{\min}\{\Lambda_2(T)\} \end{bmatrix} \tag{3.23}$$

$$pr_2\{\Lambda_2(T)\} = \begin{bmatrix} 1 \\ \frac{\alpha_2}{2y_2} \\ \frac{\beta_{21}}{2y_2} \end{bmatrix}^T \begin{bmatrix} pr_2^{\min}\{\Lambda_2(T)\} \\ \Lambda_2(T) \left(d_{\max}(T) - \frac{1}{\mu_2 m_2(T) - \Lambda_2(T)}\right) \\ pr_1\{\Lambda_1(T)\} - pr_1^{\min}\{\Lambda_1(T)\} \end{bmatrix} \tag{3.24}$$

As mentioned in (3.23) and (3.24), the price set by CSP-1, depends on the vector consisting of

- i) The cost incurred to service the demand
- ii) The product of demand serviced and the offered QoS
- iii) The profit made by the other CSP

Among these three terms, the first and the third are intuitively obvious. The second term actually suggests that the optimum price to be charged by the CSP is affected by the amount of demand serviced (which translates into market share) and the QoS offered. More importantly, the optimal price contains the product of these two parameters, which implies that the CSP can either choose to service more demand with a slightly worse QoS or less demand with a slightly better QoS. Usually, there is a defined standard set for the QoS and the CSPs compete for attracting demand. This strongly suggests that the CSP which can operate at a higher efficiency level with lower internal cost will always position itself better in the market compared to the others. At the same time, in this proposed model the cost incurred to service the demand request does not account for any profit component. The profit to the CSP is obtained from the second and the third term in the expansion of (3.23) and (3.24). These profits come from the amount of demand serviced by the CSP, the QoS provided to the customers and the pricing strategy of other CSPs. The explicit presence of the term reflecting the pricing strategy of other competing CSPs (Figure 3.1) imposes a distributed control over the market to ensure that no CSP can charge an excessive amount to the customers. Moreover, from the CSP's point of view, these pricing strategies prevent the CSP from charging the customers too much so that it does not lose its market share, as this would eventually make it uneconomical for the CSP to continue operations.

The first vector in both (3.23) and (3.24) includes the model parameters defined by the market. In an emerging market, where sufficient knowledge about the parameters is not available, these parameters will have to be learned over a period of time through the optimum trade-off curve. In our numerical simulations we have determined these parameters by trial and error.

3.4 Energy Aware Scheduling for Maximizing Renewable Energy Integration

Fast Timescale Optimization Procedure (FTOP)

Over the past few decades, a major challenge faced by mankind is the issue of environment sustainability. To tackle this issue a major topic has emerged, called *smart grid*. The high level objective of smart grid is to generate, transmit, distribute and consume electricity in a more efficient way than the traditional methods. Several crucial strategies that are prevalent in the domain of smart grid have improved the way the data centers are operated. These include demand response, distributed generation, pricing etc. The recent work in [130] presented an extensive survey on energy efficiency in communication networks, and energy efficiency in data centers [130]. The authors presented some open issues and discussed the major research challenges to make further progress in this area. Another notable work which covers the state-of-the-art theory, key strategies, protocols, and the implementation strategies in the context of various applications related to communication and networking technologies for the smart grid, is [131].

In this section, the proposed method for scheduling deferrable jobs so as to maximize the integration of renewable energy, is described. This procedure is executed multiple times within a

single *STOP* window. In *STOP*, we did not consider the scope of time scheduling of the deferrable service requests, which is quite common in the case of actual data center workload [119, 122]. A few examples of these kind of deferrable jobs are data analytics on large data sets, scientific simulation, high quality medical image processing, system backup, and various other scientific computations based services. The QoS of these types of jobs is that they have to be serviced before the user's specified deadline where the deadline would be given at the time the demand request for the service is made. These types of jobs can conceptually be thought of as consumers of energy rather than consumers of power, and as a result scheduling becomes a natural choice. In this section, we present our MPC based job scheduling algorithm to achieve higher renewable energy integration by leveraging on the deferrable demand requests; this is referred to as *Fast Timescale Optimization Procedure (FTOP)*. The theoretical foundations of MPC and the implementation algorithms may be seen in [83]. Applications of MPC based job scheduling algorithms in other domains of smart grid are available in [79]. To explicitly differentiate between *STOP* and *FTOP* we use T and t , respectively to denote a single timeslot. In *FTOP*, if T appears, it implies that the quantity in question is to be obtained from the corresponding *STOP*.

In *FTOP*, the two crucial aspects, namely the integration of renewable energy and the number of times a server is switched from OFF to ON (and vice versa) are appropriately addressed. The commercial servers typically deployed in today's data centers are not at all energy proportional. In other words, they consume a significant amount of power, even when at idle condition [9]. In this model, the timeslots in *FTOP* are indexed by t , and the optimization horizon (the window size with which the optimization problem is solved) is denoted by ω . The deferrable tasks are indexed by z . The arrival time and the deadline of the deferrable task z are

denoted by a_z and d_z , respectively. The total service requirement of deferrable task z is denoted by $\sigma^{D,z}$. The superscript D , corresponds to the fact that the job is deferrable. Number of active servers during timeslot t deployed by CSP i is denoted by $n_i(t)$. The set of all active (unfinished) jobs which require non-zero resource at time t is denoted by Z_t . The number of servers owned by CSP i is denoted by N_i , which is the maximum number of servers CSP i , can deploy to serve the customer demand at any point of time. Total amount of demand serviced by CSP i , during timeslot t is denoted by $L_i(t)$. Naturally, this total demand is the sum of the total amount of non-deferrable demand ($L_i^{ND}(t)$) and the total amount of deferrable demand ($L_i^D(t)$) serviced by CSP i , during timeslot t . Similarly, the total deferrable demand and non-deferrable demand experienced by CSP i , during the entire period T (the corresponding timeslot of the *STOP*) are denoted by $\Lambda_i^D(T)$ and $\Lambda_i^{ND}(T)$, respectively. The amounts of jobs scheduled at timeslot t , for deferrable task z , performed by CSP i is denoted by $l_i^{D,z}(t)$.

As of now, the most prominent and well accepted approach to achieve energy efficiency in data centers is to turn off the extra servers during low demand periods and to make the system power proportional [32]. The main challenge to implement this is the fact that the demand is uncertain and sometimes varies dramatically with time. As a result, switching servers ON and OFF as per the demand, may incur very frequent server switching. The adverse effects of server switching are i) delay cost, ii) energy cost, iii) increased wear and tear and iv) perceived risk of reducing server lifetime. The work [32] provides a detailed discussion on the various adverse effects of server switching.

$$J_{\varepsilon,i}(t) = (n_i(t)\gamma_i - w_i(t))^+ \quad (3.25)$$

where the brown energy used by CSP i during timeslot t , is denoted by $J_{\varepsilon,i}(t)$ and the renewable energy available to CSP i during timeslot t , is given by $w_i(t)$.

$$J_{\phi+,i}(t) = (n_i(t) - n_i(t-1))^+ \quad (3.26)$$

$$J_{\phi-,i}(t) = (n_i(t-1) - n_i(t))^+ \quad (3.27)$$

The brown energy consumption, during timeslot t , is presented in (3.25). The number of servers switched from OFF to ON (denoted by $J_{\phi+,i}(t)$) and from ON to OFF (denoted by $J_{\phi-,i}(t)$) during timeslot t , are given by (3.26) and (3.27), respectively. In *FTOP*, we minimize a linear combination of (3.25) to (3.27) with a set of constraints to ensure QoS and operational feasibility. This model is not novel and has been implemented by [32]. We augment the work in [32] by considering deferrable workload and the renewable energy integration issue. Therefore, the novelty of our work lies in the integration of *FTOP* and *STOP*, to construct a bi-level optimization procedure to implement a cloud service pricing strategy which would also promote energy efficiency in the cloud service market. The *FTOP* is given below.

$$\text{Minimize} \quad \sum_{t=1}^{\omega} (J_{\varepsilon,i}(t) + \theta_{i+} J_{\phi+,i}(t) + \theta_{i-} J_{\phi-,i}(t)) \quad (3.28)$$

Subject to:

$$\Lambda_i(T) = \Lambda_i^D(T) + \Lambda_i^{ND}(T), \quad \forall i, T \quad (3.29)$$

$$L_i^{ND}(t) + L_i^D(t) = L_i(t), \quad \forall t, i \quad (3.30)$$

$$\sum_{z \in Z_i} l_i^{D,z}(t) = L_i^D(t), \quad \forall t \in \{1, 2, \dots, \omega\} \quad (3.31)$$

$$\sum_{t=1}^{\omega} L_i(t) = \Lambda_i(T), \quad \forall t, T \quad (3.32)$$

$$\sum_{t=a_z}^{d_z} l_i^{D,z}(t) = \sigma^{D,z}, \quad \forall z \quad (3.33)$$

$$l_i^{D,z}(t) \geq 0 \quad (3.34)$$

$$n_i(t) \geq \frac{1}{\mu_i} (L_i(t) + \frac{1}{d_{\max}}), \forall i, t \quad (3.35)$$

$$0 \leq n_i(t) \leq N_i^{\max}, \forall i, t \quad (3.36)$$

$$n_i(t) \in \mathbb{N}_+ \quad (3.37)$$

The constraints in *FTOP*, are presented in (3.29) to (3.37). The total demand serviced by CSP i is the amount of demand allocated to CSP i , as obtained through *STOP*; this is the sum of both the deferrable and non-deferrable service demands. This condition is imposed through (3.29). This condition could have been incorporated in *STOP*, but since we formulate one more level of optimization we decided to incorporate everything which deals with energy efficient scheduling in *FTOP*. Similarly, in every *FTOP* operation the amount of demand serviced is the sum of both the deferrable and non-deferrable demands during each timeslot t as given by (3.30). For CSP i , during timeslot t , the amount of demand serviced for deferrable demand is equal to the aggregated sum of demand services to the individual deferrable jobs, which belongs to Z_t . This condition is imposed through (3.31). The total service demand served during $t = 1$ to ω , is equal to the total service demand received by CSP i as a result of the corresponding *STOP*; this is enforced through (3.32). This is a step where the two levels of the bi-level optimization strategy namely *STOP* and *FTOP* are directly coupled. For all deferrable jobs z , the total amount of service provided during arrival and deadline has to be equal to the size (service requirement) of the job; i.e., all the deferrable jobs are to be serviced before the respective deadline is ensured through (3.33). The amount of demand serviced in timeslot t , corresponding to service request z , has to be non-negative as imposed through (3.34). The QoS guarantee for servicing the demand is given in (3.35). This is obtained from a M/M/k queueing model, which is an appropriate model for data center systems [78]. The number of servers deployed by CSP i for the timeslot t is non negative and upper bounded by the total number of servers owned by CSP i ; this is imposed

through (3.36). This procedure builds the framework for scheduling the deferrable jobs in line with the availability of renewable energy, which is the *demand response* we implement through *FTOP*. Finally, (3.37) represents the fact that $n_i(t)$ can be a non-negative real number and not necessarily an integer. This assumption is adopted by some previous works such as [32, 67] and is justified by the fact that typically $n_i(t)$ is quite large and a round off to the nearest integer (greater than the obtained solution) does not impose any penalty. This simplification results in a convex optimization problem which is computationally tractable.

The the algorithms for *STOP* and *FTOP* are implemented in CVX (academic licensed version) [132], which is a package for solving disciplined convex optimization problems built on the MATLAB platform. We choose to solve *STOP* and *FTOP* by SDPT3 subroutines [133], which is the default solver in the academic licensed version of CVX. All simulations are performed on a HP Pro 3130 MT machine with an i5 processor (2.67 GHz) and 4.0 GB of RAM. At each instance, both *STOP* and *FTOP* takes less than 1 second to converge in MATLAB version 7.11.0. This time is negligible in comparison with the timeslot length of *FTOP*.

3.5 Performance Evaluation

This section describes the methodologies and the set up under which the performances of the proposed algorithms are evaluated. In order to examine the efficiency and effectiveness of the proposed algorithms in real systems, case studies are carried out under realistic parameter settings and the results are discussed. It is assumed that there are four CSPs in operation which are competing in the CSM as shown in Fig. 3.1. It is further assumed that each of the four CSPs

operates one data center, located in California, New York, Pennsylvania and Indiana, respectively. (It should be noted that the proposed model can indeed address the case where some of the CSPs own multiple data centers.) The electricity prices, renewable generation and PUE values are correspondingly shifted in time to emulate a practical scenerio. The electricity price data is obtained from [134]. For the *STOP*, the time interval of solving the optimization problem is 1 hour, whereas for *FTOP*, it is 10 mins. This implies that the demand distribution is performed in a one-hour timescale and the energy efficient job scheduling is performed in a 10-min timescale. The renewable energy trace is obtained from [106]. The overall service demand is collected from a commercial server cluster owned by Facebook, which was used in [106]. The PUE data is obtained from a commercial data center owned by Google, which was also used in [106]. For the deferability of the workload, the slacks are generated by following a normal distribution with a mean of 30 mins and standard deviation of 10 mins. It is assumed that 50% of the total demand is deferrable where the extent to which a demand can be deferred is provided by the customer at the time of submission of the demand request.

In the core of the proposed pricing strategy, there is a price regulator which regulates the price charged by the individual CSPs in a collaborative manner. Alternatively, the CSPs may coordinate among themselves to iteratively determine the optimal prices by considering the prices charged by the other competitor CSPs. It may be noted that, strategically charging different prices to different customers as per the probability that a particular customer may move to some other competitor, is already prevalent in the US airline industry [135]. In the context of CSM, the market is still in a nascent stage and thus we believe sophisticated and customized pricing schemes will indeed be adopted by the market in the near future, as the market becomes more matured. We refer to this as pricing through coordinated regulation (CR) where the

individual CSPs coordinate with both the regulator and the other CSPs to create a mutually beneficial economically sustainable CSM. The performance of this strategy is evaluated by comparing this with a baseline scenario without CR. The baseline scheme is realized by setting $\beta_{ij} = 0$ in (3.5) as presented above. This means that as per the baseline strategy the regulator does not consider the price charged by other CSPs operating in the CSM. Under this baseline scheme, the demand attracted by a particular CSP is a function of the price that the CSP is charging to the customers and the QoS level that it provides.

3.5.1 Benchmark strategy for comparing fairness of revenue and profit

Let $\Theta(\Omega)$ be the total revenue generated in the market (i.e., when aggregated across all the CSPs operating in the market) and $\Pi(\Omega)$ be the total profit earned aggregated across all the CSPs. Let $\Theta_i(\Omega)$ be the revenue generated by CSP i and $\Pi_i(\Omega)$ be the profit earned by CSP i . $\Theta(\Omega)$ and $\Pi(\Omega)$ are calculated by

$$\Theta(\Omega) = \sum_{i=1}^N \Theta_i(\Omega) \quad (3.38)$$

$$\Pi(\Omega) = \sum_{i=1}^N \Pi_i(\Omega) \quad (3.39)$$

The benchmark revenue for CSP i , (denoted by $\Theta_i(\Omega)$) is obtained by:

$$\Theta_i(\Omega) = \frac{\Theta(\Omega)}{\sum_{T=1}^{\Omega} \sum_{i=1}^N \Lambda_i(T)} \left(\sum_{T=1}^{\Omega} \Lambda_i(T) \right) \quad (3.40)$$

Similarly, the benchmark profit for CSP i , (denoted by $\Pi_i(\Omega)$) is obtained by:

$$\Pi_i(\Omega) = \frac{\Pi(\Omega)}{\sum_{T=1}^{\Omega} \sum_{i=1}^N \Lambda_i(T)} (\sum_{T=1}^{\Omega} \Lambda_i(T)) \quad (3.41)$$

Economic efficiency of the proposed pricing scheme is compared with this fair benchmark pricing scheme. The results demonstrate that the proposed method with *CR* substantially outperforms the scenerio with no *CR* in terms of economic efficiency and fairness. The benchmark profit and revenue are obtained by sharing the profit and revenue generated in the market as per the proportion of the demand serviced by them as given by (3.38) through (3.41). We carry out the comparison by computing the deviations of the profit and revenue earned by the CSPs from the profit and revenue under this benchmark scheme, before and after the implementation of *CR*. It is observed that there is remarkable improvement in deviation from the benchmark scheme after the implementation of *CR*. This benchmark pricing scheme is a hypothetical oracle policy since it requires prior knowledge of demand distribution, revenue and profit and would be impossible to implement in practice. Hence this comparison provides a conservative evaluation of the efficiency of the proposed strategy. Additionally, the benchmark pricing scheme is inefficient in terms of both energy efficiency and the objective of promoting an energy efficient system, which is the primary issue addressed in this chapter. However, this benchmark pricing scheme is theoretically quite straight forward, and fair in terms of revenue and profit generation from the CSP's point of view and rational in terms of the service cost incurred by the customers.

3.5.2. Results and discussions

This subsection presents the results obtained through numerical simulations to evaluate the performance of the proposed pricing strategy in terms of fairness, economic efficiency, and sustainability. Each of these has specific implications in this context. By fairness we mean how closely $\Theta_i(\Omega)$ tracks $\Theta_i(\Omega)$ and $\Pi_i(\Omega)$ tracks $\Pi_i(\Omega)$. Quantitatively the difference in percentage in revenue under the benchmark scheme and our proposed strategy, for CSP- i is determined by $\frac{100 \times (\Theta_i(\Omega) - \Theta_i(\Omega))}{\Theta_i(\Omega)}$. Similarly, the difference in profit under the benchmark

scheme and our proposed strategy, for CSP- i is obtained by $\frac{100 \times (\Pi_i(\Omega) - \Pi_i(\Omega))}{\Pi_i(\Omega)}$. These

quantities are evaluated before and after implementation of CR . These two quantities serve as good indicators of business sustainability of the operating CSPs and thereby determine the economic efficiency of the pricing scheme. As an instance, if any CSP is earning profit and revenue significantly higher than the fair profit and revenue then eventually the customers will go away from this CSP and the CSP will get extinct from the market. We demonstrate that, the proposed CR prevents the CSPs from that situation and thereby creates a sustainable market. Additionally, under the proposed pricing scheme both the revenue and profit should increase as the amount of demand serviced increases. This can be seen from Figure 3.2 which shows that the profit does increase in a normal and expected fashion. A similar trend of increasing revenue for increasing demand is also observed for the revenue, as shown in Figure 3.3. Table 3.1 presents the deviation from the fair profit and revenue earned by each of the CSPs before and after the implementation of CR to demonstrate the improvement. The results indicate that after CR the profit and revenue are significantly closer to the fair ones (refer to Figure 3.2 and Figure 3.3).

(a)

(b)

Figure 3.2: Profit earned by the CSPs before and after implementation of CR: (a) before, (b) after

Figure 3.3: Revenue earned by the CSPs before and after implementation of CR: (a) before, (b) after

We present the aggregate profit and revenue earned by the CSPs before and after implementation of *CR*, in Figure 3.4. This actually indicates which CSPs are earning more profit and revenue after applying *CR*. It is to be noted that the ones earning less revenue and profit after application of *CR*, are benefited in longer run. Otherwise this is quite evident that these CSPs will get extinct from the market due to price sensitivity of the rational customers.

Table 3.1: Deviation of profit and revenue earned by the CSPs from the benchmark pricing scheme with and without *CR* (in percentage)

	Without <i>CR</i>			
	CSP-1 over benchmark	CSP-2 over benchmark	CSP-3 over benchmark	CSP-4 over benchmark
Profit	-17.35	-9.8	20.45	11.5
Revenue	-7.53	1.96	6.63	2.47
	With <i>CR</i>			
	CSP-1 over benchmark	CSP-2 over benchmark	CSP-3 over benchmark	CSP-4 over benchmark
Profit	-0.48	-0.12	1.14	0.02
Revenue	-7.07	2.93	4.96	2.36

(a)

(b)

Figure 3.4: Comparison of total profit and revenue earned by individual CSPs before and after implementation of CR: (a) Profit, (b) Revenue

In our simulation CSP-1 is the only CSP which is making a profit much lower than the benchmark profit and serving the highest demand of the market, i.e., in terms of demand CSP-1 has the highest market share. After application of *CR* the profit earned by CSP-1 has gone up. The same thing has happened with CSP-2 as well. For CSP-3 we observe that it has earned significantly higher profit as compared to the benchmark strategy, before applying *CR* (Figure 3.2 (a)). If this continues then, eventually customers will realize this and move to some other CSPs and as a result CSP-3 will become extinct from the CSM. This is controlled by *CR* since after applying this CSP-3 has reduced the profit margin to cope up with the market competition. Same thing has occurred with CSP-4 as well, which was also charging more without *CR*.

In Figure 3.5 and Table 3.2, the cumulative profit, revenue earned by the CSPs and the demand serviced by them are provided and the average profit earned per request by the CSPs is presented. This is indeed an important parameter which should not vary too much among the CSPs in an efficient CSM. Results indicate the implementation of *CR* significantly flattens the average profit earned by the CSPs, which means *CR* creates a fair, economically efficient and sustainable CSM.

