Energy Efficient Computing, Clusters, Grids and Clouds: A Taxonomy and Survey

Muhammad Zakarya[1,2] and Lee Gillam[1]

[1]Department of Computer Science, University of Surrey, UK
[2]Abdul Wali Khan University, Mardan, Pakistan
{mohd.zakarya,l.gillam}@surrey.ac.uk

Abstract. Cloud computing continues to play a major role in transforming the IT industry by facilitating elastic on-demand provisioning of computational resources including processors, storage and networks. This is necessarily accompanied by the creation, and refreshes, of large-scale systems including cluster, grids and datacenters from which such resources are provided. These systems consume substantial amounts of energy, with associated costs, leading to significant CO_2 emissions. In 2014, these systems consumed 70 billion kWh of energy in US; this is 1.8% of the US total energy consumption, and future consumption is expected to continue around this level with approximately 73 billion kWh by 2020. The energy bills for major cloud service providers are typically the second largest item in their budgets due to the increased number of computational resources. Energy efficiency in these systems serves the providers interests in saving money to enable reinvestment, reduce supply costs and also reduces CO_2 emissions. In this paper, we discuss energy consumption in large scale computing systems, such as scientific high performance computing systems, clusters, grids and clouds, and whether it is possible to decrease energy consumption without detrimental impact on service quality and performance. We discuss a number of approaches, reported in the literature, that claim to improve the energy efficiency of such large scale computing systems, and identify a number of open challenges. Key findings include: (i) in clusters and grids, use of system level efficiency techniques might increase their energy consumption; (ii) in (virtualized) clouds, efficient scheduling and resource allocation can lead to substantially greater economies than consolidation through migration; and (iii) in clusters, switching off idle resources is more energy efficient, however in (production) clouds, performance is affected due to demand fluctuation.

1 Introduction

Large scale computing systems as observed in the top500 [1] supercomputers, clusters, grids [2], and clouds [3] consist of a large number of Information & Communication Technology (ICT) resources that are connected through a network. Supercomputers and clusters are non-distrusted systems which are used to solve large problems quickly where large mathematical calculations are involved like weather forecasting, defence & control systems etc. Distributed systems (grids and clouds) are preferred over non-distributed systems for three main reasons including reliability, distributed nature of applications and concurrent execution of applications [2]. These systems provide their services to users on either best or commercially reasonable effort policy.

Cluster, grid, cloud and datacenter service providers maintain a large pool of computational resources, that needs more energy to (i) operate properly and (ii) to cool down the heat generated. A recent report [4] show that in 2015, across the world almost 416.2 terawatt hours of energy was consumed by datacenters which is higher that the UK's total consumption. The amount of energy consumption will continue to increase with increasing capacity unless energy efficient management techniques are established and applied [5]. The resource allocation and management algorithms along with the physical infrastructure of the datacenter are needed to reduce the environmental impact (CO_2 emission) and make them more energy and cost efficient [6]. In this paper, we discuss energy consumption of clusters, grids and clouds, and whether it is possible to minimize energy consumption without detrimental impact on service quality and performance.

1.1 Clusters

A cluster (computer cluster) ties together a number of computers through a LAN, that essentially act as a single computer, that can be more cost effective than a single computer of comparable performance. Typically clusters are homogeneous, however there are certain heterogeneous HPC clusters like OSCAR [2]. A centralized job scheduler is responsible for resource allocation and placement.

1.2 Grids

"A grid is a system in which computing and data resources belonging to many enterprisers are organized into a single, virtual computing entity that can be transparently utilized to solve compute and data intensive problems" [7]. Computational grids which are based on the notion of virtual organizations [8], use high speed public networks for high availability of computational resources or multiple clusters at low-cost. Grids are more heterogeneous and geographically dispersed (distributed) than clusters. Due to multiple administrative domains, a hierarchical job scheduler (local and meta) is responsible for resource management in grids [2].

1.3 Clouds

According to National Institute of Standards and Technology's (NIST), *"cloud computing is a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction"* [3]. Compared to grids, a pricing model is associated to clouds, essentially virtualized, where the resource demand can be unpredictable. A hierarchical job scheduler [8] is used to provision resources for customers. Cloud computing is divided into three major types including Software as a Service (SaaS), Platform as a Service (PaaS) and Infrastructure as a Service (IaaS). Cloud computing provide ICT devices on lower costs to customers, which can be elastic to fulfil the varying resource demand, flexible and can provide resources on lower Total Cost of Ownership (TCO). Customers are looking for QoS and reliable services when they pay, therefore the focus of service providers is to satisfy customer needs with minimum energy consumed – performance-cost-energy model. A cloud can be public and/or private. Public clouds are operated by a third party such as Amazon[1] and Google[2], so the enterprises do not need to manage and maintain their allocated resources. Private clouds (based on the OpenStack[3]) are operated by large enterprises, who can afford the cost of operating and maintaining cloud datacenters.

Datacenters: Datacenters provide an IT backbone for cloud computing which may consist of large servers (in thousands) that process large tasks for businesses, complex scientific problems and facilitate customers to accomplish their business goals. A server can be virtualized that runs multiple VMs possibly of different instance type, for different users. Virtualization can provide opportunities for server consolidation that increases the ration of resource usage. Datacenters are IaaS, which are located in different geographical areas – for example Amazon $EC2$ regions that will have one or more virtualized datacenters. A high level system architecture for OpenStack [9] cloud computing platform is shown in Figure 1. A Global Resource Manager (GRM) is responsible for resource allocation, VM placement and initiating migration at the controller level. A Local Resource Manager (LRM) is working on each server, which monitor the resource usage, decide a host is under-loaded or over-loaded and inform the GRM to take appropriate action on resource reallocation.

[1] https://aws.amazon.com/
[2] https://cloud.google.com/
[3] https://www.openstack.org/

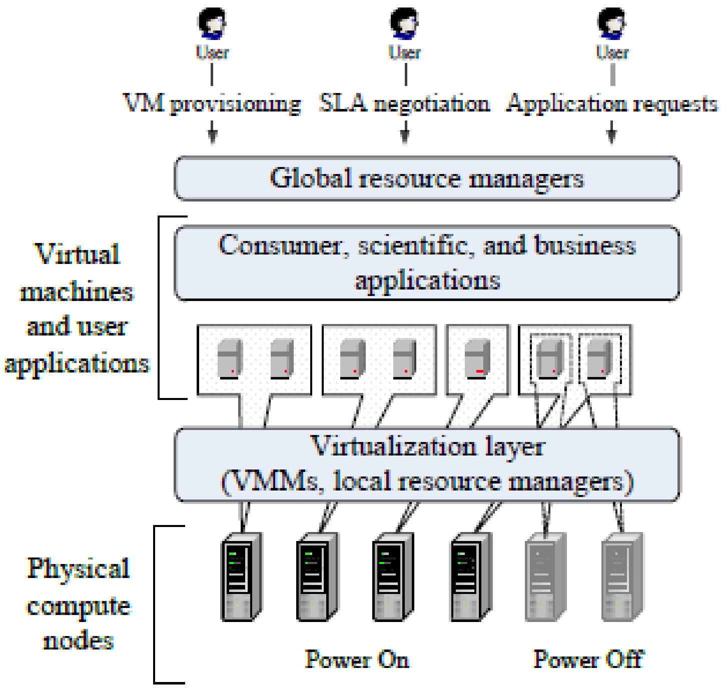

Fig. 1: A high level system architecture for cloud computing [9]

1.4 High Performance Computing (HPC)

According to NIST, *"HPC enables work on challenging problems that are beyond the capacity and capability of desktop computing resources"*. Supercomputers (such as INCITE [10]) and production clusters are considered the two most reliable HPC systems, where the focus can be to provision powerful resources for short period to complete the job quickly. These systems differs from High Throughout Computing (HTC) like HTCondor, where powerful resources are provisioned over a long period. Grids and clouds are not HPC, but they can certainly support HPC workloads as observed in TeraGrid (grid) [11] and Amazon new generation C4 instances (cloud) [12]. Despite knowing the differences between these systems [13], in the rest of document, by HPC we would mean any system amongst supercomputers, clusters, grids or clouds.

1.5 Problem

The modern computing era and the problems we are facing like global warming, increased fuel & energy cost and global crisis have driven scholars to investigate and decrease the energy demand of ICT equipment in small and large scale datacenters which are currently growing with the development of cloud infrastructure. As discussed in [14], the energy sector is one of the major Greenhouse Gas (GHG) emitters throughout the world, which produces 43% of GHGs in total. Following green computing principles, it is essential to minimize energy consumption in datacenters; ICT equipment that will lower GHG emissions. A key challenge for the owners of datacenters is the energy shortfall/outages that will continue to rise in future, due to the closure of all nuclear power plants in Germany, and reduction in coal-based power plants in the UK. In 2007, [15] the IT sector energy requirements (including PCs, cooling, servers, telephony, networks, mobiles, printers & office telecommunication) and CO_2 emissions were projected somewhat equal to that of airline industry, which is 2% of the global emissions. In 2013, US datacenters consumed an estimated 91 billion KWh of energy and were projected to be roughly consuming 140 billion KWh annually by 2020 [16]. However, due to state-of-the-art energy efficiency techniques implemented, a recent study [17] report that datacenters are accountable for 70 billion kWh of energy that is 1.8% of the total energy consumed in the US and are expected to consume approximately 73 billion kWh by 2020. Currently, the share of ICT equipment to global GHG emissions is around 1.6% and it is estimated to be around 2% by 2020 [18]. Table 1 shows the % of CO_2 emissions from ICT equipment.

The datacenter energy cost is in competition with the infrastructure cost. According to Amazon [19], the server cost in their datacenter is about 53%, while the energy cost is 42%, including 23% for servers energy, 19% for cooling and 5% for other infrastructure like lightning. Therefore, if money can be saved on the energy budget of the datacenters, this will result in greater profit and will have improved environmental sustainability [15]. A number of techniques are used to implement energy efficient mechanisms [20] in ICT equipment like processors and network cards, to make them less power hungry. These include mechanisms to reduce the energy consumption of processors especially in latest CPU architectures (multi-processors and multi-cores).

Table 1: CO_2 emissions (%) from ICT equipment (2005-2006)

PCs & Monitors	Servers	Telecom			LANs	Printers
		Fixed	Mobile	Office		
40	23	15	9	3.5	3.5	6

The focus of this paper is to discuss a number of state-of-the-art approaches, reported in the literature, that claim to improve the energy efficiency of large scale computing systems, and identify a number of open challenges. We will discuss energy efficient scheduling for single-processor, multi-processor and multi-core systems first in Section 4, and later in Section 5, we present a taxonomy to classify the existing work based on the methodology used for energy efficiency. Efforts are made to summarize datacenter level resource management techniques including virtualization, load balancing and server consolidation with migration.

1.6 Motivation

Both economical and environmental issues related to large scale datacenters motivate us for this study. With the rapid uptake of cloud datacenters to host industrial applications, reducing the operational costs of powering and cooling large scale datacenters is a major economical concern. As for every 10°C increase in temperature, the system failure rate doubles, hence reduced temperature could also improve datacenters reliability [3]. Various studies are conducted to elaborate and investigate green computing and datacenters. The purpose of this survey is to analyze the energy consumption of ICT equipment in general (including Laptops, PCs etc.) and focus on large scale HPC systems, clusters, grids and cloud datacenters. We explain a taxonomy of techniques that are proposed to enhance the energy efficiency of these systems. Conceptually clusters, grids and clouds are treated the same [2], hence considerable efforts have been made to analyze and differentiate the energy efficiency methods proposed for these systems. The contributions of the survey are as follows:

1. A taxonomy of energy efficient computing
2. System level energy efficient CPU scheduling in single systems and clusters
3. Datacenter level resource management for energy efficiency in grids and clouds

Our survey is different from those conducted in [9], [21], [14]. The techniques presented in [9], provide a taxonomy of optimizations, but the energy efficiency techniques in virtualized cloud environments have not been studied. Similarly, the surveys conducted in [21] [14], only focus on energy efficient datacenters and have ignored the energy efficiency of system level CPU and resource scheduling. These studies have also ignored the energy efficiency of multi-processors, multi-cores and clusters that form a base for grids and cloud systems. We start from the energy efficiency of a single system (its different components) and explore the energy efficiency of large scale cloud datacenters, storage systems and networking. We

believe that this survey will provide readers with an understanding of the essential concepts of the resource management techniques (to achieve energy efficiency) in cluster, grids and cloud systems. Furthermore, this survey will help researchers to identify the key and outstanding issues for further investigation.
The rest of the paper is based on the following roadmap. In Section 2, we introduce cloud datacenters and the need to make them more energy efficient. Modelling energy consumption of a non-virtualized or virtualized server is discussed in Section 3. Section 4 is devoted to system level (operating system level) green scheduling in terms of multi-processor and multi-core technologies. A brief survey on the existing scheduling algorithms is presented. Section 5 introduces the taxonomy of power management techniques in computing systems. Section 6 introduces different approaches to make large scale HPC, clusters, grids and cloud datacenters greener and more energy efficient. Application level green approaches are explained in Section 7. In Section 8 we discuss metrics, which are used to measure the energy efficiency of datacenters. Section 9 introduces new research challenges in the field of energy efficient computing and datacenters. Section 10 concludes the article, by offering recommendations.

2 Clusters, Grids, Cloud Datacenters & Energy Consumption

Clusters, grids and cloud datacenters consists of hundreds to thousands servers, processing large number of tasks for large businesses like Google [22] and Amazon, solve complex scientific problems and assist customers to achieve their business and scientific goals. These large scale cloud systems combined with clusters and grids are almost 90% idle (worldwide) [23], consumes large amount of energy and produce large amount of CO_2 emissions, which costs billions of dollars energy bills for the major service providers. Various studies [17], [21], [23], including the analysis by Natural Resources Defense Council [24] suggest that approximately 30% of the running servers in US datacenters are idle and the others are under-utilized, thus making it possible to save energy and money with state-of-the-art energy efficient methods [21]. A cloud datacenter consists of the following components:

1. Computing equipment including servers, network, and storage equipment.
2. Un-interruptible power supply components including switch gear, generators, PDU and batteries.
3. Cooling systems including Computer Room Air Conditioning units (CRACs), Heating Ventilation and Air-Conditioning (HVAC), chillers, pumps, direct expansion air handler (DX) units and cooling towers.
4. Supplementary equipment like workstations / laptops used to monitor, KVM switches.

