Survey Energy Management Approaches in Data Centres

Bouchra Morchid[1], Siham Benhadou[2], Mariam Benhadou[3], Abdellah Haddout[4], Hicham Medromi[5]
Engineering Research Laboratory (LRI), System Architecture Team (EAS) [1, 2, 5]
National and High School of Electricity and Mechanic (ENSEM) Hassan II University, Casablanca, Morocco
Research Foundation for Development and Innovation in Science and Engineering[2, 5]
Laboratory of industrial management and energy and technology of plastics and composites[3, 4]
National and High School of Electricity and Mechanic (ENSEM) Hassan II University Casablanca
Morocco Research Foundation for Development and Innovation in Science and Engineering

***Abstract*—Data centers are today the technological backbone for any company. However, the failure to control energy consumption leads to very high operating costs and carbon dioxide emissions. On the other hand, reducing power consumption in data centers can lead to a degradation of application performance and quality of service in terms of SLA Service Level Agreement. It is therefore essential to find a compromise in terms of energy efficiency and resource consumption. This paper highlights the different approaches of energy management, related studies, algorithms used, the advantages and weaknesses of each approach related to server virtualization, consolidation and deconsolidation of virtual machines.**

Keywords—Data center; power management; virtual machines; physical servers "hosts"; energy efficiency; SLA (Service Level Agreement); PUE (Power Usage Effectiveness); QoS (Quality of Service)

I. Introduction

Currently, technological advances and new business models of cloud computing have contributed to the development of data centers in number and size.

This strong growth has not only contributed to an increase in the cost of electricity and enormous energy consumption, but has also generated significant carbon dioxide emissions, thus contributing to the greenhouse effect [1]. Indeed, in addition to the significant amount of electricity consumption for IT resources, it requires such considerable energy resources to cool this infrastructure from the heat dissipated, including servers and storage racks.

Data centers' energy costs are increasing, which requires reducing their energy consumption and improving energy efficiency through the development of efficient management approaches.

The rest of this paper is organized as follows: In Section II, we describe different energy management approaches in Data center. In Section III, we present related studies especially on server virtualization, consolidation and deconsolidation of virtual machines for optimizing energy efficiency in data centers. Section IV describes discussions about some limitations of the related studies. Finally, the conclusion and perspectives of evolution of different approaches are drawn in Section V.

II. State of the Art

A. Context

Each new service to be implemented by companies often requires the acquisition of new equipment. According to many estimates, the level of server utilization is generally less than 20% [2], although a significant part of the energy consumed in the data center is used to power inactive machines [3]. An underutilized server consumes more energy and has a negative impact on energy efficiency due to the non-linear characteristics of server power proportionality at these usage levels [4].

Moreover, the energy consumption of IT resources directly impacts the cooling power required to cool and maintain the ambient temperature of the data center at a predetermined threshold value. Indeed, according to Patel et al. [1], each watt consumed by computer equipment requires an additional power of 0.5 to 1 watt to operate the cooling system.

In [5], an indicator of energy efficiency in data centers "PUE: Power Usage Effectiveness" is the usage efficiency metric, which corresponds to the ratio between the total energy consumed in the data center and the total energy consumed by IT equipment (servers, storage, networks).

From an energy point of view, a data center with a PUE equal to 1 consumes less energy and is most efficient [6].

In this regard, the approaches used to improve energy efficiency in data centers must take into account the various parameters that influence energy efficiency.

B. The Different Approaches

Several energy management approaches have been discussed to optimize energy efficiency. In [7], these techniques can be divided into static and dynamic: Static energy management and dynamic energy management (see Fig. 1):

a) Static energy management: The first approach consists in applying optimization methods when designing hardware components (circuit, logic, architectural system levels) and also software components such as OS, compilers, etc. The low-powered ARM system or atom-based processor servers that have a reduced performance capability and are more energy efficient can be used as a system device to reduce static energy consumption [1]. In order to optimize the static energy consumption of the data center, it is necessary to avoid oversizing of capacity, however, care must be taken to design a precise dimensioning while taking into account possible evolutions to avoid bottlenecks and respect SLA services.

b) Dynamic energy management: The second approach is to apply different methods of power optimization at the hardware and software level to ensure energy efficiency.

- At the hardware level: Dynamic power management (DPM), Dynamic voltage and frequency scaling (DVFS).
- At the software level: Virtualization controls, thermal controls and server heterogeneity controls, etc.