(a)

(b)

Figure 3.5: Average profit per request earned by the CSPs and amount of demand served by the CSPs before and after implementation of CR: (a) Before, (b) After

Table 3.2: Cumulative profit, revenue earned by the CSPs and the amount of demand serviced by them

	Without CR			
	CSP-1	CSP-2	CSP-3	CSP-4
Average profit	0.4228	0.4519	0.6237	0.5607
Average revenue	18.52	20.3148	21.3306	20.4227
No. of requests	18782000	15593000	9209800	13833000
Cumulative profit	7941030	7046476.7	5744152.3	7756163.1
Cumulative revenue	347885839	316768676	196450560	282507209
	With CR			
	CSP-1	CSP-2	CSP-3	CSP-4
Average profit	0.5051	0.5069	0.5133	0.5076
Average revenue	18.7842	20.7182	21.1012	20.5987
No. of requests	18645000	15037000	8934700	13977000
Cumulative profit	9417590	7622255.3	4586181.5	7094725.2
Cumulative revenue	350231409	311539573	188532892	287908030

(a)

Figure 3.6: Profit earned by the CSPs and the amount of demand (no. of requests serviced) before and after implementation of CR: (a) before, (b) after

(a)

Figure 3.7: Revenue earned by the CSPs and the amount of demand (no. of requests serviced) before and after implementation of *CR*: (a) before, (b) after

It is observed that the average profit (per job) earned by both CSP-1 and CSP-2 increase after applying *CR* while that of CSP-3 and CSP-4 decrease correspondingly. Therefore this regulation will protect CSP-3 and CSP-4 from charging more to the customers; if they do not do this then they may eventually get extinct from the market (Figure 3.5). Figure 3.6 presents how the profit grows with the amount of demand serviced before and after implementing *CR* for all the four CSPs. Similarly, the growth of revenue with the amount of demand serviced is presented in Figure 3.7. It is observed that in terms of revenue, the effect of *CR* is quite small as compared to that in terms of profit.

The results of the proposed energy efficient scheduling through *FTOP*, is presented with respect to a baseline scheduling policy. The baseline scheduling is the one where after allocation of demand (as per *STOP*), the demand is serviced uniformly among the 6 timeslots. This is a straight forward policy for scheduling and does not consider any attribute of environmental sustainability. It may be noted that this strategy is currently adopted by some CSPs [86]. On the

other hand this strategy does not perform any switching of servers during one timeslot (*STOP*) and hence no switching cost is incurred. In terms of energy cost we see that our proposed methods perform much better than this baseline scheme. The comparison results with respect to cumulative electricity cost incurred, cumulative renewable energy used, cumulative number of server switchings from OFF to ON and from ON to OFF are presented in Figure 3.8, 3.9 and in Table-3.3. The results indicate the proposed MPC based scheduling scheme *FTOP*, performs much better than the baseline scheduling scheme in terms of cumulative electricity cost incurred, cumulative renewable energy used, cumulative number of server switchings from OFF to ON and from ON to OFF.

Figure 3.8: Cumulative electricity cost incurred and cumulative renewable energy used by each CSP under our proposed technique for job scheduling *FTOP* and a benchmark scheme: a) Cumulative electricity cost incurred, b) Cumulative renewable energy used.

Figure 3.9: Cumulative number of server switchings and comparison with a baseline strategy: a) Cumulative switching from OFF to ON, b) Cumulative switching from ON to OFF.

Table-3.3: Comparison of *FTOP* with a baseline scheme

Parameter	CSP-1 over baseline	CSP-2 over baseline	CSP-3 over baseline	CSP-4 over baseline
Cumulative electricity cost	-23.3% (↓)	-20.54% (↓)	-23.24% (↓)	-26.33% (↓)
Cumulative renewable energy used	5.32% (↑)	5.23% (↑)	5.19% (↑)	5.57% (↑)
Cumulative switching: OFF to ON	-89.1% (↓)	-57.85% (↓)	-76.45% (↓)	-90.33% (↓)
Cumulative switching: ON to OFF	-50.79% (↓)	-28.66% (↓)	-42.96% (↓)	-52.27% (↓)

3.6 Summary

In this chapter we have considered a CSM, consisting of multiple CSPs which compete with each other with each aiming to maximize its revenue. However, the customers in the CSM (individuals or organizations) are price sensitive, i.e., they prefer to direct their demands to the CSP which would offer better service at a lower cost. Three different entities are considered, namely i) regulatory bodies, ii) CSPs and iii) the customers present in the CSM. Each of them has different objectives and plays a different role in the market. The main contribution of this chapter is to suggest the pricing strategy to be used by the CSPs, which would be paid by the customers to get their demands serviced. The primary objective is to create a fair, sustainable and economically efficient cloud service market. Additionally, we enable this pricing strategy to promote energy efficiency and more importantly renewable energy integration by realistically modeling the incentives prevalent in the market. The issue of economic sustainability is associated appropriately with the issue of environmental sustainability. This objective is achieved by proposing a bi-level optimization strategy, where two optimization problems are solved over different timescales. In the slow timescale, the market demands are distributed among the CSPs and the pricing related issues are addressed. An effective strategy is then proposed where the individual CSPs coordinate with the regulator to determine the service price charged by the CSPs. In the fast timescale, we have leveraged a MPC based energy-aware scheduling procedure to achieve energy efficiency by increasing the integration of renewable energy while considering of adverse effects of server switching from OFF to ON and vice-versa. We have carried out simulation with realistic cloud service demands with four CSPs and

demonstrated the efficiency of this strategy in terms of fairness, economic efficiency and business sustainability. We have integrated and aligned these two optimization problems to propose an energy efficiency aware pricing scheme for cloud market to promote and incentivize green cloud computing. To evaluate the performance of the proposed energy efficient scheduling technique we have carried out numerical studies with real life electricity price, renewable energy generation and workload data traces.

Chapter 4

Demand Response in Data Centers Through Energy Efficient Scheduling and Simple Incentivization

4.1 Introduction

This chapter focuses on a single data center to address some practical operational details. Towards the objective of improving economic sustainability, this chapter aims to, i) reduce the capital cost by reducing the number of servers, and ii) reduce the operational cost by reducing the total electricity cost incurred by the Cloud Service Provider (CSP). Towards the objective of improving environmental sustainability, this chapter aims to increase the renewable energy usage by energy aware scheduling. Furthermore an incentivization technique is developed to achieve more aggressive reduction of electricity cost and increase in renewable energy usage. The previous chapter focused on the high level dynamics of a Cloud Service Market (CSM). This chapter considers a single data center where various services are deployed in different tiers (which is typically the case for a real large scale system) so as to capture the crucial intricacies of real operational details. To tackle resource provisioning for deferrable jobs, Model Predictive Control (MPC) based scheduling is used. This chapter develops strategies to formulate and solve the optimal server provisioning problem in two different ways, namely (i) minimization of

operational cost and (ii) joint minimization of capital and operational costs, at different tiers, based on a discrete time model. Additionally, a simple and effective mechanism is designed to facilitate more renewable integration into the grid through monetary incentives on deadline deferral. It is demonstrated that the proposed method reduces the cumulative electricity cost and increases the cumulative renewable energy used, with no additional switching cost.

It may be noted that the previous chapter did not make any effort to investigate the opportunities to reduce the capital cost incurred by a CSP, which a critical contribution is made by this chapter. This chapter explores the opportunity for reduction in capital investment by reducing the peak server provisioning. To focus on operational details of a CSP this chapter considers a practical situation of a data center, where different services (applications) are provisioned in different tiers. Under this tiered architecture, the real intricacies of deferrable jobs are modeled. This chapter also considers the cost (adverse effect) of toggling a server (switching from OFF to ON and vice versa), similar to the previous chapter. Some preliminary results to highlight the usefulness of scheduling in data centers were presented in our previous work [136]. In fact, this chapter presents more rigorous analyses by providing sensitivity analysis with respect to some important model parameters, such as weightage on electricity cost in the optimization problem and the length of optimization horizon and also provides more results. Additionally, we propose here a novel and simple technique for increasing renewable energy usage by providing incentives for deadline deferral and evaluate its performance. Moreover, some additional attributes are also incorporated (briefly outline a heuristic method to improve cooling system efficiency, provide analysis on effect of electricity price forecast error etc.) to address some meaningful issues, which augment the work in [136]. The contents presented in this chapter are mostly based on the work in [137].

The rest of the chapter is organized as follows. Section 4.2 describes the model through which the resource provisioning issue is appropriately captured for a data center with tiered deployment of different functions and applications. Section 4.3 describes the model for implementing energy efficient scheduling of deferrable but deadline-sensitive jobs by accounting for the deferability (often referred to as *slack* [103] of the job request. It then proposes two ways of formulating the optimal server provisioning strategy at different tiers to provide some insights to the CSPs about the tradeoff between certain parameters.

Figure 4.1: Tiered architecture of data center

The first one (of the two formulations) considers the operational cost and switching cost. In the second one, apart from operational cost minimization, the issue of peak server provisioning is also considered and the switching cost is captured in a different way. Section 4.4 proposes a simple and easy to implement mechanism to incentivize the customers to allow further deferral

of the deadline of deferrable jobs in order to integrate more renewable energy. Section 4.5 describes the simulation set up and presents and discusses the simulation results. Section 4.6 concludes the chapter.

4.2 System Model

4.2.1 Model of the Tiered Architecture of a Data Center

It is a common practice to deploy different services in a data center in different tiers [35]. Most of the previous works for enhancing energy efficiency assumed a simple flat architecture of data center, which is not really the case in practical systems. Figure 4.1 shows the architecture of a typical data center with three tiers. A very common example of Internet service is web service. These days a very rich set of various web services are provided by the CSPs which may or may not be the owner of the data center. E-commerce such as online trading and online shopping are very important examples of web applications, which provide a lot of economic values in the business process of the company which implements and provides these [138]. Real system implementations of web applications are composed of different sub-systems, including web interfaces, logic and databases [139]. Typically, each tier in a data center is equipped with different functionalities. Most of the previous works did not consider the tiered architecture of the data centers, which would make the strategies proposed quite inappropriate to practically implement in large scale systems. The model proposed in this chapter can suitably address this issue, which makes the proposed strategies more valuable and close to practical large scale systems. The details of practical system implementation and the associated resource provisioning strategies may be found in [34]. After receiving job requests from customers, a load dispatcher

allocates the jobs to the appropriate tier depending on the requirement. Typically, in this service architecture (Figure 4.1) the web servers are responsible for displaying the results of the request on the client side, through appropriate interfaces such as web browsers. The second tier, namely the application tier consisting of application servers are typically responsible for performing the processing as per the business logic implemented in the system. The third tier shown in Figure 4.1 consists of data base servers and is responsible for data base processing and data accessing. Additionally, the model proposed here can address arbitrary number of tiers and does not sacrifice on generality, although a three-tier architecture is most common in practice [34].

This chapter studies the issue of resource provisioning for deferrable, deadline oriented jobs through energy efficient scheduling. Table 4.1 provides an overview of different types of workloads exhibited by a typical commercial data center. For the delay sensitive non-deferrable jobs, the resource provisioning is typically done by employing suitable queueing theoretic models [140]. Recently, with the tremendous developments in both hardware and software systems, computationally intensive jobs are becoming increasingly popular and very widely used. A very popular programming paradigm is MapReduce which can perform computationally intensive jobs across many servers. A pioneer work on MapReduce authored by industry practitioners from Google may be found in [141]. An open source implementation of MapReduce is Hadoop which has also become very popular among the practitioners [142], [143]. The inherent idea behind this paradigm is to parallelize a large task into a number of smaller tasks and execute them in a pool of large number of servers. The Map phase divides the task into different smaller task and generates a key-value pair corresponding to each smaller task. Then the Reduce phase executes all the smaller tasks, and more importantly combines the results of the smaller tasks to produce the end results. MapReduce is being increasingly adopted by

many companies and given the extremely valuable power of parallelizing, its use should become more widespread in the near future [144].

There has been a substantial amount of previous work to consider energy efficient optimum resource provisioning for non-deferrable jobs. Some notable references are [67], [32], [35]. For non-deferrable jobs, the SLA is commonly based on the upper bound on the expected response time, or the upper bound on some percentile of response time [32]. Compared to that, little work has been done on optimum resource provisioning for deferrable jobs.

Table 4.1

Classification of different types of workload in terms of deferability

Workload type	Nature
Business transactional application	Non-deferrable
Video on demand	Non-deferrable
Real time user requests	Non-deferrable
Scientific applications	Deferrable
Simulations	Deferrable
Large scale data analytics	Deferrable
High quality medical image processing	Deferrable

However, the work which is closest to the one presented in this chapter is [35]. The authors of [35] proposed an analytical framework for resource provisioning in a multi-tier server cluster for session based non-deferrable jobs, which is a pioneer work in this area. They applied the Mean Value Analysis (MVA) technique in queuing theory to implement an efficient iterative algorithm to address session based traffic which triggers multiple jobs at various tiers, by modeling the

system as a closed network of queues [140]. They modeled session based the internet workload, where a session typically comprise of a series of requests from the customer. Each of these requests may trigger multiple requests at different tiers and there can be think times in between which correspond to phenomena such as network delay etc. However, the objective of this chapter is to contribute more towards energy efficiency aspects. Furthermore a heuristic based strategy is designed for increasing renewable energy penetration by providing incentives for deadline deferral.

This chapter considers a discrete time model whose timeslot length matches with the timescale at which the data center updates its resource provisioning decision. Typically this ranges from a few minutes to a few hours depending on the internal management policies of the CSP [67]. To tackle the issue of resource provisioning for deferrable jobs, a technique prevalent in control engineering, called Model Predictive Control (MPC), is used. This is also known as Receding Horizon Control (RHC). A relevant reference on theoretical background of various algorithms in MPC and their implementations are found in [83], [145]. With MPC, an optimization problem is solved at each timeslot, to determine the control actions over a fixed horizon of time. The action for the first timeslot is applied to the system and the model parameters are updated. In the next timeslot the optimization problem is solved again with the modified model parameters and thereby the window moves forward. At every timeslot, the available information gets updated and thereby the problem instance changes at every iteration and the model error decreases as the window moves on. Some notable references on application of MPC based optimization techniques in resource scheduling are [80, 146].

In this chapter, the deferrable jobs are called as tasks, where each task requires a certain amount of jobs to be serviced at different tiers. Loosely, the deferrable tasks are referred to as

deferrable jobs to avoid notational complexity. However, each task is composed of a finite amount of service requirement which is equivalent to size of the task at different tiers. The size of the tasks is termed as amount of jobs, to be serviced to finish the task. A task is characterized by its arrival time, deadline and the amount of jobs at various tiers. The deferrable jobs can be conceptually considered as consumers of energy rather than consumers of power.

4.3 Problem Formulation

The model captures the cooling (and its associated costs) through a parameter called Power Usage Effectiveness (PUE). PUE is defined as the ratio of the total power consumed to the power consumed by the IT system [38]. In fact there has been considerable progress in improving the energy efficiency of data centers by improving the PUE. This is an important issue since the cooling and air distribution accounts for roughly 40~45% of the total electricity bill in a typical data center [59]. This chapter does not consider the dynamics of cooling systems in detail. Nevertheless, it may be noted that the strategies proposed here will act over all this improvements additively. In job scheduling, the macro level cooling dynamics of a data center is captured through PUE. A fairly holistic modeling framework is designed here to capture the important attributes including real time electricity price, renewable energy generation, macro level cooling dynamics etc.

It is assumed that the data center operates in the wholesale electricity market and is subjected to real time spot market pricing, which is quite common in practice [19]. Another prevalent charging plan charges a fixed price/MWh and an additional demand charge during peak hours [67], which may vary over a sliding window. The model is indeed general enough to

capture all these pricing schemes mentioned. Another assumption is that these prices are known a priori since the actual prices are very close to the forecasted ones. Similarly, for renewable energy, there are many efficient techniques to forecast the future availability a couple of hours ahead with prediction errors less than 5% [147]. Moreover, with recent developments and deeper understanding on cloud computing and in parallel with the progress in areas such as time series analysis, machine learning etc., even the future jobs can be predicted quite accurately a couple of hours ahead of time [148]. To provide some insights to the reader on optimal server provisioning, this chapter formulates the optimization problem in two different ways and investigates and analyzes the solutions. One is based on minimizing a linear combination of electricity cost and switching cost with adjustable knobs to control the relative priority between the two. A very elegant online algorithm called Lazy Capacity Provisioning (LCP) was proposed based on a similar model in [32]. However, this chapter exploits the time flexibility of the deferrable jobs, focusing more on renewable integration by energy efficient scheduling. The second procedure proposed in this chapter addresses for the first time, both the capital and operational costs under the same framework, [137]. In this MPC based optimization problem, the objective function minimizes a linear combination of electricity cost, the variance of the number of servers provisioned and the maximum number of servers provisioned at the different tiers. In the second formulation the switching cost is captured by the variance of the number of servers provisioned, over the optimization horizon (window size of the MPC problem). There is always a trade-off between catering to the varying renewables and not toggling servers too much. The aim is to strike a reasonable balance between them. Specifically, the objective is to incorporate a substantial portion of the available renewable energy, which at the same time does not incur too much switching cost, subject to SLA constraints. Through these formulations a number of crucial

issues within a data center are accurately captured. Importantly, all the strategies can apply to the tiered architecture of the data center since our model is able to discriminate between the jobs attributed to the different tiers. This is a significant advancement over previous works which were mostly based on a flat architecture model of the data center, where all the servers are assumed to be identical in terms of the way they handled different types of job requests [19, 67]. That assumption would make the model quite restrictive and substantially different from real life deployment.

All the vectors are denoted by an arrow ($\rightarrow$) above and are J dimensional column vectors unless mentioned otherwise. A J dimensional column vector whose elements are real and non-negative is denoted by $\vec{\square}_+^J$. A J dimensional row vector whose entries are all 1 is represented by $\vec{1}^T$. The price of electricity in real time market in timeslot t is denoted by $p(t)$.

4.3.1 Minimization of Operational Cost (*Min-OpC Procedure*)

It is assumed that the power consumed by an active server is proportional to the server load (which is normalized to its capacity i.e., utilization), while there is a fixed power consumption at idling condition. This assumption is quite standard and indeed very close to the actual power consumption pattern of a typical commercial server [32]. As per this affine (linear plus a constant) model, if $e(\rho)$ is the energy consumed by a server when the server load is ρ , then,

$$e(\rho) = \alpha_0 + \alpha_1 \rho; 0 \leq \rho < 1 \quad (4.1)$$

where α_0 is a constant which depends on the power consumed by a server in the idling condition and α_1 is a constant which depends on the increase in power consumption with increase in server

utilization. The authors in [32] proved that when the objective is to minimize a linear combination of energy cost, delay cost (which reflects the loss of revenue due to delay) and switching cost, then balancing the load uniformly across all the active servers becomes the optimal choice of server provisioning for non-deferrable jobs. This chapter augments the work of [32] by additionally considering aspects like renewables, brown electricity price and, most importantly, deferrable jobs.

As per the notation we adopt here, the non-deferrable jobs to be served at tier j in timeslot t , is denoted by $\lambda_j^{ND}(t)$, and the J dimensional vector whose elements are $\lambda_j^{ND}(t)$, is denoted by $\vec{\lambda}^{ND}(t)$. The number of servers provisioned for non-deferrable jobs at tier j in timeslot t is denoted by $m_j^{ND}(t)$, and the J dimensional vector whose elements are $m_j^{ND}(t)$, is denoted by $\vec{m}^{ND}(t)$. Power usage effectiveness of the data center in timeslot t is denoted by $PUE(t)$. Total renewable energy available in timeslot t is denoted by $w(t)$. The deferrable tasks are indexed by z . For each task z , the arrival time is denoted by a_z and the deadline, by which it must be completed, is denoted by d_z . The set of all unfinished (active) tasks at the beginning of timeslot t , which need some service at any tier is denoted by $Z(t)$. The service requirement (job size) for the task z at tier j is denoted as L_j^z , and the J dimensional vector whose entries are L_j^z , is denoted by $\vec{L}^z$. The amount of jobs (in size) to be scheduled at tier j in timeslot t for the active task z is denoted by $\lambda_j^{D,z}(t)$, and the J dimensional vector whose entries are $\lambda_j^{D,z}(t)$ is denoted by $\vec{\lambda}^{D,z}(t)$. The number of servers provisioned in timeslot t for deferrable task z , at tier j is denoted by $m_j^{D,z}(t)$, and the J dimensional vector whose entries are $m_j^{D,z}(t)$ is denoted by $\vec{m}^{D,z}(t)$. The total

amount of deferrable job (in size) to be scheduled at tier j in timeslot t is denoted by $\lambda_j^D(t)$, and the J dimensional vector whose entries are $\lambda_j^D(t)$ is denoted by $\vec{\lambda}^D(t)$. The total number of servers for deferrable tasks (jobs) at tier j in timeslot t for deferrable tasks is denoted by $m_j^D(t)$, and the J dimensional vector whose entries are $m_j^D(t)$ is denoted by $\vec{m}^D(t)$. The lower and upper bounds on the amount of jobs (in size) to be scheduled at tier j in timeslot t for the active task z , are denoted by $\lambda_j^{D,z,(-)}(t)$ and $\lambda_j^{D,z,(+)}(t)$, respectively. Similarly, the lower and upper bounds on the amount of jobs (in size) to be scheduled at tier j in timeslot t for the active task z , are denoted by $m_j^{D,z,(-)}(t)$ and $m_j^{D,z,(+)}(t)$, respectively. The lower and upper bounds on the total amount of jobs served at tier j in timeslot t are denoted by $\lambda_j^{(-)}(t)$ and $\lambda_j^{(+)}(t)$, respectively. Similarly, the lower and upper bounds on the total number of servers provisioned at tier j in timeslot t are denoted by $m_j^{(-)}(t)$ and $m_j^{(+)}(t)$, respectively. The total number of servers provided during timeslot t is denoted by $\vec{m}(t)$, whose entries are denoted by $m_j(t)$. Similarly, the total amount of jobs served at tier j in timeslot t is denoted by $\vec{\lambda}(t)$, whose entries are denoted by $\lambda_j(t)$. Let ψ be a function that determines the number of servers required to serve the amount of jobs given by $\vec{\lambda}^{ND}(t)$. Then

$$\psi\{\vec{\lambda}^{ND}(t)\} = \vec{m}^{ND}(t), \forall t \quad (4.2)$$

where $\psi\{.\}$ denotes a fairly general formulation. In fact it can even be an iterative algorithm as described in [35] or it may be calculated from queueing theoretic models as in [32]. It may be

noted that in this case $\psi\{.\}$ does not have to be convex since we can calculate the server requirement for non-deferrable jobs separately, outside the optimization problem.