Such components are also installed in large scale HPC clusters and grids. Various techniques are used to minimize the energy consumption of these systems. At server level, CPU is considered as a major energy consumer (Table 2), and research has focused on reducing its consumption (frequency or voltage scaling) as explained in Table 3. Other server components like network card, fan and motherboard can be either completely switched off or put into low power mode to save energy. Despite the potential benefits of switching components off, the motherboard is a high consuming component (Table 2), that can be switched off only if the whole server can [25]. Due to the nature of clusters workload, such techniques can be more energy efficient in clusters. However, due to the elastic nature of cloud datacenters and unpredictable customers resource demand, server level methods could not be used to achieve greater energy efficiency in grid and cloud systems. Other high level resource management techniques would be more helpful to maintain system performance and SLAs with customers.
Another technique is the use of virtualization technology, which allows a number of devices (servers, networks) to execute more services on the same hardware, to improve resource utilization. According to NIST *"virtualization is the simulation of the software and/or hardware upon which other software runs under the control of a hypervisor"*. Virtualization is used to achieve energy efficiency in grids and cloud datacenters through increased utilization that can be achieved with server consolidation [27] through Virtual Machine (VM) migration

Table 2: A typical server's component energy consumption [26]

Component	Peak power (W)	Count	Total power (W)	%
CPU	40	2	40	37.6
PCI slots	25	2	50	23.5
Memory	9	4	36	16.9
Motherboard	25	1	25	11.7
Disk	12	1	12	5.6
Fan	10	1	10	4.7
Total			213	

and with the use of energy efficient scheduling approaches [21]. Various studies shows that virtualization raises the utilization ratio up to 50%, save more energy[4] (through consolidation) and consequently drop the CO_2 emissions [28] [29]. Similarly, resource consolidation moves the servers around to reduce the number of active resources and can switch off some underutilized resources to save the resource idle energy. Load balancing [30] attempts to distribute the load across different servers that can be achieved through server consolidation in clouds and task consolidation [31] in clusters and grids. Load balancing can increase the servers utilization level, and can possibly minimize the total energy consumption. However, increasing the resource utilization does not always result in minimum energy consumption due to workload variations [31], hence an Energy-aware Task Consolidation (ETC) method is proposed in [31] to restrict the CPU usage below a specific threshold value. Estimating such a threshold level for a server in clouds, where the demand varies over time, is a challenging task due to workload unpredictability and heterogeneity.
Recently, most software[5] (applications) and hardware have delivered support for virtualization. We can virtualize many aspects such as applications, hardware, OS and storage[6], and easily manage them. To make use of virtualization, competent scheduling algorithms are essential, that are generally applied by cloud resource manager to provision only the required (ideal) number of resources. There are a large number of scheduling algorithms that focus on minimizing the total completion time of the tasks in distributed systems [32], [31], [33], however, these algorithms does not guarantee energy efficiency for large scale cloud datacenters – because after finishing task execution, the resources may become idle that still consumes energy. Some proposed system level techniques, scheduling, and switching off unused servers through consolidating the workload on fewer servers in datacenters, their pros and cons are discussed in Section 6.6.
Table 3 summarizes some techniques that are implemented in computing to achieve energy efficiency [21], including Static Power Management (SPM) and Dynamic Power Management (DPM). Dyanamic Voltage & Frequency Scaling (DVFS) scale the processor frequency up or down, if feasible to minimize the amount of energy consumed at the system level. The schedulers are designed to make use of such technologies at system level, to achieve energy efficiency. Similarly, Selective connectedness allows resources to go idle or low power mode for some time transparently in the network [34]. There is potential to increase the focus on these methodologies in terms of green datacenters, where it might not be possible to switch off servers due to demand variation and unpredictability. Such mechanisms including system level and high level resource management are discussed in Section 4 and 5 respectively.

3 Modelling Energy Consumption for Non-Virtualized & Virtualized Platforms

Due to the size of clouds, It is extremely difficult to conduct repeatable experiments on a real infrastructure, which is required to evaluate and compare different scheduling and resource management policies [9]. Hence, to ensure the repeatability of experiments, researcher use

[4] http://www.vertatique.com/average-power-use-server
[5] http://www.salesforce.com/
[6] https://www.techopedia.com/definition/4798/storage-virtualization

Table 3: Energy efficiency techniques to make ICT equipment greener

method	Scope	Explanation	Benefits	Shortcomings
DPM [Sec. 5]	Server	Switch ON/OFF servers according to workload [9]	More energy saving as compared to SPM	A realistic methodology is mandatory to coordinate, predict future workload
	Network	Switch off network when idle	Energy is not wasted when no traffic or packets are there	Degrades network performance
Scheduling with DVS & DVFS [Sec. 4.2, 6.1]	Server	Increase and decrease frequency and voltage according to the task criticality and priority [35], [36]	Reduced carbon emissions as servers are utilized according to voltage, frequency and utilized as needed	An approach essential to scale voltage or frequency independent of the workload. Insignificant energy savings in large systems. In heterogeneous datacenters difficult to implement
	Network	Link sleep and active, according to the packets availability through Adaptive Link Rate (ALR)	ALR achieves negligible energy savings [21]	A strategy needed to decide link rate, LPI states. ALR can achieve significant energy savings if implemented properly [37]
Virtualization [Sec. 6.6]	Server	Dynamically provision resources according to QoS requirements [38], [30]	More energy savings and high utilization ratio	Extensively used, live migration effect network performance
Protocols [Sec. 6.2]	Network	Keep minimum switches on for tolerable throughput energy efficient Transmission Control Protocol (TCP) [39]	Data packets are transferred only in an energy efficient state	Difficult to modify existing protocols to save energy
Server consolidation [Sec. 6.6]	Server	Consolidate the load on minimum servers [40] [41]	Increase utilization ratio. Have overcome the problems of server sprawl	Categorization of the servers needs proper planning and large amount of time. Failure of single consolidated server
Load Balancing [Sec. 4.3]	Server	Balance the workload among different servers to balance utilization	Equal utilization as no server is overloaded	Tough implementation in heterogeneous platform [30]
Reduce number of active servers [Sec. 6.5]	Server	Decrease the number of active servers in datacenters (consolidation) [23], [42]	Less active servers produce less CO_2 emissions and requires less energy to operate	Not essentially reduce the energy consumption in a heterogeneous cloud datacenter

simulations as a way to evaluate the performance of their proposed policies. To produce valid results that can be achievable in a real platform, plausible assumptions and models based on different characteristics of servers are still needed. As, the focus of this study is energy, thus we explain how the energy consumption of a typical real server is modelled. As shown in Table 2, the CPU is the chief energy consuming component in a typical server. Majority of the proposed energy models[7] are based on the assumption that a server total energy consumption can be modelled as the energy consumed by its CPU [26], [43]. An empherical study (based on the benchmarks) of CPU utilization and its energy consumption in [26], suggests that a linear relationship can be made. This relationship of CPU utilization and energy consumption can be used to estimate the server total energy/power consumption,

[7] https://www.infoq.com/articles/power-consumption-servers

given by Equation 1:

$$P(u) = P_{idle} + (P_{max} - P_{idle}).u \tag{1}$$

where $P(u)$ is the estimated server power consumption, P_{idle} is static power consumed by the CPU when the server is idle, P_{max} is the power consumed by the CPU when the server is fully utilized, and u is the current CPU utilization. The portion $(P_{max} - P_{idle}).u$ is called dynamic power consumption, and is treated as a linear function of utilization. The power model in Equation 1, allows an estimate of energy consumed by a server (non-virtualized) given CPU utilization and the energy consumed at idle and fully utilized state. Other models like PowerTOP and Joulemeter [25] are also accurate for non-virtualized servers. However, in virtualized servers, there is a multi-layered software stack that contains physical devices, native OS, a hypervisor and several VMs on the same hardware. The privileged hypervisor have control over virtualised server – they can get coarse-grained energy consumption of the server. However, fine-grained VM level energy consumption cannot be measured in this way [44], as the capability of server is multiplexed across several VMs, whose real energy consumption is decided by the characteristics of running application. Moreover, the virtualization layer makes it challenging to isolate application's energy consumption from server's total energy consumption.

Energy consumption of a physical server can be measured by metering the provided energy (through sensors or watt-meters), but VMs are made in software and cannot be attached with a meter, which makes it challenging to measure VMs energy consumption. The skill to measure the energy consumption of a VM or a virtualized server is significantly important for several reasons. For instance, if a VM is consuming too much energy it may be migrated to a more energy efficient server. VMs are sized such that a certain number will fit on a server. Efficiency (of a server) for a VM is going to be a factor of how many are running on that server. Assuming a constant power use for a server that is switched on with no VMs, the first VM is going to be least energy efficient – the baseline energy use gradually spreads across all VMs. Therefore new VM level energy consumption models are needed to measure accurate levels of total energy consumption across virtualized platform [44]. In our previous work [45], we proposed an energy consumption model for virtualized systems, to decide which type of host could run a VM more energy efficiently. Ibrahim at el. [46] also proposed a VM level energy-aware profiling model that attributes the host's energy consumption to VMs. Similar to our approach, they divide the host's idle (static) energy consumption evenly amongst the number of VMs running on it. However, the dynamic energy consumption of the host is divided based on the VM CPU utilisation level only. The proposed model could estimate the energy consumption of homogeneous VMs accurately, however, heterogeneity is not considered.

4 Scheduling

The idea of scheduling (system level) is to allocate jobs to processors considering job execution time, deadline (in case of real-time systems) and other characteristics. A job may contain multiple tasks that can be executed on different processors. Tasks can be independent or dependent, where the output of one task provides input for other tasks (workflow). The two terms, jobs and tasks are used interchangeably in the rest of this document. In literature [25], [47], energy aware scheduling algorithms either focus on scheduling of tasks in such a manner that their execution completes in minimum time with reduced energy consumption or tasks are balanced over a number of processors so that all processors runs with similar utilization level. For example, if the deadline of a job can be extended, then the processor frequency may be scaled down to minimize the energy consumption. An idle processor still consumes significant amount of energy that can be 60% of its peak energy consumption. However, switching off an idle processor is not feasible in a single system, so keeping processors always doing some workload would also save significant amount of energy. Other techniques like delaying the execution of a job (with the hope that another job will finish in near future) [48] and migrating a job to another system (in case some job on the system have finished execution) [25] would also save significant amount of energy for non-critical applications. In this section we discuss energy efficient scheduling techniques in single-processor, multi-processor and multi-core systems as they are very common in HPC clusters and cloud datacenters.

4.1 Single-processor scheduling

The scheduling problem has been predominantly studied on a uniprocessor system, which contains a single processing unit in which all the jobs must be executed after satisfying some scheduling constraints. The focus of the job scheduling is to allocate n jobs to m homogeneous or heterogeneous processors, such that the total makespan is reduced. The job total makespan is the length or duration of schedule when all of its tasks have finished processing. Job scheduling approaches are normally categorized as static (off-line) and dynamic (on-line) as shown in Figure 2. In static scheduling, the workload[8] size – number of required clock cycles in millions of instruction per second (MIPS), required physical resources (CPU), and job priorities are determined prior to their execution. Information about workload size, Worst Case Execution Time (WCET), job deadline and communication time is thought to be known at execution time. Min-Min (schedule small tasks for execution first) & Min-Max (schedule large task for execution first) are two common static scheduling techniques [47]. In dynamic scheduling, algorithms may change jobs priority level during execution and resources to running processes are allocated dynamically to maximize resource utilization. We know from the previous discussion (section 4) that maximizing resource utilization could execute the given workload more efficiently, however it strongly depends on the workload. The workload size and runtime is also not known in advance, which makes dynamic algorithms more challenging and complex. An on-line scheduler decides the job placement on the fly, and then the obtained results can be compared to estimate its optimality and efficiency to optimal off-line algorithm using competitive analysis [49]. Dynamic scheduling algorithms can be either (i) static priority - job priorities are fixed, or (ii) dynamic priority - jobs priorities varies over the execution time. Another category of scheduling is real time scheduling which comprise of static and dynamic priority scheduling algorithms as well, as shown in Figure 2. Similarly, scheduling can be categorized as centralized, distributed and hierarchical policies as explained in [2]. In centralized scheduling, a single scheduler is responsible for all users jobs, while in a distributed model, several schedulers cooperate with each other to satisfy the user experience. Distributed schedulers are able to schedule from the same pool. In hierarchical models the scheduling is a combination of centralized and decentralized schedulers, where a centralized scheduler is working at the top level and a decentralized scheduler is installed at the lower level [2].

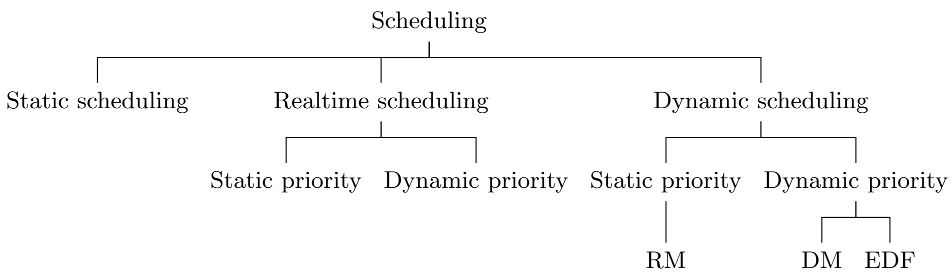

Fig. 2: Scheduling techniques for single / multi processors (RM - Rate Monotonic, DM - Deadline Monotonic, EDF - Earliest Deadline First)

[8] http://www.omgwiki.org/hpec/files/hpec-challenge/metrics.html

4.2 Multi-processor Scheduling

In multi-processor system the tasks can be executed in more than one processor (SMP[9] – simultaneous multi-processing). There are two major types of multi-processor scheduling: (i) partitioned and (ii) global scheduling. In partitioned scheduling [50], every job is allocated to a processor statically and then it is executed there without being migrated to other processors/cores. Each processor or core (based on the system architecture) have its own on ready queue. In global scheduling [51], jobs are kept in a single priority queue (shared) and the scheduler picks the job having the highest priority for execution. Unlike partitioned scheduling, jobs can migrate freely amongst different processors/cores. One of the optimal scheduling algorithms for multi-processor system is fair (Pfair) scheduling. Pfair works on the notion of "fair share of processor" where each task receives amount of CPU time slots proportional to its utilization, and hence guarantees that no deadline is missed. Assigning tasks optimally to multiple processors is a bin-packing NP-hard problem [52] and different heuristics approaches (on-line, off-line based on the nature of problem and workload) or approximation algorithms (to find approximate optimal solutions/schedule) [52] are used in the literature. Heuristics are fast enough and only provide good solutions but not optimal, however, approximation techniques are near to optimal. The main purpose of heuristics is to find out that all tasks are schedulable within the deadline (task feasibility), however they are not responsible for efficient and optimal allocation of tasks. Currently, for interested researchers, several other heuristic based multi-processor scheduling algorithms are suggested and offered in [9], [36], [53], [54] and [55].
The feasibility test of a job on a specific resource is to find out either all the tasks (real-time) in that job will meet their deadlines or not. Many parallel applications consist of several computational tasks that can be modelled as a weighted Directed Acyclic Graph (DAG). A vertex denotes a task, its weight represents task computational size and the edge (directed) shows the dependency between two tasks. An arc shows communication among two tasks, and its weight is the cost of communication [55]. While the execution of some of these tasks depends on the completion of other tasks, others can be executed at the same time, that increases parallelism of the scheduling problem. When the feasibility of such application is checked, then the communication overhead of tasks must also be considered. If all tasks are executing on the same processor, then their communication overhead is zero. However, if tasks are executing on separate processing units, then their communication overhead is positive and it is possible that the job will take longer to complete. The following postulations may be made, when designing a multi-processing unit scheduling algorithm to achieve energy efficiency:

- Task pre-emption is allowed.

When the task is running on a processor, then it could be pre-empted by some higher priority task and could be resumed later on. We assume that there is no penalty (on energy efficiency) linked with such pre-emption.

- Task migration is allowed.

The task that is executed on a processor may be pre-empted by another high priority task, and can resume its execution on another processor or even on the same processor later. And there is no penalty allied with such migration and pre-emption.

- Task parallelism is prohibited.

That is, it can only execute on one processor at a particular instant of time.
Those scheduling algorithms that are implemented for energy optimizations on uni-processor systems need to be redesigned for multi-processor platforms. Several scheduling techniques including Rate Monotonic (RM), Deadline Monotonic (DM) and Earliest Deadline First (EDF) are already optimized to achieve energy efficiency in multi-processor systems [47].

[9] http://superuser.com/questions/214331/what-is-the-difference-between-multicore-and-multiprocessor

The communication overhead must be considered while implementing such techniques as it might reduce the system performance and increase the energy consumption in case of workflow scheduling.