The DPM method consists of reducing or eliminating static server consumption by either disabling the server or switching it to a low-power standby state when it is not used on the basis of the workload.

Workload is a factor used to determine when a server can be disabled or transferred to standby mode or turned off.

This method saves energy on certain types of workloads [8]. However, additional configuration time for the transition of the server from a low-power state (Standby or off) to an operational state or vice versa can negatively affect performance and power consumption. The effectiveness of this method is based on effective workload management to meet SLAs and a rapid transition to and from standby mode.

The DVFS method consists in reducing the number of instructions that a processor can generate over a given time interval. It is about:

- Decrease the frequency and/or voltage of the processor when it is not fully used;
- Increase the processor frequency and/or voltage if performance should be improved.

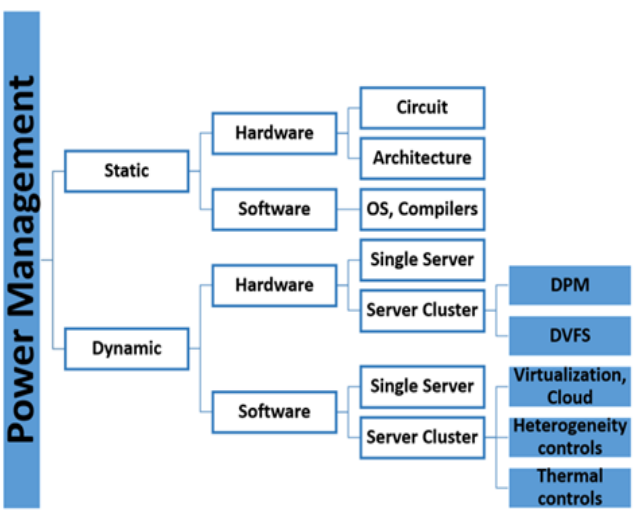

Fig. 1. Energy Management Techniques.

The DVFS method is an efficient way to manage the CPU's power consumption. However, the decrease in frequency and/or voltage has a negative impact on performance, which contributes to the increase in the time required to perform workload tasks. In addition, the increase in frequency and/or voltage contributes to much higher energy consumption.

The method based on the controls of consolidation and deconsolidation of virtual machines consists in consolidating workloads to fewer physical servers and disabling inactive servers or putting them in a low-power state. The advantage of this approach is that it is very appropriate in Cloud Data Centers and contributes to improving energy efficiency through server consolidation and disabling unused servers. Data center location can also be taken into account to migrate virtual machines to servers hosted at low-power data centers. However, this method faces constraints such as:

- A degradation in application performance due to aggressive consolidation;
- The increase in the temperature of the server hosting the virtual machines migrated to it and the creation of thermal hot spots, which also increases the cooling power required to dissipate the additional heat generated.

The method based on thermal controls consists in maintaining the thermal state of IT resources within an acceptable operating range. The increase in IT energy consumption has a direct impact on cooling power. The advantage of this method is that it allows the maintaining of the appropriate thermal profile for all IT resources in the data center, which ensures maximum reliability, longevity and return on investment. However, the dependence (linear and non-linear) between IT and non-IT class must be studied together to ensure that the total energy consumption of the data center is optimized.

The method based on server heterogeneity controls allows the workload to be allocated to the server with the most energy-efficient architecture. This is a technique used for a clustered server architecture.

This technique makes it possible to use the best features of each server to achieve energy efficiency at all levels of use. In addition, the heterogeneity of a server is considered according to the performance of the processor and the power consumption for a range of workloads. In some cases [1], such as an Infrastructure as a service (IaaS) data center, the knowledge of application workloads and the time during which virtual machines are used is unknown, which helps to depend on processor performance.

III. Related Studies

In this section we will review some related studies that focus on the use of server virtualization, virtual machine consolidation and deconsolidation approaches to optimize energy efficiency in data centers. It also reviews the different algorithms and techniques used to select virtual machines to migrate, and to select the machines to activate or deactivate.

A. The Concept of Server Virtualization and Consolidation

Virtualization has become a very common technique in modern IT architectures. This technique allows several virtual machines to coexist on the same physical server to increase its utilization rate (see Fig. 2).