The total server energy consumption for timeslot t at tier j (denoted by $Pow_j^{Serv}(t)$) is given by:

$$Pow_j^{Serv}(t) = m_j(t)e^{\frac{\lambda_j(t)}{m_j(t)}} = \alpha_0 m_j(t) + \alpha_1 \lambda_j(t), \quad \forall t, j \quad (4.3)$$

To incorporate the energy required by the cooling system $PUE(t)$ is considered. For the whole data center considering all the tiers, the total energy consumption for timeslot t (denoted by $Pow^{Tot}(t)$) is given by:

$$Pow^{Tot}(t) = (\alpha_0 \vec{1}^T \vec{m}(t) + \alpha_1 \vec{1}^T \vec{\lambda}(t)) PUE(t) \quad (4.4)$$

Current strategies for integrating renewable energy to power the data centers differ widely from one CSP to another [16]. In this work, we assume that the CSP owns a renewable generation infrastructure. Based on that the cost of electricity incurred by the CSP during timeslot t (denoted by $Cost^E(t)$) is given by:

$$Cost^E(t) = p(t)\{Pow^{Tot}(t) - w(t)\}^+, \text{ where } \{x\}^+ = \max(0, x). \quad (4.5)$$

This is justified by the fact that the cost related to renewable energy is largely dominated by the capital expenditure and its operational cost is very low [80], [90]. However this model can be broadened to consider a more general scenario where renewable is supplied by another generation company through a long term contract. Since these days the issue of environmental sustainability has become quite prominent all the responsible corporate citizens are demonstrating their commitments towards carbon footprint reduction. In some cases the CSP

owns the renewable energy generation infrastructure. In that case, the generalized incentive based model given in (4.6) can be used:

$$Cost^E(t) = \{p(t)Pow^{Tot}(t) - \beta w(t)\}^+, Pow^{Tot}(t) \geq w(t) \quad (4.6)$$

where β is the equivalent monetary incentive received by the CSP for integrating one unit of renewable energy, which can be appropriately chosen to consider the different cases as per the business practices adopted by different CSPs. It is to be noted that $Pow^{Tot}(t)$ consists of the sum of the brown energy supplied by the grid and the renewable energy supplied by the generator (utility provider) the CSP has contract with. This chapter focuses on the case where the CSP owns the renewable generation infrastructure in the rest of this work. The switching cost for one timeslot is:

$$Cost^S(t) = \vec{1}^T \{\vec{m}(t) - \vec{m}(t-1)\}^+ \quad (4.7)$$

Details about the switching cost due to server turning ON and OFF, can be found in [32]. In order to control the relative priority between electricity cost and switching cost, we introduce weighting factors k and γ . In *Min-OpC* procedure the following optimization problem is formulated and solved at each timeslot:

$$\text{Minimize} \quad k \sum_{t=1}^T Cost^E(t) + \gamma \sum_{t=1}^T Cost^S(t) \quad (4.8)$$

subject to:

$$\lambda_j^{D,z,(-)}(t) \leq \lambda_j^{D,z}(t) \leq \lambda_j^{D,z,(+)}(t); \forall t, j \text{ and } \forall z \in Z_t \quad (4.9)$$

$$\sum_{t=a_z}^{d_z} \lambda_j^{D,z}(t) = L_j^z, \forall z, j \quad (4.10)$$

$$\sum_{z \in Z_t} \lambda_j^{D,z}(t) = \lambda_j^D(t), \forall t, j \quad (4.11)$$

$$\sum_{z \in Z_t} \lambda_j^{D,z,(-)}(t) = \lambda_j^{D,(-)}(t), \forall t, j \quad (4.12)$$

$$\sum_{z \in Z_t} \lambda_j^{D,z,(+)}(t) = \lambda_j^{D,(+)}(t), \forall t, j \quad (4.13)$$

$$\lambda_j^{D,(-)}(t) \leq \lambda_j^D(t) \leq \lambda_j^{D,(+)}(t), \forall t, j \quad (4.14)$$

$$m_j^D(t) \geq 1.1\lambda_j^D(t), \forall t, j \quad (4.15)$$

$$\sum_{z \in Z_t} m_j^{D,z}(t) = m_j^D(t), \forall t, j \quad (4.16)$$

$$\sum_{z \in Z_t} m_j^{D,z,(-)}(t) = m_j^{D,(-)}(t), \forall t, j \quad (4.17)$$

$$\sum_{z \in Z_t} m_j^{D,z,(+)}(t) = m_j^{D,(+)}(t), \forall t, j \quad (4.18)$$

$$m_j^{D,z,(-)}(t) \leq m_j^{D,z}(t) \leq m_j^{D,z,(+)}(t), \forall t, j \text{ and } \forall z \in Z_t \quad (4.19)$$

$$m_j^{D,(-)}(t) \leq m_j^D(t) \leq m_j^{D,(+)}(t), \forall t, j \quad (4.20)$$

$$m_j^{ND}(t) + m_j^D(t) \leq M_j, \forall t, j \quad (4.21)$$

$$\lambda_j(t) = \lambda_j^{ND}(t) + \lambda_j^D(t), \forall t, j \quad (4.22)$$

$$m_j(t) = m_j^D(t) + m_j^{ND}(t), \forall t, j \quad (4.23)$$

The objective function of the optimization problem is presented in (4.8), which is a linear combination of electricity cost and switching cost. The decision variable are $\vec{m}(t)$ and $\vec{\lambda}(t)$. The constraints are listed from (4.9) to (4.23). The upper and lower bounds on the amount of jobs (in size) that can be executed for a particular task z at each tier and every timeslots are enforced through (4.9). The lower bound $\lambda_j^{D,z,(-)}(t)$ is typically set to zero to ensure operational feasibility, since a negative amount of jobs is not possible. However, if there are any additional service requirements for any task requiring a certain minimum execution in any timeslot then that can

also be conveniently addressed through this formulation. On the other hand, the upper bound $\lambda_j^{D,z,(+)}(t)$, the maximum amount of jobs belonging to task z , which can be executed, is derived from the maximum number of servers that can be provided, $m_j^{D,z,(+)}(t)$ for task z at tier j in timeslot t . Through this maximum number of servers some more subtle operational constraints can be addressed. Some tasks may have the requirement of execution within a specific subset of servers. This may be due to additional security issues (for sensitive jobs) or data locality issues (in case of jobs with special storage requirements) which limit the maximum number of servers for that particular task. The most important aspect of SLA, which is the completion of the deferrable task z by the given deadline, is enforced through (4.10). This ensures that all the active tasks are finished before their respective deadlines (i.e., for $\forall z \in Z(t)$, for $\forall t$) and that they receive the required amount of services at all the tiers. The total amount of deferrable jobs executed is equal to the sum of jobs executed across all the active tasks at each tier and is enforced through (4.11). The sum of the lower bounds on the amount of jobs that can be executed for each task belonging to the set of active tasks is computed through (4.12). This represents the minimum amount of total deferrable jobs (summed across all the active tasks) that has to be executed at timeslot t at various tiers in order to comply with the SLA. Similarly, (4.13) gives the maximum amount of total deferrable jobs which can be executed at tier j in timeslot t . Constraint (4.14) ensures that the amount of deferrable jobs scheduled (summed across all the active tasks) is feasible and within the limits. As for the relationship between the number of servers provisioned and the amount of jobs, it may be noted that our model can handle any function (f) of the form:

$$f : \Lambda \rightarrow \square_+^J; \Lambda = \{\vec{\lambda}(t) \mid \vec{\lambda}(t) \in \square_+^J\} \quad (4.24)$$

which maps the amount of jobs (in size) to the number of servers for each tier, as long as f is convex in $\vec{\lambda}(t)$. The convexity criterion is satisfied by most of the queueing theoretic models which have a closed form expression and does not compromise on generality [32]. In this chapter, the number of servers to be provided is computed by following a simple linear relationship between the number of servers to be provisioned and the amount of jobs to be executed as in (4.15). In order to deal with uncertainties from various sources such as breakdown, component failure and prediction uncertainty, an over-provisioning by 10% as given in (4.15), is done. This also ensures that the servers are not overloaded.

Constraints (4.16) - (4.18) ensure similar conditions on the number of servers provisioned as were done by constraints (4.11) - (4.13) for the jobs scheduled for timeslot t . Constraint (4.19) ensures that the number of servers provisioned to serve active task z at tier j , in timeslot t , is feasible and lies between the stipulated bounds. The bounds on total number of servers for deferrable jobs is ensured through (4.20) by summing across all the active tasks. Constraint (4.21) ensures that the sum of the number of servers provisioned for deferrable jobs and non-deferrable jobs (i.e., the total number of servers) at each tier does not exceed the total number of servers at that tier. It is to be noted that $m_j^{D,(+)}(t)$ can be less than $M_j - m_j^{ND}(t)$ since there may be specific server requirements for some tasks(s) due to additional security issues or data locality issues. The total amount of jobs served at tier j , in timeslot t , is the sum of the deferrable and non-deferrable jobs and is given by (4.22). Similarly, the total number of servers provisioned is the sum of servers provisioned for non-deferrable jobs and the servers provisioned for deferrable jobs, for all t , is ensured by (4.23).

In this optimization problem (from (4.8) to (4.23)) the decision variables are $\vec{m}^D(t)$ and $\vec{\lambda}^D(t)$. Once $\vec{m}^D(t)$ and $\vec{\lambda}^D(t)$ are computed as per the above procedure, the servers are allocated to the jobs following *Earliest Deadline First* (EDF) scheduling policy, i.e., the job with earliest deadline is served first. It can be shown that under EDF policy the deadline for all the jobs are met if the resource requirement does not exceed the processing capacity [80], which is indeed a strong service level guarantee associated with EDF.

The algorithm for solving the above problem and provisioning the servers at different tiers based on that will be referred to as *Min-OpC*, since it primarily aims to minimize the operational cost for energy and the number of switchings. It is worth mentioning that both the objective function and the feasible set are convex in decision variables and hence the optimization problem is convex. Therefore, efficient polynomial time reliable algorithms to solve it do exist [116]. At every timeslot, we solve a convex optimization problem. We implement the solution for the first timeslot and then update the problem instance for the next timeslot. It may be noted that, leveraging the convexity of the underlying problem, a large class of optimization problems can be solved dramatically fast using freely available software packages, such as CVX [149]. (Note that in case the original problem is not convex, it may still be possible to transform it into a convex optimization problem with acceptable relaxation [116].) In this case, since the above optimization problem is indeed convex, the solution can be obtained by expanding as per KKT (Karush–Kuhn–Tucker) conditions to obtain an optimal solution [116].

The work in [129] presents the developments in convex optimization over past few decades and gives this topic a fairly rigorous treatment. In particular there are some open source packages to generate the codes (in a couple of high level languages) for solving MPC based problem, if the

optimization problem solved at each timeslot is convex. The work in [145] describes advances that facilitate the development of custom MPC solvers, which can be incorporated in commonly used computing platform. By combining a high level specification language for optimization and code-generation, reliable and fast custom codes can be generated to solve MPC based problems. This may serve as a valuable tool to implement the real deployment is a large scale system where the problem size is enormously huge.

4.3.2 Joint Minimization of Capital and Operational Cost (*Min-CapOpC* Procedure)

This chapter now proposes a more effective method for making the server provisioning decision. In this case, the variance and peak of the number of servers provisioned at the individual tiers are incorporated into the objective function. Engineers working with Google reported that, as per a thumb rule, the approximate capital cost typically ranges between \$10~20/Watt [38], although this may change with time. This thumb rule provides a rough idea about the capital cost incurred to build a data center with a given peak power consumption. This is just to appreciate the motivation behind reducing peak power consumption. In fact, the peak power plays a crucial role in budgeting and serves as an important parameter during the project planning phase. The optimization problem proposed here differs from the previous one in the formulation of the switching cost. In this case, we capture the switching cost as the variance of the number of servers provided in the various tiers and scale it by a factor η .

As in the case of the *Min-OpC* procedure given earlier, the operational cost of electricity is:

$$Cost^E(t) = p(t)\{Pow^{Tot}(t) - w(t)\}^+ \quad (4.25)$$

In this procedure, the adverse effect of server toggling is captured through the variance of the number of active servers employed across the different time slots within the time window T . The total cost of switchings across the prediction horizon is given by:

$$Cost_{Total}^{Var} = \frac{1}{T} \sum_{j=1}^J \sum_{t=1}^T \{m_j(t) - \frac{1}{T} (\sum_{\tau=1}^T m_j(\tau))\}^2 \quad (4.26)$$

This is in fact the variance of the number of servers provisioned at each of the tiers, which involves the second moment and the mean. It turns out that capturing the switching cost through the variance imposes a tight control over switchings and strikes a good trade-off between electricity cost reduction, renewable incorporation and switching cost.

In order to exclusively focus on peak reduction (the total number of servers at different tiers), another extra term in the objective function is incorporated. This is captured in the form of the cost for peak provisioning at each tier and summed up across all the tiers. The peak provisioning cost is denoted by $Cost_{Total}^{Peak}$, which is expressed as:

$$Cost_{Total}^{Peak} = \sum_{j=1}^J \max_t \{m_j(t)\}, t \in \{1, 2, ..T\} \quad (4.27)$$

The optimization problem is formulated as:

$$Minimize: k \sum_{t=1}^T Cost^E(t) + \eta Cost_{Total}^{Var} + \delta Cost_{Total}^{Peak} \quad (4.28)$$

subject to the same set of constraints as *Min-OpC*, enumerated from (4.9) to (4.23). With the appropriate choices of k, η and δ , the relative priority among the power cost, variance (of the number of servers provisioned) and peak server provisioning capital cost is controlled. The purpose is to get more insights on how to design the optimal server provisioning strategy. The algorithm for solving the above problem and provisioning the servers at different tiers is referred

to as *Min-CapOpC*. Apart from reducing the operational cost, this aims to minimize peak server provisioning cost (capital cost) as well. The parameter choices for this are discussed in detail in Section 4.6. This optimization problem is also convex since the objective function and the feasible set are convex and is therefore also computationally tractable [116].

4.4 Heuristic Algorithm to Improve Cooling System Energy Efficiency

In *Min-OpC* and *Min-CapOpC* procedures the efficiency of the cooling system and the energy consumed by the cooling system is captured through *PUE*. Since cooling system efficiency is not the major focus of this chapter, the details of cooling systems are not specifically addressed. Nevertheless, the proposed framework can also include some heuristic methods to achieve more aggressive cooling system efficiency. As a concrete example, a threshold based policy can be implemented. As per the prediction, when the renewable is higher than an appropriately chosen threshold value, and also during the next timeslot the renewable is less than another threshold value, then the system can be cooled below a particular temperature. This will enable the cooling system to consume more energy to lower the temperature during the current timeslot when renewable generation is high, and in the following timeslot when the renewable energy generation is low, the cooling system would consume less energy. This is an area where there have been tremendous improvements over the past few decades. The various different strategies proposed in the literature can be integrated with the strategies proposed in this chapter to make the system more efficient.

4.5 Incentivization for Deadline Deferral (*IDD*)

In the smart grid arena, at macro level, the techniques to control the load according to the supply will fall into one of two categories namely, *direct load control* (DLC) and *indirect load control* (ILC) [94]. With DLC, the electricity service provider possesses the right to control the loads in a pre-specified manner. The work presented in [95] proposed a DLC based adaptive strategy to improve the operations of ancillary services. In contrast, when ILC is used, the utility provider does not have a direct control over the loads. In this case the utility provider controls the loads through various pricing signals. The various different pricing schemes include time-of-use pricing [96], coincident peak pricing [97], critical peak pricing [98], real time pricing [99], [100]. The general idea behind all these pricing strategies is to ensure that the utility providers can reflect and recover, as a function of time, the true marginal cost for producing one extra unit of electricity. There are some fundamental issues in implementing DR through dynamic pricing based approaches. Functionally, the time varying price acts like a signal to have a decentralized control on the individual consumers. However, there would be significant uncertainty in the load response for different customers which may make the effectiveness of these strategies questionable. In addition, the consumers are also exposed to the *volatility* of the electricity price. This section develops a simple, effective and *novel* mechanism to increase the renewable penetration into the grid through *incentivization on deadline deferral (IDD)*. The main idea here is to provide an economic motivation for the customers to defer the stipulated deadline of the submitted jobs till such time when more renewable is available. Typically the customers of these kinds of services are corporate clients largely guided by objectives of cost savings or monetary benefits. In this scenario, direct monetary incentive seems to be the most powerful way to attract

customers to participate in this program. Another advantage is that customers do not experience any risk in this proposed scheme as the quantity of electricity to be delivered by the new deadline (after deferral) is guaranteed. We believe that the simplicity of the proposed mechanism would facilitate its large scale adoption and the direct monetary reward provided would also motivate customers to participate in this program.

Table 4.2: Notations for incentivized deadline deferral model

Notation	Meaning
$f(\vec{L}^z)$	Number of servers required to serve the deferrable task z .
h	Number of timeslots by which the deadline of a task is shifted.
$\theta(\vec{L}^z, h)$	A function which computes the incentive when task z deferred by h timeslots.
H_{++}	A set of strictly positive integers by which the deadline deferral can be allowed.

The general theoretical framework of our proposed scheme is given next. Some additional notations are given in Table 4.2. The proposed method of incentivization is called as Incentivized Deadline Deferral (*IDD*). This explores if a particular task can be shifted profitably and the extent by which the deadline should be shifted to give maximum benefit for the CSP. The difference in cost after shifting the deadline of task z by h timeslots is:

$$\left. \begin{aligned} & p(t)\{(\alpha_0 \vec{1}^T f(\vec{L}^z) + \alpha_1 \vec{1}^T \vec{L}^z)PUE(t) - w(t)\}^+ \\ & - [p(t+h)\{(\alpha_0 \vec{1}^T f(\vec{L}^z) + \alpha_1 \vec{1}^T \vec{L}^z)PUE(t+h) - w(t+h)\}^+ + \theta(\vec{L}^z, h)] \end{aligned} \right\} \quad (4.29)$$

In this chapter the following is assumed:

$$\theta(\vec{L}^z, h) = \xi \vec{1}^T \vec{L}^z h \quad (4.30)$$

where ξ is a constant, independent of $\vec{L}^z$ and h . This implies that the incentive paid by the CSP to the customer is the product of a constant, the task size and the timeslot by which the deadline is deferred. The deadline shifting may be profitable if (4.29) is positive. The negative of the expression in (4.29) is minimized to determine which timeslot is best for shifting, considering all the above mentioned parameters. Though this is not a sufficient condition for profitable deferral, it does capture the situation where there is some rapid increase in renewable within a few adjacent timeslots. The problem is formulated to determine the profitable deferral of the deadline of the given task, as follows:

Minimize (by determining the optimal value of h)

$$\left. \begin{aligned} & p(t+h)\{(\alpha_0 \vec{1}^T f(\vec{L}^z) + \alpha_1 \vec{1}^T \vec{L}^z)PUE(t+h) - w(t+h)\}^+ \\ & + \theta(\vec{L}^z, h) - p(t)\{(\alpha_0 \vec{1}^T f(\vec{L}^z) + \alpha_1 \vec{1}^T \vec{L}^z)PUE(t) - w(t)\}^+ \end{aligned} \right\} \quad (4.31)$$

subject to:

$$\left. \begin{aligned} & \{p(t+h)\{(\alpha_0 \vec{1}^T f(\vec{L}^z) + \alpha_1 \vec{1}^T \vec{L}^z)PUE(t+h) - w(t+h)\}^+ + \theta(\vec{L}^z, h)\} \\ & - p(t)\{(\alpha_0 \vec{1}^T f(\vec{L}^z) + \alpha_1 \vec{1}^T \vec{L}^z)PUE(t) - w(t)\}^+ < 0 \\ & h \in H_{++} \end{aligned} \right\} \quad (4.32)$$

where $H_{++} = \{1, 2, 3, \dots | H_{++}|\}$.