4.3 Multi-core Scheduling

A Multi-core processor consists of multiple execution cores which performs different arithmetic and logical operations (CMP[10] – chip-level multi-processing). These systems provide similar performance to multi-processor systems where multiple processors (single core) work in parallel[11], but probably at lower cost. Theoretically, adding an extra core to the same chip doubles the performance of the chip. However, in practice performance of each core is slower than a single core processor. The communication amongst cores and main memory is achieved in two different ways: (i) through a single communication bus, which is also known as a shared memory model and (ii) through an interconnected network approach which is also known as a distributed memory model. The multi-core processor provides 4 times greater bandwidth, decreased energy consumption and running at lower frequency with the same voltage as a single core processor [56], [47].
Increasing the frequency of a single-core processor also increases the energy consumption due to the linear relationship of frequency to energy consumption [9]. For that reason, vendors adopted an approach of multiple cores on a single chip against the single core processor to increase the performance with lower energy consumption. There are several major challenges involved in using multi-core architectures to minimize energy usage [47]. *"Multi-cores have the potential to be more energy efficient, however, they does not save energy unless you simplify each individual core to make it more energy efficient"*[12]. Several industry organizations have been making enhancement in solving those challenges with streamlining applications and splitting them amongst the cores. The idea to divide the task set $t_1, t_2, t_3, ..., t_n$ into k subsets where every subset is feasible on a single core is called partitioning. To use multiple cores energy efficiently, the partitioning problem is an important concern. Multi-core hardware can increase the running application performance and minimize the energy consumption even more, if all the cores stay similarly active [47], [56]. If the cores are not equally loaded, it can waste CPU cycles and can increase message passing amongst different cores which reduces the running applications performance [50] and increases energy consumption. There are several basic steps involve in designing parallel applications for these systems. In the first step of designing algorithm, designers need to find out the opportunities for parallel execution of the application (as an application might consist of several tasks). The main purpose of this, is to define number of small tasks that can run in parallel on different cores. The output of one task may provide input for other task (workflow), so the data must be transferred between tasks to continue processing normally. This message passing or data transfer between tasks is specified in the communication phase [40]. In the last stage, designers needs to identify where tasks will execute. The load balancers [56] which are used to balance the workload amongst available servers in a compute cluster, to increase utilization level and minimize the chances of server overloading – such techniques are also used in multi-core processors.
Apart from these, different techniques have been also implemented at processor level to save energy. For example, processor speed controlling using DVFS (low utilization) [57] [58], a possibility of on/off switching (power cycling) [59] can save more energy but still considerable efforts are needed to find a conciliation between keeping a core idle, in on state or switch it off [60]. The deadline of a job must also be kept in mind, when scaling down the frequency or voltage of the core, as the job might miss the deadline due to increase in execution time while running on lower frequency [56]. Switching on/off a core also cost more energy (reconfiguration cost) and can degrade the system performance which increases

[10] https://software.intel.com/en-us/blogs/2008/04/17/the-difference-between-multi-core-and-multi-processing

[11] http://forum.cakewalk.com/Difference-between-3939multicore3939-and-3939multiprocessor3939-m1098367.aspx

[12] http://www.futurechips.org/chip-design-for-all/a-multicore-save-energy.html

energy consumption. Such techniques at server level can be more energy efficient, which are further explained in Section 6.5.

5 Taxonomy of Energy Efficient Computing

With the prompt growth of people experience in Internet, volume of mobile devices and size of datacenters, there is more demand to green ICT devices for sustainable environment and to reduce their energy usage. Largely, datacenters are overwhelmed with many thousands of servers consuming enormous amount of energy, remains almost 30% idle [24] and are hardly utilized only 10% to 40% [29]. The scheduling and load balancing methods which are suggested for system level energy efficiency [section 4], are also used at high level resource management in compute clusters, grids and large scale cloud datacenters to manage the computational resources more efficiently. Some energy management techniques (system and high level) suggested so far in the literature are discussed in this section.
In computing equipment, energy saving techniques have been categorized as Static Power Management (SPM) and Dynamic Power Management (DPM) [61]. SPM falls under the category of hardware level techniques and are most efficient at single system while DPM are application level resource management methods, which are more energy efficient in large systems. Figure 3 outlines a fleeting overview of various energy management schemes in computing equipment [9]. Most of the work in SPM techniques are related to hardware level efficiency, for example low power consumption circuit designing. DPM techniques are mostly implemented in software or on network layer, for example protocol design and algorithms. Energy aware scheduling like DVFS [62], energy efficient routing [63], ALR [64], fan control [65], load balancing, virtualization [6], resource consolidation and migration [27] are mostly studied by a number of researchers in energy efficient computing [Table 3]. Some work on selective connectedness (turn off the network devices) when they are not in use for some period, are also proposed [61]. The problem is that high availability, QoS and performance guarantee are still ignored, which is most desirable in such distributed environments as the customers pay for their provisioned resources. If QoS and expected performance level is not achievable, then the customers would not even pay or may alternately switch to other similar service providers [2]. In some environment where mobility is involved (mobile cloud computing), energy issues are critical and needs to be managed properly as power batteries are not much reliable [66]. There are also application/software level methods to limit the amount of energy used by the applications though green compilers and robust programming [67]. However, application level efficiency techniques are not the focus of this document. Interested readers are recommended to read [67], [68]. In next sections we discuss hardware level and high level resource management techniques to achieve energy efficiency at single system, clusters, grids and cloud datacenters. Frequency scaling [62] and ALR [69] are hardware level efficiency techniques and are discussed in Section 6.1 and Section 6.2 respectively. The system level scheduling for single processor, multi-processors and multi-core systems were discussed in Section 4, that make use of certain hardware applicabilities to save energy. In Section 6, we limit our study to energy efficient resource scheduling, communication, storage and resource management for large scale clusters and grids. A complete section i.e. Section 6.6 is devoted for datacenter level energy efficiency techniques. Cloud systems implements a hybrid architecture from clusters and grids, hence those methods which are also applicable to these systems, are clearly indicated where appropriate.

6 Energy Efficiency Techniques in Clusters, Grids and Clouds

In the literature [9], [14], [25], [68], [70] energy efficiency of computing devices is considered at three different levels including efficiency of hardware, resource management systems and applications as shown in Figure 3. In this section we have categorized our study in different sections including energy efficient CPU scheduling (system level technique), communication, storage and resource management (cluster/grid/datacenter level techniques including server consolidation, migration and power cycling). We have discussed latest techniques which are

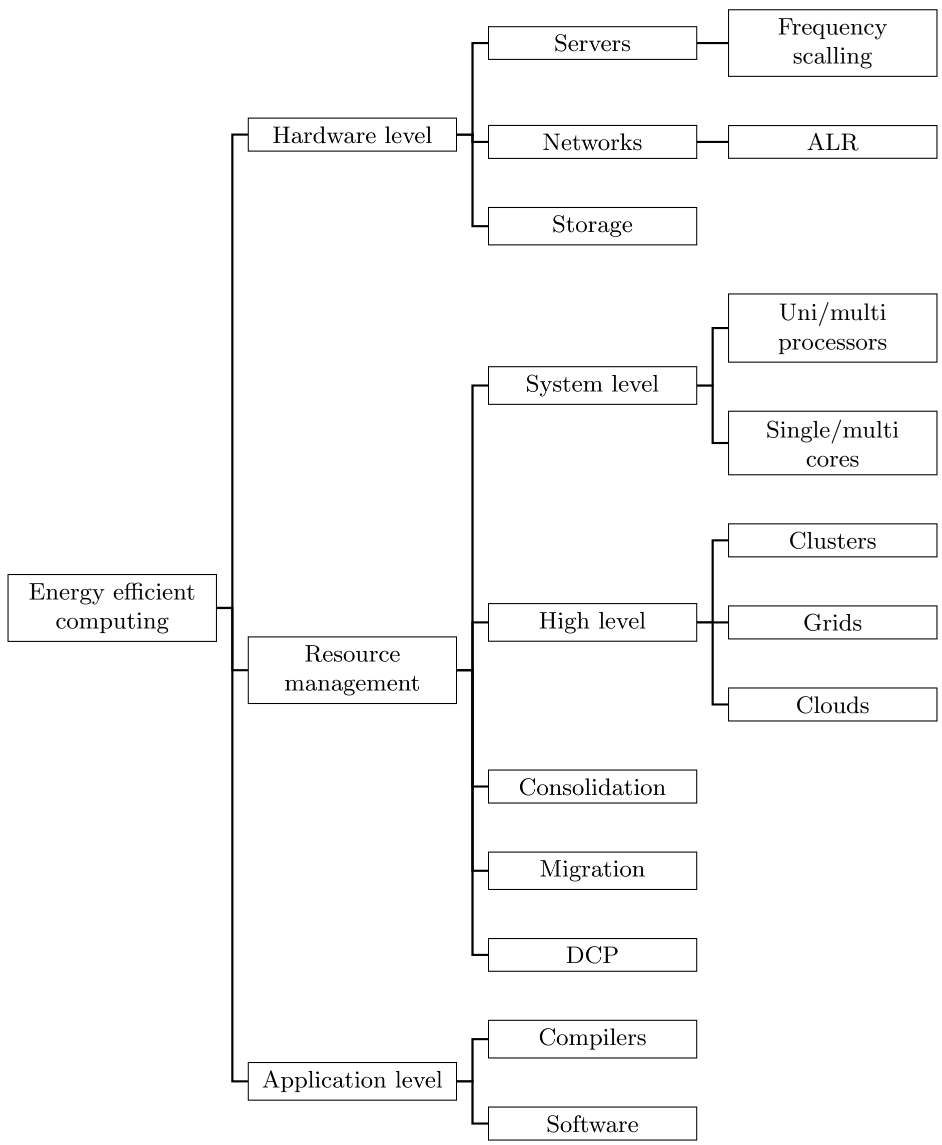

Fig. 3: Taxonomy of energy management techniques in computing systems

claimed to make large scale datacenters more energy efficient. As discussed before, a cloud datacenter consist of different ICT equipment which consumes enormous amount of energy [9], [71]. Table 4 shows the energy usage of different components in a typical datacenter. It is clear that datacenters and servers energy consumption are in competition for infrastructure. Therefore, the main focus of researchers is to decrease the energy consumption of VMs and physical servers (IaaS) in datacenters. Some methods like (i) shut down idle servers to save the idle power which is almost 70% of the peak and (ii) execution on those servers that are powered by Renewable Energy Sources (RES), are also proposed in the literature [9], [14] to minimize datacenters energy consumption. We will shortly discuss several state-of-the-art approaches in Section 6.6 for green and energy efficient datacenters. According to Environmental Protection Agency (EPA) energy efficiency report on US datacenters [72], about 70% energy saving can be achieved by implementing modern state-of-the-art efficiency techniques [9] at the cooling, power delivery & management levels of the datacenters. Table 5 shows the total energy consumption of datacenters, in 2011, along with the savings that can be achieved with state-of-the-art energy management techniques [21]. Datacenters are also accountable for GHGs productions because the energy production units (i.e energy grids)

Table 4: Energy usage of a typical datacenter (2006)

ICT Component	energy usage (TWh)	total usage %
Infrastructure	30.7	50.0
Network devices	3.0	5.0
Storage	3.2	5.0
Servers	24.5	40.0
Overall	61.4	100.0

Table 5: Energy savings from using state-of-the-art methods (2011) [21]

Module	energy usage (billion kWh)	energy usage with state-of-art (billion kWh)
Infrastructure	42.1	18.1
Network devices	4.1	1.7
Storage	4.1	18
Servers	33.7	14.5
Overall	84.1	36.1

which are used to operate these datacenters, release a large amount of CO_2 when fossil fuels including coal, oil or natural gases are used to generate electricity. Similarly, datacenter cooling equipment also emit CO_2 during operation. In 2006, the ICT sector was estimated to produce 2% of the worldwide CO_2 emissions. It has been shown that datacenters due to their large size and energy requirements are one of the key contributors to global GHGs emission that further increases global warming [73]. Hence there is a solid need to address this issue for environmental sustainability as well.

6.1 High Level Scheduling

The most considerable units of a system (or datacenter) where there is more space for reducing energy consumption are computation, communication, cooling and storage units. According to Moore's law, the performance per watt ratio is increasing, but the total energy drawn by the machine is scarcely decreasing [5]. If this tendency continues, the energy cost of a server during its lifespan will beat the hardware cost in near future [73]. The problem is even more worse for large scale clusters, grids and clouds which consumed energy of $4.5b in 2006 [9]. As stated before that system level approaches might not be energy efficient in large systems due to (i) energy consumed in idle state is much more than dynamic energy consumption, (ii) large number of idle resources, and (iii) varying demand and performance loss (heterogeneity of resources and workload). Therefore, other resource management systems are required to monitor the system state and take appropriate energy efficiency action in large systems. The term scheduling, for large systems is very similar to resource allocation and VM placement. VM placement [6] is treated as a bin packing optimization problem in the literature, and several heuristics have been proposed to allocate the provisioned resources more efficiently [74]. In [15], an energy consumption model is suggested for the private cloud datacenter based on three steps, (i) optimization, (ii) monitoring and (iii) reconfiguration. The entire state of the datacenter is continuously monitored by the optimization module to find out energy efficient software alternatives and service deployment configurations. Once it detects suitable energy efficient configuration, the deployment is reconfigured for energy efficiency. A 20% decrease in energy consumption is claimed for a private cloud datacenter. The proposed model is generic and can be implemented in traditional, super or cloud computing for energy minimization.

A brief survey on resource allocation (high level scheduling) problem in clusters, grids and cloud computing systems is conducted in [2]. Clusters support centralized job management and scheduling, grids follows a decentralized resource management system and clouds are based on hybrid architecture of clusters and grids. As clusters did not support virtualization, hence load balancing and job/process migration are more feasible to save energy in these

systems. In [20], the authors have briefly surveyed the energy efficiency in clusters using load balancing, power scalable processors (DVFS) and low power memory chips. Another work in [57], scales up and down the processor frequency based on the current demand (DVFS), to minimize the CPU energy consumption. For systems like clusters and grids, where the demand does not frequently change, DVFS can be efficient. However in clouds, as the demand varies over time, it cannot be feasible to switch processor frequency frequently. It should be kept in view, that the performance of the processor is affected when running on lower frequency and the task would take longer to complete – more energy is consumed. Thus, for application that needs a hard deadline met (hard real-time systems), it is not feasible to reduce the clock frequency. Grids are very similar to clouds, and previous studies [75] have shown that the grid nodes are mostly idle; thus consuming significant energy which can be saved by switching them off. Machine learning based load prediction techniques can be more helpful to decide when to switch on nodes again to accommodate the varying demand. Switching nodes on and off (power cycling) also cost more energy and affect the system performance, customer SLAs, resource reliability and lifetime [76].
In [29], virtualization is briefly discussed to achieve energy efficiency in computational grids and large scale clouds using server consolidation with migration technique. In virtualized platforms, server consolidation increases the utilization ratio up-to 50%, and can save large amount of energy by turning off the idle servers. Server consolidation ensures that IT setup subsidizes as little as possible to CO_2 emissions, and recover cooling & power capacity which exponentially drop energy costs. However, we found in our previous work [45], that consolidation is some time more expensive and cost extra energy due to migrations. Largely, efficient scheduling (resource management) techniques are more cost-energy efficient than migration based methods. In Section 6.6, we will briefly discuss several recently proposed energy efficient techniques in cloud datacenters.

