

A SURVEY ON DYNAMIC ENERGY MANAGEMENT AT VIRTUALIZATION LEVEL IN CLOUD DATA CENTERS

T.Veni and S. Mary Saira Bhanu

Department of Computer Science and Engineering
National Institute of Technology Tiruchirappalli
Tiruchirappalli-620 015, India
{406111001,msb}@nitt.edu

ABSTRACT

Data centers have become indispensable infrastructure for data storage and facilitating the development of diversified network services and applications offered by the cloud. Rapid development of these applications and services imposes various resource demands that results in increased energy consumption. This necessitates the development of efficient energy management techniques in data center not only for operational cost but also to reduce the amount of heat released from storage devices. Virtualization is a powerful tool for energy management that achieves efficient utilization of data center resources. Though, energy management at data centers can be static or dynamic, virtualization level energy management techniques contributes more energy conservation than hardware level. This paper surveys various issues related to dynamic energy management at virtualization level in cloud data centers.

KEYWORDS

Data Center, Virtualization, Energy Management, VM Consolidation, VM placement, VM Migration

1. INTRODUCTION

Data Centers have emerged as a back-bone infrastructure, housing large number of IT equipments such as servers, data storage, network devices, power and cooling devices etc. that facilitate the development of wide variety of services offered by the cloud [1] Currently, several service providers such as Amazon, Google, Yahoo, Microsoft, IBM and Sun, have their own data centers to provide a variety of scalable services dynamically to a large customer base [2-3]. With the rapid development of IT industry and increasing demand for cloud services, the number of data centers have increased. This has resulted in increased energy consumption. According to McKinsey report [4], the energy consumption of data centers is \$11.5 billion in 2010 and it doubles every five years. Gartner [5] also estimated that world wide IT infrastructures are responsible for 2% of

global CO₂ emissions and energy related costs account for the 12% of the total economical expenditures. The excessive energy consumption at data centers leads to high operational cost, large amount of CO₂ emission and falling reliability of hardware equipments. Hence, it is necessary to develop energy-efficient technologies in data centers not only for reducing operational cost but also for reducing its impact on the environment and system's reliability.

Virtualization technology is a powerful technology that can be used in data centers to achieve better resource utilization for energy conservation [6]. This technology allows multiple virtual machines (VMs) to share the resources on a single physical machine (PM). The features such as VM isolation and VM migration along with dynamic resource provisioning can be used either to consolidate virtual machines on fewer physical servers or to balance the load across physical servers in data centers, thereby ensuring applications' performance.

Apart from virtualization technology, the data centers use low-power hardware components and hardware level energy management techniques for energy conservation. However, managing energy at virtualization level rather than hardware level contributes a major portion of energy conservation in data centers.

Energy management techniques at the data centers can be static or dynamic. This paper surveys the various issues related to dynamic energy management at virtualization level in cloud data centers.

The rest of this paper is organized as follows. Section 2 explains taxonomy of energy management mechanisms in virtualized data centers; section 3 presents an overview of hardware level energy management techniques; section 4 presents the details study of virtualization level energy management techniques and finally, section 5 presents the conclusion.

2. TAXONOMY OF ENERGY MANAGEMENT TECHNIQUES IN DATA CENTERS

It is necessary to understand the background terminologies involved in the context of energy management techniques. The energy consumption (E) at a data center is defined as a total amount of power (P) consumed over a period of time (T) while performing the work [7], i.e.,

$$E = P * T \quad (1)$$

Thus, the energy conservation at the data centers can be achieved by minimizing average power consumption (P). As per taxonomy [7], the energy management approaches in the virtualized data center can be classified into static and dynamic energy management techniques as shown in Fig.1