For every task this optimization problem is solved to check if there is any chance of energy cost saving and increasing renewable incorporation, by deferring the deadline of that task. This is a reasonable heuristic to detect situations where renewable generation would be higher than in the previous timeslot(s). The suffix “++” denotes that tasks can only be post scheduled. The

solution process is further simplified and some analytical insights are built by considering three different cases.

Case-I:

$$(\alpha_0 \vec{1}^T f(\vec{L}^z) + \alpha_1 \vec{1}^T \vec{L}^z) PUE(d_z + h) > w(d_z + h) \text{ and } (\alpha_0 \vec{1}^T f(\vec{L}^z) + \alpha_1 \vec{1}^T \vec{L}^z) PUE(d_z) > w(d_z) \quad (4.33)$$

For some $h \in H_{++}$. With algebraic simplification, the optimization problem becomes:

Minimize

$$\left. \begin{aligned} & (\alpha_0 \vec{1}^T f(\vec{L}^z) + \alpha_1 \vec{1}^T \vec{L}^z) \{p(d_z + h) PUE(d_z + h) - p(d_z) PUE(d_z)\} \\ & + p(d_z) w(d_z) - p(d_z + h) w(d_z + h) + \theta(\vec{L}^z, h) \end{aligned} \right\} \quad (4.34)$$

subject to:

$$\left. \begin{aligned} & (\alpha_0 \vec{1}^T f(\vec{L}^z) + \alpha_1 \vec{1}^T \vec{L}^z) \{p(d_z + h) PUE(d_z + h) - p(d_z) PUE(d_z)\} \\ & + p(d_z) w(d_z) - p(d_z + h) w(d_z + h) + \theta(\vec{L}^z, h) < 0 \\ & h \in H_{++} \end{aligned} \right\} \quad (4.35)$$

Case-II:

$$(\alpha_0 \vec{1}^T f(\vec{L}^z) + \alpha_1 \vec{1}^T \vec{L}^z) PUE(d_z + h) < w(d_z + h) \text{ and } (\alpha_0 \vec{1}^T f(\vec{L}^z) + \alpha_1 \vec{1}^T \vec{L}^z) PUE(d_z) > w(d_z) \quad (4.36)$$

Hence, the optimization problem becomes:

$$\text{Minimize } \theta(\vec{L}^z, h) - p(d_z) \{(\alpha_0 \vec{1}^T f(\vec{L}^z) + \alpha_1 \vec{1}^T \vec{L}^z) PUE(d_z) - w(d_z)\} \quad (4.37)$$

subject to

$$\theta(\vec{L}^z, h) - p(d_z) \{(\alpha_0 \vec{1}^T f(\vec{L}^z) + \alpha_1 \vec{1}^T \vec{L}^z) PUE(d_z) - w(d_z)\} < 0, \quad h \in H_{++} \quad (4.38)$$

Case-III:

$$(\alpha_0 \vec{1}^T f(\vec{L}^z) + \alpha_1 \vec{1}^T \vec{L}^z) PUE(d_z + h) > w(d_z + h) \text{ and } (\alpha_0 \vec{1}^T f(\vec{L}^z) + \alpha_1 \vec{1}^T \vec{L}^z) PUE(d_z) < w(d_z) \quad (4.39)$$

It is obvious that, in Case-III, the deferral is not profitable since the renewable energy in current timeslot is more than the forecasted renewable in future timeslots. In the simulations, the function in (4.15) is applied for computing the number of servers from the amount of jobs for deferrable tasks and for the purpose of evaluating the effectiveness of the proposed mechanism. The above problem (4.33) to (4.39) is solved for all the deferrable tasks to determine if the job can be shifted profitably and then *Min-OpC* and *Min-CapOpC* are executed.

4.6 Simulation Setup and Results

4.6.1 Model Parameters

This section describes the simulation experiments performed to evaluate the proposed methods and algorithms. The duration of each time slot is 10 min, which is same as in [32]. With regard to the server power consumption, the following assumptions are made, $\alpha_0 = 0.6$ and $\alpha_1 = 0.4$, since it is reported that typically an idle server consumes 60% of its maximum power and that this grows almost linearly with the load [29]. In the simulation, a data center is considered to have three tiers with 10%, 40% and 50% servers in tiers 1, 2 and 3, respectively, with a total of 120,000 servers. The purpose of incorporating all these parameters k, γ, η and δ is

to make each of the terms in the objective function to be in comparable range so as to ensure that the objective function is sensitive with respect to each term. The different parameters are chosen to be $\gamma = 6$, $\eta = 1.2 \times 10^{-2}$ and $\delta = 100$. The authors of [32] mentioned that the cost of one OFF to ON transition for one server is almost equivalent to powering a server for one hour, which would be 6 timeslots in our model. The parameter η controls the weight of variance of the number of servers provisioned by *Min-CapOpC*, as in (4.28). To find out an appropriate η , different values of η , from 10^{-4} to 0.1 is explored. It is found that, a number of the order of 10^{-3} gives a good trade-off between electricity cost reduction, renewable usage increase and the number of switchings. For δ , which determines the weight of peak server reduction (4.28), 100 turns out to be a good choice for striking a reasonable trade-off between electricity cost and number of switchings. Simulations were performed with three values of k namely, 10, 100 and 1000 for both *Min-OpC* and *Min-CapOpC*. The results are provided in Figure 4.2. With $k=10$, there is a substantial increase in electricity cost with a moderate saving in switching cost. On the other hand with $k=1000$, there is substantial increase in switching cost with a moderate saving in electricity cost. It is observed that the sensitivity of the amount of renewable energy used with respect to k is relatively low. In both *Min-OpC* and *Min-CapOpC*, the renewable used is highest for $k=100$ and lowest for $k=1000$. However, in the case of *Min-CapOpC*, the renewable usages are almost the same for $k=10$ and 100 (Fig 4.2.b). The optimal operating point by choosing appropriate k depends on the individual CSPs. However, for the purpose of this work we choose $k=100$ in all our simulations, as it strikes a reasonable trade-off between electricity and switching cost. The data trace for electricity price is obtained from [150], which is the real time market in Singapore. To understand the effect of the optimization window size of the MPC algorithm, a simulation is performed for 100 timeslots with two different window sizes (i.e. 12, 24 timeslots)

and compared with the hypothetical case of perfect future knowledge. In Figure 4.3 (a), it is observed that, under *Min-OpC* the cumulative electricity cost, decreases, quite consistently, with increase in the window size.

Figure 4.2: The effect of scaling factor of electricity cost in terms of: (a) Cumulative electricity cost, (b) Cumulative green energy used and (c) Cumulative switchings

It may be noted that, in case of *Min-CapOpC* this trend is not very prominent. This is because of the fact that in *Min-CapOpC*, a higher priority is assigned for the switching cost, which is also a crucial parameter to implement this power saving strategy. The cumulative renewable energy used is increased quite prominently with the increase in window size for *Min-CapOpC*, whereas this remains almost unchanged in the case of *Min-OpC* (Figure 4.3 (b)). Finally, the cumulative number of switchings follows a consistent trend and both *Min-OpC* and *Min-CapOpC* performs better as the window size increases (Figure 4.3 (c)). The window size of 24 timeslots is chosen for the simulation and results are discussed in the rest of this chapter.

Figure 4.3: Sensitivity of the proposed optimization procedures with respect to the window size of the MPC. (a) Cumulative electricity cost, (b) Cumulative renewable energy used, (c) Cumulative no. of switchings

Figure 4.4 (a) shows the traces of the jobs at different tiers as well as the trace of the total jobs in the data center. Figure 4.4 (b) shows the traces of the deferrable task sizes at different tiers as well as the trace for the total size. Figure 4.4 (c) shows the distribution of deadlines of the deferrable tasks through a histogram. Taking the job trace of [106], the job statistics were generated for the simulations. For the non-deferrable jobs, it is assumed that the job size (service requirement) is exponentially distributed with mean of 1.

Figure 4.4: Workload trace: (a) Non-deferrable workload, (b) Deferrable workload, (c) Distribution of deadlines of deferrable loads

4.6.2 Comparison between *Min-OpC*, *Min-CapOpC* and Their Respective Offline Counterparts

This subsection presents the performance results in terms of cumulative electricity cost, cumulative green energy consumption and cumulative number of switchings under the proposed

algorithms, namely *Min-OpC* and *Min-CapOpC*. These are examined and discussed to guide designers in making optimum choices. It is observed that in Figure 4.5 that *Min-CapOpC* achieves 13.8% peak (maximum number of servers) reduction as compared to *Min-OpC*. As per our empirical calculation, this would translate into a capital cost reduction of US\$170 million to US\$340 million, assuming the capacity of the data center to be 120,000 servers. The proposed *Min-OpC* and *Min-CapOpC* are also analyzed by comparing with their respective offline optima. The offline optimum is the one which has perfect knowledge of the future and will therefore always be superior to its online MPC counterpart. The work in [151] aimed to understand the improvement of dynamic resource provisioning algorithms with the increase in forecast window size for the workload. They concluded that in the settings considered by them, the increase in the prediction horizon of the workload does not significantly improve the performance of the proposed algorithm, beyond a certain window size. Even though an oracle-type offline algorithm would be practically infeasible to implement, a comparison with this would enable us to perform a conservative evaluation of the proposed algorithms. It is observed that *Min-CapOpC* is quite conservative in toggling servers from ON to OFF and hence *Min-OpC* can provide a better match with the renewables by performing more frequent switching. Our strong conviction is that the proposed methods produce encouraging results considering both these aspects together. Figure 4.5 shows the number of servers provisioned under the two algorithms *Min-OpC* and *Min-CapOpC* and their offline counterparts. In terms of cumulative electricity cost there is approximately 20.6% saving for *Min-CapOpC* over *Min-OpC* (Figure 4.6). However these are achieved at the cost of less renewable integration (Figure 4.7). In Figure 4.6 it is observed that the cumulative electricity cost incurred by *Min-OpC* does not increase between timeslots 222 to 250. The reason is that the comparison is under the condition that the same amount of jobs is

executed for different algorithms, and *Min-OpC* schedules and executes the deferrable jobs earlier as compared to *Min-CapOpC*. It turns out that after 222 timeslots brown energy consumed by *Min-OpC* is practically zero (only green energy is used). In terms of green energy usage we see *Min-CapOpC* uses 11.9% less green energy than *Min-OpC* (Figure 4.7). In this case it is observed that unlike cumulative energy cost, green energy is continuously being used under both the algorithms. In terms of cumulative switching cost (measured in number of times a server turned from OFF to ON) the performances of these two algorithms are dramatically different. It is observed that *Min-CapOpC* performs only 30.1% (close to one third) cumulative number of switchings as compared to *Min-OpC* (Figure 4.8). The most interesting finding is that all the parameters of *Min-CapOpC* are less sensitive to window size than *Min-OpC*. This indicates that *Min-CapOpC* provides more robustness against forecast errors. It may be noted that the cumulative electricity cost incurred by offline *Min-OpC* is 34% less than that of its online variant. In contrast, for *Min-CapOpC*, the online algorithm incurs about 3% less electricity cost than its offline counterpart.

Figure 4.5: Comparison of no. of servers provisioned under online and offline cases under different schemes

This occurs because *Min-CapOpC* is more conservative in switching and its offline variant achieves 56% improvement in switching cost as compared to the online counterpart. For *Min-OpC*, increasing the window size dramatically improves the switching performance which is even better than what can be achieved with the offline version of *Min-CapOpC*. As shown in Figure 4.8, the offline counterpart of *Min-OpC* performs only 6.7% number of switchings as compared to its online version. This again confirms that *Min-OpC* algorithm is quite sensitive to the size of the prediction horizon. In summary, it is observed that in this setting, *Min-CapOpC* gives a reasonable balance between switching and electricity cost, renewable integration and is robust on model parameters. Our conviction is that this robustness is a crucial consideration since a robust system will be more dependable in actual implementation. We believe the findings would provide crucial insights both on the importance of carefully formulating the optimization problem and the information that can be obtained from it for efficiently operating the system over time.

Figure 4.6: Comparison of cumulative electricity cost under different schemes

Figure 4.7: Comparison of cumulative green energy consumption under different schemes

Figure 4.8: Comparison of cumulative no. of switchings under different schemes

4.6.3 Improvements after Applying *IDD*

This subsection presents the results obtained after the application of *IDD*. It is assumed that $\xi = 0.1$. This is a reasonable assumption which implies that if the service of one server is deferred by one timeslot from the original deadline then the amount paid to the customer is 10 cents. The respective deadlines of the tasks (by respective amounts) which satisfy the mentioned conditions are shifted and thereafter *Min-OpC* and *Min-CapOpC* algorithms are applied for scheduling of the deferrable tasks. This simple method is quite effective in incorporating more renewables without incurring additional switching cost. In simulations, the H_{++} is assumed to be $H_{++} = \{1, 2..6\}$ i.e. the deadline of a deferrable job can be deferred up to a maximum of 6 timeslots. Although a maximum shift of 6 timeslots is allowed, but it is noticed that in practice there is profitable deferral up to at most 3 timeslots. The reduction in electricity cost for *Min-OpC* and *Min-CapOpC*, after *IDD* is 4.5% and 3.1%, respectively, as shown in Figure 4.9 (a). Additionally, there is an increase in cumulative renewable integration under both *Min-OpC* and *Min-CapOpC*, after *IDD* is applied. In case of *Min-OpC*, renewable usage increased by 7.8% and in case of *Min-CapOpC*, it increased by 5% as shown in Figure 4.9 (b). The desirable feature of this is that this extra renewable integration and electricity cost reduction is not realized at the cost of additional switchings. In fact, it turns out that *IDD* decreases the cumulative number of switchings for *Min-OpC* and *Min-CapOpC* by 7.6% and 4.2% respectively, as shown in Figure 4.9 (c).

Figure 4.9: Performance of the proposed *IDD* (a) Cumulative electricity cost, (b) Cumulative renewable energy used, (c) Cumulative switchings

This reemphasizes the fact that the improvement through this mechanism is purely a result of this strategic *IDD*. For cumulative renewable integration and switching, no prominent sluggishness is observed after application of *IDD*. The summary of the results are presented in Table 4.3.

Investigation is carried out to measure the effect of error in electricity price forecast on the cumulative electricity cost. An error factor following Gaussian distribution with parameters, mean $\mu = 0.05$ and variance $\sigma = 0.1$, is generated to carry out the sensitivity analysis. To carry out a conservative evaluation the error factor is multiplied with the predicted price to generate the actual price and compute the cumulative electricity cost with this actual price (Figure 4.10). The results indicate that the proposed algorithms are quite robust in terms of prediction error in electricity price as long as the mean and the variance of the distribution of the error are small.

Table 4.3: Summary of results in terms of percentage

		Cost	Renewable	Switchings
<i>Min-OpC</i>	Offline over Online	34% (↓)	2.17% (↑)	93.3% (↓)
	After IDD	4.5% (↓)	7.8% (↑)	7.6% (↓)
<i>Min-CapOpC</i>	Offline over Online	-3% (↓)	12.7% (↑)	56.3% (↓)
	After IDD	3.1% (↓)	5% (↑)	4.2% (↓)

Figure 4.10: Effect of electricity price forecast error in the cumulative electricity cost

4.7 Summary

This chapter has considered the issue of determining the optimal strategy for server provisioning in data centers to reduce the capital and operational cost and to increase renewable energy usage. The main objective is to reduce peak server provisioning cost (capital cost) in addition to reducing the operational cost. Importantly, this chapter has also considered the switching cost, for toggling a server in the resource scheduling problem. Procedures are designed based on convex optimization problems to tackle these issues, for which there exist efficient polynomial time robust algorithms. Investigations are carried out to understand the performance of two different ways (formulations), namely, Minimization of Operational Cost (*Min-OpC*) and Minimization of Capital and Operational Cost (*Min-CapOpC*), for optimal server provisioning with respect to electricity cost, renewable energy used and the number of switchings. Under *Min-CapOpC* the switching cost is captured through the variance of the number of servers during different timeslots in the optimization horizon. Simulations are carried out with real data traces

of electricity price, renewable energy generation profile, workload, PUE etc. Results have shown that there is a 13.8% reduction in peak server usage, under *Min-CapOpC*, which would translate into significant saving in capital cost to construct the data center. Moreover, *Min-CapOpC* resulted in 20.6% saving in electricity cost over *Min-OpC* which is also significant. Importantly, under *Min-CapOpC* only 30% of the cumulative number of switchings are performed as compared to that under *Min-OpC*. Additionally, investigations are carried out to understand the sensitivity of the proposed algorithms with respect to the size of the prediction horizon. Furthermore, a simple strategy, namely, Incentivized Deadline Deferral (*IDD*) for achieving more renewable penetration is proposed. This strategy is enabled by providing monetary incentives for deferring the deadline until more favorable conditions are available in the system in terms of renewable availability and electricity price. The analytical framework for the proposed method is provided and the proposed algorithms are evaluated under real data traces. After implementation of *IDD* the electricity cost is reduced by 4.5% under *Min-OpC* and 3.1% under *Min-CapOpC*, accompanied by increase in renewable energy usage and decrease in cumulative number of switchings.

Chapter 5

Energy Efficiency Aware Load Distribution and Electricity Cost Volatility Control for Cloud Service Providers

5.1 Introduction

This chapter considers the case of a Cloud Service Provider (CSP) who owns multiple geographically distributed data centers, with collocated sources of renewable energy. Investigations are carried out to understand the performance of two slightly different load distribution strategies to minimize electricity cost and increase renewable incorporation considering the adverse effects of switching the servers, subject to compliance with Service Level Agreement (SLA).

Today's large to medium sized CSPs typically have numerous data centers located at different geographic locations across the globe. For these CSPs, the various services are replicated in different data centers to offer short response times, robust performance, enhanced security and operational redundancy [86]. They typically have front end servers to collect the job

requests from their customers. The issue of distributing the incoming jobs among the data centers is called *Geographical Load Balancing* (GLB) [67, 90]. In this chapter, the performances of different algorithms for GLB are compared in terms of electricity cost, renewable energy integration and number of server switchings. It is demonstrated that instead of modeling switching cost through a linear function, the proposed technique of modeling switching cost through variance achieves a better tradeoff between some important parameters.

Since the three major input parameters - electricity price, renewable energy and number of job requests - vary over time, the average cost of electricity per job request may also exhibit dramatic fluctuations. This chapter proposes to tackle this volatility by controlling the average cost of electricity per job request through leveraging contracts in the forward electricity market. The proposed strategy determines the optimal amount of electricity to be procured in the forward electricity market by performing a statistical analysis on the time series data. It is demonstrated that the proposed strategy substantially reduces the variance of the average cost of electricity per job and that this price risk mitigation is in fact achieved with a decrease in the cumulative electricity cost.

Contrary to traditional systems, nowadays a new start-up company does not have to incur a huge capital expenditure (Capex) to enter into the market, since most of the services are available from the existing CSPs in a pay-as-you-go manner. To support the ever increasing customer demands for IT services, today's CSPs incur a huge electricity cost to support their day to day operations. In fact it is reported that a significant proportion of the total operational expenditure (Opex) of a typical CSP goes towards electricity costs [152]. This chapter considers the case of a CSP which owns multiple data centers at various geographic locations. This chapter presents a holistic framework for modeling and distributing the jobs among geographically distributed data

centers considering the electricity price in the wholesale market, local renewable availability, cooling related conditions (at a macroscopic level), carbon emission effects and, most crucially, the adverse effects of server switching. Based on the model, investigations are carried out to understand various trade-offs so as to have more clarity and insight on the optimal choice for GLB. It may be noted that the previous chapters 3 and 4 did not consider the issue of mitigating of price risk that the CSPs are typically exposed to. This is indeed a crucial issue associated with economic sustainability. With the increase in renewable energy usage the volatility in electricity price would be further aggravated. This chapter addresses the issue of electricity price volatility by proposing a novel method to tackle this using the forward electricity market [113]. To emulate the complexity of the real system, extensive simulations are performed with real data traces on electricity price, renewable generation, cooling related conditions, workload, and carbon emission portfolio at ten locations across the globe. The contents of the present chapter builds on the initial work in [153].

It may be noted that, this chapter focuses on multiple data centers that are owned by a single big organization such as Google, Microsoft, etc, and considers non-deferrable workload. Our conviction is that Chapters 3, 4 and 5 together cover a fairly general scenario for the practitioners and would increase the scope of this thesis. The rest of this chapter is organized as follows. Section 5.2, elaborates the model and proposes two optimization problems with different structures of the objective function and solves them. Section 5.3 introduces a strategy to perform hedging on cost volatility of electricity by buying optimal quantity of electricity in the forward market. Section 5.4 describes the simulation setup and the data traces used in the study. Section 5.5 discuss the results obtained through simulations and provides some useful insights to

guide the CSPs on the deployment of an appropriate load distribution strategy. Section 5.6 concludes the chapter.

5.2 System Model of a CSP with Multiple Distributed Data Centers

This section first describes the operations of a CSP with multiple geographically distributed data centers. It then describes and investigates two different formulations of the optimization procedure in terms of electricity cost, server switching cost, service delay cost and carbon emission cost subjected to the SLA and operational constraints. These two formulations are presented in subsections 5.2.1 and 5.2.2, respectively. The objective is to provide more insights on how the formulation of the optimization strategy affects the performance of the corresponding load distribution procedure, in terms of various parameters.