6.2 Energy Efficient Communication

With the rising electrification and connectedness of the society including smart phones, laptops, ad-hoc and Wi-Fi networks, a report [73] has suggested that 1% to 3% of US electricity use comes from datacenters. This figure is even larger if energy use also accounts for the consumer devices and networks (fixed and wireless). In [77], it is estimated that access networks including consumer devices use as much energy as datacenters, and these have a faster growth rate of usage. Table 6 shows the network devices estimated energy consumption in Italy by 2020 – where due to scaling effect, reducing the least consuming equipment energy consumption can still lead to large savings [78]. The study in [77], [21] also signifies the need of energy efficient techniques to reduce the energy consumed by access networks and servers in cloud datacenters. In [21], it is estimated that a typical datacenter network accounts for 30% of the total energy consumed – including 15% for access switches, 10% for aggregate switches and 5% for core switches. In state-of-the-art literature [64], [77], [78], [21], server

Table 6: Networks energy consumption forecast 2015–2020 [78]

Network type	Energy consumption (W)	No of devices	Energy consumption (GWh/year)
Home	10	17,500,000	1,533
Access	1,280	27,344	307
Transport	6,000	1,750	92
Core	10,000	175	15
Total			1,947

communication is one of the prime energy consumers where energy optimization must deal with performance, QoS and energy savings trade-offs. There are network hardware that offers different features, creating a chance for energy-efficient operations including turning off network interfaces when there is no traffic and the media is idle. Using such techniques,

significant amount of energy can be saved [78], however despite its benefits, these techniques raises other problems including loss of connectivity and long re-synchronization periods. Additionally, constantly power cycling network devices can be more energy consuming than keeping them on all the time. Therefore, the underlying communication protocols must be optimized and re-designed to enhance the energy-efficient operations of the network [79].
To enhance the network performance with minimal energy consumption, other techniques like ALR [69], [64], [80] slowdown speed of the communication media whenever there is no or less data packets to transfer. As stated in [25], the difference between an idle and a fully utilized Ethernet link is negligible, hence the energy savings with ALR are not much significant. Similarly, the congestion control mechanism should be re-designed to minimize data transfer in such a way that the throughput is never affected; during a congested network. Energy efficient TCP [39] has overcome the two major problems of simple TCP, (i) the acknowledgement scheme that does not provide sufficient information about the state of the destination server to the sending server, and (ii) the window management that makes TCP aware of the burst errors. These two optimizations cause a 75% reduction in energy overhead. Such type of energy efficient protocols can reduce the energy consumption of datacenter networks [34] as shown in Table 7. The problems with ALR implementation include queue management,

Table 7: Energy consumption of datacenter network devices

LAN Switches	Hubs	Routers	WAN Switches
54%	26%	18%	2%

increased packet delay, packet loss ratio and QoS, which effect the network performance and throughput. Low power idle (LPI) mode [25] for Ethernet devices are defined by IEEE 802.3az standard, which can save at least 50% of the network controller energy consumption in the US, which consumed an estimated 5.3 terawatts hours in 2005. Switching off network devices, adapting the speed to the need (ALR), or a combination of both is more energy efficient but are worst in transfer time and performance. In clusters and grids, for certain workloads & QoS constraints, this delay can be acceptable or shortened with an accurate traffic prediction technique. However, in clouds, where the demand varies in an unpredictable way and the customers pay for their expected service time – such techniques can be less energy efficient or may result in SLA violation and performance degradation. To minimize the energy consumption of the global networks, the work in [63] has considered the problem as a NP-hard bin-packing problem. The objective function is to minimize the number of network elements needed for the communication and an Integer Linear Programming (ILP) formulation (heuristic) is proposed to solve it. In large scale datacenter networks with thousands of devices, ILP is not an efficient and faster approach to achieve the objective [27]. For such systems, high level energy efficient techniques like Green TCP, network algorithms, Clean-slate approaches and energy-aware framework (network virtualization) [25] can be more performance-energy efficient.

6.3 Energy Efficient Storage

Industry reviews [81], [17] have illustrated that US datacenters energy cost had increased by 15% per year since 2011. Along with numerous modules installed in a datacenter, storage devices are one of the prime energy consumers. One of these survey report has further suggested that storage module consumes approximately 27% energy in datacenters. Increasing demand for performance had introduced storage devices with higher energy needs; and this trend is continuously rising annually by 60% [82]. Given the well-known growth in datacenter TCO, several solutions that can decrease the energy consumed by storage module, and keeping customers data highly available, are proposed in the literature [25]. On one side, efforts are continued to make the drives more energy efficient. On the other side, efficient algorithms are proposed to decrease the energy consumption.
Using High Capacity (HC) drives can considerably change the storage energy consumption.

Classic Serial Advanced Technology Attachment (SATA) drives available in the market consume 50% less energy than capacity fiber channel drives per TB. Having the highest storage density per drive, HC drives minimize storage energy consumption. Table 8 shows the potential savings from HC drives [82]. The article [82] has discussed several findings to make the storage more energy efficient including consolidation of storage module, usage of HC drives, protection against storage module failure, migration of data to more energy efficient module, increase in disk utilization, backup of data, elimination of storage overhead and measure energy efficiency of the storage module & datacenter on a regular basis. In [83]

Table 8: Savings from greater volume drives [82]

	Old System	New System	Improvements
Number of systems	11	1-FAS 3020 with	
Systems details :	3-F810	3 disk shelves	
	4-F880		
	2-F820		
	1-F840		
	1-F825		
Energy savings without cooling (KWh)	113651	20919	81% decrease
Space (Cubic Feet)	63.0	4.3	93% decrease
Capacity (GBs)	9776	14000	16% increase

an efficient buffer-disk scheduling algorithm is suggested which is 38% more energy efficient and keep idle disks in sleep mode to save energy. In [84] a new architecture is proposed for energy efficient Redundant Array of Independent Disks (RAID) system to save energy using RAID redundant information. New controller-level cache management & I/O scheduling techniques are proposed for energy efficient RAID-5 and RAID-1 respectively. The experimental study for single speed disks show that RAID-1 can save 30% energy and would save 11% more along with RAID-5. Similarly, for multi speed disks RAID-1 can save 22% energy, and would save 11% additional energy with RAID-5. Also, energy-aware cache management systems have been suggested in the literature [85], [86] to save more energy. Belady's Offline Power-aware Greedy (OPG) algorithm [86] has minimized cache misses and the idea is extended for greater efficiency to an on-line energy-aware cache replacement method in [85]. Although their results are interesting however, OPG is designed for single disk and the online algorithm is considered for several disks. It can be more efficient if the proposed on-line algorithm can be implemented on a single disk. Similarly, the storage cache is active all the time and both techniques have ignored its energy consumption.

Pergamum [87] is a network of smart storage modules that stores data energy efficiently and reliably. At present, Massive Arrays of Idle Disks (MAID) systems retain storage modules idle to save energy. However, Pergamum complements Non-Volatile Random-Access Memory (NVRAM) at each node to store data signatures, raw data and allowing data requests to be completed while the storage module is switched off to minimize energy consumption. Several other techniques have been proposed to minimize the energy consumption in storage module include DPM schemes [88], cache management [85], pre-fetching [89], software directed power management [89], redundancy [89] and multi-speed setting [90]. However, the study of energy efficient parallel and distributed storage module in large scale systems is still in its infancy and need more attention of the research community.

6.4 Renewable Energy Sources (RES)

Renewable Energy Sources (RES) like wind, solar etc. also play a major part in decreasing clusters, grids and clouds energy consumption. As grids and datacenters are distributed in multiple regions, that are powered by different energy sources. Routing more user requests to the region which is powered by cheaper production technology (RES), can save enough energy. However, RES are intermittent and require policies to tackle the variability in supply. There are at least three benefits: (i) all oversupply of renewables allows for more energy to be provided back to the electricity grid; (ii) low supply of renewables means lessened demand on non-renewable sources from the electricity grid; (iii) lessened reliance on means to store renewable energy reduces the costs of management and replenishment of storage mechanisms, such as batteries, and extending the life of these mechanisms. RES region can also provide other benefits to on-cite local clusters and datacenters including (i) free cooling with outside air (as in Sweden) and/or (ii) cooling with sea water (as in Google Finland datacenter). Cooling costs almost 33% of the total datacenter energy consumption, hence significant energy can be saved that can be up-to 40% for large production clouds such as Google[13].

6.5 Resource Management

Resource management is significant to the victory of large scale computing, as it regulates the efficiency with which resources are used and that QoS is provided to the customers. Today's computing systems deliver simply a best-effort or reasonable (in clouds) service to their consumers. There are certain complex applications which involve immense computational power, strict delay and other best effort services; failure to provide the required performance makes consumers unwilling to pay. In order to provide QoS and guaranteed performance to the customers there is a strong need for efficient performance based resource management and scheduling techniques. Performance based resource management provides assured services to premium consumers (SLA) and reasonable services to other consumers. Such approaches are able to allocate resources (VM placement) near-optimally, in view of task characteristics and performance requirements [2]. Modern methodologies of energy efficient resource management for datacenters typically model this interesting problem as a bin-packing optimization problem with the objective to decrease the number of required servers to accommodate current demand (VMs). However, decreasing the number of servers may not essentially reduce the energy consumption of a heterogeneous cluster due to varying demand and resources reconfiguration cost.

In [91], energy saving in clusters is discussed in terms of two methods, (i) local and (ii) cluster-wide. Local methods focus on decreasing energy consumption of a single server either with (i) reduced processor clock speed/voltage (DVFS) and/or (ii) energy savings in network modules (switches) etc. DVFS, Dynamic Link Scaling (DLS) and request batching are the common examples of such implementation. Cluster-wide methods aggregate total workload on minimum number of servers and savings are achieved by switching off the idle servers in the cluster. The energy saved from slowing down the processor speed and then scale it back, is far less than the savings with switching off idle servers (idle processor consumes more energy). Therefore, cluster wide techniques like Independent Voltage Scaling (IVS), Coordinated Voltage Scaling (CVS) [61] etc. are used to save more energy in clusters. These homogeneous systems are for specific applications and sometime go to failure state for reallocation of the workload. The proposed stable framework runs in an energy efficient way, and have little influence on system performance. Inner job and outer job blocking problems [91] are mostly influencing the QoS, performance and throughput of a cluster. Solution to the former one is to balance the jobs among servers; however it require that the CPU and the memory requirement are known in advance. The second problem happens due to a large task having incredible requirement that cannot be fulfilled. MAGNET [91] has tried to solve these two problems and thus save more energy with improved performance score.

In [23], an energy efficient resource management system for virtualized datacenters is proposed which decreases operational expenses (OPEx) without affecting the QoS and response time. Energy savings are achieved by VMs consolidation that matches to current resource utilization, inter-networking topology and thermal states of servers. A similar approach is also presented in [42] where a workload power-aware on-line provisioning technique reduces the energy consumption by turning off subsystems that are not needed by VMs. Other resource management schemes, for example [92], [93], [76] have also discussed energy efficiency

[13] https://deepmind.com/blog/deepmind-ai-reduces-google-data-centre-cooling-bill-40/

but have ignored the heterogeneity of clusters and workloads.
Reducing the number of servers involve consolidation of current demand on fewer servers, and switching off the idle servers to save energy [76]. The idle servers also consume significant amount of energy, however a researcher from a cloud provider states that it is infeasible to switch off idle servers in their datacenter [94]. Also consolidation involves migration which have energy overhead, and sometime it is more energy efficient to run the VM without migrating it to a less efficient target server. It is also possible that the VM got terminated before or just after its migration is completed, hence the migration effort is wasted. Migration would be more feasible if the VM is migrated to a more energy efficient target server where the VM will be able to recoup back its migration energy overhead. A comprehensive research is required to know which VMs (of which runtime) should be able to recoup back their migration energy and how much energy savings can be made if the VM is subsequently running more efficiently [45].
We found in a simulated experiment, that server consolidation is more energy efficient if the idle servers are switched off, as they are consuming more energy. However, the tasks arrival and unpredictable resource demand are challenging to estimate when and which servers to be switched on/off. A server consumes more energy during the boot and shut-down process (server reconfiguration costs), that should be kept into consideration with increased failure rate. According to [95], the transition energy spent when switching a server from sleep state to on state is 4260 Joules, and that when switching a server from on state to sleep state is about 5510 Joules. If a server takes longer to boot (setup delay) [76], and if the demand suddenly rises, then the performance (in terms of scheduling delay [42]) and QoS is also affected. Efficient scheduling and resource management techniques are required to minimize the number of used servers without affecting the system performance [76]. The problem of this type - scaling the resources dynamically to meet the current resource demand, is called Dynamic Capacity Planning and is vastly studied by the research community in [96], [97].
Additionally, a major role in energy consumption is also played by the end user equipment. Although these equipment are rarely used in datacenters but still it should be noted that they also affect the environmental sustainability and increases global warming. Table 9 shows the energy consumed by end user desktops and monitors in three different states [98].

Table 9: Energy demand of user monitors [98]

	Energy consumption (Wh)		
Type	On	Sleep	Standby
Desktop	74.16	4.60	2.80
CRT	61.39	3.68	2.03
LCD	35.51	1.15	0.96

6.6 Datacenter Level Energy Efficiency Techniques

Largely, the energy efficiency techniques in datacenters fall into three broad categories; (i) Dynamic Capacity Planing (DCP) which allows hosts to switch on/off dynamically to save energy, (ii) Dynamic Voltage Frequency Scaling (DVFS) where a host adjusts its operating frequency (voltage) to lower power mode, dynamically, and (iii) Dynamic Power Management (DPM) schemes (resource management), which allows host's components to be in sleep state (to save energy) and decides when and for how long the component should be put to sleep [99]. Lim [100] proposed a hybrid scheduling scheme consisting of all these three approaches and demonstrated through extensive simulations that the proposed approach is 50% more energy efficient.
Figure 4 shows a taxonomy of different approaches that are proposed in the literature to achieve energy efficiency in cloud datacenters [9], [25], [100]. Virtualization allows different VMs running on a single server, thus creating more opportunities to consolidate the workload on lessened servers to save energy in datacenters. DCP [96] allows switching on/off the

available resources to meet only the current demand, that could save more energy costs for production clouds - where the resource demand is low [22]. In [38], the authors have suggested that light sensors could be installed inside a datacenter to observe environmental influences to subordinate the energy costs and increase the energy efficiency.

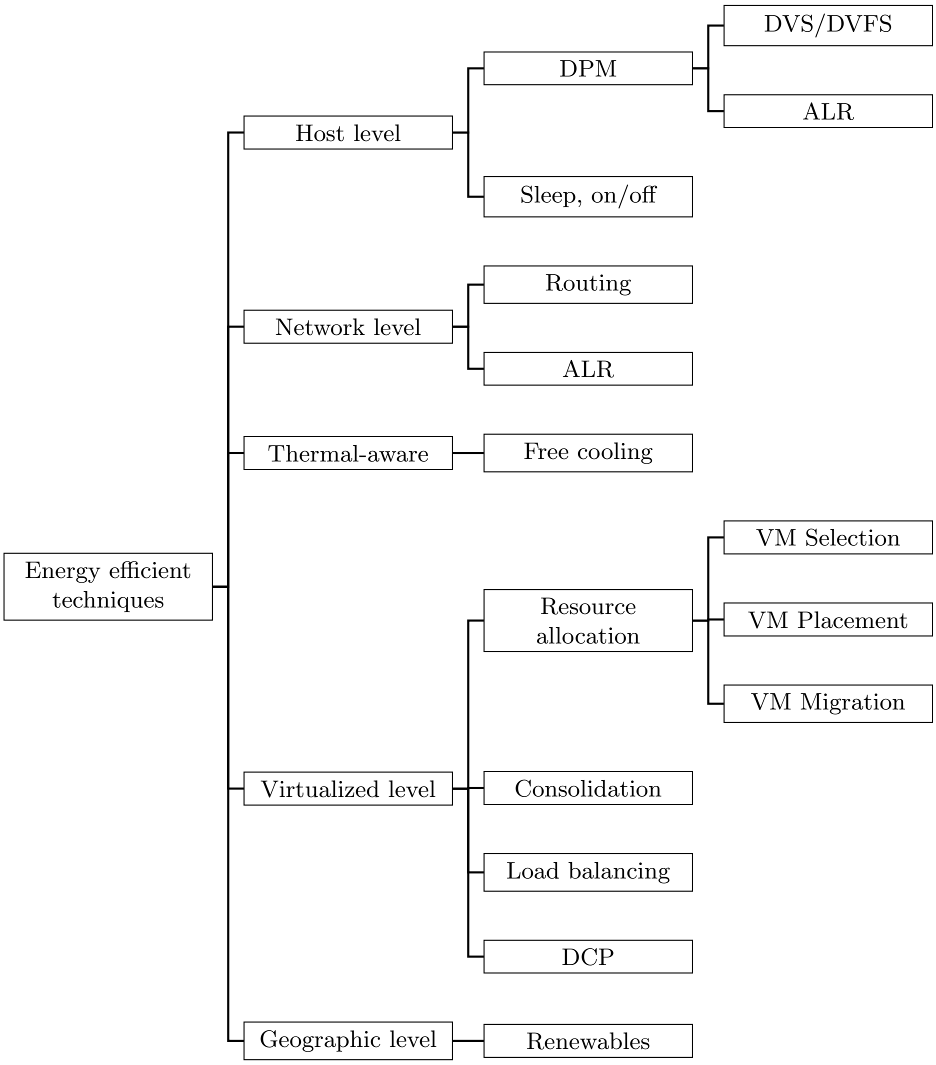

Fig. 4: Different approaches to make datacenters greener

Table 10 summarizes some techniques implemented in current cloud datacenters to make them greener and more energy efficient. There are several other ways to reduce the energy use of these large systems, which we will discuss in following subsections.