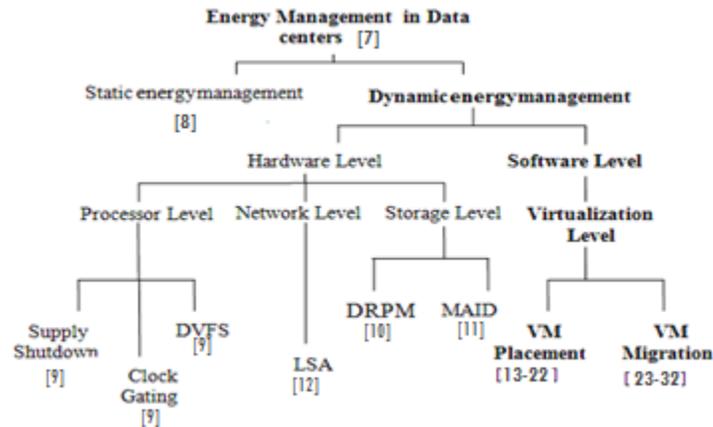


Fig.1. Taxonomy of energy management techniques in data centers

Static Energy Management (SEM) technique [8] uses low-power hardware components at the data centers for energy savings. The request for the cloud services running in the data center is volatile in nature. This technique fails to address the runtime adaptation of data centers according to the service request level to avoid resource wastage and it is also a very expensive technique for energy conservation.

3. OVERVIEW OF DYNAMIC ENERGY MANAGEMENT TECHNIQUES IN DATA CENTERS

Dynamic energy management (DEM) techniques dynamically reconfigure the system based on current resource requirements of the requested services. DEM utilizes the power scalable hardware components and software approaches to optimize the energy consumption at data centers. Hence, the large scale computing systems focus on DEM techniques for energy conservation. DEM techniques are classified into hardware and virtualization assisted techniques based on the level of their applicability.

Hardware Level Energy Management Techniques

The energy management techniques need to be applied on each hardware components such as the processors, network equipments, and storage devices to reduce overall hardware level energy consumption at data centers.

Processor Level

Energy consumption of a processor consists of two components such as static power and dynamic power. The static power is the power drawn even if the processor is idle whereas dynamic power depends on the usage of various resources at server [7]. The dynamic power consumption of a CMOS-based processor is given by

$$\text{Dynamic Power} = A * F * C * V^2 \quad (2)$$

Where A is the percentage of active gates, F is clock frequency of the processor, C is capacitance load of the processor, and V is the voltage supplied. Benini et al. [9] classified the energy management in processors into

1. **Network level Clock Gating** – Here, the supplied voltage or the clock frequency is reduced for idle components
2. **Dynamic Voltage and Frequency Scaling (DVFS)** – In DVFS, the voltage and frequency of the processor can be scaled dynamically depending on the request that is being processed
3. **Supply Shutdown** – In this case, the idle components are powered off to decrease power dissipation as there is a leakage of current even if all the clocks are halted

Network Level

As the size of the data center and network infrastructure is exploding, it is necessary to address the energy consumed by the network devices such as routers, switches, NIC, etc. **Link State Adaptation (LSA)** [12] is a technique that changes the link speed of network dynamically to reduce energy consumption in networks.

Storage Level

The storage devices constitute a significant fraction of the overall energy budget. The storage level energy management is accomplished by the following techniques:-

- **Dynamic Rotations Per-Minute (DRPM):** In [10], DRPM is proposed which dynamically changes the speed of disk rotation per minute (RPM) , thereby controlling the power explored in the spindle motor driving the platters
- **Massive Array of Idle Disks (MAID):** In [11], MAID is proposed to place copies of required files on a sub set of files called “cache disks” and non-cache disks are spin down for energy conservation

The hardware techniques do not contribute much in total energy conservation at data centers. Therefore, it is essential to apply energy management techniques at virtualization level to minimize energy consumption at data centers.