This chapter assumes that these front end servers are collocated with the data centers. All the data centers are connected through a high bandwidth communication network (often optical) that offers high availability and reliability. Figure 5.1, presents a schematic diagram of the interconnection of multiple geographically distributed data centers connected with a reliable and high performance communication network, owned by a single CSP. It is assumed that each data center is equipped with a renewable energy source, which can supply green energy only to its corresponding data center without any transmission and distribution loss. On site renewable sources are quite common among some CSPs contributing towards their environmental sustainability [16, 104]. In this system, every data center collects customer requests (in the form of job requests) and services them in the same or different location based on their internal

strategies. As per current industry practices, either some appropriate load distribution algorithm is employed to minimize the response time [86] or it is simply done based on local geographical proximity.

Usually the task of load distribution and resource allocation is modeled through a discrete time model. The length of the timeslot at which resource provisioning decision is updated typically ranges from a few minutes to a few hours [67]. This chapter assumes that in a single timeslot, the number of servers provided at different locations remains constant. The objectives are to reduce electricity cost and increase green energy usages while at the same time reduce the number of server switchings. The strategy proposed in this chapter leverages spatial diversity in electricity price and the spatial variation of renewable energy generation, whereas the focus of Chapter 4 was to understand the operational intricacies to a reasonable extent. To tackle the issue of resource provisioning, a technique prevalent in control engineering, called Model Predictive Control (MPC), is applied to design the proposed algorithm.

Figure 5.1: Operation of a CSP with multiple geographically distributed data centers

The reason for choosing the MPC based framework is to appropriately address the issue of server switching and the adverse effects associated with it. In this model the objective function of the optimization problem has three major goals. Firstly, it minimizes the electricity cost incurred by the CSP by leveraging the spatial diversity of the electricity cost. Secondly, it maximizes the renewable energy usage at all locations. Thirdly, it minimizes the number of switchings of the servers. The first two can be merged by appropriately pricing the green energy. However, the third factor is not aligned with the first two. In fact, very often, there is a trade-off between the first two and the third. The problem is complicated further by the fact that we need to deal with many such locations. On one hand, the CSP can implement a greedy strategy for load

distribution first to consume the entire renewable energy available and then start loading data centers in the ascending order of electricity price. Since, electricity price and renewable energy generation both exhibit dramatic fluctuations, this greedy load distribution can result in arbitrarily many switchings. On the other hand, if the CSP prefers least number of switchings, some heuristic algorithm can be deployed to maintain a near constant number of active servers at every location, irrespective of the electricity price and renewable energy. Apart from these, the optimization problem also takes a linear combination of the adverse effects of carbon emission and response delay, which are incorporated in the objective function of the optimization problem.

The difference between the two optimization problems lies in the way the switching cost is captured. In the first formulation, we capture the switching cost by a linear function. In the second formulation we address the switching cost by considering the variance of the number of servers provisioned across all the timeslots (within T) and then sum up across all the locations. This captures the second moment of the number of servers provisioned at all locations. It turns out that capturing the switching cost through the second moment (variance) imposes a tighter control over switching and strikes a better trade-off between electricity cost reduction, renewable incorporation and switching cost. This first algorithm is called as *LB-L* (Load balancing with linearized effect of switching). The second algorithm is named as *LB-SM* (load balancing by considering the second moment effect of switching). The former is similar to the strategies proposed in [32, 33], but the latter is novel and is proposed for the first time in [153]. However, it may be noted that both the works [33] and [32] considered the case of a single data center, whereas we address the case of a CSP owning multiple data centers that are geographically distributed at different locations. Specifically, in [33], the authors investigated the server

provisioning by leveraging deferrable workload, which are to be served before the predefined deadline. In [32] the authors investigated the effects of dynamic provisioning of servers as per the demand and presented an online server provisioning algorithm which is provably optimal. The focus of this chapter is more on i) electricity expense incurred by the CSP and, ii) renewable energy integration. In this chapter, the works in [32, 33] are augmented by additionally considering the electricity expense, the renewable energy integration and exclusively accounting for the carbon emission cost as shown in (5.8) and (5.9). A geographic region is typically served by a few electricity providers. The generation portfolio and thereby the emission caused by one unit of energy consumption, varies across geographic locations and across time, e.g. the generation of energy in the state of Washington is much cleaner than that in the state of Texas. It is also reported, e.g. in [88] that this difference can be quite significant. We believe that the holistic model we consider here takes these factors into account and would provide a reasonable amount of confidence to the CSPs to practically implement the proposed algorithms.

5.2.1. Optimization with Linear Switching Cost (*LB-L Procedure*)

In this subsection, the first formulation of the optimal load distribution problem is discussed. Given below is the notation used in this chapter. Data center locations are indexed by i and the timeslots are indexed by t . The window size of the MPC algorithm is denoted by T . Price of electricity in the real time wholesale market at location i in timeslot t is denoted by $p(i, t)$. The amount of renewable energy available at location i in timeslot t is denoted by $w(i, t)$. The Power Usage Effectiveness (PUE) at location i in timeslot t is denoted by $PUE(i, t)$. PUE is the ratio of the total power consumed to the power consumed by the IT systems. The service rate of

a server at location i is denoted by $\mu(i)$. Total number of servers at location i is denoted by $M(i)$. The Carbon emission at location i when one server is active for one timeslot, in g/Kwh is denoted by $c(i)$. Average delay experienced by a job while transferring from location i to location j in timeslot t is denoted by $d(i, j, t)$. The amount of jobs transferred from location j to location i in timeslot t is denoted by $l(i, j, t)$. The amount of jobs originally arriving at location i in timeslot t , before the load distribution is denoted by $L(i, t)$. The constants to control the priorities for switching cost, the queueing delay at a data center (after arriving at the destination data center), the network delay (during transmission) and the carbon emission cost are denoted by β , γ , δ and ε respectively. The model addresses the heterogeneity of power characteristics of servers at different locations. The total power consumption at location i in timeslot t is denoted by $Pow(i, t)$. The total power consumed by the servers at location i in timeslot t is denoted by $Pow^{(Serv)}(i, t)$. The cost of electricity summed across all the locations at timeslot t is denoted by C_t . The switching cost summed across all the locations in timeslot t under $LB-L$ is denoted by S_t^{LB-L} . The response time after reaching the destination data center (service time + waiting time) summed across all the locations in timeslot t is denoted by $D_{t,Q}$. The network delay in timeslot t is denoted by $D_{t,NW}$. The carbon emission cost summed across all the locations in timeslot t is denoted by CE_t .

As in Chapter 4, it is assumed that the power consumed by an active server is an affine function in server load (normalized to its capacity, i.e. utilization). This is modeled as in [32] to be:

$$e_i(\rho) = \alpha_0(i) + \alpha_1(i)\rho ; \quad 0 \leq \rho < 1 \quad (5.1)$$

where $\alpha_0(i)$ is a constant which depends on the idling power consumption of a server and $\alpha_1(i)$ is a constant which depends on how power consumption increases with the load (utilization) of a server at location i . It was proved in [32] that when the objective is to minimize a linear combination of energy cost, delay cost (reflecting the loss of revenue due to delay) and switching cost, balancing the load across all the active servers uniformly becomes the optimal strategy. Applying this strategy in our model, we get:

$$Pow^{(Serv)}(i,t) = m(i,t)e_i\left(\frac{\lambda(i,t)}{m(i,t)}\right) = \alpha_0(i)m(i,t) + \alpha_1(i)\lambda(i,t) \quad (5.2)$$

Data centers spend a significant part of their energy expenditure in cooling and air distribution systems (to maintain the temperature within the allowed range defined for optimum performance of the servers). To appropriately address this, we choose to consider the most prominent parameter which summarizes the power consumed by the cooling system, called Power Usage Effectiveness (*PUE*) [38]. Due to variation in outside temperature, humidity and other weather factors, the *PUE* exhibits significant variation over time even in a single data center [106]. The total power consumed at location i , during timeslot t is then given by:

$$Pow(i,t) = \{\alpha_0(i)m(i,t) + \alpha_1(i)\lambda(i,t)\}PUE(i,t) \quad (5.3)$$

It is assumed that the CSP owns the renewable energy source and therefore there is no marginal cost for renewable energy usage. This is also justified by the fact that the cost related to renewable energy is largely dominated by the capital expenditure and its operational cost is very low [79, 90]. Hence, the total cost of electricity summed across all the locations in timeslot t , is obtained as:

$$C_t = \sum_{i=1}^N p(i,t) \{Pow(i,t) - w(i,t)\}^+ \quad (5.4)$$

where $\{a\}^+ = \max(a, 0)$.

There have been substantial amount of previous works showing that the cost for switching the server from OFF to ON state is dominant and the energy consumed for the ON to OFF transition is negligible [32, 154]. Additionally, the cost for transitioning ON to OFF can be incorporated in the next power up operation [32]. The total switching cost (across all the locations) for timeslot t is:

$$S_t^{LB-L} = \sum_{i=1}^N \{m(i,t) - m(i,t-1)\}^+ \quad (5.5)$$

The adverse effects associated with switching are: (i) energy consumed during transition (since the server does not perform any work during the switching time), (ii) delay experienced in replicating and migrating data through virtualization based techniques, (iii) unfavorable effects on server lifetime due to frequent switching etc. The potential adverse effects of server switching are discussed in detail in [32]. The most common way to capture the above mentioned adverse effects of switching is by representing these adverse effects (cost) as an increasing function of the number of servers transitioned from OFF to ON and vice-versa. This chapter focuses on the server switching issues and does not explicitly model the other aspects of a data center. Moreover, the other adverse effects of server switching (such as on the routers and network links) scale quite linearly with the number of servers switched. In this chapter, we consider only the cost incurred during OFF to ON transition.

Furthermore, to control the effect of switching cost with respect to the other attributes (such as electricity cost, service delay cost, carbon emission cost), the two parameters β and η are

used, in case of *LB-L* and *LB-SM*, as given by (5.9) and (5.18), respectively. If β and η are increased then the switching cost will have more priority over other costs in the overall objective function. As a result, there will be less number of server switchings, at the cost of sacrificing the other aspects. On the other hand, if β and η are decreased, then the switching cost will have less priority. In that case, the number of switchings would be high and at that cost, there will be improvements in terms of the other aspects.

The response time $D_{t,q}$ consists of waiting time and the execution time. For $m(i,t)$ queues, each queue observes a job arrival rate $\frac{\lambda(i,t)}{m(i,t)}$. Following [32, 67] and assuming a M/G/1, processor sharing (PS) model, the average response time is given by:

$$\frac{1}{\mu(i) - \lambda(i,t) / m(i,t)}.$$

$$D_{t,Q} = \sum_{i=1}^N \sum_{j=1}^N l(i,j,t) \frac{m(i,t)}{\mu(i)m(i,t) - \lambda(i,t)} = \sum_{i=1}^N \frac{m(i,t)\lambda(i,t)}{\mu(i)m(i,t) - \lambda(i,t)} \quad (5.6)$$

$$D_{t,NW} = \sum_{i=1}^N \sum_{j \in \{1,2..N\} - \{i\}} d(i,j,t) l(i,j,t) \quad (5.7)$$

$$CE_t = \sum_{i=1}^N c(i) \{Pow(i,t) - w(i,t)\}^+ \quad (5.8)$$

The objective function of the optimization problem is then,

Minimize

$$\sum_{t=1}^T (C_t + \beta S_t^{LB-L} + \gamma D_{t,Q} + \delta D_{t,NW} + \varepsilon CE_t) \quad (5.9)$$

In this case a linear combination of electricity cost, switching cost, response time cost (after reaching the destination data center), network delay cost (migrating the jobs from source to destination data center) and carbon emission cost is minimized subjected to a set of constraints to

ensure the compliance to SLA and the operational feasibility constraints. Through the constants β , γ , δ , and ε the respective weightage for switching cost, response time cost, network delay cost and carbon emission cost are controlled, respectively. By increasing the respective constants, more weightage can be assigned on the corresponding objective which would result in more aggressive reduction of the corresponding cost. The aim of this optimization procedure is to strike a reasonable trade-off between all the performance parameters and make the server provisioning decision accordingly.

It may be noted that the combination of C_t and CE_t can be expressed in a more compact form as follows:

$$C_t + \varepsilon CE_t = \sum_{i=1}^N [(p(i,t) + \varepsilon c(i)) \times \{Pow(i,t) - w(i,t)\}^+] \quad (5.10)$$

Subject to:

$$\mu(i)m(i,t) \geq 1.1\lambda(i,t), \forall i,t \quad (5.11)$$

$$\lambda(i,t) = \sum_{j=1}^N l(i,j,t), \forall i,t \quad (5.12)$$

$$\sum_{i=1}^N L(i,t) = \sum_{i=1}^N \lambda(i,t) = \sum_{i=1}^N \sum_{j=1}^N l(i,j,t), \forall t \quad (5.13)$$

$$\lambda(i,t) \geq 0, \forall i,t \quad (5.14)$$

$$l(i,j,t) = 0, \text{ for some } i,j,t \quad (5.15)$$

$$m(i,t) \leq M(i), \forall i,t \quad (5.16)$$

The constraints in (5.11) to (5.16) ensure the compliance to SLA and operation feasibility. The stability of queues at each of the locations is ensured through (5.11). In order to deal with situations like server break down, component failure and communication link failure, 10% server

overprovisioning is done since resource under provisioning and violation of SLA will have an extremely detrimental effect on the financial revenue of the CSP. Through (5.12) it is enforced that the total amount of workload serviced at each location i is the sum of the amount of workloads coming from all locations, including the workload coming from i itself. Through (5.13) it is enforced that the equilibrium condition that the total amount of incoming jobs (customer requests) summed across all the locations is the total amount of jobs serviced summed across all the locations for every timeslot t . This condition ensures that all the jobs coming at every location are serviced for every timeslot t . The amount of jobs serviced at each location i is non-negative and is enforced through (5.14). The issues like, the topology of the data center interconnection network, network congestion and link breakdown are appropriately addressed through (5.15). This formulation can dynamically react to such scenarios and restrict the data transfer between some locations. This can also address the scenario when certain jobs are to be executed at certain locations. The proposed algorithm can set corresponding $l(i, j, t) = 0$, for those jobs which are not to be executed at location i . On the other hand if there is congestion in some link, or link breakdown then the corresponding $l(i, j, t) = 0$, for the time during the congestion or breakdown period. Through (5.16) it is enforced that the number of servers provided is less than the number of servers physically available, for all i and t .

The decision variables in this optimization problem are $m(i, t)$ and $\lambda(i, t)$. The solutions to the above optimization problem from (5.2) to (5.16) are denoted by $m^{LB-L}(i, t)$ and $\lambda^{LB-L}(i, t)$. Here the output parameters are $m^{LB-L}(i, t)$ and $\lambda^{LB-L}(i, t)$, and the rest are the input parameters.

The objective function of the above optimization problem (5.2) to (5.16) is convex in decision variables and the solution space (defined by the intersection of the constraints) is also a convex

set; therefore the problem is a *convex optimization* problem. It is well established that there are efficient polynomial time algorithms to solve convex optimization problem very reliably and fast [116]. Under the MPC framework, in each timeslot, the proposed strategy solves a convex optimization problem and implements the solution for the current timeslot and update the problem instance for the next timeslot. There are many standard solvers available, which work over different computing platforms and have convenient user friendly interfaces. The literature on convex optimization problem is quite developed [116, 129]. In case the original problem is not convex it can still be transformed into a convex optimization problem by performing appropriate relaxation and translation [116]. In this case, since the optimization problem is convex, the solutions can be obtained by expanding as per KKT (Karush–Kuhn–Tucker) conditions since the KKT conditions are necessary and sufficient for the solution to be optimal [116].

5.2.2. Optimization Based on Second Moments of the Switching Cost (*LB-SM* Procedure)

This subsection discusses another formulation of the optimal load distribution strategy. The purpose is to investigate how the different formulations affect the performance of the optimization strategy in terms of cumulative electricity cost, cumulative renewable energy used and cumulative number of switchings. In this case, the adverse effect of server toggling is captured and quantified through the variance of the number of active servers employed across the different timeslots within the time window T . The total cost of switchings across the prediction horizon is given by:

$$S_{Tot}^{LB-SM} = \sum_{i=1}^N \frac{1}{T} \sum_{t=1}^T \left\{ m(i,t) - \frac{1}{T} \sum_{\tau=1}^T m(i,\tau) \right\}^2 = \frac{1}{T} \sum_{i=1}^N \sum_{t=1}^T (m(i,t) - E[m(i,t)])^2 \quad (5.17)$$

where $E[Z]$ denotes the average of the random variable Z .

This form of switching cost is novel in capturing the adverse effect of server switching. In this case, the average is computed over all the timeslots and then summed across all the locations. This reflects the variance of the number of servers provisioned at one location over different timeslots and is then summed across all the locations. It turns out that capturing the switching cost through the variance imposes a tighter control over switching and strikes a better trade-off between electricity cost reduction, renewable energy incorporation and server switching cost. In this procedure, the priority for switching cost is controlled through a constant η . A high value of η would assign higher weightage on switching cost and thereby tend to reduce the number of switchings. Conversely, a low value of η would perform more frequent switching and will in turn achieve better performance in terms of other objectives such as electricity cost and renewable energy integration. In this case the optimization problem becomes:

Minimize

$$\sum_{t=1}^T (C_t + \gamma D_{t,Q} + \delta D_{t,NW} + \varepsilon C E_t) + \eta S_{Tot}^{LB-SM} \quad (5.18)$$

subject to the constraints mentioned under *LB-L*, from (5.11) to (5.16), presented above. The decision variables in this optimization problem are $m(i,t)$ and $\lambda(i,t)$. The solutions to the above optimization problem from is denoted by $m^{LB-SM}(i,t)$ and $\lambda^{LB-SM}(i,t)$.

This optimization problem is also convex since the objective function and the feasible set, are both convex in the decision variables $m(i,t)$ and $\lambda(i,t)$. Hence, the solutions can be obtained by expanding as per KKT (Karush–Kuhn–Tucker) conditions since the KKT conditions are necessary and sufficient for a solution to be optimal.

5.3 Electricity Cost Variability Control (*ECVC*)

In this section, a novel strategy is proposed to mitigate the risk associated with the huge variability of the average cost to serve a single job request, as incurred by the CSP. This strategy is formulated by using a simple derivative product, namely forward contract. This strategy is applied on the top of both *LB-L* and *LB-SM* procedures and the results are presented in Section 5.5.

Following the deregulation of electricity market, various new market participants came up to replace the conventional vertically integrated structure [108, 109]. In this new paradigm, there are five categories of market participants within the end to end process of generation, transmission and distribution of electricity, namely i) generation companies, ii) load serving entities, iii) power marketers, iv) exchanges and market makers, and v) independent system operators [113]. This new paradigm enables electricity to be traded in the commodity market. Moreover with high renewable integration, the uncertainty in the system is aggravated since generation of renewable exhibits substantial amount of unpredictability, uncontrollability and intermittency which altogether results in substantial amount of randomness [25]. The major economic issue with this type of market is that the price varies over time. This is quite similar to real time price of commodities such as coal, gold etc. However, a major difference between

inherent nature of the electricity market with other commodity markets is that electricity is not storable, so the consumers do not have the option of operating with an appropriately chosen level of inventory. The various models for constructing different derivatives in the commodity market with storable commodities may be found in [110].

CSPs operating in multiple geographical locations are excellent candidates for leveraging the spatial diversity of renewable generation. However, the price volatility is a major concern of the CSP operating in this framework. Additionally, electricity trading is not mature enough and, most importantly, today's paradigm does not offer any mechanism for mitigating this price variability in a way which is particularly suitable for data centers. Recently, the authors of [114] proposed a mechanism to deal with this uncertainty. Under a stochastic optimization framework they proposed a strategy leveraging the electricity price diversity in the day ahead and in real time electricity markets. The work which is most relevant to ours is [115], which deals with the electricity price risk the CSP is exposed to. However, the model presented in that work did not consider switching costs, which would be of major concern to a CSP for implementing any power saving strategy. Moreover, it also did not consider the issue of integrating variable renewable energy sources.

In finance, the strategies to mitigate risk associated with price volatility are referred to as *hedging*. Hedging is an efficient tool to deal with the uncertainty associated with a wide array of parameters whose future realization is not known in advance. Hedging may be exercised through various kinds of tools such as forward contracts, options, futures and other derivatives [110]. A fairly rigorous mathematical background of design of these derivative products may be found in [111]. Here the hedging is exercised through a simple derivative, namely forward contract. Forward contracts are contracts which specify future delivery of a given quantity

(amount) of a particular commodity at a given fixed price (known at the beginning of the contract), before the expiry. In this context, the CSP executes a trading through forward contract with the electricity provider at all the locations, at each timeslot. The strategy for determining the optimal amount of electricity to be traded through the forward market is discussed in this section. It may be noted that the market for electricity trading is still in its nascent phase. However, as renewable penetration increases, the need for applying tools to mitigate the operational risk will become more prominent.