Virtualization provides the following advantages [9]:

- High flexibility in server management, administration and occupancy.
- Allows users and administrators to create, save (control point), copy, share, migrate, read, modify and cancel the execution state of the machine.
- Reduction of costs.
- Reduction of energy costs.
- Reduction of CO2 emissions.

Concerning the consolidation of virtual machines, it is a technique that eliminates the energy consumption of underutilized servers. It consists in migrating virtual machines from an underutilized physical server to another physical server. This technique makes it possible to disable underutilized physical servers or put them in low-power mode, and to increase the utilization rate of the physical server containing the migrated virtual machines.

In addition, the main factors driving the consolidation of virtual machines are the impact on performance and energy consumption. In fact, a massive migration leads to a degradation of performance and additional energy consumption. This is why it is important to study when virtual machines can be consolidated or deconsolidated? To which physical servers should the virtual machine be migrated to, which will ensure energy savings? What is the migration scenario to adopt? What are the migration costs?

B. pMapper

pMaper is an application mapping framework (see Fig. 3), in a cluster environment of heterogeneous virtualized servers, that dynamically places them on physical hosts in order to minimize energy consumption while respecting performance guarantees and so that the cost/power trade-off is optimized.

The pMapper framework is based on the use of three different managers, with one arbitrator to ensure consistency between the three managers [10]:

Performance Manager: It provides an overview of the application in terms of QoS, SLA and performance, and communicates software actions such as virtual machine resizing and/or resting. In the case of heterogeneous platforms, the performance manager has a knowledge base to determine the performance of an application when a virtual machine is migrated from one platform to another.

Power Manager: It initiates power management on a hardware layer; it examines current power consumption and can suggest a limitation by applying a technique that adjusts the voltage/frequency dynamically or explicitly limits the processor.

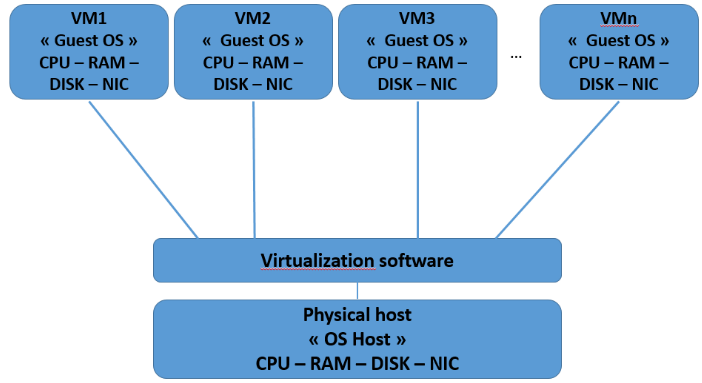

Fig. 2. Virtual Machine Diagram.

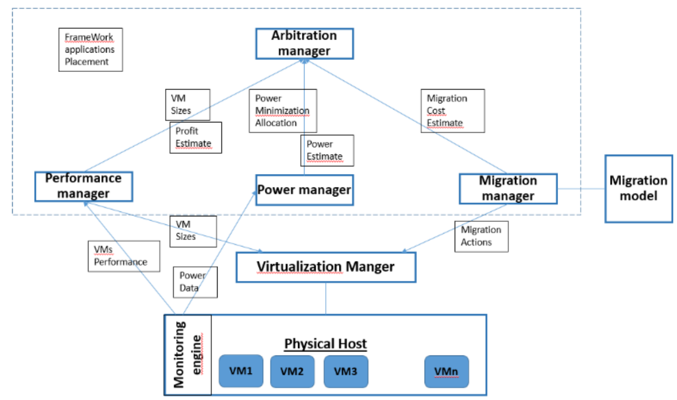

Fig. 3. pMapper Architecture for the Placement of Applications[10].

Migration Manager: Executes direct migration of virtual machines to consolidate workload. The migration manager estimates the cost of moving from one location to another and uses the migration model to make the estimate.

Arbitration Manager: It explores the configuration space of eligible virtual machine sizes and locations and implements an algorithm to calculate the best virtual machine locations and sizes, based on estimates received from Performance, Power and Migration managers.

