Experimental Analysis of Application Specific Energy Efficiency of Data Centers with Heterogeneous Servers

Grace Metri, Soumyasudharsan Srinivasaraghavan, Weisong Shi, and Monica Brockmeyer
Department of Computer Science
Wayne State University
Detroit, Michigan
{gmetri, sudharsan, weisong, mbrockmeyer}@wayne.edu

***Abstract*— Energy efficiency is an important issue for data centers given the amount of energy they consume yearly. However, there is still a gap of understanding of how exactly the application type and the heterogeneity of servers and their configuration impact the energy efficiency of data centers. To this end, we introduce the notion of Application Specific Energy Efficiency (ASEE) in order to rank energy efficiency of heterogeneous servers based on the hosted applications. We conducted extensive sets of experiments using three benchmarks: TPC-W, BS Seeker, and Matrix Stressmark. We observed that each server has different ASEE value based on the type of application running, the size of the virtual machine, the application load, and the scalability factor. In some cases, we witnessed 70% of ASEE improvement by changing the virtual machine size within the same node while keeping an identical load. In different cases, we witnessed up to 86% of ASEE improvement by running the same application with the same load within the same size of virtual machine but on different nodes. Our observation has many implications which include but are not limited to improving virtual machine scheduling based on the ASEE rank of the node. Another implication stresses on the importance of accurate prediction of application load and selecting the appropriate virtual machine size in order to improve the ASEE.**

Keywords: Virtualization, Cloud Computing, Power Management of Data Centers, Energy Efficiency.

I. INTRODUCTION

Cloud computing is a popular computing paradigm enabled by the large number of data centers and the advantages of virtualization. Data centers, driven by the economy of scale, are mainly built out of cheap, unreliable, heterogeneous interconnected commodity components [1]. Most of today's data centers are composed of heterogeneous servers because many misbehaved or failed servers get replaced with different ones [2]. The heterogeneity also stems from the fact that when a data center undergoes any upgrades, it is technically impossible to replace all the servers due to the size of the centers. Thus, new servers with different specification are added to the old ones. Many of today's data centers are considered mega data centers [3-5] because they house over tens of thousands of servers consuming tens of mega-watts of energy during peak hours adding up to 9.3 million dollars a year which is actually a small fraction when compared to the 7.2 billion consumed by the all servers around the world in 2005. As a result, improving energy efficiency within a datacenter will have a huge positive financial impact and can significantly reduce their carbon footprint.

With the extensive use of the cloud, there is a large spectrum of applications running on heterogeneous nodes where each application exhibits diverse utilization of resources and a broad range of workload. Toward this end, there is still a lack of understanding of how the application type and its workload can affect the energy consumed by the servers.

Traditionally, manufacturers of hardware components have established the energy efficiency of each component separately as its performance per watt and then compute the overall energy efficiency to the entire server. However, a common scenario within a datacenter is to have components of a single node replaced as they fail, resulting in varying degrees of energy efficient components within the same server. Since each application uses different percentages of the server components such as CPU or memory and since each component can have different energy efficiency values, it becomes necessary to calculate the application specific energy efficiency of each server.

In this paper, we focus on a black box technique to profile the energy efficiency of each server based on the application type. We examine the energy consumed and application performance of three different types of benchmarks, TPC-W [6], BS Seeker [7] and Matrix Stressmark [8] where each type of benchmark utilizes different proportions of hardware resources.

Toward this end, we provide the following contributions:

- We propose the concept of Application Specific Energy Efficiency (ASEE) where each server can have a different energy efficiency value based on the hosted type of application.
- Through detailed measurement based analysis on a heterogeneous cluster, we demonstrate that energy consumed to run a benchmark using the same input varies depending on the node type. We also exhibit a relationship between the size of the running VM and the overall energy efficiency of the server. We also present a strong relationship between the application type and energy efficiency of the server as the degree of scalability changes. The pattern of ASEE drastically varies as we vary the application type.

- Based on our observation, we derived a list of implications which can be used to increase the overall energy efficiency of a data center.

The rest of the paper is organized as follows. In Sections II and III, we define application specific energy efficiency (ASEE) and evaluation methods. Section IV gives the experimental results. Section V shows the implication of our results. Related work is presented in Section VI. Section VII concludes this paper.