In the proposed model, there are three major sources of uncertainty, including electricity price, renewable energy and the number of job requests. Typical electricity prices in the United States and some other countries range between \$30 and \$60 per megawatt hour. There are, however, instances when the price increased up to \$7,000 and even \$10,000 per megawatt hour and in some cases persisted at \$1,000 per megawatt hour for even as long as a couple of days [117]. In this section the aim is to mitigate the volatility associated with the electricity cost, incurred by the CSP. More specifically, the goal is to minimize the variance of the average cost of electricity per job request by leveraging forward contract in the wholesale electricity market. The notations used in this section are explained and described in Table 5.1, presented below.

Table 5.1: Description of notations used in *ECVC*

Notation	Meaning
$g(i, t)$	Cost of serving all the job requests at location i during timeslot t .
$g(i, t)_\lambda$	Average cost of serving one job request at location i during timeslot t .
$g^*(i, t)$	Cost of serving all the job requests at location i during timeslot t , with <i>ECVC</i> .
$g^*(i, t)_\lambda$	Average cost of serving one job request at location i during timeslot t , with <i>ECVC</i> .
$G(t)$	Cost of serving all the job requests across all the locations during timeslot t .
$G_\lambda(t)$	Average cost of serving one job request, across all the locations during timeslot t .
$G^*(t)$	Cost of serving all the job requests across all the locations during timeslot t after applying <i>ECVC</i> .
$G_\lambda^*(t)$	Average cost of serving one job request during timeslot t , after <i>ECVC</i> .
$f_i(t, \tau)$	Price of electricity in the forward market during timeslot t , with contract expiry at τ .
$q(i)$	Quantity of electricity traded in forward market at location i .
$q^*(i)$	Optimum quantity of electricity traded in forward market at location i .
$Var[x]$	Variance of x .
$\Delta Var(G_\lambda)$	Percentage reduction in variance of $G_\lambda(t)$ after applying <i>ECVC</i> .

The aim is to determine the optimum contract in terms of the quantity of electricity to be traded in the forward market with known and fixed price at the beginning of contract. The cost of serving all the job requests at location i during timeslot t , with *ECVC* is obtained by:

$$g^*(i, t) = g(i, t) - q(i)(p(i, t) - f_i(t, \tau)) \quad (5.19)$$

The forward price $f_i(t, \tau)$ is known at t and the contract is up to time τ . In our model $g(i, t)$ is a function of $p(i, t), m(i, t), w(i, t)$ and $\lambda(i, t)$.

$$g^*(i, t)_\lambda = \frac{g(i, t)}{\lambda(i, t)} - \frac{q(i)(p(i, t) - f_i(t, \tau))}{\lambda(i, t)} = g(i, t)_\lambda - \frac{q(i)(p(i, t) - f_i(t, \tau))}{\lambda(i, t)} \quad (5.20)$$

The variance of the cost of serving all the job requests at location i during timeslot t , with *ECVC* is obtained by:

$$\left. \begin{aligned} Var[g^*(i, t)_\lambda] &= Var\left[g(i, t)_\lambda - \frac{q(i)(p(i, t) - f_i(t, \tau))}{\lambda(i, t)}\right] \\ &= Var[g(i, t)_\lambda] + (q(i))^2 Var\left[\frac{(p(i, t) - f_i(t, \tau))}{\lambda(i, t)}\right] - 2q_i Cov\left[g(i, t)_\lambda, \frac{(p(i, t) - f_i(t, \tau))}{\lambda(i, t)}\right] \end{aligned} \right\} \quad (5.21)$$

Differentiating this with respect to $q(i)$ and equating the first order derivative to zero, the value of the amount of electricity to be traded is obtained. This is denoted by $q^*(i)$ as the optimal value of $q(i)$ which minimizes the variance of cost of electricity.

$$q^*(i) = \frac{Cov\left[g(i, t)_\lambda, \frac{(p(i, t) - f_i(t, \tau))}{\lambda(i, t)}\right]}{Var\left[\frac{(p(i, t) - f_i(t, \tau))}{\lambda(i, t)}\right]} \quad (5.22)$$

Substituting this value of $q^*(i)$ into (21), we have:

$$Var[g^*(i, t)_\lambda] = Var[g(i, t)_\lambda] - \frac{(Cov\left[g(i, t)_\lambda, \frac{(p(i, t) - f_i(t, \tau))}{\lambda(i, t)}\right])^2}{Var\left[\frac{(p(i, t) - f_i(t, \tau))}{\lambda(i, t)}\right]} \quad (5.23)$$

The second order derivative of $Var[g^*(i,t)_\lambda]$ with respect to $q(i)$ is non-negative (zero in the trivial case when the variance is zero), confirming that point to be the minima. Now the denominator of the second term in (5.23) is always non-negative and when the covariance is non-zero, i.e. when the two random variables $g(i,t)_\lambda$ and $\frac{(p(i,t) - f_i(t, \tau))}{\lambda(i,t)}$ are not independent, then the second term takes a strictly positive value. This analytically justifies the decrease in variance of price and hence reducing the risk. In fact the proposed algorithm achieves more reduction in risk as the covariance between the two random variables increase, that is, when the two random variables are more correlated.

It may be noted that, the covariance between, $g(i,t)_\lambda$ and $\frac{(p(i,t) - f_i(t, \tau))}{\lambda(i,t)}$ can be negative if they are negatively correlated, i.e., if increase in one results in decrease in the other. The intuition behind this is that if the average electricity price per job increases in spite of decrease in difference between real time price and forward market price, then that implies that the forward market price is getting high and it is better to sell off some amount of electricity in the forward market instead of buying. On the other hand if $g(i,t)_\lambda$ decreases with increase in $\frac{(p(i,t) - f_i(t, \tau))}{\lambda(i,t)}$, then that indicates more renewable energy is available and it is worth selling energy in the forward market at the forward contract price. The quantity $q^*(i)$ is determined as per (5.22) for each location individually and the performance of the proposed strategy is analyzed considering the overall system by summing it over all the locations.

$$G_\lambda(t) = \frac{\sum_{i=1}^N g(i,t)}{\sum_{i=1}^N \lambda(i,t)} \text{ and } G_\lambda^*(t) = \frac{\sum_{i=1}^N g^*(i,t)}{\sum_{i=1}^N \lambda(i,t)} \quad (5.24)$$

Percentage reduction in variance of average cost of electricity per job request served is obtained by:

$$\Delta Var(G_{\lambda}) = \frac{100 \times (Var[G_{\lambda}(t)] - Var[G_{\lambda}^*(t)])}{Var[G_{\lambda}(t)]} \quad (5.25)$$

The previous work in [115] proposed an idea along similar lines. However, the cost function designed here $g(i, t)_{\lambda}$ is substantially different from that in their model since we have considered the issue of server switching. Additionally the strategy proposed here incorporates the issue of renewable energy integration and one of our objectives is to increase renewable energy penetration. To the best of our knowledge, the work in [153] is the first to propose a scheme for electricity cost variability control for renewable energy powered geographically distributed data centers, with reasonable number of server switchings. Since the amount of renewable energy at different locations can be modeled as a random variable, this brings in a major source of variability in the price per unit workload which makes hedging more meaningful.

Our strong conviction is that the proposed strategy will act as a bridge between the communities working on stochastic electricity pricing models to apply various tools [112, 113, 117, 118] prevalent in quantitative finance and the engineering community aiming at power cost reduction for the CSP. The effectiveness of this time series analysis depends on the stationarity of the distributions of the associated random variables. In this chapter, the random variables are mapped into a general wide sense stationary process. Even though this is a simplification of a process which could be more complex, the numerical simulation shows that the proposed method is nevertheless quite effective in reducing the variance of the average cost of electricity per job request. As a future extension of this work, the plan is to do more work on the stationarity analysis of these distributions [155, 156].

5.4 Performance Evaluation

This section describes the methodologies and the set up under which the performance of the proposed algorithms is evaluated. In order to examine the efficiency and effectiveness of the proposed algorithms in real systems, case studies are carried out under realistic parameter settings. Real workload traces are collected, which serve as the incoming workload in the simulation experiments. Then the proposed algorithms *LB-L* and *LB-SM* are applied for distributing the incoming workload among all the data centers. The amount of electricity to be traded through the forward contract, are computed as per the proposed *ECVC* strategy. Furthermore the performance of the proposed *LB-L* and *LB-SM* algorithms are compared with a general baseline strategy whereby the workload is uniformly distributed among the data centers.

5.4.1 Locations and Interconnection Topology

It is assumed that the CSP has one data center in each of the following ten locations including Paris, Germany, Switzerland, Netherlands, California, Indiana, Connecticut, New York, Pennsylvania and Singapore. These locations are selected because Google has data centers at these locations. Figure 5.2 represents the interconnection network between the data centers. Most of the previous works considered that all data centers are fully connected to each other and thereby the practical issues related to constraints of the underlying connection topology and congestion issues remain unaddressed. Although it is possible to route the traffic through alternative path(s), this may result in unacceptable network delay given the massive amount of

traffic. This chapter therefore considers the restriction that jobs can only be transferred between the locations through direct link. That is, jobs can only be transferred directly from source to the destination data center with no intermediate data center in-between. This would eliminate the situation where some jobs will experience high network delay. The duration of each timeslot in our simulation is 10 minutes, which is in line with current industry practices [67]. In the MPC problem, the optimization horizon size for both *LB-L* and *LB-SM* is 9 timeslots, which corresponds to one and half hours. This is a reasonable assumption since all the parameters in our model, such as electricity price, renewable energy and workload can be accurately predicted a couple of hours ahead of time [147, 148].

Figure 5.2: Topology of the distributed data center network used in the simulation

5.4.2 Electricity Price, Renewable Energy Generation and Job Request Distribution over Time

It may be noted that electricity price, renewable energy and the number of job requests exhibit a strong diurnal pattern i.e., they vary across the day, and between one day and another day the pattern turns out to be quite similar. So for each parameter, the data are shifted as per the local time. The simulation is started when local time at California is (00:00) and the respective local times at other 9 locations are considered with respect to that and the data are shifted in time accordingly. Electricity price data is collected from [157, 158]. For the sake of brevity the time series data of electricity price (in USD/Mwh) and renewable energy are reported for three locations (out of ten) namely, Paris, California and Singapore in Figure 5.3 and 5.4, respectively. Different wind and solar portfolio at different locations are considered to emulate the practical scenerio. As an instance, in Paris, 50% wind, 50% solar is assumed, whereas in California 20% solar and rest wind, is assumed. We believe that through this we can capture the realistic renewable generation process. For the other seven locations, similar method is deployed to generate the renewable energy trace. The service rate of servers at different locations is given in Table 5.2. In this chapter it is assumed that the service rate of all the servers is the same at a given location. The workload trace is taken from a commercial server cluster owned by Facebook and we shift it in time as per the local time in the locations considered in the simulations to emulate the real scenario faced by a CSP. The carbon emission rate at different locations is presented in Figure 5.5, which is used in the simulation. It is computed from the energy generation portfolio at the respective locations and considering the carbon emission of the corresponding source of generation, in the unit of g/kWh. The weightage of carbon emission on the load distribution strategy can be controlled by changing ϵ , in (5.9) and (5.18). A prior work

taking the carbon emission as a contributing factor for load distribution is presented in [88]. For the locations in US the electricity generation portfolio data is collected from [158] and for the other locations the data is collected from [159, 160].

Figure 5.3: Electricity price in three different locations

Figure 5.4: Renewable energy generation at three different locations (in percentage of the total capacity)

In *ECVC* the optimal amount of electricity to be procured is computed as per (5.22). This minimization of the variance in average cost (of electricity) to serve a job request is carried out over a moving window of 75 timeslots, which is equivalent to 750 min or 12.5 hours in the simulation. This is determined through trial and error method. The window lengths of 50, 75 and 100 timeslots are considered and it is observed that the window length of 75 timeslots results in the highest reduction in the variance of the average electricity cost per job request for the data set used.

5.4.3 Implementation of *LB-L* and *LB-SM*

The algorithms, *LB-L* and *LB-SM* are implemented in CVX [161] which is a package for solving disciplined convex optimization problems, built upon MATLAB platform. The SDPT3 solver [162] is leveraged which is the default solver in CVX. All simulations are performed on a HP Pro 3130 MT machine with an i5 processor (2.67 GHz) and 4.0 GB of RAM. At each timeslot *LB-L* takes approximately 1.5 seconds and *LB-SM* takes approximately 2 seconds to converge in MATLAB 7.11.0, under the above mentioned settings.

Table 5.2: The service rate of servers at different locations

Location	Service rate
Paris	1.25
Germany	1
Switzerland	1.25
Netherlands	1.5
California	1.25

Chicago	1
Connecticut	1
New York	1.25
Pennsylvania	1.25
Singapore	1.2

Figure 5.5: Carbon emission rate at different locations used in the simulation

5.5 Results

In this section the results obtained through the simulations are presented and discussed under the settings described in Section 5.4. The simulations are performed for two days, i.e. over 288 timeslots.

First the performance of the proposed algorithms is compared with a general baseline algorithm of near uniform load balancing (*ULB*). Under this baseline strategy, the incoming job requests are distributed across all the data centers as uniformly as possible. Although this may be

inefficient in terms of electricity cost, renewable energy consumption and server switchings, it would be easy to implement and would be intuitively expected to perform well in terms of network congestion and delay; it may be noted that *ULB* is indeed still implemented by some CSPs in practice [86]. This comparison will indeed quantitatively indicate the system performance improvement of the proposed algorithms as compared to the baseline algorithm *ULB* in terms of cumulative electricity cost, cumulative renewable energy used and cumulative number of switchings. The observations on the results of all these comparisons are given next.

Figure 5.6: Comparison of cumulative electricity cost

Figure 5.7: Comparison of cumulative renewable energy used

It is observed that both the algorithms *LB-L* and *LB-SM* proposed here perform substantially better than the baseline algorithm *ULB*. For the sake of improved visibility the y-axis is drawn in logarithmic scale for Figures 5.6, 5.7 and 5.8. In terms of cumulative electricity cost the percentage reduction of *LB-L* over *ULB* is 94.88% and *LB-SM* over *ULB* is 94.68 (Figure 5.6). In terms of renewable energy usage the *LB-L* consumes 3.45% less and *LB-SM* consumes 1.85% more renewable energy as compared to *ULB* (Figure 5.7). In terms of cumulative number of switching *LB-L* results into 85.99% and *LB-SM* results into 91.41% reduction as compared to *ULB* (Figure 5.8). However, the objective of this chapter is to understand how the two slightly different strategies for optimal load distribution perform in terms of cumulative electricity cost, cumulative renewable energy usage and cumulative number of switchings. Now the performance results obtained from *LB-L* and *LB-SM* are compared. For the sake of improved visibility results under *LB-L* and *LB-SM* are plotted in a linear scale. In terms of electricity cost, the performance of *LB-L* is slightly better than *LB-SM* (Figure 5.9). The

cumulative electricity cost incurred by *LB-SM* is 3.47% higher than that of *LB-L*. In terms of renewable energy integration, *LB-SM* outperforms *LB-L* by 5.2% (Figure 5.10).

Figure 5.8: Comparison of cumulative number of switchings

This difference is because of the fact that, under *LB-SM* the number of servers provided is slightly higher than that of *LB-L*. As a result of that, *LB-SM* achieves a better performance, in terms of percentage of jobs queued, which is discussed later in this section. In terms of electricity price and renewable incorporation, the performances of the two algorithms are quite close. However, in terms of the cumulative number of switchings, *LB-SM* performs much better than *LB-L*. This makes the former superior to the latter when all the three aspects are considered together, i.e. cumulative electricity cost, cumulative renewable energy used and the cumulative number of switchings. The server switching (OFF to ON and vice versa) incurs various types of adverse effects on energy consumption and overall life cycle of the server. These adverse effects are, i) energy used during the process of turning on, ii) delay for migrating jobs, iii) increased wear and tear, and finally, iv) the risk associated with the server switching. The work [32]

presents a detailed discussion on these issues. However, it is hard to establish the monetary cost of one switching operation, since the adverse effects of switching is not very measurable and quantifiable in this context. It is observed that *LB-SM* incurs 63.03% less switching as compared to *LB-L* (Figure 5.11).

Figure 5.9: Comparison of cumulative electricity cost under *LB-L* and *LB-SM*

It may be noted that the switching cost is fundamentally different from energy cost and renewable energy incorporation in that the future knowledge improves the performance of switching cost. Intuitively, if the algorithm knows that in the next timeslot the number of servers provisioned is going to be high then the algorithm may choose not to turn-off too many servers. This may be done to avoid high switching cost at the cost of very nominally higher electricity cost. In terms of cumulative carbon emission *LB-L* performs 5.7% better than *LB-SM*.

Figure 5.10: Comparison of cumulative renewable energy used under *LB-L* and *LB-SM*

Figure 5.11: Comparison of cumulative carbon emission under *LB-L* and *LB-SM*

Figure 5.12: Comparison of cumulative no. of switchings

To understand how the performance of *LB-L* and *LB-SM* improves with increasing window size, the switching cost incurred under the proposed algorithms is compared with respect to the *offline optima*. Offline counterpart of an algorithm is the one with future knowledge of all the timeslots, as used in the simulation. This is impossible to implement in practice, since it is infeasible to predict the model parameters beyond certain window with acceptable precision. However, this type of comparison will provide a conservative estimate on the performance of the proposed algorithms. This comparison also provides an insight on the extent to which performance improves by increasing the prediction window size.

It is observed that in the case of *LB-L*, the offline counterpart achieves 9.22% improvement in the cumulative number of switchings, whereas in case of *LB-SM* this is only 1.69% (Figure 5.12). This indicates that *LB-SM* provides more robustness compared to *LB-L*. This is indeed a crucial factor to be considered by the CSP since robustness plays an important role when there is a possibility of prediction error.

Now the result of the queueing theoretic analysis which is carried out is discussed. This analysis is carried out with the given number of servers provisioned and the average job requests. A unit of 5 seconds i.e., is considered in 10 minutes slot, where there are 120 such intervals. For each 5 seconds, the incoming workload is generated with an arbitrary distribution and the percentage of jobs that had to queue is tracked.

Figure 5.13: Percentage of jobs queued

It is to be noted that if there are locations with faster servers, eventually more jobs will be allocated there for higher efficiency of the overall system. However, those locations will then have a higher percentage of jobs which will have to queue, with roughly the same average response time. This is intuitive, since if we can reduce the service time (by employing faster service rate), then we can afford more queueing delay to achieve the same response time. This simulation is run for the entire duration, and the results are reported, in Figure 5.13. It is observed that the number of servers provided under *LB-SM* is slightly higher as compared to *LB-*

L , so the percentage of jobs queued is less with $LB-SM$. Nevertheless, this improved service quality is attained by incurring more electricity cost (Figure 5.9).

Now the effectiveness of the proposed method for Electricity Cost Variability Control ($ECVC$) is presented. The electricity price risk mitigation is in its germinating phase and the CSPs require more sophisticated and efficient strategies to handle it. We believe our work would facilitate the understanding of application tools and methods prevalent in quantitative finance which can be adapted for use in the domains of energy efficiency of geographically distributed data centers owned by a CSP.

The average costs of electricity per job request under $LB-L$ and $LB-SM$, before and after applying $ECVC$, are presented in Figure 5.14 and Figure 5.15, respectively, to visually depict how the variability of the average cost of electricity per job reduces after applying $ECVC$. As per the proposed $ECVC$ strategy, the optimum amount of electricity to be procured as per (5.22) minimizes the variance in average cost (of electricity) to serve a job request. This computation is carried out over a moving window of 75 timeslots, which is equivalent to 750 min or 12.5 hours in our simulation. It is to be noted that, from timeslot 0 to 75, we collect the data to compute the optimum amount of electricity to be procured. The first time the trading is executed is at the 76th timeslot and thereafter (as a result of $ECVC$) the average electricity cost to serve a job request changes. The effectiveness of the proposed $ECVC$ strategy is evaluated by quantifying the reduction in variance of the average cost of electricity per job request in aggregate (summed across all the locations), as mentioned in (5.25), over the entire simulation period.

Figure 5.14: Comparison of average cost per request before and after *ECV* for *LB-L*

Figure 5.15: Comparison of average cost per request before and after *ECV* in *LB-SM*

It is observed that the results of the simulation with real data sets are in line with the theoretical analysis presented in (5.19) through (5.25). In the case of *LB-L*, after applying *ECV*, the variance of average electricity cost per job is reduced by 15.12% and in the case of *LB-SM*

the improvement is 20.74%. This leads us to conclude that *ECVC* is quite effective in mitigating the risk pertaining to electricity price volatility, which is of paramount importance from the financial perspective of the CSP.

Figure 5.16: Comparison of cumulative electricity cost before and after *ECVC*

The cumulative electricity cost under both algorithms, before and after *ECVC*, is presented in Figure 5.16. It is encouraging to note that this risk mitigation is achieved with additional reduction in cumulative electricity cost. For *LB-L* the reduction in cumulative electricity cost is 1.59% and for *LB-SM* the reduction is 4.71% after applying *ECVC*. Given the large electricity costs incurred by a big CSP, this reduction in cumulative electricity cost will result in a considerable amount of cost savings.