Virtualization Level Energy Management Techniques

The energy consumption problem has been addressed partially by making improvements in the physical infrastructure of modern data centers. According to Open Compute project report [33], 93% of the data center’s energy consumption depends upon efficient utilization of computing resources at data centers. Virtualization is a key technology that facilitates the better use of available data center resources using the technique called **VM consolidation** [13]. This technique involves gathering of several virtual machines into single physical server thereby increasing the resource utilization. It also helps to reduce the energy consumption by putting idle physical servers into sleep mode or power save mode. The VM consolidation can be performed either statistically or dynamically.

- In **static VM consolidation**, the virtual machine manager (VMM) allocates the physical resources to VMs based on peak load demand (overprovision). This leads to resource wastage because the workloads are not always at peak.
- In **dynamic VM consolidation**, the VMM changes the VM capacities according to the current workload demands. This helps in utilizing the data centers resources efficiently

The steps for VM consolidation are as follows:

1. Gathering the requirements of VM and placing the VM on PM (initial placement)
2. Monitoring and profiling the resource utilization for hotspot detection (overloaded and under loaded PMs)
3. Resize and remap the VM to another PM.

The VM placement and VM migration act as backbone to VM consolidation process. The issues such as heterogeneity and scalability of physical resources, volatile workloads, and migration cost make the VM consolidation process difficult. The challenges for efficient VM consolidation are:-

1. Finding the proper VM- to-PM mappings to minimize the acceptable cost function
2. Dynamic hotspot detection
3. Performing VM migration process with minimal service downtime and resource consumption during migration process

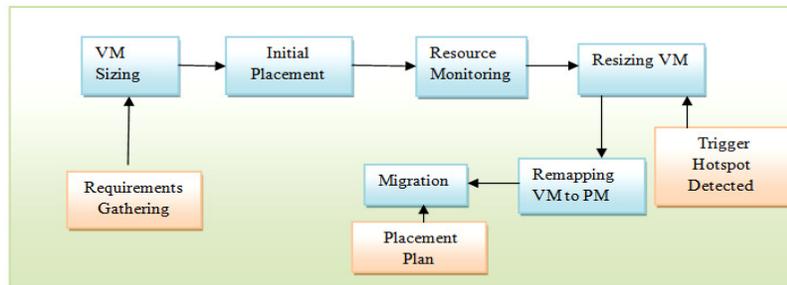


Fig. 1. Virtual Machine consolidation steps

Virtual Machine Placement

VM placement [14] is the process of mapping VMs to PMs. The placement approach should consider multiple resources such as CPU, memory, disk storages and network bandwidth to reduce the energy consumption at data centers and also maintain the energy performance tradeoff.

The virtual machine placement involves two main steps:

- Provisioning of resources for the virtual machines according to the capacity requirements of corresponding applications (VM sizing)
- Actual placement of VMs onto PMs.

The goal of virtual machine placement is to do VM consolidation either for power savings or load balancing to deliver best possible QoS to the applications running in VMs. VM placement problem is a NP –Hard problem [34] and no optimal solution exists for it. A set of heuristics is used to solve VM placement problem.

4. STATE-OF-THE ART IN VIRTUAL MACHINE PLACEMENT

The energy consumption at data centers increases with increase in the number of VMs on the PM and the type of applications running on them. Therefore, optimal placement of the VMs on the PM is crucial for energy conservation.

Availability aware virtual machine placement

The main challenge of the VM placement algorithms is to identify suitable VM-PM mapping so as to satisfy workload demands of various applications. Once the hotspot is detected in source PM, the VM placement algorithms find the resource demands of each VM and identifies the target PM based on the resource availability of the PMs [13-17].

Affinity aware VM placement

Affinity aware VM placement is based on relationship among VMs. This not only improves the response time of the applications but also optimizes the resources used for VM consolidation. If a hotspot has occurred because of network resource crunch among PMs, then placing two communicating VMs on the same PM will reduce the network overload. In [15] [20], the intra and inter-PM network traffic and also network topology has been considered for VM placement. In [17] [19], the inter memory sharing VMs are identified and placed on a single PM. This reduces the number of servers and memory foot print used.