II. Definition of Application Specific Energy Efficiency

As the energy consumption of data centers keeps on doubling every five years, energy efficiency of those data centers becomes a very important topic. Energy efficiency is defined as performance per watt which is calculated by running a benchmark and getting its performance score which gets divided then by the average system power usage [9]. Since each computer component consumes different energy values and since data centers are hosting heterogeneous types of applications each requiring a different combination of system usage, we are looking at energy efficiency from a global perspective where each server can have their energy efficiency compared to one another based on the type of application running. Thus, we define Application Specific Energy Efficiency (ASEE).

$$ASEE = \frac{\textit{Load of the application}}{\textit{Energy consumed by the application}}$$

The definition of *Load* can be different based on the type of application. In the case of web applications such as TPC_W, the *Load* is defined as the throughput. In the case of applications focusing on processing data such as BS Seeker, the *Load* is defined as the size of data processed. In the case of arithmetic operations such as Matrix Stressmark, the *Load* is the number of operations, e.g., iterations.

III. Evaluation Methods

We evaluated ASEE by creating a heterogeneous cloud and running experiments using three applications: TPC-W, BS Seeker, and Matrix Streesmark. The goal is to observe how ASEE changes based on server type, virtual machine (VM) size, load size, and scaling factor.

A. Experimental Setup

We created a heterogeneous cloud using Eucalyptus [10]. Our cloud contains the nodes as described in Table 2. Node-1 is our head node. Our cloud supported VMs as described in Table 3. Due to the size of Node-2, Node-3, and Node-4, they can only support up a (L) VM. The nodes and VMs ran CentOS 5.5. We used Apache Tomcat version 5.5.20 as the application server and MySQL 5.0.77 as the database server.

We did all our power consumption measurements using an electronic watt meter manufactured by Electronic Educational Devices Inc, Denver, CO [11]. The model used is Wattsup?/PRO/ES/.Net. The voltage is 120 VAC, 60 HZ and the max wattage is 1800 Watts. The outlet rating is 120 VAC/15 amps. The measurement accuracy is +/- 1.5 % and the selected interval of time between records is one second.

We collected the performance metrics via Oprofile [12], a system-wide profiler for Linux systems capable of profiling hardware performance counters of the CPU at low overhead.

Energy consumed can also depend on the air inlet temperature due to the power consumption of fans and temperature leakage. Our nodes are placed in the same location, making the latter a constant environmental variable that we did not need take into account. We calculated the energy consumed by calculating the average power consumed during the interval run time of the benchmark. Then, we multiplied the average power by the duration of the execution time of the benchmark.

$$\textit{Energy Consumed} = \textit{Average of power} * \textit{Duration of test}$$

B. TPC-W

TPC-W is an online bookstore serving as a transactional web e-commerce benchmark. We deployed a Java TPC-W implementation [13] based on TPC-W specification 1.0.1.

TPC-W consists of a client-server architecture where traffic generated by customers is emulated via remote Emulated Browsers (EBs). Each EB session consists of a series of sequential interactions such as searching for a product, browsing the list of products, adding items to the shopping cart, cart check out, and so on.

TPC-W contains three mixes as described in table 1. Browsing requests consist of checking the home page, new products, best sellers, search requests, and search results. Order requests consist of checking shopping cart, registration, buy request, order inquiry, order display, and so on. Since browsing requests are composed of requests that put pressure on the database server, mix 1 has mainly disk accesses whereas mix 3 has the least [14].

All our tests included 10,000 items in the database. For each of the mixes, we ran TPC-W with four different numbers of concurrent clients: 250, 500, 750, and 1000 EBs which we controlled via the Remote Browsing Emulator (RBE) which was shipped with TPC-W. We calculated the energy consumed by each test, the throughput which is the total number of web interactions requested and completed successfully, and the ASEE.

C. BS Seeker

BS Seeker is a CPU intensive Bio-informatics application designed for mapping bisulfite-treated reads in genome-wide measurements of DNA methylation at single nucleotide resolution. We deployed a Python BS Seeker [15] which takes an input file containing the genome reference and then converts them to a three-letter alphabet and uses Bowtie [16] to align the reads to reference genome. All of our testing with BS Seeker, we gave it as an input the same file of size 100k.

TABLE 1: REQUEST COMPOSITION OF TPCW-W

	Mix 1: Browsing	Mix 2: Shopping	Mix 3: Ordering
Browsing request	95%	80%	50%
Ordering request	5%	20%	50%

TABLE 2: CLOUD NODE TYPES

Server	Architecture & CPU type	Core	RAM	Cache	HD	Speed per Core
Node 1	64 bit Intel (R) XEON (TM) E5620 i386/ arch_perfmon	16	12GB	L1 8KB L2 1024KB L3 12MB	855 GB	2.4 GHz
Node 2 & 3	64 bit Intel (R) XEON (TM) i386/p4-ht	1	2GB	L1 8KB L2 1024KB	28 GB	2.8 GHz
Node 4	64 bit Intel (R) XEON (TM) i386/p4-ht	1	2GB	L1 8KB L2 1024KB	97 GB	2.8 GHz
Node 5	64 bit Intel (R) XEON (TM) E5620 i386/ arch_perfmon	16	12GB	L1 8KB L2 1024KB L3 12MB	855 GB	2.4 GHz