Virtualization: In terms of cloud computing and datacenters, virtualization [3] is considered as the most promising approach to save energy, which increases the resource utilization ratio. For different types of workload scheduled on physical and virtualized servers, the study in [38] has suggested that two virtualized servers would save 51.7% more energy as compared to two physical servers treating the same type of workload. If the electricity cost is 10p per kWh, a server that operates 24/7 will save £259.99 per year using server virtualization technique. There are many idle servers which waste a lot of energy, while do not process any important information [24]. Therefore, many researchers are looking for a novel comput-

Table 10: Current approaches to make datacenters energy efficient

Technique	Explanation	Benefits	Shortcomings
Virtualization	Dynamically provision resources [38], [101], [102]	Efficient energy saving	Widely used, VM live migration effect network performance
Server consolidation & encapsulating application	Reduces active servers by consolidating the workload of multiple servers [102]	Increases the utilization ratio of servers, reduce SLA violation ratio [102]	Categorization of servers failure of single consolidated server [103]
DCP	Adjust the available resources to current demand [100]	More energy efficient [96], [97]	DCP involved additional cost of switching on/off the resources that could violate customers SLA's [99]
Load Balancing	Balance the workload among different servers to fully and average utilization [6], [104]	Equal utilization	Tough job to implement in a heterogeneous platform [103]
Scheduling & VMs Placement	Place VMs onto a suitable (most energy efficient) servers [105], [106]	Server and Communication system energy efficient	Planning and live migration SLA violation [91], [107]
Live migrations	Migrate VMs from over-utilized & under-utilized to more efficient servers	Less energy consumption	Service level of running application affected
Renewables	Migrate VMs to servers operated by renewable energy sources	More energy efficient and economical	Renewables are intermittent, involves migrations that cost extra energy

ing prototype to realize green computing infrastructure. Some centralized-based computing models based on virtualization including clouds, began to advance the quality, efficiency and high availability of IT resources. Although these approaches are abolishing the old "one server, one application" model and it becomes a trend that different VMs run on a server. Therefore, the security issue has become a barrier of virtualization in such open Internet platform [101].

In cloud computing the communication amongst consumers and service providers is dependant on brokerage that provides negotiation techniques. It also manages available system resources to understand the goals of both communicating parties. The work in [101] has focused on building brokerage and related services which results in growth of clouds. SLA also plays a major role in such activities. Bin Li [108] has analyzed current issues in SLA and has categorized necessary characteristics required in cloud SLAs. Computing as a utility is discussed and it also explores the SLA frameworks with contemplations for building cloud SLAs. It recommends cloud brokerage as a suitable option for providing management in dynamic cloud platforms and presents autonomic SLA (ASLAs) framework for utility purpose. The authors have compared Amazon EC_2 (public cloud), Eucalyptus (private cloud) and HTCondor (grid) [2] from which they find that public cloud are faster than a supercomputer for some specific applications with well-known requirements and performance.

Consolidation: In 2010, the energy consumption by datacenters was expected to be in the range of 1.1% to 1.5% of the global energy use and is likely to rise more in the near future. The goal of consolidation (server/task) is to increase the computing resources utilization and diminish energy consumption under workload independent QoS constraints [49]. Energy consumption of the physical machines is decreased by vigorously activating and disabling (power cycling) them to fulfill current resource demand (DCP). The technique proposed in [49] is scalable, distributed and enable service providers to offer elastic resource provisioning with minimum energy, OPEx and CO_2 emissions. The extraordinary heat produced from large processing at datacenters leads to a number of issues like reduced reliability, availability and hardware life time. The released heat must be dissipated properly to keep the system modules operating safely, avoiding system failures and crashes. To reduce the cooling cost, the thermal state of each server is continuously observed and VMs are reallocated to other servers (server consolidation [27]) in case if the server is overheated.

The consolidation problem can be divided into four sub-problems including (i) server under-load detection, (ii) server over-load detection, (iii) VMs selection for migration and (iv) placement of the selected VMs on servers. The Global Resource Manager (GRM) decides VM allocation and placement. The burden on the GRM is condensed as it only sorts VM placement nominated for migration [107]. Furthermore, the GRM can be duplicated to eliminate the single point of failure and making the system entirely decentralized & distributed. Experimental results have shown that dynamic VM consolidation can minimize 30% of the total energy consumed (in case of no consolidation), having a minor impact on system performance [107]. Advanced distributed VM placement algorithms, VM network topologies, exploiting VM resource usage patterns, thermal-aware dynamic VM consolidation, dynamic and heterogeneous SLAs, power capping, VM consolidation algorithm analysis and replicated GRMs are the main issues that need researchers attention to further advance this area of interest.

In [102] an energy efficient approach is presented where the volume of resources assigned to a VM is adjusted based on application utilities, existing resources and energy costs. A min, max and share parameter concept is presented. Many proposed techniques like [109], [41] analyzed server utilizations which can result in a consolidation plan to virtualize servers. However, all these fails to yield benefit of the min, max and shares parameters. Although these three parameters are valuable only under high load circumstances as shown in [102] and it is of worth that high load circumstances will occur frequently in recent virtualized datacenters.

Dynamic VM consolidation lets cloud service providers to improve resource usability and minimize energy consumption. Each server is continuously monitored, and an appropriate action is taken, if certain threshold value (lower/upper) is reached. For example, if a server is underutilized, the current accommodated demand on it can be migrated to another server, and it is switched off to save energy. Setting proper threshold values is difficult and wrong values might reduce the cloud energy efficiency and performance. Hence, in [107] as shown in Table 11, a new idea of estimating the server's threshold values adaptively is introduced that use statistical guesses based on VM resource usage. The upper and lower utilization threshold values are dynamically determined according to the server's workload. The model consists of a local manager, a global manager and a VM monitor. The VM monitor perform actual VMs resizing, VMs migration and also monitor changes in the energy states of the servers. For efficient VM placement all VMs are sorted in decreasing order of their utilization values and assign each VM to a server having a minimum growth of energy consumption. The article concludes that it is essential to reduce number of VM migrations which may originate certain SLA violations and performance degradation. The proposed methodology beats other strategies (migration-aware) in total number of SLA violations, which are less than 1% and the amount of VMs migrated still providing comparable energy consumption.

Unsuitability between servers design and customer requirements in datacenters clues to issues like reduced load-balancing and large amount of energy consumption. In [104], the authors have described a power-aware load-balancing policy based on adaptive VM migration technique. If the current demand is greater or lesser than upper & lower utilization threshold values respectively, some VMs are migrated to increase resource utilization and thus minimize total energy consumption. The authors have also reduced the total number of migrations using minimum migration time strategy [49] that results in load balancing and meeting SLA requirements. The data collected from more than 5000 servers in 6 months shows that server's utilization hardly approaches 100%. Similarly, idle servers still consumes 70% of their peak energy. So underutilized servers are enormously fruitless from an energy consumption point of view.

Table 11: Basic architecture proposed in [110]

GM (master node) gathers information from the LM to maintain the total sight of resource utilization and issues commands to optimize VM placement	Global Manager (GM)
LM (each node) module of VMM continuously monitor utilization of a node, resize VMs according to the current resource requirements and decide VM migrations	Local Manager (LM) VM Monitor (VMM)
Virtualization	VM | VM | VM | VM Physical Node

VM Placement: How to place VM requests into available servers concerning to power intake, has turn out to be a crucial research problem [74], [111]. The placement problem is NP-hard and a number of heuristics like first fit, fill-up etc, are proposed to solve it [30], [105] as a bin-packing problem. Pietri et al. [70] has conducted a systematic review of VM placement techniques in clouds, however, energy efficiency is not discussed. EnaCloud [103] has allowed VMs placement onto servers in an energy efficient way. To save energy, VM encapsulates application that supports live migration to decrease the amount of active servers. The tactic is implemented as an advantageous module in internet-Oriented Virtual Computing (iVIC) platform. The results have shown that a single desktop machine having 2 cores idle, consumes 85W, and the consumption is even double when the machine is fully utilized. In distributing systems like clusters, tasks scheduling focus on how to allocate independent tasks to balance the workload amongst different servers to maximize throughput. In clouds, the placement problem is more challenging due to heterogeneity of resources, applications, and especially with the focus how to run an individual VM more efficiently.
There are a number of approaches [25] to balance the load amongst different servers, however these approaches are not suitable for cloud systems. Different techniques like DVS, DVFS, switching off display, sleep mode etc. are only suitable for a single server. They cannot be used to attain enough energy savings in HPC systems like grids and clouds datacenters because energy saved by decreasing the processor speed is less than switching off a server. Current approaches have implemented workload-aware (server/vm) consolidation, but still in some situations the associated SLA is violated [103]. As consolidation can be more expensive for certain workload, efficient placement policies are more cost-energy efficient. In cluster and grids, a server efficiency factor like the one used by [112], might be of significant importance to run the task on most efficient server. However, in virtualized platform, as the efficiency of a server is dependent on how many VMs are running on it (the baseline idle power in divided amongst all running VMs), new metrics like the one proposed in [45] are sill required.

Dynamic Capacity Planing (DCP): *"Capacity planning is a process through which the procurement of IT resources, infrastructure and services are planned over a specific period of time. It is an IT management practice to predict and forecast the future requirements of an enterprise IT environment and its associated essential entities/services/components"*[14]. The purpose of the DCP is to plan so well that new resources are added dynamically (just in time) to meet the expected resource demand but not so early that resources go unused for longer periods. The successful capacity planning and management is one that makes the trade-offs between the present and the future that overall prove to be the most cost-efficient[15]. The question that why would the service providers switch off a host, is also studied in the liter-

[14] https://www.youtube.com/watch?v=WivORAcBhV8
[15] http://searchenterprisewan.techtarget.com/definition/capacity-planning

ature [113]. As discussed before, an idle host still consumes approximately half the amount of energy consumed when running at peak load. Most datacenters are utilized in the range of 20 to 40% [114], which means there are a lot of hosts idling (not being used, but consuming energy), that can be switched off to save energy. However, what type of hosts can be switched off [96], because it is impractical to switch off servers (like web servers) that should be running 24 hours a day, 7 days a week.
Zhang et al. [96] proposed a framework for DCP and analyzed it with extensive simulations using real workload traces from Googles compute clusters. Their framework consist of five different modules – (i) a scheduler which assign tasks to active hosts, (ii) a monitoring module which is responsible for collecting CPU and memory usage statistics, (iii) a prediction module which estimate the future resource usage, (iv) a controller to control the resources considering the reconfiguration costs and (v) a capacity provisioning module which identifies which hosts should be switched off or switched on. They suggest that DCP could save significant amount of energy and hence cost, while maintaining an acceptable average scheduling delay (the performance objective in term of SLAs). However, their findings are only limited to homogeneous hosts with identical resource capacities. In [113], it is claimed that switch off technique can save 8% energy at a cluster operated by Cornell University by switching off 16% servers over a period of six months. Zhang et al. extends their own work in [96] with HARMONY [97], a heterogeneity-aware framework that dynamically adjusts the number of hosts to strike a balance between energy savings and scheduling delay (SLA), while considering the hosts reconfiguration cost.

Migration: Greater performance, fault tolerance and enriched resource manageability (server or task consolidation) are some of the benefits of virtualization technology. VM live migration in clouds and task or process migration in clusters lets the current demand to be consolidated on fewer servers with small service downtime. Service levels of active applications are negatively affected during a migration, therefore, there is a need to better understand its effects on system performance and energy efficiency. In [115], the authors have discussed such an important study area inside Xen VMs. In general, the migration overhead is tolerable but cannot be ignored in datacenters where service availability is managed by strict SLAs and the customer pay for their services. Beloglazov et al. [110] have discussed a distributed resource allocation & management policy with focus on live VM migration approach. To avoid performance degradation VMs are migrated from overloaded servers; and to increase the resource utilization and decrease the amount of energy consumed, VMs are migrated from under-loaded servers to switched them off. Which, when and where to migrate VMs are the main points that have distinguished this work from earlier techniques.
In [116] authors have presented a bi-phase optimization technique that decreases energy consumption by decreasing the total number of VM migrations while maintaining QoS to its best level. In the first phase i.e. VM selection; determine the list of all VMs to migrate from over-utilized & under-utilized servers and add them to the VMs pool. In the second phase i.e. VM placement; all VMs available in the pool are placed on suitable servers subject to particular heuristic functions. The problem, to find which server is over utilized and which one is underutilized is solved by defining some lower and upper utilization threshold values. As discussed before, estimating these threshold values are difficult and are not guaranteed to produce correct measurement if the workload is more diverse and/or is memory intensive. When migrating VMs amongst different servers, the migration energy overhead [115] must also be kept in view. As discussed earlier, migration would be more costly if the VM (in migration) got terminated during or just after the migration process is completed. A VM migration would be more economical and energy efficient if the migration is performed to a more energy efficient target server, enabling the VM to recover back its migration cost and runs subsequently more efficiently on the target server [45].
Table 12 summarizes different approaches (discussed in this survey), based on system characteristics (clusters, grids, clouds) and the approach (scheduling, virtualization, consolidation etc.) used for energy efficiency. It will help the research community to find the gap for further research, that still exists to make HPC systems and large scale cloud datacenters more energy efficient.