Workload aware VM placement.

The efficient VM consolidation can be accomplished by considering workload interference between co-located VMs to avoid performance degradation. In workload aware VM placement mechanism, the different combination of workload mix such as CPU, I/O intensive, etc., are investigated in order to find optimal VM placement to make better utilization of data center resources [21] [22].

Virtual Machine Migration

VM migration [23] has become an indispensable tool that facilitates resource management at data centers for a wide variety of key scenarios such as load balancing, fault- tolerance, hot spot mitigation and server consolidation. The VM migration can be either offline or live. The offline VM migration follows traditional suspend and resume mechanism which increases service down time which leads to SLA violation. Live migration relocates the VM from source to target PM, without disturbing the applications running in the VM which results in minimal downtime. In live migration process, the virtual disks are stored on some shared storage such as Network File Systems, which is accessible from entire network. Hence, virtual disks are not transferred. The steps of Live migration are as follows (Fig. 3):-

- Once migration process is triggered, source PM transfers the VM's memory pages to target PM without suspending the VM. The memory pages are dirtied during the migration process. When this dirtied memory pages are transferred to target PM, it leads to multiple iterations.
- This iterative process ends and enters into suspend -and- copy phase when either one of the following conditions occurs:
 - The numbers of pre-copy iterations are greater than the threshold value (n)
 - The network traffic generated is greater than the multiples of VM size
 - The memory page dirty rate is greater than a memory transmission rate
 - The number of memory pages dirtied during current iteration becomes smaller than a predefined threshold

- In suspend-and-copy phase, the VM being migrated is suspended in the source PM and the remaining dirty memory pages as well as the state of the CPU register are transferred to the target PM. Once the migration process is completed, the target PM resumes the VM.

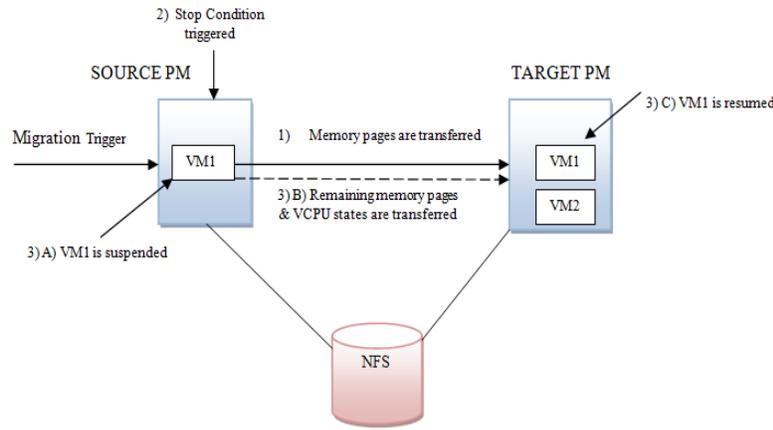


Fig. 2. Live Virtual Machine Migration Mechanism

5. STATE-OF-THE ART IN LIVE MIGRATION RESEARCH

It is hard to provide a quick live VM migration with less resource overhead because large amount of data transfer during migration results in performance degradation of VM applications. The response time and throughput of the data center applications heavily depends on the VM migration process. Hence, it is essential to improve the performance of migration process to facilitate efficient resource management at data centers.

Optimizations of live migration.

Many methods have been proposed to improve live migration and optimize its performance metrics such as total migration time and downtime to provide uninterruptible services to applications in VMs.

- **Memory compression:** Live migration performance is improved by minimizing the amount of data transfer to the destination using the technique called memory page compression [24]. However, performing compression and decompression in source and target PMs introduces significant overhead that results in performance degradation. Therefore, it is necessary to ensure trade-off between resource overhead and compression effects to achieve efficient migration.
- **Delta page transfer:** It is a technique [25] [32] that reduces the network bandwidth consumption by maintaining a cache of previously transferred memory pages. This optimizes the transmission of dirtied pages by sending the difference between cached page content and page content that is going to be transferred. This improves the live migration process with reduced risk of service interruption
- **Dynamic self ballooning (DSB):** It is a technique [26], that enables the VM OS (called guest OS) to return its free memory back to the VMM without explicit human intervention.