TABLE 3: TYPES OF AVAILABLE INSTANCES

Virtual Machine	Number of Cores	RAM	Hard Disk
Small (S)	1 core	128 MB	10 GB
Medium (M)	1 core	256 MB	10 GB
Large (L)	1 core	512 MB	10 GB
XLarge (XL)	2 cores	2 GB	40 GB
XXLarge (XXL)	4 cores	4 GB	50 GB

D. *Matrix Stressmark*

The Matrix Stressmark is part of DARPA Data Intensive Systems (DIS) Stressmark suite. It represents a sparse matrix with vector multiplication where memory accesses are irregular with a mixed level of reuse. It generates data at runtime. Then, it performs a number of iterations in order to calculate the equation A • x = b where A is a sparse nxn matrix and x and b are vectors with n elements each. This type of memory access patterns and arithmetic operations on the data is common with scientific applications [9].

For our experiments, the dimension n and vectors x and b were set to 3,000 and the number of non-zero items were 1,000,000 items. The data collected upon completion of the test were the number of iterations and total time to execute.

IV. EXPERIMENTAL RESULTS

We conducted a large number of experiments to investigate the relationship between application type and energy consumption of servers. We also observed the impact of VM size on the application performance and the energy consumption of the node. Then, we calculated the ASEE in order to establish a ranking of nodes based on their energy efficiency while running a specific application.

A. *TPC-W Results*

The first set of experiments consists of running TPC-W for each mix on the physical node. We ran the experiments for 250, 500, 750, and 1000 EBs. The throughput, energy consumed, and ASEE are displayed in Fig. 1, Fig. 2 and Fig. 3 respectively. We noticed that the identical nodes had similar throughput and energy consumption values. The small variations are due to the fact that the EBs requests are randomly generated and may be accessing different locations on disk. We also noticed that Node-5 had the highest throughput and the least amount of energy consumed when compared to the other nodes. In addition, as the number of EBs increase, the energy consumption increase as well but the increase in energy consumption is not proportional to the increase in throughput. For instance, when running mix 3 on Node-2 and changing the number of EBs from 250 to 750, the throughput increased by 135.5% and the energy consumed increased by 16%.

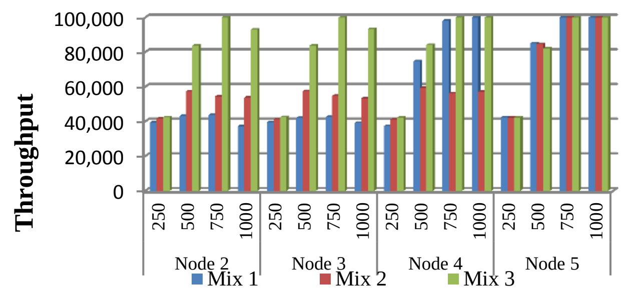

Figure 1: Throughput for TPC-W on the physical nodes.

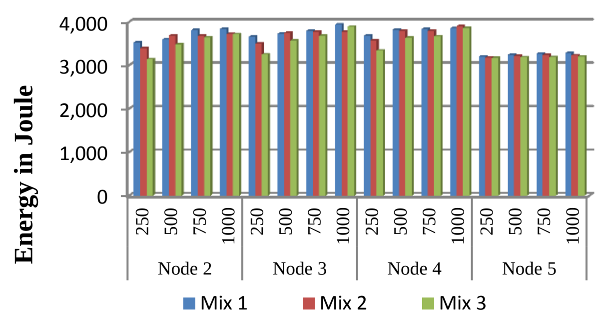

Figure 2: Energy consumption for TPC-W on the physical nodes.

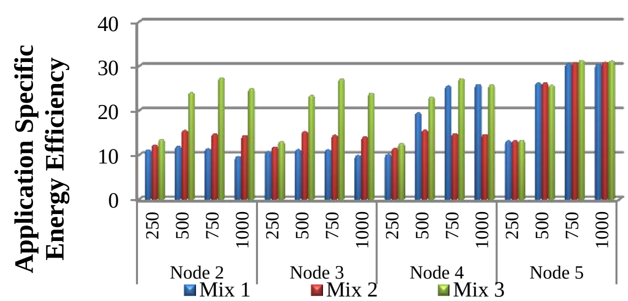

Figure 3: Application Specific Energy Efficiency when running TPC-w on the physical nodes.