5.6 Summary

This chapter has considered the energy cost and the associated cost volatility issues faced by a cloud service provider having multiple data centers located across the globe, with collocated renewable energy sources. The idea is based upon formulating a dynamic load distribution strategy to exploit the spatial variation in electricity price and renewable energy generation. Based on this idea, two MPC based algorithms are designed for geographical load balancing in order to reduce the cost of electricity and to increase renewable energy penetration. The performances of the proposed algorithms are evaluated and compared in terms of cumulative energy cost, cumulative renewable integration and cumulative number of switchings. It is concluded that the algorithm *LB-SM* performs better than *LB-L*. The proposed *LB-L* and *LB-SM* are also compared with a baseline strategy for load distribution and it is observed that both of them perform much better than the baseline algorithm. This chapter has also addressed the issue of volatility of cost of electricity, which is uncertain and varies with time. A method is proposed to minimize the variation in the average cost of electricity per job request by borrowing optimal contracts in the forward electricity market. The optimal quantity which minimizes the variance of average cost of electricity per job request, is analytically determined. Extensive simulation experiments have been conducted to evaluate the performance of the proposed cost variability control strategy. The simulation results have shown a 15.12% reduction in variance of average cost of electricity in the case of *LB-L* and 20.74% reduction in the case of *LB-SM*. We believe this is a meaningful contribution towards mitigating the electricity price risk that the CSP is exposed to. It is also encouraging to note that this risk mitigation is achieved along with a slight

reduction in cumulative electricity cost. Given the large electricity costs incurred by a big CSP, this reduction in cumulative electricity cost will result in a considerable amount of cost savings.

Chapter 6

Conclusions and Future Research

6.1 Conclusions

This chapter presents the conclusions of this thesis. The objectives of this thesis were to jointly address the issues of economic sustainability and environmental sustainability associated with cloud computing system. Economic sustainability is closely associated with the operational cost and the capital cost incurred by the *Cloud Service Provider* (CSP) to provide services to the customers. This thesis considered the holistic view of the entire ecosystem of a *Cloud Service Market* (CSM) in Chapter 3 and developed strategies to compute the optimal prices charged by the individual CSPs operating in the competitive CSM. In this market there are multiple CSPs, price sensitive rational customers and a regulatory body and each of them have their own objectives. The CSPs are competing among themselves to attract customer demands. There is a tradeoff between the long term economic sustainability (market share) and the short term economic sustainability (earning profit), from the perspective of a CSP. The proposed pricing strategy achieved an optimal tradeoff between these two aspects. The issue of operational cost control and reduction by reducing the electricity cost incurred by the CSPs was addressed

appropriately. At the same time, strategies were formulated to increase the renewable energy integration which is the major aspect in terms of environmental sustainability. In the study, both non-deferrable and deferrable workloads were considered. The two issues namely, economic sustainability and environmental sustainability were coupled together through a bi-level optimization strategy where two optimizations were carried out in two different timescales. In the slower timescale the prices charged to the customers were computed and in the faster timescale the energy efficient job scheduling was carried out. In summary, in the slower timescale the issue of economic sustainability was addressed and in the faster timescale the issue of environmental sustainability was addressed. The effectiveness of the proposed bi-level optimization strategy was demonstrated by carrying out simulations with real data traces of electricity price, renewable energy generation, workload, PUE etc. Evaluations showed that the proposed strategy was able to introduce a fair and sustainable charging system in the competitive CSM, with price sensitive rational customers. For operational cost and renewable energy usage, it was demonstrated that the proposed strategy achieved substantial reduction in electricity cost and significant increase in renewable energy integration. Importantly, it considered the adverse effects of frequent server switching and investigated the implications of incorporating the switching through linear and nonlinear cost functions. The proposed model is tunable through some parameters and thereby in this modeling paradigm the individual CSPs may compute the optimal operating points by choosing the set of parameters according to its own preferences.

To consider the specific operational points, Chapter 4 focused on specific details by considering a single data center where different applications were deployed in different tiers. In this tiered architecture, both deferrable and non-deferrable workloads were considered, incorporating practical application specific details. Under this framework an energy efficient job

scheduling strategy was proposed to reduce the electricity cost (incurred by the CSP) and increase the renewable energy integration in the grid. In the case of a single data center, investigations were carried out on two ways for capturing the server switching cost, and some useful insights on various performance parameters were provided. It may be noted that, the aspect of electricity cost (incurred by the CSP) reduction is closely associated with economic sustainability and the aspect of renewable energy integration is closely associated with environmental sustainability in the CSM. Thereafter, a strategy for providing monetary incentive to the customers was designed which allowed some reasonable deferral of jobs with respect to the original deadline of completion. Results of extensive simulations demonstrated that the proposed strategy is effective in achieving aggressive electricity cost reduction and renewable energy integration. The simulations were carried out with real data traces of electricity price, renewable energy generation, workload, PUE etc.

Chapter 5 considered the case of a CSP having multiple geographically distributed data centers across the globe. This is the case of a medium to large sized CSP in practice. Under this framework, investigations were carried out to gain some insights on the performance of *Geographical Load Balancing* (GLB) strategies, how the different formulations of the optimization strategy affect the cumulative electricity cost, cumulative renewable energy used and cumulative number of switchings. It considered more implementation level practical issues like the topological details of the interconnections between the data centers. Furthermore, the issue of risk associated with electricity cost volatility was addressed, which is again associated with economic sustainability of the CSP. Based on a simple derivative product, namely forward contract, a strategy for trading was designed in the wholesale electricity market to reduce the electricity cost volatility. The effectiveness of the proposed strategies was evaluated and

demonstrated through simulations with real data traces of electricity price, renewable energy generation, workload, PUE etc. This thesis comprehensively addressed the crucial aspects of economic sustainability and environmental sustainability for future cloud computing systems.

6.2 Future Work

Recently, many start-ups are offering innovative products and services to their clients (end customers) by taking the role of a Software-as-a-Service (SaaS) provider. Being small to medium size organizations, they typically implement their services and applications in the cloud by leasing the resources (such as servers, storage etc) from large organizations which owns facilities and infrastructures. These big organizations play the role of an Infrastructure-as-a-Service (IaaS) or Platform-as-a-Service (PaaS) provider in the market. In such a marketplace, SaaS providers charge the end customers for the services they provide, while the IaaS/PaaS providers charge the SaaS providers for the infrastructure and computing resources obtained from them. This means the CSM is no longer flat. Instead, the market is inherently tiered where there are end customers who have contracts with the SaaS providers, who in turn have contracts with the respective IaaS/PaaS providers. The present thesis considers a flat CSM with multiple CSPs and many end customers, but ignores any further hierarchical details.

A useful extension of the work reported in this thesis will be to develop efficient and economically sustainable pricing strategies for a tiered CSM as described above, which would ensure that the consumers are fairly charged for the services they received and there is no monopoly in the market. In the case of monopoly, the powerful CSP will have the market power to strategically manipulate the service prices in the market. In this context, economic

sustainability would mean that the market mechanism should on one hand prevent the players from losing their market share or losing profit, and on the other hand promote innovation with high QoS offered to consumers.

Another direction is to incorporate energy efficiency into the pricing strategies. Effective strategies need to be designed to incentivize renewable energy incorporation through an appropriate pricing mechanism. The significant novelty associated with this attribute is that the existing models are based on assumptions which are probably too simplistic (and often unrealistic) for practical application in a tiered CSM. Additionally, in such a market there has to be well chosen financial products to mitigate the risk of price volatility due to the complicated dynamics of the very large CSM. It may be possible that the most powerful way to realize the goal of green cloud computing would be through appropriate pricing and incentivization techniques. In summary, the future work plan includes three aspects given below:

- 1) To develop tractable models with reasonable generality to capture the price sensitive complex tiered markets with many end customers, several IaaS/PaaS, SaaS providers and regulatory bodies.
- 2) To formulate effective strategies for promoting energy efficiency, such as renewable energy integration.
- 3) To propose techniques to mitigate the risk associated with the price volatility for services offered in the cloud service market, in this three tier market architecture.

In fact the only existing work explicitly considering the three tier CSM with reasonable rigor is [163]. This paper presented some insightful results and discussed some potentially valuable research directions. However, they did not consider the energy efficiency issue which is a major focus of our future work. It may be particularly meaningful to establish a strong

connection between pricing and energy efficiency, which would then be an extremely powerful method for achieving various crucial objectives in the smart grid arena.

Author's Publications

Journal Papers

1. D Paul, W.D. Zhong, S.K. Bose, "Energy efficient cloud service pricing: A two-timescale optimization approach," Elsevier Journal of Network and Computer Applications, Accepted. (related to chapter 3)
2. D Paul, W.D. Zhong, S.K. Bose, "Demand response in data centers through energy efficient scheduling and simple incentivization," IEEE Systems Journal, vol. PP, DOI: 10.1109/JSYST.2015.2476357. (related to chapter 4)
3. D Paul, W.D. Zhong, S.K. Bose, "Energy efficiency aware load distribution and electricity cost volatility control for cloud service providers," Elsevier Journal of Network and Computer Applications, vol. 59, pp. 185-197, 2016, DOI: 10.1016/j.jnca.2015.08.012. (related to chapter 5)
4. D Paul, W.D. Zhong, S.K. Bose, "Energy Aware Pricing in a Three Tiered Cloud Service Market," MDPI, Electronics, Special issue on Data Centre Energy Efficiency (Accepted)

Conference Papers

1. D Paul, W.D. Zhong, "Price and renewable aware geographical load balancing technique for data centres," in *Proc. International Conference on Information, Communications and Signal Processing (ICICS)*, pp. 1-5, 2013, DOI: 10.1109/ICICS.2013.6782783.
2. D Paul, W.D. Zhong, S.K. Bose, "Energy efficient scheduling in data centers," in *Proc. International Conference on Communications (ICC)*, pp. 5948-5953, 2015, DOI: 10.1109/ICC.2015.7249270. (related to chapter 4)
3. D Paul, W.D. Zhong, S.K. Bose, "Energy efficient cloud service pricing," in *Proc. International Conference on Information and Communication Systems (ICICS)*, 2015. (related to chapter 3)

Bibliography

- [1] E. J. Umble, R. R. Haft, and M. M. Umble, "Enterprise resource planning: Implementation procedures and critical success factors," *European journal of operational research*, vol. 146, pp. 241-257, 2003.
- [2] M. D. Assunção, R. N. Calheiros, S. Bianchi, M. A. Netto, and R. Buyya, "Big Data computing and clouds: Trends and future directions," *Journal of Parallel and Distributed Computing*, vol. 79, pp. 3-15, 2015.
- [3] N. Marz and J. Warren, *Big Data: Principles and best practices of scalable realtime data systems*: Manning Publications Co., 2015.
- [4] K. Lee, M. Lam, R. Pedarsani, D. Papailiopoulos, and K. Ramchandran, "Speeding Up Distributed Machine Learning Using Codes," *arXiv preprint arXiv:1512.02673*, 2015.
- [5] X. Wu, X. Zhu, G.-Q. Wu, and W. Ding, "Data mining with big data," *IEEE Trans. Knowledge and Data Engineering*, vol. 26, pp. 97-107, 2014.
- [6] V. Pejovic and M. Musolesi, "Anticipatory mobile computing: A survey of the state of the art and research challenges," *ACM Computing Surveys (CSUR)*, vol. 47, pp. 1-29, 2015.
- [7] A. Zanella, N. Bui, A. Castellani, L. Vangelista, and M. Zorzi, "Internet of things for smart cities," *IEEE Internet of Things Journal*, vol. 1, pp. 22-32, 2014.
- [8] C. Rottondi, "Greenpeace report "Clicking Green: A Guide to Building the Green Internet"," 2015.
- [9] M. Armbrust, A. Fox, R. Griffith, A. D. Joseph, R. Katz, A. Konwinski, *et al.*, "A view of cloud computing," *Communications of the ACM*, vol. 53, pp. 50-58, 2010.
- [10] B. C. Tak, B. Urgaonkar, and A. Sivasubramaniam, "To move or not to move: The economics of cloud computing," in *Proc. USENIX Conference on Hot Topics in Cloud Computing*, pp. 1-6, 2011.
- [11] T. Navaro, "Global cloud market will be worth \$121bn by next year, report finds.," Available: <http://www.cloudcomputing-news.net/news/2014/mar/05/global-cloud-market-will-be-worth-121bn-next-year-report-finds/>, 2014.
- [12] "Gramatics. Inc," <http://www.graymatics.com/>, 2014.

- [13] F. Wen, "Market power in electricity supply," *IEEE Trans. Energy Conversion*, vol. 16, pp. 352-360, 2001.
- [14] "Annual report, 2014," *Google Inc.*, 2014.
- [15] J. Koomey, "Growth in data center electricity use 2005 to 2010," *A report by Analytical Press, completed at the request of The New York Times*, p. 9, 2011.
- [16] G. Cook, "Clicking clean:how companies are creating the green Internet," *Technical report, Greenpeace, Washington, DC*, 2014.
- [17] G. Ghatikar, "Demand response opportunities and enabling technologies for data centers: Findings from field studies," 2014.
- [18] J. Glanz, "Google details, and defends, its use of electricity," *The New York Times*; http://www.nytimes.com/2011/09/09/technology/google-details-and-defends-its-use-of-electricity.html?_r=3&, 2011.
- [19] A. Qureshi, R. Weber, H. Balakrishnan, J. Gutttag, and B. Maggs, "Cutting the electric bill for internet-scale systems," *ACM SIGCOMM Computer Communication Review*, vol. 39, pp. 123-134, 2009.
- [20] "Energy Flow Chart," *Lawrence Livermore National Laboratory*; <https://flowcharts.llnl.gov/>, 2015.
- [21] N. Stern, "The economics of climate change," *The American Economic Review*, pp. 1-37, 2008.
- [22] D. L. Hawkins, J. Blatchford, and Y. V. Makarov, "Wind integration issues and solutions in California," in *Proc. IEEE Power Engineering Society General Meeting*, 2007.
- [23] H. Farhangi, "The path of the smart grid," *IEEE Power and Energy Magazine*, vol. 8, pp. 18-28, 2010.
- [24] E. Y. Bitar, "Selling Random Energy," *PhD thesis, University of California, Berkeley*, 2011.
- [25] D. E. Bakken, A. Bose, K. M. Chandy, P. P. Khargonekar, A. Kuh, S. H. Low, *et al.*, "GRIP-Grids with intelligent periphery: Control architectures for Grid2050," in *Proc. IEEE International Conference on Smart Grid Communications (SmartGridComm)* pp. 7-12, 2011.
- [26] F. Rahimi and A. Ipakchi, "Demand response as a market resource under the smart grid paradigm," *IEEE Trans. Smart Grid*, vol. 1, pp. 82-88, 2010.
- [27] P. Siano, "Demand response and smart grids—A survey," *Renewable and Sustainable Energy Reviews*, vol. 30, pp. 461-478, 2014.
- [28] V. Shah, J. Booream-Phelps, and S. Min, "2014 Outlook: Let the Second Gold Rush Begin," *Deutsche Bank, Markets Research*, 2014.
- [29] L. A. Barroso and U. Hölzle, "The case for energy-proportional computing," *IEEE computer*, vol. 40, pp. 33-37, 2007.

- [30] G. Linden, "Make data useful," *Presentation, Amazon*, 2006.
- [31] A. Gandhi, M. Harchol-Balter, and I. Adan, "Server farms with setup costs," *Performance Evaluation*, vol. 67, pp. 1123-1138, 2010.
- [32] M. Lin, A. Wierman, L. L. Andrew, and E. Thereska, "Dynamic right-sizing for power-proportional data centers," *IEEE/ACM Trans. on Networking (TON)*, vol. 21, pp. 1378-1391, 2013.
- [33] M. A. Adnan, R. Sugihara, Y. Ma, and R. Gupta, "Dynamic deferral of workload for capacity provisioning in data centers," *arXiv preprint arXiv:1109.3839, Nov 2012*, 2012.
- [34] D. Huang, B. He, and C. Miao, "A survey of resource management in multi-tier web applications," *Communications Surveys & Tutorials, IEEE*, vol. 16, pp. 1574-1590, 2014.
- [35] B. Urgaonkar, G. Pacifici, P. Shenoy, M. Spreitzer, and A. Tantawi, "An analytical model for multi-tier internet services and its applications," in *Proc. ACM SIGMETRICS Performance Evaluation Review*, vol. 33, pp. 291-302, 2005.
- [36] L. Sun, H. Dong, F. K. Hussain, O. K. Hussain, and E. Chang, "Cloud service selection: State-of-the-art and future research directions," *Journal of Network and Computer Applications*, vol. 45, pp. 134-150, 2014.
- [37] J. Huang, G. Liu, and Q. Duan, "On modeling and optimization for composite network-Cloud service provisioning," *Journal of Network and Computer Applications*, vol. 45, pp. 35-43, 2014.
- [38] L. A. Barroso and U. Hölzle, "The datacenter as a computer: An introduction to the design of warehouse-scale machines," *Synthesis lectures on computer architecture*, vol. 4, pp. 1-108, 2009.
- [39] S. Anand and A. Sen, "Human development and economic sustainability," *World development*, vol. 28, pp. 2029-2049, 2000.
- [40] R. M. Solow, "Sustainability: an economist's perspective," *Economics of the environment*, vol. 1, pp. 131-141, 1991.
- [41] C. S. Yeo, S. Venugopal, X. Chu, and R. Buyya, "Autonomic metered pricing for a utility computing service," *Future Generation Computer Systems*, vol. 26, pp. 1368-1380, 2010.
- [42] Amazon, "Amazon EC2," <http://aws.amazon.com/ec2/spot/>, 2014.
- [43] H. Xu and B. Li, "Dynamic cloud pricing for revenue maximization," *IEEE Trans. Cloud Computing*, vol. 1, pp. 158 - 171, 2013.
- [44] C. Joe-Wong and S. Sen, "Mathematical frameworks for pricing in the cloud: Revenue, fairness, and resource allocations," *arXiv preprint arXiv:1212.0022*, 2012.
- [45] I. Menache, A. Ozdaglar, and N. Shimkin, "Socially optimal pricing of cloud computing resources," in *Proc. International ICST Conference on Performance Evaluation Methodologies and Tools*, pp. 322-331, 2011.

- [46] R. Pal and P. Hui, "Economic models for cloud service markets," *Distributed Computing and Networking*, pp. 382-396, 2012.
- [47] S. S. Manvi and G. Krishna Shyam, "Resource management for Infrastructure as a Service (IaaS) in cloud computing: A survey," *Journal of Network and Computer Applications*, vol. 41, pp. 424-440, 2014.
- [48] L. Zuo, L. Shu, C. Zhu, and Z. Zhou, "A resource evaluation model based on entropy optimization toward green cloud," in *Proc. International Conference on Semantics, Knowledge and Grids (SKG)*, 2013; p. 74-81, 2013.
- [49] K. RahimiZadeh, M. AnaLoui, P. Kabiri, and B. Javadi, "Performance modeling and analysis of virtualized multi-tier applications under dynamic workloads," *Journal of Network and Computer Applications*, vol. 56, pp. 166-187, 2015.
- [50] B. Jennings and R. Stadler, "Resource management in clouds: Survey and research challenges," *Journal of Network and Systems Management*, pp. 1-53, 2014.
- [51] R. T. Kaushik, P. Sarkar, and A. Gharaibeh, "Greening the compute cloud's pricing plans," in *Proc. Workshop on Power-Aware Computing and Systems*, pp. 1-6, 2013.
- [52] Y. Li, D. Chiu, C. Liu, L. T. Phan, T. Gill, S. Aggarwal, *et al.*, "Towards dynamic pricing-based collaborative optimizations for green data centers," in *Proc. IEEE 29th International Conference on Data Engineering Workshops (ICDEW)*, pp. 272-278, 2013.
- [53] D. Niyato, L. Xiao, P. Wang, D. I. Kim, and Z. Han, "Economics of Internet of Things (IoT): An Information Market Approach," *To appear in IEEE Wireless Communications; arXiv:1510.06837*, 2016.
- [54] S. Sen, C. Joe-Wong, S. Ha, and M. Chiang, "A survey of smart data pricing: Past proposals, current plans, and future trends," *ACM Computing Surveys (CSUR)*, vol. 46, pp. 1-37, 2013.
- [55] S. Kaxiras and M. Martonosi, "Computer architecture techniques for power-efficiency," *Synthesis Lectures on Computer Architecture*, vol. 3, pp. 1-207, 2008.
- [56] S. Herbert and D. Marculescu, "Analysis of dynamic voltage/frequency scaling in chip-multiprocessors," in *Proc. ACM/IEEE International Symposium on Low Power Electronics and Design (ISLPED)*, pp. 38-43, 2007.
- [57] M. Dayarathna, Y. Wen, and R. Fan, "Data center energy consumption modeling: A survey," *IEEE Communications Surveys & Tutorials*, vol. PP, 2015.
- [58] A. Beloglazov, R. Buyya, Y. C. Lee, and A. Zomaya, "A taxonomy and survey of energy-efficient data centers and cloud computing systems," *Advances in Computers 2011*, vol. 82, pp. 47-111, 2011.
- [59] M. Pedram, "Energy-efficient datacenters," *IEEE Trans. on Computer-Aided Design of Integrated Circuits and Systems*, vol. 31, pp. 1465-1484, 2012.
- [60] M. Pore, Z. Abbasi, S. K. Gupta, and G. Varsamopoulos, "Techniques to achieve energy proportionality in data centers: A survey," *Handbook on Data Centers*, pp. 109-162, 2015.