Table 12: Techniques for energy efficient computing and cloud datacenters

WORK		PLT			SCD		DCP	DPM	VRT		CNSL		DVS	LB	NET	STR	COL	RES	GHGs
	Cluster	Grid	Cloud	System	High	App			Server	Net	Server	Task							
[35], [92], [101]			✓		✓				✓	✓	✓		✓	✓					
[6], [36], [109]	✓			✓										✓					
[103], [116], [117]			✓		✓				✓	✓	✓								
[49], [91], [102]			✓	✓	✓														
[21], [72], [67], [75], [76], [73]	✓			✓		✓									✓				
[31], [32], [38], [57]	✓	✓		✓									✓		✓				
[20], [61]		✓	✓						✓		✓				✓				
[38], [57], [60]						✓									✓				
[2], [3], [30], [64]	✓	✓	✓		✓				✓					✓					
[27], [29]	✓	✓		✓								✓							
[21], [51], [50]	✓	✓		✓							✓			✓	✓			✓	
[53], [98]											✓						✓	✓	✓
[47], [52], [54]													✓						
[79], [69], [80], [63], [83], [84], [85], [86]																✓			
[82], [87]	✓	✓		✓	✓					✓		✓							✓
[95], [101]	✓	✓				✓							✓						✓
[107], [104]		✓	✓								✓	✓		✓					
[106], [115], [110], [118], [33], [119], [120]			✓		✓				✓		✓	✓				✓			
[14], [15], [17], [19]			✓		✓				✓						✓		✓		✓
[34], [9]	✓			✓									✓						
[30], [74], [121]			✓		✓						✓								
[41], [42]			✓		✓				✓		✓			✓					
[97], [122], [100]	✓		✓		✓		✓		✓		✓								
[99], [96]			✓		✓		✓	✓	✓		✓								✓

Acronym used in the table are as given below:
LB – Load Balancing, DVS – Dynamic Voltage & Frequency Scaling, PLT – Platform, SCD – Scheduling approach, DPM – Dynamic Power Management
VRT – Virtualization, CNSL – Consolidation technique using power cycling, STR – Storage system, COL – Cooling system, NET – Networks
RES – Renewable Energy Sources, GHGs – Greenhouse Gases, DCP – Dynamic Capacity Planing

7 Green Compilers and Applications

Apart from all these techniques, applications and compilers can also be designed in such a way that they can run with minimum energy consumption. Programming techniques like active polling and waiting loop[16], frequently wake up the CPU and could waste significant amount of energy. The authors in [25], have discussed several examples of applications that wake up the CPU hundreds times/second unnecessarily. Similarly, Universal Serial Bus (USB) adapter also takes time to initialize and cost more energy, that can be avoided in HPC systems where they are used hardly. In [67] an energy efficient compiler is proposed which executes instructions with minimum energy consumption. As software drives the bare hardware; hence design and development phase decisions will have momentous control on energy consumption of the server [123]. The proposed solution [67] have focused on power management measures for software level and their utilization in scheduler and compiler. A hardware independent Distributed Green Compiler (DGC) is presented that distributes software source code over a grid, redesigns binary code after applying green tactics which give green recommendation to programmer for energy savings. Performance evaluation of DGC showed that it preserves 40% energy clock cycles. Similarly, in clouds, similar applications would have different runtimes due to the performance variations (heterogeneity of resource and instances) and would cause different energy consumption as; Zhang et al. demonstrated in [122]. Therefore, it is of worth to know different applications (workloads) and their energy consumption, to schedule and run them more efficiently. Piraghaj et al. [68] have briefly discussed application level energy efficiency techniques for PaaS.

8 Metrics used to measure Energy Efficiency

Energy efficiency metrics are used to analyse and conclude if a datacenter can be boosted up before a new energy efficient datacenter is desirable. A typical datacenter consumes substantial volume of energy and results in enormous amounts of CO_2 emissions. Additionally, the humidity in a datacenter can origin hardware failures and increases cooling costs. A number of metrics like Power Usage Effectiveness (PUE) and Data Center Efficiency (DCE) are proposed in the literature [118]. Figure 5 summarizes a few of the most widely used metrics to measure datacenter efficiency. Methods like (i) reduce datacenter temperature [33], [119], (ii) increase server utilization [120], [124] and (iii) decrease the energy consumption of computational resources [125], [117] have their contributions towards green datacenters, however these articles lacks measurement for green quality i.e. how much energy is consumed, how much useful work done and how much CO_2 is produced etc.
PUE and DCE will facilitate service providers to guess the energy requirements and energy usage of their datacenter equipment. This will also enable them to match their obtained outcomes with outcomes obtained from other datacenters, which will help them to easily decide, are energy efficiency improvements needed in their datacenters or not [126]. PUE was given proper importance in the literature, however, DCE is not enough successful to measure energy efficiency. PUE is given by equation 2:

$$PUE = \frac{P_{IT}}{P_{total}} \tag{2}$$

Where P_{total} is total facility power measured through a utility meter and P_{IT} is the energy consumed by IT equipment in a datacenter. DCE actually shows the percentage of IT equipment energy efficiency, which is given by following equation 3:

$$DCE = \frac{1}{PUE} = \frac{P_{IT}}{P_{total}} \tag{3}$$

In [127], a new metric Data Center Infrastructure Efficiency (DCiE) is proposed, which is given by equation 4 below:

$$DCiE = \frac{1}{PUE} = \frac{P_{IT}}{P_{total}} * 100\% \tag{4}$$

[16] https://software.intel.com/en-us/articles/benefitting-power-and-performance-sleep-loops

The above equation 2 shows that larger PUE values represent less efficiency and vice versa. Therefore, PUE benchmarks the amount of energy deployed usefully (used by IT equipment) and the amount of energy wasted (used by other facilities) in datacenters.

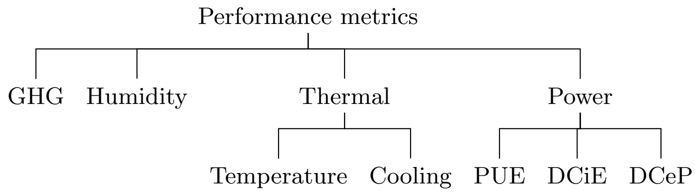

Fig. 5: Performance metrics

Table 13 shows the PUE values based on several experiments which were performed in a small datacenter [128]. The authors demonstrated that overall performance of the datacenter is very poor in terms of energy efficiency which have a PUE of 3.2. The study in [128] also suggested that the main reason of this inefficiency is that; the datacenter contains 150 racks and 85% of these racks are underutilized which consume significant amount of energy without doing any useful work.

Table 13: PUE efficiency values [128]

PUE	DCE (%)	Level of efficiency
3.0	33	Very inefficient
2.5	40	Inefficient
2.0	50	Average
1.5	67	efficient
1.2	83	very efficient
1.1	87	standard

These energy efficiency metrics are energetic tools for service providers to use when judging their services performance and deciding which resource should be considered for enhancement in terms of energy efficiency. Unfortunately, PUE is just a measure of how much energy is consumed by the compute equipment in a datacenter; it does not describes how the energy is used i.e. how much the equipment are energy efficient[17]. The energy sources powering a datacenter, the amount of useful work done and the underlying network used to move data around are also important, as they govern the amount of energy consumed and CO_2 emitted. A better Data Center Productivity (DCP) metric has been identified by TGG [118] as $DCeP$ which is given by following equation 5.

$$DCeP = \frac{W}{E_{use-dc}} \tag{5}$$

Where W is the useful amount of work done and E_{use-dc} is the total amount of energy consumed by datacenter to produce W. The key problem with this efficiency metric is calculating an accurate measure of the quantity in numerator [118]. A green datacenter must maintain the requirements of achieving good computational performance and hence cost w.r.t customers and minimum energy (more profit) w.r.t providers. Other metrics like performance per watt (PPW) and performance to power ratio (PPR) are used by the well known

[17] http://www.datacenterknowledge.com/archives/2011/11/15/pue-is-dead-the-case-for-performance-per-watt/

GREEN500[18] List to rank energy efficient supercomputers worldwide and SPECpower[19] benchmarks organization to estimate the server energy efficiency, respectively. Similar performance based metrics are still needed that mirrors how datacenters performance will perform, after the implementation of green approaches.

9 Research Challenges

System level approaches: Current approaches proposed in the literature like [96], [99], [9] expects homogeneity of the physical resources in a datacenter. Only a few researchers like [97], [122] have addressed solutions for heterogeneous datacenters. Furthermore, most of these methodologies have considered the server processor as the only resource that contributes greater to the energy consumption. Other important system components such as memory, fan and disks that also consume significant energy, are ignored. Runtime energy reduction techniques can decrease the energy consumption in large systems to some extent. However, the system level techniques like DVFS and ALR are not able to save more energy due to the cost of running an idle server [68]. Additionally, in systems with hard deadlines, performance [129] and QoS is affected with running processors at lower frequencies [47]. Hence, there is a need for other higher level resource management methods for energy savings in HPC and datacenters.

Virtualization: A simple way to reduce the hardware cost is to use virtualization technology (see Table 14). Virtualization can be useful to numerous resource types i.e. hardware or software including network links, software resources and storage modules etc. A classic illustration of virtualization is to share servers that reduces hardware cost, improves energy management, reduces cooling cost and carbon footprints in datacenters [34]. Due to sharing a single server, the security issue has become a barrier of virtualization in open Internet platform. The authors in [101] have highlighted several crucial security issues in green cloud computing platform, and suggest a secure virtualization model for energy efficiency. Secondly, as virtualization allows multiple VMs of different capacities to share a host resources (co-location or sibling VMs) that comes at a price of contention for the available resources – which leads to high variation in performance and hence cost [100].

Consolidation: Resource allocation & placement, VM selection, consolidation with migration and replacement algorithms have greater research potentials. In server consolidation, which VM, when and where (which server) to migrate the VM, are the basic questions that needs more research. These procedures are strongly dependent on (i) the system performance in case if servers are overloaded (SLA violation may occur) and (ii) on energy efficiency if several servers are under loaded (the demand can be consolidated on fewer servers to switch-off some servers to save idle energy consumption). Another important issue is the VM performance that could lead to more cost or SLA violation; and needs to be considered when a VM is migrated from one server to another target server [129].
Network consolidation [121] could save considerable energy, however they are habitually treated as too expensive in large systems where the objective is to deliver extreme throughput and QoS at slightest latency. Similarly, consolidating traffic and switching off network devices have effects on network performance, QoS, latency, throughput, packets overcrowding and the underlying communication protocols [21]. There are still certain gaps, such as when to switch-on the network switches back (sleep state delay) and the size of the queue (used to hold the requests when the device is in sleep state), that are keys to meet the performance-cost-energy model.

DCP: Effective capacity planning and management for cloud allows the service providers to predict the system performance (to avoid SLAs with customers) and enables them to

[18] https://www.top500.org/green500/
[19] https://www.spec.org/power_ssj2008/

take important business decisions for cost saving purposes at capital (how much hardware resources to buy) and operational (energy and licences costs) levels. Cloud capacity management is claimed as an underrated problem that still need to address[20]. To be cost effective, it is very important to study different types of cloud workload and the future demand (workload estimation models) to avoid under and over provisioning of resources [130]. Also, DCP might not be feasible in large production clouds like Google[21], involves hosts reconfiguration cost and could affect the hosts lifetime [100].

Migration: As discussed in Section 6.5, server consolidation involves migration that have energy cost. It is possible that the efforts for migration are wasted if a VM terminates during or just after the migration process is completed. A migration would be more energy efficient if it can recoup back its migration cost and subsequently run on a more energy efficient target server [45]. Knowing which VMs of runtimes, would recoup back their migration energy cost is an interesting research problem. An analysis of real workload traces from public cloud providers [22] can be more helpful is such energy efficient migration decisions. Currently majority of cloud researchers rely on parallel systems workload, which is different from cloud workload. Public cloud data [22], would be more realistic and accurate to study the performance of different scheduling and migration techniques.

Performance: Several performance studies like [129], [131] conducted on EC2 demonstrates that VM runtimes over different CPU models backing a single instance class will have consistent, i.e. largely predictable, performance variations with respect to the CPU model. The distribution of runtimes of application benchmarks within the VMs shows a lognormal distribution (multi-modal) with positive skewness - which describes the variation in performance. There are at least two benefits in understanding such performance variation for different CPU models in a cloud platform:

1. knowing the relationship between CPU model and runtime means decisions over scheduling can be made to optimize overall energy efficiency;
2. migration of a VM to a host which leads to an increase in runtime would increase costs for the customer, and as a consequence may violate SLAs.

In respect to (1), if a VM can be migrated to a better performing host, for its workloads, better energy efficiencies may be achieved. But, in respect to (2) if migration would make performance worse, not migrating may be better. Hence it is important to study the variation in energy efficiencies in heterogeneous clusters with different hosts having the abilities to run different kinds of workload. Ideally, a cloud provider should look to the best trade-off between performance, cost and energy requirements. However, a realistic and in practice approach for resource provisioning and consolidation is needed such that the energy requirements to run the service and the expected performance/price goals can be met.

Renewables: The current trends of cloud service providers towards using renewable energy sources that may operate intermittently [132], and hence necessitate falling back to the energy grid, also implies a need for consolidation policies to be able to effectively switch between the available energy sources, as well as to reduce the replacement cycle of renewable capture and storage equipment. With 640 datacenter outages in the UK alone in 2015 and outages expected to be more common in near future [133], there is a need at least for proper capacity planning, consolidation of workloads onto servers powered by renewables, and migration of workloads when it is most energy, and therefore cost, efficient, to safeguard supply and reduce the drain on renewable generation and storage equipment.

[20] http://enterprise.netscout.com/cio_brief/capacity-management-cloud-underrated-problem-you-need-address

[21] https://www.youtube.com/watch?v=7MwxA4Fj2l4

Real-time Services: As Cloud computing becomes growing for Anything as a Service (XaaS) model, modern real-time cloud services including gaming applications, flight-control systems and image processing are also becoming popular and accessible over the cloud [134]. Real-Time Services (RTSs) are those whose precision depends not only on the logical results but also on the time in which these results are made [47]. RTSs need large volume of computational resources to scale user requests and satisfy timely deadlines simultaneously. An usual RTS involves numerous Real-Time Applications (RTAs) that are further divided into subtasks. As long as a group of applications or tasks for a given RTS meet all their deadlines, the service achieves the QoS settled with the customers. Cloud computing model can provide this scalability within the timing constraints to these RTSs but is more challenging in case if energy efficient scheduling and consolidation approaches are considered[22]. For example, as discussed earlier, scaling down the processor frequency or even migrations can increase application runtime, which can result in deadline miss. Hence, it is very important to have a look on application types and their deadlines before the implementation of green techniques to save energy.

Datacenter Costs: Total Cost of Ownership (TCO) of a typical datacenter (as shown in Table 14) includes cost of IT equipment, infrastructure, energy and other operating cost like management and maintenance. Reducing TCO including CapEx (capital expenses) and OpEx (operational expenses) is an active research issue. Infrastructure costs includes cooling, power distribution, backup power, and power conditioning costs and needs to be considered for reduction. Other momentous costs like architectural and engineering, land and property, IT build-out rack, electric wiring, routers, bridges, switches, wide area networking & communications, electricity, safety, operations & maintenance must also be considered [5] for possible decrease.

Table 14: Total Cost of a datacenter ownership

Infrastructure	ICT equipment	Energy	Operating
31%	47%	10%	12%

Datacenter Simulators: Due to time and the amount of resources available to the researchers in academia, it is often infeasible for them to conduct extensive and repeatable experiments in order to validate and verify their research findings [107], [111], [49]. Therefore majority of the researchers use different kinds of simulator and then, based on several plausible assumptions, generalize their findings [135]. However, simulators are not guaranteed to produce accurate and verifiable results due to the lake of mathematical proofs and real testbeds. Therefore, concrete studies like the one presented in [136], are needed to facilitate the researchers and even service providers to validate their hypotheses before they put it on the cloud.

10 Concluding Remarks

Large systems like clusters, grids and datacenters energy costs can be divided into two essential types; (i) energy consumption of ICT equipment and (ii) infrastructure level energy consumption like servers cooling etc. A recent study [17] shows that in 2014, the US datacenters almost consumed 70 billion kWh of energy that is 1.8% of the total consumption and is expected to reach 73 billion kWh by 2020. Similarly, the current share of ICT equipment to global GHG emissions is around 1.6% and it is estimated to be around 2% by 2020 [18]. It has been reported that a typical datacenter energy consumption accounts for more than

[22] https://aws.amazon.com/it/hpc/

12% of monthly operational expenditures. For large industries like Google and Amazon, a 3% reduction in energy cost can translate into over a million dollars in cost savings [96]. As, the typical datacenter energy consumption has increased significantly since 2006 [137] and is expected to increase more in near future, the survey provides a detailed comparison and description of the energy efficient techniques in three broad categories of distributed systems namely clusters, grids, and cloud datacenters. In this review article we studied the energy efficiency of these systems at three levels i.e. (i) hardware, (ii) resource management and (iii) applications.

We found that for certain kinds of workload, the system level efficiency techniques might increase cluster energy efficiency with some performance loss, however in grids, scheduling and efficient resource allocation are more efficient than system level methods. Similarly, in virtualized clouds, efficient scheduling and resource allocation is more economical than consolidation with migration technique, for certain types of workload (application). From a datacenter perspective, the two major points of energy efficient techniques are: (i) reduce the energy consumption of ICT equipment and (ii) minimize CO_2 emissions for environmental sustainability. To meet the challenges of today's elastic cloud systems and unpredictable customers workload, efficient scheduling techniques are still required as this would be more economical and energy efficient to implement as compared to server consolidation and VM migration techniques. The survey will help the readers to analyse the gap between what is already available in existing systems and what is still required, so that outstanding research issues can be identified.