The second set of experiments consists of running the three types of mixes for 250, 500, 750, and 1000 EBs within the supported VMs on all the nodes. The throughput for mix 1 and mix 3 are displayed in Fig. 4 and Fig. 5 respectively. The energy consumption for mix 1 and mix 3 are displayed in Fig 6. and Fig 7. respectively. The ASEE for mix 1 and mix 3 are displayed in Fig. 8 and Fig. 9 respectively. (Mix 2 exhibits similar patterns as Mix 3 but the results are not displayed for space constrains). In addition, we included Table 4 as the ASEE for all the nodes for mix 2.

Note: Node-2, Node-3, and Node-4 constantly failed within the (S) VMs even with 250 EBs. Based on TPC-W evaluation performed by Garcia et. al. [17], they found that with 256 Mbytes of memory which is the memory size set for our (S) VM, TPC-W suffers severe memory starvation. However, due to the presence of L3 cache in Node-5, the experiments did not fail.

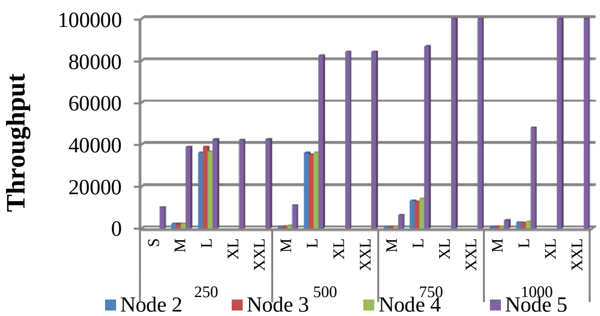

Figure 4: Comparing Mix 1 throughput for heterogeneous nodes.

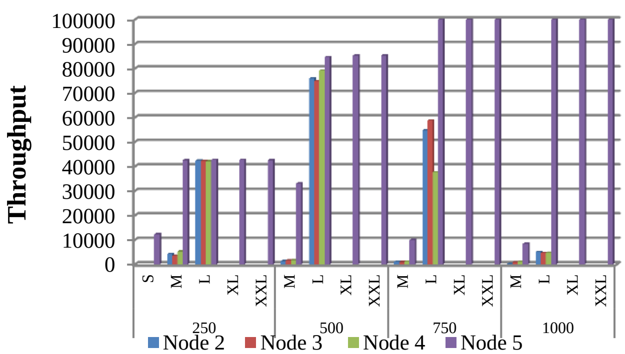

Figure 5: Comparing mix 3 throughput for heterogeneous nodes.

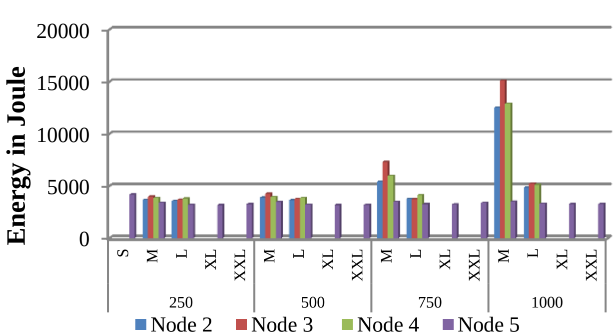

Figure 6: Comparing mix 1 energy consumption for heterogeneous Nodes.

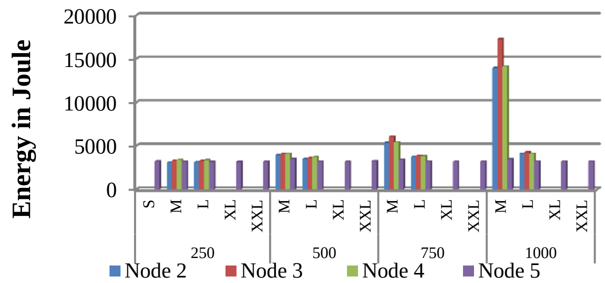

Figure 7: Comparing mix 3 energy consumption for heterogeneous nodes.

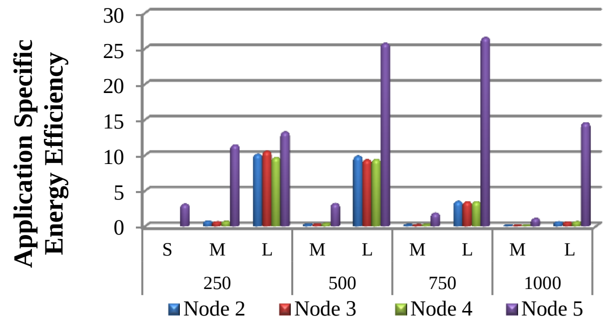

Figure 8: Comparing mix 1 Application Specific Energy Efficiency.