This approach uses application placement algorithms to minimize the overall cost of "power and migration" while respecting performance constraints:

- The Min Power Parity (mPP) algorithm: It allows to place virtual machines on all servers in a way that the overall power consumed by all servers is minimized. However, the last configuration is not taken into account by this algorithm, which can lead to large-scale migrations and therefore a high overall cost (power + migration).
- The Min Power Placement with History (mPPH) algorithm: It works in the same way as mPP except that it takes into account the history (the previous placement is also taken into account). The mPPH algorithm tries to minimize migrations by relocating a minimum number of virtual machines, while migrating to the new optimal use of server targets.

- The pMaP algorithm: It is placed at the arbitration manager level to optimize the compromise of power migration. However, when the load is unbalanced to minimize power, pMaP leads to a large number of abandoned requests. The decrease in applications is considered as penalty of placement.
- The pMaP +: It takes into account the penalty and only selects intermediaries with a penalty below a predefined threshold.

The pMaP and mPPH algorithms are the two best performing algorithms, even with an increased number of servers.

The mPP and mPPH dynamic algorithms are capable of saving about 25% power from static virtual machine locations and with balanced load.

In view of the weaknesses of this solution, the effects of migration, in particular on total energy consumption and performance are not taken into account. When the load is unbalanced, a large number of requests are abandoned or delayed by applying the pMaP algorithm.

C. Kusic's Approach: Limited Anticipation Control (LLC)

The approach adopted consists in minimizing energy costs in a virtualized server cluster environment and under conditions of workload uncertainty, using sequential optimization using limited anticipation control (LLC). The purpose of this control is to minimize energy consumption and SLA contract violations in order to maximize the profits that may be lost while waiting for a virtual machine and its host to be activated, usually between three and four minutes. Pragmatically, revenues can be generated through response times, when they are below a predetermined threshold value, they give rise to a reward for the service provider, but if they exceed this threshold and contribute to violations of the SLA contract, they result in the payment of a penalty by the provider.

So, to meet this objective of profit maximization, the online controller determines the number of physical and virtual machines to allocate to each service for which virtual machines and their hosts are enabled or disabled based on workload demand, as well as the part of the CPU to allocate to each virtual machine.

The LLC approach modelizes the cost of control, i.e. the transition costs from an inactive to an active state when provisioning virtual machines or vice versa from an active to an inactive state. The LLC approach explicitly encodes the risk associated with procurement decision-making, since in an operating environment with a highly variable workload, an aggressive transition of virtual machines can occur and can therefore reduce benefits. Also, the LLC approach uses a hierarchical LLC structure, which breaks down the control problem into a set of smaller sub-problems and is solved cooperatively by several controllers, to achieve faster operation when workload intensity changes very quickly. This decomposition into sub-controllers will allow the main controller to adapt to these variations and plan resources over short periods of time, generally in the range of 10 seconds to a few minutes.

The Kalman filter is used to estimate the number of future requests, thus allowing to predict the future state of the system and consequently makes reallocations of CPU parts and virtual machine hosting mappings.

This approach has been implemented in a virtualized server cluster environment. The LLC method allows multi-objective optimization "in terms of energy reduction, reduction of SLA contract violations, increase in profits" under explicit operating constraints [2].

In terms of results, the server cluster, managed according to the LLC approach, maintains, on average, 22% of the cost of energy consumption over a 24-hour period compared to a system that operates without dynamic control while meeting QoS objectives.

The controller's execution time is relatively short, which does not pose any constraints in a high-demand operating environment.

Taking into account the notion of risk inherent in procurement decision-making has effects on control performance because a controller who takes this risk factor into account reduces the number of SLA violations by 35% compared to a controller with no risk.

The weaknesses of this approach mean that DVS controls are not taken into account.

This model requires supervised learning for application-specific adjustments, the number of virtual machines contributes more to energy consumption than resource use aspects. On the other hand, due to the complexity of the model, the execution time of the optimization controller reaches 30 minutes, even for 15 nodes, which is not suitable for large real systems.

D. Migration based on utilization "UMA"

The Migration based on utilization is a technique that consists of migrating virtual machines to stable physical servers in order to effectively reduce migration time and energy consumption (see Fig. 4).