Acknowledgement

This work is supported by Department of Computer Science, University of Surrey, UK and Abdul Wali Khan University, Mardan, Pakistan. The authors are thankful to Dr. Joseph Chrol-Cannon and Santosh Tirunagari from Department of Computer Science, University of Surrey, UK for their review, valuable comments, and suggestions for technical improvement of this work in hand.

References

1. Jack J Dongarra, Hans W Meuer, and Erich Strohmaier. Top500 supercomputer sites, 1994.
2. Hameed Hussain, Saif Ur Rehman Malik, Abdul Hameed, Samee Ullah Khan, Gage Bickler, Nasro Min-Allah, Muhammad Bilal Qureshi, Limin Zhang, Wang Yongji, Nasir Ghani, et al. A survey on resource allocation in high performance distributed computing systems. *Parallel Computing*, 39(11):709–736, 2013.
3. Rajkumar Buyya, Chee Shin Yeo, and Srikumar Venugopal. Market-oriented cloud computing: Vision, hype, and reality for delivering it services as computing utilities. In *High Performance Computing and Communications, 2008. HPCC'08. 10th IEEE International Conference on*, pages 5–13. Ieee, 2008.
4. Global warming: Data centres to consume three times as much energy in next decade, experts warn. http://www.independent.co.uk/, 2016. [Online; accessed 3-Aug-16].
5. Jonathan G Koomey et al. Estimating total power consumption by servers in the us and the world, 2007.
6. Anton Beloglazov, Jemal Abawajy, and Rajkumar Buyya. Energy-aware resource allocation heuristics for efficient management of data centers for cloud computing. *Future generation computer systems*, 28(5):755–768, 2012.
7. Christopher Dabrowski. Reliability in grid computing systems. *Concurrency and Computation: Practice and Experience*, 21(8):927–959, 2009.
8. Yanmin Zhu and Lionel M Ni. A survey on grid scheduling systems. *Department of Computer Science, Hong Kong University of science and Technology*, page 32, 2003.
9. Anton Beloglazov, Rajkumar Buyya, Young Choon Lee, Albert Zomaya, et al. A taxonomy and survey of energy-efficient data centers and cloud computing systems. *Advances in computers*, 82(2):47–111, 2011.
10. INCITE. http://www.doeleadershipcomputing.org, US Department of Energy. [Online; accessed 19-August-16].

11. Peter H Beckman. Building the teragrid. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 363(1833):1715–1728, 2005.
12. Randy Bias. Grid, cloud, hpc ... what's the diff? http://cloudscaling.com/blog/cloud-computing/grid-cloud-hpc-whats-the-diff/. [Online; accessed 19-August-16].
13. Jason Cope, Kamil Iskra, Dries Kimpe, and Robert Ross. Grids and hpc: not as different as you might think?, 2010.
14. Tarandeep Kaur and Inderveer Chana. Energy efficiency techniques in cloud computing: A survey and taxonomy. *ACM Computing Surveys (CSUR)*, 48(2):22, 2015.
15. Robert Basmadjian, Hermann De Meer, Ricardo Lent, and Giovanni Giuliani. Cloud computing and its interest in saving energy: the use case of a private cloud. *Journal of Cloud Computing*, 1(1):1–25, 2012.
16. NRDC. Scaling up energy efficiency across the data center industry: Evaluating key drivers and barriers. *NRDC Issue Brief, IP:14-08-A*, pages 1–35, 2014.
17. A Shehabi, SJ Smith, N Horner, I Azevedo, R Brown, J Koomey, E Masanet, D Sartor, M Herrlin, and W Lintner. United states data center energy usage report. *Lawrence Berkeley National Laboratory, Berkeley, California. LBNL-1005775 Page*, 4, 2016.
18. Accenture. Smarter2030: Ict solutions for 21st century challenges online:. http://smarter2030.gesi.org/downloads/Full_report2.pdf, 2015.
19. James Hamilton. Cooperative expendable micro-slice servers (cems): low cost, low power servers for internet-scale services. In *Conference on Innovative Data Systems Research (CIDR09)(January 2009)*. Citeseer, 2009.
20. Giorgio Luigi Valentini, Walter Lassonde, Samee Ullah Khan, Nasro Min-Allah, Sajjad A Madani, Juan Li, Limin Zhang, Lizhe Wang, Nasir Ghani, Joanna Kolodziej, et al. An overview of energy efficiency techniques in cluster computing systems. *Cluster Computing*, 16(1):3–15, 2013.
21. Junaid Shuja, Sajjad A Madani, Kashif Bilal, Khizar Hayat, Samee U Khan, and Shahzad Sarwar. Energy-efficient data centers. *Computing*, 94(12):973–994, 2012.
22. Charles Reiss, John Wilkes, and Joseph L Hellerstein. Google cluster-usage traces: format+ schema. *Google Inc., Mountain View, CA, USA, Technical Report*, 2011.
23. Anton Beloglazov and Rajkumar Buyya. Energy efficient resource management in virtualized cloud data centers. In *Proceedings of the 2010 10th IEEE/ACM international conference on cluster, cloud and grid computing*, pages 826–831. IEEE Computer Society, 2010.
24. NRDC. America's data centers are wasting huge amounts of energy: critical action needed to save billions of dollars and kilowatts. *NRDC Issue Brief, IB:14-08-A*, pages 1–6, 2014.
25. Anne-Cecile Orgerie, Marcos Dias de Assuncao, and Laurent Lefevre. A survey on techniques for improving the energy efficiency of large-scale distributed systems. *ACM Computing Surveys (CSUR)*, 46(4):47, 2014.
26. Xiaobo Fan, Wolf-Dietrich Weber, and Luiz Andre Barroso. Power provisioning for a warehouse-sized computer. In *ACM SIGARCH Computer Architecture News*, volume 35, pages 13–23. ACM, 2007.
27. Tiago C. Ferreto, Marco A S Netto, Rodrigo N. Calheiros, and César A F De Rose. Server consolidation with migration control for virtualized data centers. *Future Generation Computer Systems*, 27(8):1027–1034, 2011.
28. Mueen Uddin, Muhammad Talha, Azizah Abdul Rahman, Asadullah Shah, Jameel Ahmed, and Jamshed Memon. Green information technology (it) framework for energy efficient data centers using virtualization. *International Journal of Physical Sciences*, 7(13):2052–2065, 2012.
29. Mueen Uddin and Azizah Abdul Rahman. Server consolidation: An approach to make data centers energy efficient and green. *arXiv preprint arXiv:1010.5037*, 2010.
30. Xin Li, Zhuzhong Qian, Sanglu Lu, and Jie Wu. Energy efficient virtual machine placement algorithm with balanced and improved resource utilization in a data center. *Mathematical and Computer Modelling*, 58(5):1222–1235, 2013.
31. Ching-Hsien Hsu, Kenn D Slagter, Shih-Chang Chen, and Yeh-Ching Chung. Optimizing energy consumption with task consolidation in clouds. *Information Sciences*, 258:452–462, 2014.
32. Sujit Tilak and Dipti Patil. A survey of various scheduling algorithms in cloud environment. *International Journal of Engineering Inventions*, 1(2):36–39, 2012.
33. Qinghui Tang, Sandeep Kumar S Gupta, and Georgios Varsamopoulos. Energy-efficient thermal-aware task scheduling for homogeneous high-performance computing data centers: A cyber-physical approach. *Parallel and Distributed Systems, IEEE Transactions on*, 19(11):1458–1472, 2008.
34. Aruna Prem Bianzino, Claude Chaudet, Dario Rossi, and Jean-Louis Rougier. A survey of green networking research. *Communications Surveys & Tutorials, IEEE*, 14(1):3–20, 2012.

35. C Yaashuwanth and Dr R Ramesh. A new scheduling algorithms for real time tasks. *arXiv preprint arXiv:0912.0606*, 2009.
36. Jagbeer Singh, Bichitrananda Patra, and Satyendra Prasad Singh. An algorithm to reduce the time complexity of earliest deadline first scheduling algorithm in real-time system. *International Journal of Advanced Computer Science & Applications*, 1(2):31–37.
37. Madhu Gupta and Sushil Singh. Dynamic ethernet link shutdown for energy conservation on ethernet links. In *Communications, 2007. ICC'07. IEEE International Conference on*, pages 6156–6161. IEEE, 2007.
38. Lu Liu, Osama Masfary, and Nick Antonopoulos. Energy performance assessment of virtualization technologies using small environmental monitoring sensors. *Sensors*, 12(5):6610–6628, 2012.
39. Raffaele Bolla, Roberto Bruschi, Franco Davoli, and Chiara Lombardo. Fine-grained energy-efficient consolidation in sdn networks and devices. *Network and Service Management, IEEE Transactions on*, 12(2):132–145, 2015.
40. Alexandra Fedorova, Margo Seltzer, and Michael D Smith. Cache-fair thread scheduling for multicore processors. *Division of Engineering and Applied Sciences, Harvard University, Tech. Rep. TR-17-06*, 2006.
41. Gunjan Khanna, Kirk Beaty, Gautam Kar, and Andrzej Kochut. Application performance management in virtualized server environments. In *Network Operations and Management Symposium, 2006. NOMS 2006. 10th IEEE/IFIP*, pages 373–381. IEEE, 2006.
42. Ivan Rodero, Juan Jaramillo, Andres Quiroz, Manish Parashar, Francesc Guim, and Stephen Poole. Energy-efficient application-aware online provisioning for virtualized clouds and data centers. In *Green Computing Conference, 2010 International*, pages 31–45. IEEE, 2010.
43. Christoph Möbius, Waltenegus Dargie, and Alexander Schill. Power consumption estimation models for processors, virtual machines, and servers. *IEEE Transactions on Parallel and Distributed Systems*, 25(6):1600–1614, 2014.
44. Mascha Kurpicz, Anita Sobe, et al. How much does a vm cost? energy-proportional accounting in vm-based environments. In *2016 24th Euromicro International Conference on Parallel, Distributed, and Network-Based Processing (PDP)*, pages 651–658. IEEE, 2016.
45. Muhammad Zakarya and Lee Gillam. An energy aware cost recovery approach for virtual machine migration. In *Lecturer Notes in Computer Science*. GECON, 2016.
46. Ibrahim Alzamil, Karim Djemame, Django Armstrong, and Richard Kavanagh. Energy-aware profiling for cloud computing environments. *Electronic Notes in Theoretical Computer Science*, 318:91–108, 2015.
47. Nasro Min-Allah, Samee Ullah Khan, and Wang Yongji. Optimal task execution times for periodic tasks using nonlinear constrained optimization. *The Journal of Supercomputing*, 59(3):1120–1138, 2012.
48. Georges Da Costa, Marcos Dias De Assuncao, Jean-Patrick Gelas, Yiannis Georgiou, Laurent Lefèvre, Anne-Cécile Orgerie, Jean-Marc Pierson, Olivier Richard, and Amal Sayah. Multi-facet approach to reduce energy consumption in clouds and grids: the green-net framework. In *Proceedings of the 1st international conference on energy-efficient computing and networking*, pages 95–104. ACM, 2010.
49. Anton Beloglazov and Rajkumar Buyya. Optimal online deterministic algorithms and adaptive heuristics for energy and performance efficient dynamic consolidation of virtual machines in cloud data centers. *Concurrency and Computation: Practice and Experience*, 24(13):1397–1420, 2012.
50. Karthik Lakshmanan, Ragunathan Raj Rajkumar, and John P Lehoczky. Partitioned fixed-priority preemptive scheduling for multi-core processors. In *Real-Time Systems, 2009. ECRTS'09. 21st Euromicro Conference on*, pages 239–248. IEEE, 2009.
51. Shinpei Kato, Akira Takeda, and Nobuyuki Yamasaki. Global rate-monotonic scheduling with priority promotion, 2010.
52. Guolong Lin and Rajmohan Rajaraman. Approximation algorithms for multiprocessor scheduling under uncertainty. *Theory of Computing Systems*, 47(4):856–877, 2010.
53. Yoshifumi Manabe and Shigemi Aoyagi. A feasibility decision algorithm for rate monotonic scheduling of periodic real-time tasks. In *Real-Time Technology and Applications Symposium, 1995. Proceedings*, pages 212–218. IEEE, 1995.
54. Sanjoy Baruah and Joël Goossens. Deadline monotonic scheduling on uniform multiprocessors. In *Principles of Distributed Systems*, pages 89–104. Springer, 2008.
55. Moïse W Convolbo and Jerry Chou. Cost-aware dag scheduling algorithms for minimizing execution cost on cloud resources. *The Journal of Supercomputing*, 72(3):985–1012, 2016.

56. Nasro Min-Allah, Hameed Hussain, Samee Ullah Khan, and Albert Y Zomaya. Power efficient rate monotonic scheduling for multi-core systems. *Journal of Parallel and Distributed Computing*, 72(1):48–57, 2012.
57. Frances Yao, Alan Demers, and Scott Shenker. A scheduling model for reduced cpu energy. In *Foundations of Computer Science, 1995. Proceedings., 36th Annual Symposium on*, pages 374–382. IEEE, 1995.
58. Susanne Albers, Fabian Müller, and Swen Schmelzer. Speed scaling on parallel processors. *Algorithmica*, 68(2):404–425, 2014.
59. Susanne Albers and Antonios Antoniadis. Race to idle: new algorithms for speed scaling with a sleep state. *ACM Transactions on Algorithms (TALG)*, 10(2):9, 2014.
60. Xingyu Yang, Weiguo Zhang, Yong Zhang, and Weijun Xu. Optimal randomized algorithm for a generalized ski-rental with interest rate. *Information Processing Letters*, 112(13):548–551, 2012.
61. Wissam Chedid and Chansu Yu. Survey on power management techniques for energy efficient computer systems. *Laboratory Report. Mobile Computing Research Lab*, 2002.
62. Tom Guérout, Thierry Monteil, Georges Da Costa, Rodrigo Neves Calheiros, Rajkumar Buyya, and Mihai Alexandru. Energy-aware simulation with dvfs. *Simulation Modelling Practice and Theory*, 39:76–91, 2013.
63. Frédéric Giroire, Dorian Mazauric, Joanna Moulierac, and Brice Onfroy. Minimizing routing energy consumption: from theoretical to practical results. In *Green Computing and Communications (GreenCom), 2010 IEEE/ACM Int'l Conference on & Int'l Conference on Cyber, Physical and Social Computing (CPSCom)*, pages 252–259. IEEE, 2010.
64. Chamara Gunaratne, Ken Christensen, Bruce Nordman, and Stephen Suen. Reducing the energy consumption of ethernet with adaptive link rate (alr). *Computers, IEEE Transactions on*, 57(4):448–461, 2008.
65. Daniel Guimaraes do Lago, Edmundo RM Madeira, and Luiz Fernando Bittencourt. Power-aware virtual machine scheduling on clouds using active cooling control and dvfs. In *Proceedings of the 9th International Workshop on Middleware for Grids, Clouds and e-Science*, page 2. ACM, 2011.
66. Mohammad S Obaidat, Alagan Anpalagan, and Isaac Woungang. *Handbook of green information and communication systems*. Academic Press, 2012.
67. Faiza Fakhar, Barkha Javed, Raihan ur Rasool, Owais Malik, and Khurram Zulfiqar. Software level green computing for large scale systems. *Journal of Cloud Computing: Advances, Systems and Applications*, 1(1):1, 2012.
68. Sareh Fotuhi Piraghaj, Amir Vahid Dastjerdi, Rodrigo N Calheiros, and Rajkumar Buyya. A survey and taxonomy of energy efficient resource management techniques in platform as a service cloud. *Handbook of Research on End-to-End Cloud Computing Architecture Design*, page 410, 2016.
69. Baoke Zhang, Karthikeyan Sabhanatarajan, Ann Gordon-Ross, and Alan George. Real-time performance analysis of adaptive link rate. In *Local Computer Networks, 2008. LCN 2008. 33rd IEEE Conference on*, pages 282–288. IEEE, 2008.
70. Ilia Pietri and Rizos Sakellariou. Mapping virtual machines onto physical machines in cloud computing: A survey. *ACM Computing Surveys (CSUR)*, 49(3):49, 2016.
71. NRDC. Scaling up energy efficiency across the data center industry: Evaluating key drivers and barriers. *NRDC Issue Brief, IP:14-08-A*, pages 1–35, 2014.
72. Richard Brown et al. Report to congress on server and data center energy efficiency: Public law 109-431. *Lawrence Berkeley National Laboratory*, 2008.
73. Jonathan Koomey. Growth in data center electricity use 2005 to 2010. *A report by Analytical Press, completed at the request of The New York Times*, 9, 2011.
74. Norman Bobroff, Andrzej Kochut, and Kirk Beaty. Dynamic placement of virtual machines for managing sla violations. In *2007 10th IFIP/IEEE International Symposium on Integrated Network Management*, pages 119–128. IEEE, 2007.
75. Anne Cécile Orgerie, Laurent Lefèvre, and Jean Patrick Gelas. Save watts in your grid: Green strategies for energy-aware framework in large scale distributed systems. In *Parallel and Distributed Systems, 2008. ICPADS'08. 14th IEEE International Conference on*, pages 171–178. IEEE, 2008.
76. Anshul Gandhi, Varun Gupta, Mor Harchol-Balter, and Michael A Kozuch. Optimality analysis of energy-performance trade-off for server farm management. *Performance Evaluation*, 67(11):1155–1171, 2010.
77. Steven Lanzisera, Bruce Nordman, and Richard E Brown. Data network equipment energy use and savings potential in buildings. *Energy Efficiency*, 5(2):149–162, 2012.