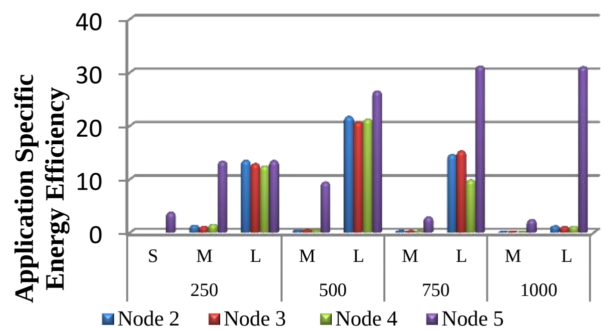

Figure 9: Comparing mix 3 Application Specific Energy Efficiency.

We analyzed our collected data by comparing the following metrics:

1) **Instance Efficiency:** Keeping the node type and EB size constant, we compared the throughput, energy consumption and ASEE based on VM size.
2) **Load Efficiency:** Keeping the node type and VM size constant, we compared the throughput, energy consumed and ASEE based on EB size.
3) **Mix Efficiency:** For each node, keeping the node type, the instance size, and EB size constant, we compared the throughput and energy consumption based on the mix type.
4) **Node Efficiency:** For all the nodes, keeping the EB size, instance size, and mix type constant, we observed how their ASEE compare to one another.

1) Instance Efficiency:

When keeping the node type and EB size constant, we noticed that as we increase the instance size from (S) to (L), their throughput values increased consistently whereas their energy consumption decreased. Intuitively, the more resources used the greater the energy consumption. But in the case of TPC-W, the smaller the instance the higher the disk accesses due to the thrashing of cache which leads to increase in energy consumption. On Node-5, when we increased the VM size from (L) to (XL), the throughput increased and the energy consumption decreased. For instance, in mix 1 with 1000 EBs, the throughput increased by 107%, the energy consumption decreased by 1%. Even though, the (XL) VM has two cores instead of one, the extra energy consumption by the additional core was cancelled out by the much larger added memory which reduced the number of accesses to the disk. In addition, the extra core should have reduced the latency time which in turn increased the throughput which directly affects the ASEE. On the other hand, when we compared the (XL) to the (XXL) VM, we noticed that the throughput was either slightly increased or it remained the same. However, the ASEE of an (XL) VM is better than the ASEE of the (XXL). The (XL) has two cores whereas the (XXL) has four cores. Having four cores would have reduced the latency to approach zero, however, the throughput gets to a point where it levels off because the user think time becomes dominant of the possible request generation rate [18].

2) Load Efficiency:

By keeping the node type and VM size constant, and comparing the throughput and ASEE based on the EB size, the throughput increases until the VM reaches its capacity of peak throughput value and then it sharply drops after the peak is reached. The peak throughput is highly dependent on the node type and VM size. We observed that each node type and corresponding VM size consistently reached the same peak value. For smaller instances ((S), (M), and (L)), when the peak throughput value is reached, that is when the VM has the highest value of throughput and it is the most ASEE. However, once the number of concurrent connections exceeds the peak, not only does the number of throughput sharply drop due to dropped connections, but also the energy consumption rises sharply leading to a very low ASEE VM. For larger instances ((XL) and (XXL)), we noticed that once we change the EB size from 750 to 1000, the throughput remains high, however the ASEE decreases.

3) Mix Efficiency:

After keeping the node type, the VM size, and EB size constant, and then comparing the throughput and energy consumption based on the mix type, we notice that mix 1 has the lowest throughput when compared to the other two mixes and mix 3 has the highest throughput when compared to the other two mixes. Since mix 1 has a larger number of disk and memory accesses as opposed to mix 3 due to its composition, it takes longer to complete mix 1 requests compared to mix 2 and mix 3. As a result, ASEE of mix 1 is lower than the other two especially for mix 3.

4) Node Efficiency:

After keeping the EB size, instance size, and mix type constant, and then comparing how their ASEE compare to one another, we noticed that Node-5 is more ASEE than all the remaining nodes which is due to the presence of L3 cache. This proves that having the same VM type running on different nodes can have different ASEE even when running the same application with the same input. Therefore, the hypothesis considering scheduling the VMs based on their ASEE in order to reduce the overall energy consumption of a data center and increasing its energy efficiency is a valid hypothesis.