- **Data deduplication:** It is a technique [31-32] that uses hash based finger print technology that finds the similar and identical data inside the memory and disk of a single VM and optimizes the live migration process. This technique requires transferring memory data from source to target PM only once.
- **Post-Copy Approach:** It is a technique [26] that postpones the memory state transfer phase from the source to target PM after the VM's CPU state has been transferred to the target PM and resumed there. This approach ensures that each memory page is transferred at most once, hence avoiding the duplicate transmission overhead of pre-copy approach. However, if the destination PM crashes, VM cannot restart in the source PM. Post Copy approach does not have the reliability as Pre-Copy approach.
- **Hybrid pre and post copy:** It is a technique [27] that performs single round of pre-copying which precedes the virtual CPU state transfer. This is followed by post-copying of the remaining dirty pages from the source PM. Therefore, it gets the benefits of both approaches and improves live migration process.

Multiple VM Migrations

Whenever there is a need to shift entire virtual clusters to different locations, it requires large amount of data to be transferred over the network which results in network and CPU overhead. This leads to performance degradation of applications running in the PM. So identical contents of co-located VMs in the physical machine can be pro-actively tracked and transferred only once to the target PM. This improves the performance of multiple VM migration. This is known as “Gang Scheduling” of virtual machines. It optimizes both memory and network overhead of migration. The works [27-32] gang scheduling strategies.

Migration in network.

Several enterprises have data centers that are spread across the world to offer efficient services to worldwide users. The network can be LAN or WAN. Live migration of VMs across worldwide data centers [29-30] [32] entails transferring virtual CPU, memory state and disk state to target PM. Also it necessitates network re-configuration as the VMs shifts into a new subnet where a new IP address is assigned to the VM.

6. ANALYSIS

The live migration approach consumes resources such as CPU, memory and network which can seriously impact the performance of both the VM being migrated as well as other VMs in the PM. The performance of the migration approach depends on the parameters such as VM size, memory page dirty rate, network traffic and amount of available CPU resource.

Reference	Techniques used	Parameters considered				Performance metrics accounted			Single VM/ Multiple VM	WAN		
		VM size	Page dirty rate	N/W traffic	CPU cycles	Total mig- down time	N/W BW	CPU util		Y/N	n / w	stor age
clark et al [23]	Pre copy		✓	✓		✓			Single VM	No		
Jun et al [24]	Precopy+memory compression		✓			✓			Single VM	No		
Hacking et al [25]	Pre copy+ delta page transfer		✓	✓		✓	✓		Multiple VM	Yes		✓
Himes et al [26]	Post copy+ adaptive pre-paging+dynamic self ballooning		✓			✓			Single VM	No		
Sahni et al [27]	Precopy+post copy		✓			✓			Single VM	No		
Deshpande et al [28]	Precopy+multiple migrations		✓	✓		✓	✓		Multiple VM	No		
Riteau et al [29]	Precopy+content hashed hashing		✓			✓			Multiple VM	yes		✓
Bradford et al [30]	Write throttling+ DNS and IP tunneling		✓	✓		✓	✓		Multiple VM	yes	✓	✓
Wood et al [32]	Delta page transfer+ data deduplication+ VPN configuration		✓	✓		✓	✓		Multiple VM	yes	✓	✓
Zhang et al [31]	Precopy+ data duplication		✓			✓			Single VM	No		

Table 1. Comparison of state-of-art existing live migration approaches

From Table 1, it is observed that the two important parameters viz., VM size and available CPU cycles, have not been taken into consideration during migration. For efficient migration, all the mentioned parameters should be taken into account and also the performance of the migration algorithm needs to be measured in terms of performance metrics as shown in Table 1.