- [61] R. Bolla, R. Bruschi, F. Davoli, and F. Cucchietti, "Energy efficiency in the future internet: a survey of existing approaches and trends in energy-aware fixed network infrastructures," *Communications Surveys & Tutorials, IEEE*, vol. 13, pp. 223-244, 2011.
- [62] G. Chen, W. He, J. Liu, S. Nath, L. Rigas, L. Xiao, *et al.*, "Energy-aware server provisioning and load dispatching for connection intensive internet services," in *Proc. USENIX Symposium on Networked Systems Design and Implementation*, vol. 8, pp. 337-350, 2008.
- [63] A. Gandhi, V. Gupta, M. Harchol-Balter, and M. A. Kozuch, "Optimality analysis of energy-performance trade-off for server farm management," *Performance Evaluation*, vol. 67, pp. 1155-1171, 2010.
- [64] S. Albers, "Energy-efficient algorithms," *Communications of the ACM*, vol. 53, pp. 86-96, 2010.
- [65] G. Wang and T. E. Ng, "The impact of virtualization on network performance of amazon ec2 data center," in *Proc. INFOCOM*, pp. 1-9, 2010.
- [66] H. Goudarzi and M. Pedram, "Achieving Energy Efficiency in Datacenters by Virtual Machine Sizing, Replication, and Placement," *Advances in Computers*, 2015.
- [67] Z. Liu, M. Lin, A. Wierman, S. H. Low, and L. L. Andrew, "Greening geographical load balancing," in *Proc. of the ACM SIGMETRICS joint international conference on Measurement and modeling of computer systems*, pp. 233-244, 2011.
- [68] J. Luo, L. Rao, and X. Liu, "eco-idc: Trade delay for energy cost with service delay guarantee for internet data centers," in *Proc. IEEE International Conference on Cluster Computing (CLUSTER)*, pp. 45-53, 2012.
- [69] Y. Yao, L. Huang, A. Sharma, L. Golubchik, and M. Neely, "Data centers power reduction: A two time scale approach for delay tolerant workloads," in *Proc. IEEE INFOCOM*, , pp. 1431-1439, 2012.
- [70] L. Parolini, B. Sinopoli, B. H. Krogh, and Z. Wang, "A cyber-physical systems approach to data center modeling and control for energy efficiency," *Proceedings of the IEEE*, vol. 100, pp. 254-268, 2012.
- [71] C. Bash and G. Forman, "Cool Job Allocation: Measuring the Power Savings of Placing Jobs at Cooling-Efficient Locations in the Data Center," in *Proc. USENIX Annual Technical Conference*, vol. 138, p. 140, 2007.
- [72] J. D. Moore, J. S. Chase, P. Ranganathan, and R. K. Sharma, "Making Scheduling" Cool": Temperature-Aware Workload Placement in Data Centers," *General Track, USENIX Annual Technical Conference*, , pp. 61-75, 2005.
- [73] L. A. Johnson and D. C. Montgomery, *Operations research in production planning, scheduling, and inventory control* vol. 6: Wiley New York, 1974.
- [74] M. L. Pinedo, *Scheduling: theory, algorithms, and systems*: Springer Science & Business Media, 2012.
- [75] A. Wierman, "Scheduling for today's computer systems: Bridging theory and practice," *PhD. Thesis, Carnegie Mellon University, Pittsburgh*, , 2007.

- [76] B. A. Shirazi, K. M. Kavi, and A. R. Hurson, *Scheduling and load balancing in parallel and distributed systems*: IEEE Computer Society Press, 1995.
- [77] H. El-Rewini and T. G. Lewis, "Scheduling parallel program tasks onto arbitrary target machines," *Journal of parallel and Distributed Computing*, vol. 9, pp. 138-153, 1990.
- [78] M. Harchol-Balter, *Performance Modeling and Design of Computer Systems: Queueing Theory in Action*: Cambridge University Press, 2013.
- [79] A. Subramanian, M. Garcia, A. Dominguez-Garcia, D. Callaway, K. Poolla, and P. Varaiya, "Real-time scheduling of deferrable electric loads," in *Proc. of American Control Conference (ACC)*, pp. 3643-3650, 2012.
- [80] A. Subramanian, M. J. Garcia, D. S. Callaway, K. Poolla, and P. Varaiya, "Real-time scheduling of distributed resources," *IEEE Transactions on Smart Grid* 2013; , vol. vol. 4, pp. p. 2122-2130, 2013.
- [81] C. L. Liu and J. W. Layland, "Scheduling algorithms for multiprogramming in a hard-real-time environment," *Journal of the ACM (JACM)*, vol. 20, pp. 46-61, 1973.
- [82] S. Baruah and J. Goossens, "Scheduling real-time tasks: Algorithms and complexity," *Handbook of scheduling: Algorithms, models, and performance analysis*, vol. 3, 2004.
- [83] E. F. Camacho and C. B. Alba, *Model predictive control*: Springer Science & Business Media, 2013.
- [84] G. Hug-Glanzmann, "Predictive control for balancing wind generation variability using run-of-river power plants," in *Proc. IEEE Power and Energy Society General Meeting*., pp. 1-8, 2011.
- [85] W. Xiao, S. Zhang, J. Lin, and C. K. Tham, "Energy-efficient adaptive sensor scheduling for target tracking in wireless sensor networks," *Journal of Control Theory and Applications*, vol. 8, pp. 86-92, 2010.
- [86] P. Wendell, J. W. Jiang, M. J. Freedman, and J. Rexford, "Donar: decentralized server selection for cloud services," *ACM SIGCOMM Computer Communication Review* 2010;, vol. 40, pp. 231-242, 2010.
- [87] C. Stewart and K. Shen, "Some joules are more precious than others: Managing renewable energy in the datacenter," in *Proc. Workshop on Power Aware Computing and Systems*, pp. 1-5, 2009.
- [88] P. X. Gao, A. R. Curtis, B. Wong, and S. Keshav, "It's not easy being green," *ACM SIGCOMM Computer Communication Review*, vol. 42, pp. 211-222, 2012.
- [89] K. Le, R. Bianchini, T. D. Nguyen, O. Bilgir, and M. Martonosi, "Capping the brown energy consumption of internet services at low cost," in *Proc. International Conference of Green Computing*, pp. 3-14, 2010.
- [90] Z. Liu, M. Lin, A. Wierman, S. H. Low, and L. L. Andrew, "Geographical load balancing with renewables," *ACM SIGMETRICS Performance Evaluation Review*, vol. 39, pp. 62-66, 2011.

- [91] L. Rao, X. Liu, L. Xie, and W. Liu, "Coordinated energy cost management of distributed internet data centers in smart grid," *IEEE Trans. Smart Grid*, vol. 3, pp. 50-58, 2012.
- [92] D. P. Bertsekas and J. N. Tsitsiklis, *Parallel and distributed computation: numerical methods*: Prentice-Hall, Inc., 1989.
- [93] A. Rahman, X. Liu, and F. Kong, "A survey on geographic load balancing based data center power management in the smart grid environment," *Communications Surveys & Tutorials, IEEE*, vol. 16, pp. 214-233, 2014.
- [94] M. Alizadeh, X. Li, Z. Wang, A. Scaglione, and R. Melton, "Demand-side management in the smart grid: Information processing for the power switch," *Signal Processing Magazine, IEEE*, vol. 29, pp. 55-67, 2012.
- [95] K.-Y. Huang and Y.-C. Huang, "Integrating direct load control with interruptible load management to provide instantaneous reserves for ancillary services," *IEEE Trans. Power Systems*, vol. 19, pp. 1626-1634, 2004.
- [96] Y. Jia-hai, "Customer response under time-of-use electricity pricing policy based on multi-agent system simulation," in *Proc. Power Systems Conference and Exposition, (PES);* , pp. 814-818, 2006.
- [97] "Utility: Electricity tariffs," *The Official Website of the City of Fort Collins, Colorado USA*; <http://www.fcgov.com/utilities/business/rates/electric/coincident-peak>, 2013.
- [98] M. A. Piette, D. Watson, N. Motegi, S. Kiliccote, and P. Xu, "Automated Critical Peak Pricing field tests: Program description and results," *Lawrence Berkeley National Laboratory*, 2006.
- [99] R. Sioshansi, "Evaluating the impacts of real-time pricing on the cost and value of wind generation," *IEEE Trans. Power Systems*, vol. 25, pp. 741-748, 2010.
- [100] A.-H. Mohsenian-Rad and A. Leon-Garcia, "Optimal residential load control with price prediction in real-time electricity pricing environments," *IEEE Trans. Smart Grid*, vol. 1, pp. 120-133, 2010.
- [101] M. Irlbeck, D. Bytschkow, G. Hackenberg, and V. Koutsoumpas, "Towards a bottom-up development of reference architectures for smart energy systems," in *Proc. International Workshop on Software Engineering Challenges for the Smart Grid (SE4SG)*, 2013.
- [102] G. Fridgen, L. Häfner, C. König, and T. Sachs, "The Value of IS-Enabled Flexibility in Electricity Demand—a Real Options Approach," in *Proc. International Conference on Wirtschaftsinformatik*, 2015.
- [103] A. Krioukov, C. Goebel, S. Alspaugh, Y. Chen, D. E. Culler, and R. H. Katz, "Integrating renewable energy using data analytics systems: challenges and opportunities," *IEEE Data Eng. Bull.*, vol. 34, pp. 3-11, 2011.
- [104] W. Deng, F. Liu, H. Jin, B. Li, and D. Li, "Harnessing renewable energy in cloud datacenters: opportunities and challenges," *IEEE Network*, vol. 28, pp. 48-55, 2014.
- [105] X. Dong, T. El-Gorashi, and J. M. Elmirghani, "Green IP over WDM networks with data centers," *Journal of Lightwave Technology*, vol. 29, pp. 1861-1880, 2011.

- [106] Z. Liu, I. Liu, S. Low, and A. Wierman, "Pricing data center demand response," in *Proc. of ACM international conference on Measurement and modeling of computer systems*, pp. 111-123, 2014.
- [107] A. Wierman, Z. Liu, and H. Mohsenian-Rad, "Opportunities and challenges for data center demand response," in *Proc. IEEE International Green Computing Conference (IGCC)*, pp. 1 - 10, 2014.
- [108] S. Stoft, *Power system economics: Designing markets for electricity*, 2002.
- [109] S. L. Puller, "Pricing and firm conduct in California's deregulated electricity market," *The Review of Economics and Statistics*, vol. 89, pp. 75-87, 2007.
- [110] J. Hull, *Options, Futures, and Other Derivatives, 7th edition*. Pearson Education India;, 2010.
- [111] P. Wilmott, S. Howison, and J. Dewynne, *The mathematics of financial derivatives: a student introduction*: Cambridge University Press, 1995.
- [112] S. S. Oren, "Integrating real and financial options in demand-side electricity contracts," *Decision Support Systems*;, vol. 30, pp. 279-288, 2001.
- [113] P. L. Skantze and M. Ilic, *Valuation, hedging, and speculation in competitive electricity markets: a fundamental approach*, Springer; , 2001.
- [114] M. Ghamkhari, H. Mohsenian-Rad, and A. Wierman, "Optimal risk-aware power procurement for data centers in day-ahead and real-time electricity markets," in *Proc. IEEE Computer Communications Workshops (INFOCOM Workshops)*, , pp. 610 - 615, 2014.
- [115] L. Rao, X. Liu, L. Xie, and Z. Pang, "Hedging against uncertainty: A tale of Internet data center operations under smart grid environment," *IEEE Trans. Smart Grid*; , vol. 2, pp. 555-563, 2011.
- [116] S. P. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge university press, 2004.
- [117] Y. Oum and S. S. Oren, "Optimal static hedging of volumetric risk in a competitive wholesale electricity market," *Decision Analysis*;, vol. 7, pp. 107-122, 2010.
- [118] M. Coulon, W. B. Powell, and R. Sircar, "A model for hedging load and price risk in the Texas electricity market," *Energy Economics*, vol. 40, pp. 976-988, 2013.
- [119] A. Krioukov, S. Alspaugh, P. Mohan, S. Dawson-Haggerty, D. E. Culler, and R. H. Katz, "Design and evaluation of an energy agile computing cluster," *EECS Department, University of California, Berkeley, Tech. Rep. UCB/EECS-2012-13*, 2012.
- [120] M. H. Albadi and E. El-Saadany, "A summary of demand response in electricity markets," *Electric Power Systems Research*, vol. 78, pp. 1989-1996, 2008.

- [121] Í. Goiri, K. Le, M. E. Haque, R. Beauchea, T. D. Nguyen, J. Guitart, *et al.*, "Greenslot: scheduling energy consumption in green datacenters," in *Proc. International Conference for High Performance Computing, Networking, Storage and Analysis*, pp. 1-11, 2011.
- [122] Z. Liu, Y. Chen, C. Bash, A. Wierman, D. Gmach, Z. Wang, *et al.*, "Renewable and cooling aware workload management for sustainable data centers," *ACM SIGMETRICS Performance Evaluation Review*, vol. 40, pp. 175-186, 2012.
- [123] D. Paul, W. Zhong, and S. K. Bose, "Energy Efficient Cloud Service Pricing: A Two-Timescale Optimization Approach," *Journal of Computer and System Sciences*, *Accepted*.
- [124] B. Urgaonkar, G. Kesidis, U. V. Shanbhag, and C. Wang, "Pricing of service in clouds: optimal response and strategic interactions," *ACM SIGMETRICS Performance Evaluation Review*, vol. 41, pp. 28-30, 2014.
- [125] Amazon, "EC2 prices, Amazon Web Services," Available: <http://aws.amazon.com/ec2/pricing/>, 2015.
- [126] S. S. Manvi and G. K. Shyam, "Resource management for Infrastructure as a Service (IaaS) in cloud computing: A survey," *Journal of Network and Computer Applications*, vol. 41, pp. 424-440, 2014.
- [127] R. Buyya, C. S. Yeo, S. Venugopal, J. Broberg, and I. Brandic, "Cloud computing and emerging IT platforms: Vision, hype, and reality for delivering computing as the 5th utility," *Future Generation computer systems*, vol. 25, pp. 599-616, 2009.
- [128] B. P. Sharma, R. K. Thulasiram, P. Thulasiraman, and R. Buyya, "Clabacus: A Risk-Adjusted Cloud Resources Pricing Model using Financial Option Theory," *IEEE Trans. on Cloud Computing (TCC)*, vol. 3, pp. 332 - 344, 2015.
- [129] A. Ben-Tal and A. Nemirovski, *Lectures on modern convex optimization: analysis, algorithms, and engineering applications* vol. 2: SIAM, 2001.
- [130] M. Erol-Kantarci and H. T. Mouftah, "Energy-efficient information and communication infrastructures in the smart grid: A survey on interactions and open issues," *IEEE Communications Surveys & Tutorials*, vol. 17, pp. 179-197, 2015.
- [131] E. Hossain, Z. Han, and H. V. Poor, *Smart grid communications and networking*: Cambridge University Press, 2012.
- [132] M. Grant, S. Boyd, and Y. Ye, "CVX: Matlab software for disciplined convex programming," ed, 2008.
- [133] K.-C. Toh, M. J. Todd, and R. H. Tütüncü, "SDPT3—a MATLAB software package for semidefinite programming, version 1.3," *Optimization methods and software*, vol. 11, pp. 545-581, 1999.
- [134] "<http://www.eia.gov/electricity/wholesale/>," US energy information administration.
- [135] S. Borenstein and N. L. Rose, "Competition and price dispersion in the US airline industry," *National Bureau of Economic Research*, 1991.

- [136] D. Paul, W. Zhong, and S. K. Bose, "Energy efficient scheduling in data centers," in *Proc. IEEE International Conference on Communications (ICC)*; DOI: 10.1109/ICC.2015.7249270, pp. 5948 - 5953, June, 2015.
- [137] D. Paul, W.-D. Zhong, and S. K. Bose, "Demand Response in Data Centers Through Energy-Efficient Scheduling and Simple Incentivization," *IEEE Systems Journal*; DOI: 10.1109/JSYST.2015.2476357, vol. PP, 2015.
- [138] L. Rao, "Global E-Commerce Revenue To Grow By 19 Percent In 2011 To \$680B," *J.P. Morgan*; Available online: <http://techcrunch.com/2011/01/03/j-p-morgan-global-e-commerce-revenue-to-grow-by-19-percent-in-2011-to-680b/>, 2011.
- [139] M. Endrei, J. Ang, A. Arsanjani, S. Chua, P. Comte, P. Krogdahl, *et al.*, *Patterns: service-oriented architecture and web services*: IBM Corporation, International Technical Support Organization, 2004.
- [140] S. K. Bose, *An introduction to queueing systems*: Springer Science & Business Media, 2013.
- [141] J. Dean and S. Ghemawat, "MapReduce: simplified data processing on large clusters," *Communications of the ACM*, vol. 51, pp. 107-113, 2008.
- [142] T. White, *Hadoop: The definitive guide*: "O'Reilly Media, Inc.", 2012.
- [143] A. Rabkin, C. Reiss, R. Katz, and D. Patterson, "Using clouds for MapReduce measurement assignments," *ACM Trans. On Computing Education (TOCE)*, vol. 13, p. 2, 2013.
- [144] H. Yang, A. Dasdan, R.-L. Hsiao, and D. S. Parker, "Map-reduce-merge: simplified relational data processing on large clusters," in *Proc. ACM SIGMOD International Conference on Management of Data*, pp. 1029-1040, 2007.
- [145] J. Mattingley, Y. Wang, and S. Boyd, "Receding horizon control," *IEEE Control Systems Magazine*, vol. 31, pp. 52-65, 2011.
- [146] Z. Ma, D. S. Callaway, and I. A. Hiskens, "Decentralized charging control of large populations of plug-in electric vehicles," *IEEE Trans. Control Systems Technology*, pp. 67-78, 2013.
- [147] P. P. Varaiya, F. F. Wu, and J. W. Bialek, "Smart operation of smart grid: Risk-limiting dispatch," *Proceedings of the IEEE*, vol. 99, pp. 40-57, 2011.
- [148] N. Singh and S. Rao, "Ensemble Learning for Large-scale Workload Prediction," *IEEE Trans. Emerging Topics in Computing*, vol. 2, pp. 149 - 165, 2014.
- [149] M. Grant, S. Boyd, and Y. Ye, "CVX: Matlab software for disciplined convex programming," version 2.1 ed, 2014.
- [150] EMC. (2014). *Market clearing price of electricity, Singapore*. Available: <https://www.emcsg.com/marketdata/priceinformation>
- [151] T. Lu, M. Chen, and L. L. Andrew, "Simple and effective dynamic provisioning for power-proportional data centers," *IEEE Trans. Parallel and Distributed Systems*, vol. 24, pp. 1161-1171, 2013.

- [152] A. Greenberg, J. Hamilton, D. A. Maltz, and P. Patel, "The cost of a cloud: research problems in data center networks," *ACM SIGCOMM Computer Communication Review*, vol. 39, pp. 68-73, 2008.
- [153] D. Paul, W. Zhong, and S. K. Bose, "Energy efficiency aware load distribution and electricity cost volatility control for cloud service providers," *Journal of Network and Computer Applications*, vol. 59, pp. 185-197, 2016.
- [154] A. Gandhi, M. Harchol-Balter, and M. A. Kozuch, "Are sleep states effective in data centers?," in *Proc. International Green Computing Conference (IGCC)*, pp. 1-10, 2012.
- [155] S. Lloyd, "A sampling theorem for stationary (wide sense) stochastic processes," *Transactions of the American Mathematical Society*, vol. 92, pp. 1-12, 1959.
- [156] N. Wiener, *Extrapolation, interpolation, and smoothing of stationary time series* vol. 2: MIT press Cambridge, MA, 1949.
- [157] EEX. (2014). *European Energy Exchange*. Available: <https://www.eex.com/en/>
- [158] EIA. (2014). *Energy Information Administration, Department of Energy US; 2014*
- [159] EMA. (2014). *Energy Market Authority, Singapore*.
- [160] Wikipedia: https://en.wikipedia.org/wiki/Energy_in_Germany;
https://en.wikipedia.org/wiki/Electricity_sector_in_the_Netherlands; https://en.wikipedia.org/wiki/Energy_in_France;
https://en.wikipedia.org/wiki/Electricity_sector_in_Switzerland; , 2014.
- [161] M. Grant, S. Boyd, and Y. Ye, "CVX: Matlab Software for Disciplined Convex Programming, version 2.1," 2014.
- [162] K.-C. Toh, M. J. Todd, and R. H. Tütüncü, "SDPT3—a MATLAB software package for semidefinite programming, version 1.3," *Optimization methods and software 1999*; , vol. vol. 11, pp. p. 545-581, 1999.
- [163] J. Anselmi, D. Ardagna, J. C. Lui, A. Wierman, Y. Xu, and Z. Yang, "The economics of the Cloud: Price competition and congestion," *Submitted to Operations Research*, 2015.