The next set of experiments is to determine how the scalability of a node can affect its throughput, power consumption, and ASEE. The throughput, energy consumed and ASEE of mix 2 on Node-5 are displayed in Fig. 10, Fig 11, and Fig. 12. As we increase from one VM to more VMs running on the same node, the ASEE and throughput decreases slightly as the scalability degree increases until it reaches a scalability threshold where the energy consumption spikes, in addition the throughput and ASEE drops sharply. In addition, the smaller the VM, the lower is its scalability threshold and vice versa. This is evident when you compare one and eight (S) VMs with one and eight (L) VMs. Since Eucalyptus network model forwards all the traffic from the VM on the cluster to the cluster controller, having multiple network traffic across many VMs can saturate the cluster controller network bandwidth which lead to dropped connections [19].

We also noticed a reverse relationship between throughput and energy variation. Fig. 13 displays the relationship between throughput and energy for mix 1 with 250 EBs (similar results were witnessed with combination of mixes and concurrent requests). For clarity purposes of the graph, we reduced the values of throughput by 10% for all of our VM test cases. Based on the graph, it is apparent that when the throughput values drop as a result of failed requests, the energy consumption spikes. However, as the number of throughput remains consistent, so does the energy consumption.

B. BS Seeker Results

We first ran BS Seeker on the bare metal and then on all the nodes with different VM sizes and collected the energy consumed. Fig. 14 and Fig. 15 display the energy consumed and ASEE respectively. Based on our results, Node-5 is the most energy efficient whether BS seeker is running on the bare metal or within the same VM type as the other nodes within our cloud. In addition, we noticed that when running our benchmark in a (M) VM on Node-2, Node-3, Node-4, and Node-5 we significantly improve the ASEE by 71%, 71%, 72%, and 4% respectively when compared to running it in a (S) VM on the same nodes. Also when running the benchmark on a (L) instance as opposed to a (M) instance, we can improve the ASEE between 1% and 7%.

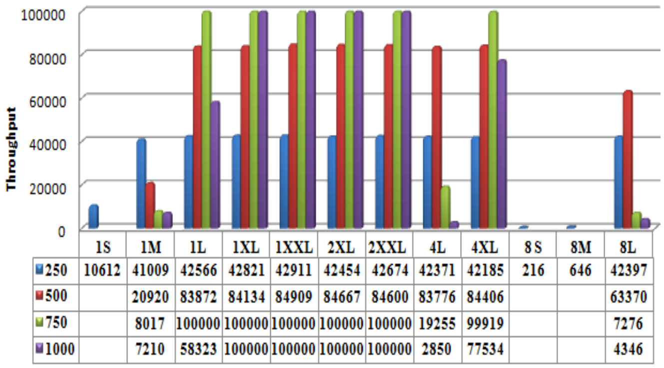

	1S	1M	1L	1XL	1XXL	2XL	2XXL	4L	4XL	8 S	8M	8L
250	10612	41009	42566	42821	42911	42454	42674	42371	42185	216	646	42397
500		20920	83872	84134	84909	84667	84600	83776	84406			63370
750		8017	100000	100000	100000	100000	100000	19255	99919			7276
1000		7210	58323	100000	100000	100000	100000	2850	77534			4346

Figure 10: Mix 2 throughput based on node scalability.

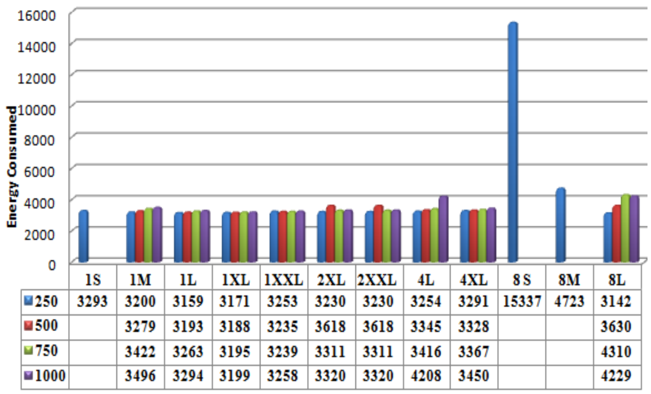

	1S	1M	1L	1XL	1XXL	2XL	2XXL	4L	4XL	8 S	8M	8L
250	3293	3200	3159	3171	3253	3230	3230	3254	3291	15337	4723	3142
500		3279	3193	3188	3235	3618	3618	3345	3328			3630
750		3422	3263	3195	3239	3311	3311	3416	3367			4310
1000		3496	3294	3199	3258	3320	3320	4208	3450			4229

Figure 11: Mix 2 energy consumed based on node scalability.