78. Raffaele Bolla, Franco Davoli, Roberto Bruschi, Ken Christensen, Flavio Cucchietti, and Suresh Singh. The potential impact of green technologies in next-generation wireline networks: Is there room for energy saving optimization? *IEEE Communications Magazine*, 49(8):80–86, 2011.
79. Andreas Berl, Erol Gelenbe, Marco Di Girolamo, Giovanni Giuliani, Hermann De Meer, Minh Quan Dang, and Kostas Pentikousis. Energy-efficient cloud computing. *The computer journal*, 53(7):1045–1051, 2010.
80. Maruti Gupta and Suresh Singh. Using low-power modes for energy conservation in ethernet lans. In *INFOCOM*, volume 7, pages 2451–2455, 2007.
81. NRDC. America's data centers are wasting huge amounts of energy: critical action needed to save billions of dollars and kilowatts. *IB:14-08-A*, pages 1–6, 2014.
82. Brett Battles, Cathy Belleville, Susan Grabau, and Judith Maurier. Reducing data center power consumption through efficient storage. *Network Appliance, Inc*, 2007.
83. Ziliang Zong, Matt Briggs, Nick O'Connor, and Xiao Qin. An energy-efficient framework for large-scale parallel storage systems. In *Parallel and Distributed Processing Symposium, 2007. IPDPS 2007. IEEE International*, pages 1–7. IEEE, 2007.
84. Dong Li and Jun Wang. Eeraid: energy efficient redundant and inexpensive disk array. In *Proceedings of the 11th workshop on ACM SIGOPS European workshop*, page 29. ACM, 2004.
85. Qingbo Zhu, Francis M David, Christo F Devaraj, Zhenmin Li, Yuanyuan Zhou, and Pei Cao. Reducing energy consumption of disk storage using power-aware cache management. In *Software, IEE Proceedings-*, pages 118–118. IEEE, 2004.
86. Laszlo A. Belady. A study of replacement algorithms for a virtual-storage computer. *IBM Systems journal*, 5(2):78–101, 1966.
87. Mark W Storer, Kevin M Greenan, Ethan L Miller, and Kaladhar Voruganti. Pergamum: Replacing tape with energy efficient, reliable, disk-based archival storage. In *Proceedings of the 6th USENIX Conference on File and Storage Technologies*, page 1. USENIX Association, 2008.
88. Kester Li, Roger Kumpf, Paul Horton, and Thomas E Anderson. A quantitative analysis of disk drive power management in portable computers. In *USENIX winter*, pages 279–291, 1994.
89. Seung Woo Son and Mahmut Kandemir. Energy-aware data prefetching for multi-speed disks. In *Proceedings of the 3rd conference on Computing frontiers*, pages 105–114. ACM, 2006.
90. Sudhanva Gurumurthi, Anand Sivasubramaniam, Mahmut Kandemir, and Hubertus Franke. Drpm: dynamic speed control for power management in server class disks. In *Computer Architecture, 2003. Proceedings. 30th Annual International Symposium on*, pages 169–179. IEEE, 2003.
91. Liting Hu, Hai Jin, Xiaofei Liao, Xianjie Xiong, and Haikun Liu. Magnet: A novel scheduling policy for power reduction in cluster with virtual machines. In *2008 IEEE International Conference on Cluster Computing*, pages 13–22. IEEE, 2008.
92. Hady S Abdelsalam, Kurt Maly, Ravi Mukkamala, Mohammad Zubair, and David Kaminsky. Analysis of energy efficiency in clouds. In *Future Computing, Service Computation, Cognitive, Adaptive, Content, Patterns, 2009. COMPUTATIONWORLD'09. Computation World:*, pages 416–421. IEEE, 2009.
93. Young Choon Lee and Albert Y Zomaya. Energy efficient utilization of resources in cloud computing systems. *The Journal of Supercomputing*, 60(2):268–280, 2012.
94. https://www.youtube.com/watch?v=7MwxA4Fj2l4. [Online; accessed 3-Oct-15].
95. Luiz André Barroso and Urs Hölzle. The case for energy-proportional computing. *Computer*, 40(12):33–37, 2007.
96. Qi Zhang, Mohamed Faten Zhani, Shuo Zhang, Quanyan Zhu, Raouf Boutaba, and Joseph L Hellerstein. Dynamic energy-aware capacity provisioning for cloud computing environments. In *Proceedings of the 9th international conference on Autonomic computing*, pages 145–154. ACM, 2012.
97. Qi Zhang, Mohamed Faten Zhani, Raouf Boutaba, and Joseph L Hellerstein. Harmony: Dynamic heterogeneity-aware resource provisioning in the cloud. In *Distributed Computing Systems (ICDCS), 2013 IEEE 33rd International Conference on*, pages 510–519. IEEE, 2013.
98. Paul Anderson, Gaynor Backhouse, Daniel Curtis, Simon Redding, and David Wallom. Low carbon computing: a view to 2050 and beyond. *JISC Technology & Standards Watch (TechWatch)*, 2009.
99. Seung-Hwan Lim, Bikash Sharma, Byung Chul Tak, and Chita R Das. A dynamic energy management scheme for multi-tier data centers. In *Performance Analysis of Systems and Software (ISPASS), 2011 IEEE International Symposium on*, pages 257–266. IEEE, 2011.
100. Seung-Hwan Lim. *Managing performance and energy in large scale data centers*. PhD thesis, The Pennsylvania State University, 2012.

101. Jianxin Li, Bo Li, Tianyu Wo, Chunming Hu, Jinpeng Huai, Lu Liu, and KP Lam. Cyberguarder: A virtualization security assurance architecture for green cloud computing. *Future Generation Computer Systems*, 28(2):379–390, 2012.
102. Michael Cardosa, Madhukar R Korupolu, and Aameek Singh. Shares and utilities based power consolidation in virtualized server environments. In *Integrated Network Management, 2009. IM'09. IFIP/IEEE International Symposium on*, pages 327–334. IEEE, 2009.
103. Bo Li, Jianxin Li, Jinpeng Huai, Tianyu Wo, Qin Li, and Liang Zhong. Enacloud: An energy-saving application live placement approach for cloud computing environments. In *Cloud Computing, 2009. CLOUD'09. IEEE International Conference on*, pages 17–24. IEEE, 2009.
104. Khushbu Maurya and Richa Sinha. Energy conscious dynamic provisioning of virtual machines using adaptive migration thresholds in cloud data center. *Int. J. Comput. Sci. Mob. Comput*, 3(2):74–82, 2013.
105. Grant Wu, Maolin Tang, Yu-Chu Tian, and Wei Li. Energy-efficient virtual machine placement in data centers by genetic algorithm. In *Neural Information Processing*, pages 315–323. Springer, 2012.
106. Rajkumar Buyya, Anton Beloglazov, and Jemal Abawajy. Energy-efficient management of data center resources for cloud computing: a vision, architectural elements, and open challenges. *arXiv preprint arXiv:1006.0308*, 2010.
107. Anton Beloglazov and Rajkumar Buyya. Adaptive threshold-based approach for energy-efficient consolidation of virtual machines in cloud data centers. In *Proceedings of the 8th International Workshop on Middleware for Grids, Clouds and e-Science*, volume 4. ACM, 2010.
108. Bin Li, Lee Gillam, and John OLoughlin. Towards application-specific service level agreements: Experiments in clouds and grids. In *Cloud Computing*, pages 361–372. Springer, 2010.
109. Akshat Verma, Puneet Ahuja, and Anindya Neogi. pmapper: power and migration cost aware application placement in virtualized systems. In *Middleware 2008*, pages 243–264. Springer, 2008.
110. Anton Beloglazov and Rajkumar Buyya. Openstack neat: a framework for dynamic and energy-efficient consolidation of virtual machines in openstack clouds. *Concurrency and Computation: Practice and Experience*, 27(5):1310–1333, 2015.
111. Mehiar Dabbagh, Bechir Hamdaoui, Mohsen Guizani, and Ammar Rayes. Release-time aware vm placement. In *Globecom Workshops (GC Wkshps), 2014*, pages 122–126. IEEE, 2014.
112. https://http://www.top500.org. [Online; accessed 9-Apr-16].
113. Dave Ohara. Sustainable computing is it time to turn off your servers? https://technet.microsoft.com/en-us/magazine/cc700341.aspx. [Online; accessed 26-August-16].
114. Charles Reiss, a Tumanov, and Gr Ganger. Towards understanding heterogeneous clouds at scale: Google trace analysis. *... Center for Cloud ...*, 2012.
115. William Voorsluys, James Broberg, Srikumar Venugopal, and Rajkumar Buyya. Cost of virtual machine live migration in clouds: A performance evaluation. In *IEEE International Conference on Cloud Computing*, pages 254–265. Springer, 2009.
116. Mohammad Moein Taheri and Kamran Zamanifar. 2-phase optimization method for energy aware scheduling of virtual machines in cloud data centers. In *Internet Technology and Secured Transactions (ICITST), 2011 International Conference for*, pages 525–530. IEEE, 2011.
117. Nanyan Jiang and Manish Parashar. Enabling autonomic power-aware management of instrumented data centers. In *Parallel & Distributed Processing, 2009. IPDPS 2009. IEEE International Symposium on*, pages 1–8. IEEE, 2009.
118. Lizhe Wang and Samee U Khan. Review of performance metrics for green data centers: a taxonomy study. *The Journal of Supercomputing*, 63(3):639–656, 2013.
119. Daniel C Vanderster, Amirali Baniasadi, and Nikitas J Dimopoulos. Exploiting task temperature profiling in temperature-aware task scheduling for computational clusters. In *Advances in Computer Systems Architecture*, pages 175–185. Springer, 2007.
120. Sameep Mehta and Anindya Neogi. Recon: A tool to recommend dynamic server consolidation in multi-cluster data centers. In *Network Operations and Management Symposium, 2008. NOMS 2008. IEEE*, pages 363–370. IEEE, 2008.
121. Raffaele Bolla, Roberto Bruschi, Franco Davoli, and Andrea Ranieri. Energy-aware performance optimization for next-generation green network equipment. In *Proceedings of the 2nd ACM SIGCOMM workshop on Programmable routers for extensible services of tomorrow*, pages 49–54. ACM, 2009.
122. Zhuoyao Zhang, Ludmila Cherkasova, and Boon Thau Loo. Exploiting cloud heterogeneity to optimize performance and cost of mapreduce processing. *ACM SIGMETRICS Performance Evaluation Review*, 42(4):38–50, 2015.

123. Candy Pang, Abram Hindle, Bram Adams, and Ahmed E Hassan. What do programmers know about software energy consumption? *IEEE Software*, 33(3):83–89, 2016.
124. Aameek Singh, Madhukar Korupolu, and Dushmanta Mohapatra. Server-storage virtualization: integration and load balancing in data centers. In *Proceedings of the 2008 ACM/IEEE conference on Supercomputing*, page 53. IEEE Press, 2008.
125. Rajarshi Das, Jeffrey O Kephart, Charles Lefurgy, Gerald Tesauro, David W Levine, and Hoi Chan. Autonomic multi-agent management of power and performance in data centers. In *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems: industrial track*, pages 107–114. International Foundation for Autonomous Agents and Multiagent Systems, 2008.
126. Green Grid Industry Consortium et al. Green grid metrics describing data center power efficiency. *Technical Committee White Paper, available at http://www. thegreengrid. org/˜/media/WhitePapers/Green_Grid_Metrics_WP. pdf*, 2007.
127. Christian Belady, Andy Rawson, John Pfleuger, and Tahir Cader. Green grid data center power efficiency metrics: Pue and dcie. Technical report, Technical report, Green Grid, 2008.
128. Mueen Uddin, Asadullah Shah, Raed Alsaqour, and Jamshed Memon. Measuring efficiency of tier level data centers to implement green energy efficient data centers. *Middle-East Journal of Scientific Research*, 15(2):200–207, 2013.
129. John OLoughlin and Lee Gillam. Performance evaluation for cost-efficient public infrastructure cloud use. In *International Conference on Grid Economics and Business Models*, pages 133–145. Springer, 2014.
130. Rafidah Pakir Mohamad, Dimitrios S Kolovos, and Richard F Paige. Cloud computing workload and capacity management using domain specific modelling. https://www-users.cs.york.ac.uk/dkolovos/research/publications/pdf/cloud-mde-2012.pdf. [Online; accessed 19-Aug-16].
131. John OLoughlin and Lee Gillam. Sibling virtual machine co-location confirmation and avoidance tactics for public infrastructure clouds. *The Journal of Supercomputing*, 72(3):961–984, 2016.
132. Christopher Stewart and Kai Shen. Some joules are more precious than others: Managing renewable energy in the datacenter. In *Proceedings of the Workshop on Power Aware Computing and Systems*, pages 15–19. IEEE, 2009.
133. http://www.telegraph.co.uk/finance/newsbysector/energy/11923465/Blackout-risk-rises-as-UK-energy-crisis-deepens.html. [Online; accessed 21-July-16].
134. Kyong Hoon Kim, Anton Beloglazov, and Rajkumar Buyya. Power-aware provisioning of virtual machines for real-time cloud services. *Concurrency and Computation: Practice and Experience*, 23(13):1491–1505, 2011.
135. Rodrigo N Calheiros, Rajiv Ranjan, César AF De Rose, and Rajkumar Buyya. Cloudsim: A novel framework for modeling and simulation of cloud computing infrastructures and services. *arXiv preprint arXiv:0903.2525*, 2009.
136. Mohamed Abu Sharkh, Ali Kanso, Abdallah Shami, and Peter Öhlén. Building a cloud on earth: A study of cloud computing data center simulators. *Computer Networks*, 108:78–96, 2016.
137. Íñigo Goiri, Kien Le, Md E Haque, Ryan Beauchea, Thu D Nguyen, Jordi Guitart, Jordi Torres, and Ricardo Bianchini. Greenslot: scheduling energy consumption in green datacenters. In *Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis*, page 20. ACM, 2011.