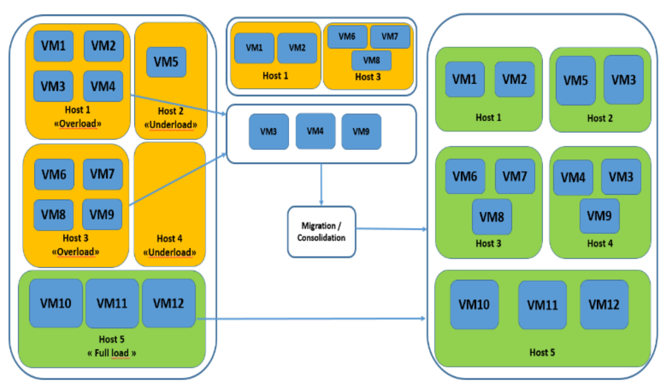

Fig. 4. VM Consolidation Framework based on Utilization [11].

This approach is based on a workload detection module that categorizes hosts into three classes: overloaded, fully loaded and underloaded. The migration probability is calculated for virtual machines hosted at the overloaded hosts and thus queue candidates waiting for migration are identified. Virtual machines hosted at the underloaded host level are consolidated to improve host utilization. No changes are made to the hosts at full load. With this solution, hosts can achieve optimal performance while offering high energy efficiency.

This approach uses Tabu Search to find the optimal solution when some inactive hosts are enabled to host virtual machines to migrate.

This algorithm uses the BFD algorithm first for a pre-migration for all virtual machines. Thus, a mapping table of virtual machines and hosts is created. After that, the virtual machine is optimized using the Tabu Search algorithm. In the Tabu search algorithm, if a possible migration has already occurred in a short period of time, this migration is placed in the tabu list, to avoid that the algorithm considers this possibility several times. Then after a number of iterations, the optimal solution is obtained.

The results of a study are carried out on four available hosts with uses of 50%, 60%, 65% and 0% respectively in a homogeneous environment. Five virtual machines are waiting for migration with MIPS requests for 20%, 20%, 20%, 15%, 15%, 15% and 10% processor utilization. The Tabu list was used, the UMA algorithm can achieve the optimal overall solution to the problem and which allowed a use of 70%, 60%, 65%, and 60% respectively for all hosts, which is closer to the total load [11].

UMA technology reduces about 77.5-82.4% of virtual machine migrations and saves up to 39.3-42.2% in energy consumption compared to MinPower policy "mPP". In terms of the number of active servers, this technique reduces the number of active hosts from 39.2% to 45.7% compared to the MinPower policy "mPP". The resource exploitation rate is between 70% and 90%.

For the SLA violation, the results show that UMA slightly mitigates the SLA violation. Compared to MinPower "mPP", UMA works a little better than the MinPower "mPP" algorithm when the overload threshold is equal to 0.8, and very similar when the overload threshold is equal to 0.9 in terms of SLA violations.

The weaknesses of the UMA technique are:

The power transition of the server state and the costs associated to the transition latency are not taken into account.

The impact of migration, in particular on total energy consumption and performance, is not also taken into account [1].

E. Modified Best Fit Decreasing Adjustment Approach "MBFD"

The best modified decreasing adjustment "MBFD" is a technique that consists in placing virtual machines in such a way that the most used virtual machine is migrated to the physical server that provides the lowest power consumption. The virtual machines are ordered in decreasing order according to CPU usage, and each virtual machine is then assigned to a host that allows a small increase in power consumption using the MBFD algorithm.

In addition, selecting the machine to migrate is an essential step in optimizing the allocation of virtual machines. To this end, selection strategies are applied to determine when and which virtual machines should be selected for migration. This consists of defining two upper and lower CPU usage thresholds for each host and maintaining total CPU usage by all virtual machines between these two thresholds. When a host has CPU usage below the lower usage threshold, all virtual machines must be migrated from this host and put into standby mode. If a host has CPU usage above the upper usage threshold, some virtual machines must be migrated from this host to reduce resource usage and prevent SLA violations. The strategies for selecting virtual machines to be migrated from a host whose CPU usage has exceeded the upper usage threshold are [12]:

- Minimization of Migration Time (MMT): The algorithm sorts the list of virtual machines in decreasing order of CPU usage. Then, this algorithm performs several iterations to identify a virtual machine and remove it from the list until the host CPU usage is below the maximum usage threshold. The complexity of the algorithm is proportional to the number of overused hosts and the number of virtual machines allocated to these hosts.
- Highest Potential Growth (HPG): This strategy consists of migrating virtual machines with lower CPU usage relative to the usage capacity defined in the virtual machine settings, in order to minimize the potential increase in host usage and avoid SLA violations.
- Random Choice (RC): This strategy of random choice (RC) is based on a random selection of a number of virtual machines necessary to reduce the CPU usage of a host that has exceeded the maximum usage threshold.