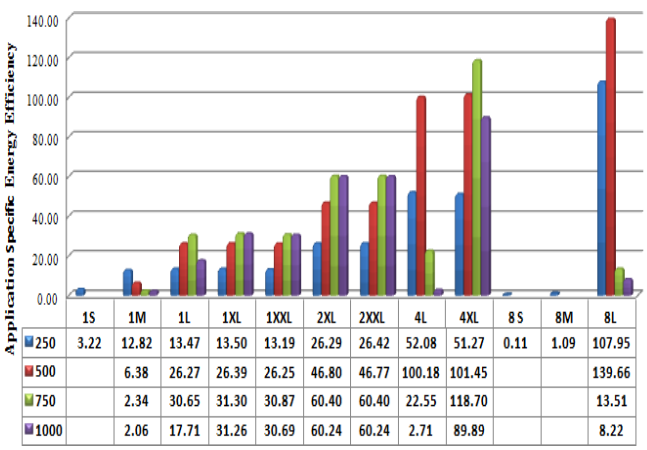

	1S	1M	1L	1XL	1XXL	2XL	2XXL	4L	4XL	8 S	8M	8L
250	3.22	12.82	13.47	13.50	13.19	26.29	26.42	52.08	51.27	0.11	1.09	107.95
500		6.38	26.27	26.39	26.25	46.80	46.77	100.18	101.45			139.66
750		2.34	30.65	31.30	30.87	60.40	60.40	22.55	118.70			13.51
1000		2.06	17.71	31.26	30.69	60.24	60.24	2.71	89.89			8.22

Figure 12: Mix 2 Application Specific Energy Efficiency based on node scalability

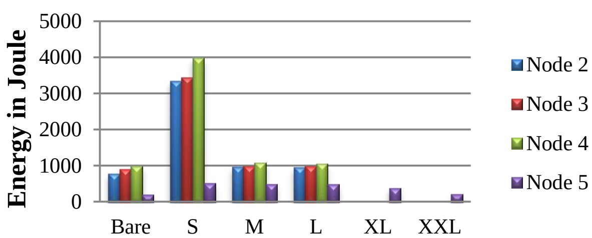

Figure 14: Energy consumption for BS Seeker

TABLE 4: APPLICATION SPECIFIC ENERGY EFFICIENCY FOR MIX 2

		Node 2	Node 3	Node 4	Node 5
250	S				3.22
	M	0.98	1.07	1.27	12.82
	L	11.90	12.36	11.70	13.47
	XL				13.50
	XXL				13.19
500	M	0.30	0.27	0.29	6.38
	L	14.03	13.15	13.58	26.27
	XL				26.39
	XXL				26.25
750	M	0.20	0.19	0.20	2.34
	L	2.45	2.06	5.96	30.64
	XL				31.30
	XXL				30.87
1000	M	0.07	0.06	0.06	2.06
	L	0.98	1.38	1.38	17.70
	XL				31.26
	XXL				30.70

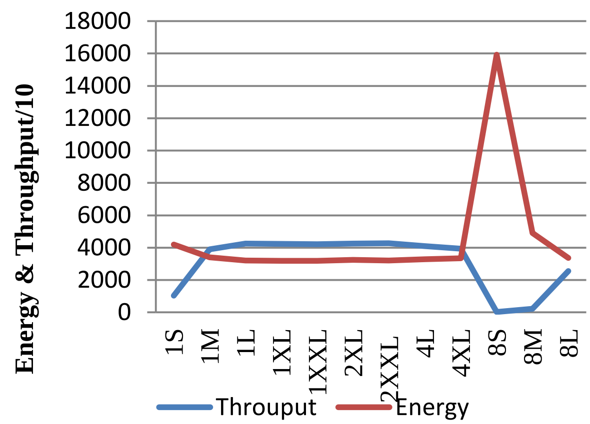

Figure 13: Relationship between throughput and energy consumed for mix 1 with 250 EBs.

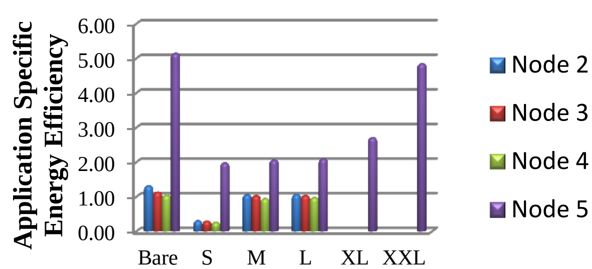

Fig 15: Application Specific Energy Efficiency of BS Seeker.

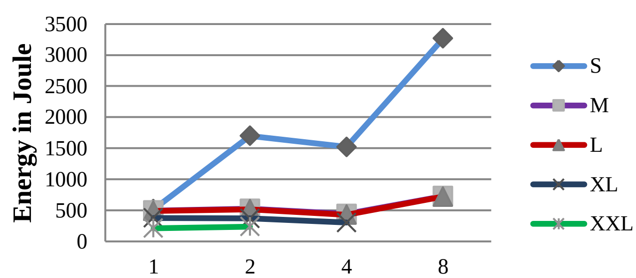

Figure 16: Energy consumption for BS Seeker based on virtual machine scaling.