7. CONCLUSION

For dynamic energy management at virtualization level, VM placement and VM migration needs to be done. In this paper, the existing approaches for VM placement and VM migration have been discussed and analyzed. Issues such as intelligent workload mixing during VM placement, efficient WAN migration, etc. needs to be addressed to provide continuous and efficient energy management in virtualized data centers.

REFERENCES

- [1] Kant., K.: Data Center Evolution A Tutorial on State of the Art, Issues, and Challenges. J. Computer Networks. 53, 2939–2965 (2009)
- [2] Amazon Elastic Compute Cloud (Amazon EC2), <http://aws.amazon.com/ec2/>
- [3] Microsoft's Cloud Platform (Windows Azure), <http://www.windowsazure.com/>
- [4] Kaplan, J., Forrest, W., Kindler, N.: Revolutionizing Data Center Energy Efficiency. Technical report , McKinsey & Company(2008)
- [5] Gartner Says Energy-Related Costs Account for Approximately 12 Percent of Overall Data Center Expenditures, <http://www.gartner.com/it/page.jsp?id=1442113> (2011)
- [6] Barham, P., Dragovic, B., Fraser, K., Hand, S., Harris, T., Ho, A., Neugebauer, R., Pratt, I., Warfield, A. : Xen and the Art of Virtualization. In: 19th ACM Symposium on Operating Systems Principles, pp. 164-177, USA(2003)