The MBFD approach reduces energy consumption by 77% compared to the NPA policy and by 53% compared to the DVFS approach with 5.4% of SLA violations.

The weaknesses of this solution are the fact that this work uses different virtual machine migration models, actions to reconfigure reactive virtual machines instead of proactive, and time of transition and consumption to turn servers on and off and vice versa.

IV. Discussion

In this paper we have presented different approaches to power management in data centers and related studies including virtualization, consolidation and deconsolidation of virtual machines with some associated algorithms, as well as their advantages and weaknesses.

We find that most approaches use virtual machine placement algorithms to minimize energy costs. In addition, the costs of migrating virtual machines and the impact on performance are not always taken into account and therefore it is difficult to quantify a net benefit in terms of total energy consumption.

V. Conclusion and Perspectives

In this paper we focused on dynamic energy management at the software level by applying virtualization, consolidation and deconsolidation techniques to virtual machines. This technique can lead to the creation of hot spots on the server hosting several virtual machines.

We started by making a comparative study of the different power management approaches within data centers. Then we identified the weaknesses of some approaches in virtualization, consolidation and deconsolidation of virtual machines.

There are still challenges to be addressed in adopting a virtual machine consolidation approach to optimize energy efficiency improvements without impacting performance. It is interesting to note that these approaches studied may include many extensions, the main ones are as follows: Consideration of migration costs, dynamic workload and also the non-IT components (cooling equipment) to avoid the creation of hot spots.

We intend to do a detailed comparative study of the different architectures in terms of virtualization, consolidation and deconsolidation the virtual machines to draw an optimal architecture that aims to optimize the energy efficiency of the data centers while taking into account the aforementioned extensions.

References

[1] V.K. MohanRaj, R.Shriram, 22 April 2016, Power management in virtualized datacenter – A survey.

[2] Dara Kusic, Jeffrey O. Kephart, James E. Hanson, Nagarajan Kandasamy, Guofei Jian, September 21, 2008, Power and Performance Management of Virtualized Computing Environments Via Lookahead Control.

[3] Andrew Krioukov, Prashanth Mohan, Sara Alspaugh, Laura Keys, David Culler, Randy Katz, NapSAC : Design and Implementation of a Power-Proportional Web Cluster, Computer Science Division, University of California, Berkeley.

[4] Annavaram,M., Wong,D., 2012. KnightShift: scaling the energy proportionality wall through server-level heterogeneity .In : Proceedings of the Annual IEEE /ACM International Symposiumon Micro architecture,pp.119–130.

[5] Seyed Morteza Nabavinejad, Maziar Goudarzi, Communication-Awareness for Energy-Efficiency in Datacenters Department of Computer Engineering, Energy Aware Systems Lab, Sharif University of Technology, Tehran, Iran, p. 203.

[6] [6] Fabien Douchet, Juin 2016, Optimisation énergétique de data centers par utilisation de liquides pour le refroidissement des baies informatiques, Thèse / Université de Bretagne-Sud.

[7] Beloglazov,A. , 2013, Energy-Efficient Management of Virtual Machines in Data Centers for Cloud Computing (Ph.D.Thesis). University of Melbourne.

[8] Gandhi,A., Kozuch,M., Harchol-Balter,M., 2012, Are sleep states effective in data centers? In: Proceedings of the International Green Computing Conference and Workshops,pp.1–10.

[9] Ankita Jiyani, 2016, Energy Efficient Techniques In Virtualized Data Center : A Survey, Research Scholar : Computer Science & Engineering Dept. Swami Keshvanad Institute of Technology Jaipur, India, p. 145.

[10] Akshat Verma, Puneet Ahuja, Anindya Neogi, pMapper : Power and Migration Cost Aware Application Placement in Virtualized Systems, IBM India Research Lab, IIT Delhi.

[11] Qi Chen, Jianxin Chen, Baoyu Zheng, Jingwu Cui, Yi Qian, 2015, Utilization-based VM Consolidation Scheme for Power Efficiency in Cloud Data Centers.

[12] Anton Beloglazov, Jemal Abawajy, Rajkumar Buyyaa, 2011, Energy-aware resource allocation heuristics for efficient management of data centers for Cloud computing.