- [7] Beloglazov, A., Buyya, R., Lee, Y., Zomaya, A. : A Taxonomy and Survey of Energy Efficient Data Centers and Cloud Computing. *J. Advances in Computers*. 82, 47-111(2011)
- [8] Luigi, G., Lassonde, W., Khan, S., Valentini, G., et al : An Overview of Energy Efficiency Techniques in Cluster Computing Systems. *Cluster computing*, Springer, 2011
- [9] Benini, L., Bogliolo, A., Micheli, G. : A Survey of Design Techniques for System-Level Dynamic Power Management. In : *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*. 8, 299-316 (2000)
- [10] Gurumurthi, S., Sivasubramaniam, A., Kandemir, M. et al : DRPM: Dynamic Speed Control for Power Management in Server Class Disk. In: *30th Annual International Symposium on Computer Architecture*, pp 169–181. USA (2003)
- [11] Colarelli, D., Grunwald, D.: Massive Arrays of Idle Disks for Storage Archives. In: *Proceedings of the ACM/IEEE Conference on SuperComputing*. pp 1–11. USA(2002)
- [12] Gunaratne, C., Christensen, K., Nordman, B., et al : Reducing the Energy Consumption of Ethernet with Adaptive Link Rate (ALR).*J. IEEE Trans Computer* 57, 448-461(2008)
- [13] Srikantiah, S., Kansal, A.,Zhao, F. : Energy Aware Consolidation for Cloud Computing. In: *Proceedings of the 2008 Conference on Power Aware Computing and Systems*. USA(2008)
- [14] Hyser, C., McKee, B., Garner, R., Watson, B. :Autonomic Virtual Machine Placement in the Data Center. HP laboratories, HPL-2007-189, 2008
- [15] Bobroff, N., Kochut, A., Beaty, K.: Dynamic Placement of Virtual Machines for Managing SLA violations, IBM T.J Watson Research Center, USA.
- [16] Beloglazov, A., Buyya, R., : Energy Efficient Resource Management in Virtualized Cloud Data Centers. In: *International Conference on Cluster,Cloud and Grid Computing*, pp. 826-831,2010.
- [17] Wood , T., Levin, G., Shenoy,P., Desnoyres, P., Cecchet,E., Corner, M.: Memory Bud-dies: Exploiting Page Sharing for Smart Co-location in Virtualized Data Centers: In: *International Conference on Virtual Execution Environments*, pp 31-40, USA(2009)
- [18] Meng, X, Pappas, V., Zhang, L. : Improving the Scalability of Data Center Networks with Traffic Aware Virtual Machine Placement, In:29th International Conference on Information Communications, pp. 1154-1162, USA(2010)
- [19] Sindelar, M., Staraman, R., Shenoy,P.: Sharing Aware Algorithms for Virtual Machine Co-location”. In: *23rd ACM Symposium on Parallelism in Algorithms and Architectures*, pp 367-378, USA(2011)
- [20] Reutiman, R., Sonnek,J, Greensky, J, Chandra A.: Starling Minimizing Communication Overhead in Virtualized Computing Platforms using Decentralized Affinity Aware Migration. Technical Report, University of Minnesota, 2009
- [21] Corradi, A., Fanelli,M., Foschini, L.: VM consolidation: A Real Case Based on Open Stack Cloud, In: *J. Future generation of computer system*
- [22] Kousiouris, U., Cacinotta, T., Varvarigou, T.: The Effects of Scheduling, Workload type and Consolidation Scenarios on Virtual Machine Performance and their Prediction through Optimized Artificial Neural Networks, *J. of system software* 84, 1270-1291(2011)
- [23] Clark,C., Fraser, K., Hand, S., Hnasen, J., Jul, E., Limpach,C., Pratt,I., Warfield, A.: Live Migration of Virtual Mahines. In:NSDI'05 Proceedings of the 2nd Conference on Symposium on Networked Systems Design and Implementation , PP. 273-286, USA(2005)
- [24] Jin, H., Deng, L., Wu,S., Shi, X., Pan,X.: Live Virtual Machine Migration with Adaptive Memory Compression. In: *IEEE CLUSTER*, New Orleans, LA(2009)
- [25] Hacking, S., Hudzia, B.: Improving the Live Migration Process of Large Enterprise Applications. In: *3rd International Workshop on Virtualization Technologies in Distributed Computing*, pp 51-58, USA(2009)
- [26] Hines , M., Gopalan, K., : Post-copy Based Live Virtual Machine Migration Using Adaptive Pre-paging and Dynamic self ballooning. In: *International Conference on Virtual Execution Environments*, PP. 51-60. USA(2009)
- [27] Sahni, S., Varma,V,: A Hybrid Approach to Live Migration of Virtual Machines . In: *International Conference on Cloud Computing in Emerging Markets*, India (2012)
- [28] Deshpande, U., Wang, X., Gopalan, K.: Live Gang Migration of Virtual Machines. In: *International Conference on High Performance Distributed Computing*, pp. 135-146 USA(2011)

- [29] Riteau, P., Morin, C., and Priol, T.: Shringer.: Efficient Wide Area Live Virtual Machine Migration Using Distributed Content Based Addressing. In: 17th International Conference on Parallel Processing, pp. 431-442. Heidelberg(2011)
- [30] Bradford , R, Kotsovinos , E., Feldmann, A., Schioberg, H., Live Wide Area Migration of Virtual Machine including Local Persistent State. In : 3rd International Conference on Virtual Execution Environments, pp. 169-179. USA(2007)
- [31] Zhang, X.,Huo, Z., Ma., J, Meng, Dan: Exploiting data duplication to accelerate live virtual machine migration
- [32] Wood, T.,Shenoy, p., Ramakrishnan k.k., Merve, J.: Cloudnet: Dynamic Pooling of Cloud Resources by Live WAN Migrations of Virtual Machines. In: 7th ACM SIGPLAN/SIGOPS International Conference on Virtual Execution Environments, pp. 121-132, USA(2011)
- [33] Open compute project, <http://opencompute.org/2012/11/08/building-a-better-server-chassis>
- [34] Skiena, S.: The Algorithm Design manual, ISBN 0-387-94860-0ssss