The other noteworthy observation is as the instance size increases, its ASEE increases as well. Since BS Seeker is highly CPU intensive, the improvement in ASEE between (S), (M), and (L) instances is not significant since they all have the same number of cores. However, there is a significant improvement when we change the VM size from (L) to (XL) and from (XL) to (XXL) because the number of cores increases, resulting in faster computation, and thus shorter runtime, and higher ASEE.

We also performed tests on Node-5 in order to compare the scalability effect on the energy represented in Fig 16. We made a counter intuitive observation when scaling from two VMs to four VMs, the energy consumption actually decreased by up to 16%. Then, the energy consumption increased after scaling from four VMs to eight VMs by up to 114% when comparing the (S) VMs and up to 66% when comparing the (M) and (L) VMs.

C. Matrix Stressmark Results

We followed the same methodology and the test plan for the Matrix Stressmark. Unlike BS seeker results where the energy consumed dropped significantly when we ran the benchmark on Node-2, Node-3, or Node-4 compared to Node-5, in the case of the Matrix Stressmark, the improvement was at most 20% as shown in Fig. 17. Unlike TPC-W results where the ASEE significantly improved in many cases when the VM size increased, the VM size did not impact the ASEE in the case of the Matrix Stressmark. The difference was ± 0.01 which is a negligible difference.

On the other hand, another interesting pattern emerged. When we scaled the number of VMs as shown in Fig. 18, the ASEE improved proportionally to the scaling degree. The ASEE is not the only improvement, the completion time of each benchmark decreased as the scaling increased. For instance, in the case of (S) instances, when we scaled from 1(S) to 2(S), 2(S) to 4(S), and 4(S) to 8(S), the completion time improved by 2.2%, 13.93%, and 8.21% respectively. In order to understand this behavior, we used Oprofile to inspect the last level of cache misses LLC_MISSES. As we increased the number of VMs, the percentage of LLC_MISSES decreased leading to the improvement in completion time. When there are many VMs running, then the allocation of memory to each VM is stricter. Therefore, having a smaller memory means the data location is less sparse leading to increased chances of subsequent data accesses to be available in cache.

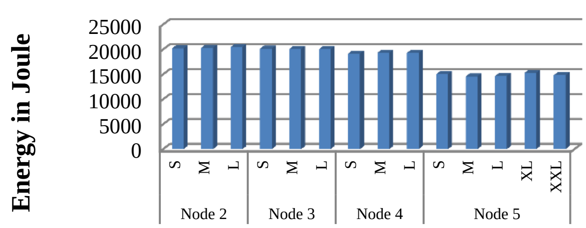

Figure 17: Energy consumption for Matrix Stressmark.

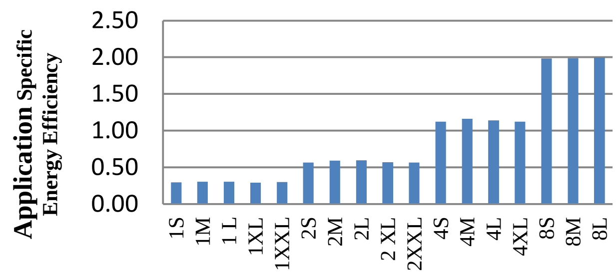

Figure 18: Application Specific Energy Efficiency of running Matrix Stressmark based on virtual machine scaling.

V. IMPLICATIONS

Based on our observations, we derived a list of implications which can be used to increase the overall energy efficiency of a datacenter.

1. Each node within a cloud has different ASEE. This observation can be useful to optimize the VM scheduler within data centers where the assignment of VMs to the available nodes can be based on their ASEE rank.
2. The more resources used by a single application, the better its ASEE and its performance. However, based on the type of application, after reaching a certain VM size, the law of diminishing returns applies where the addition of resources can negatively affect the ASEE. Due to the fact that the number of physical resources is finite, it will be interesting when building ASEE models for scheduling resources, to develop an algorithm which develops equilibrium between available resources and ASEE. In addition, since accurate prediction of resources needed per application load can directly affect the energy efficiency of the data centers, research related to this topic will have a stronger impact because they will not simply improve the application's performance but they can increase the energy efficiency as well.
3. Another valuable observation is the importance of accurate predictability of resource needs especially for e-commerce applications. Having Service Level Agreements (SLAs) which guarantee low number of failed requests is not only in the best interest of the clients but also for the data center service provider. Having high requests failure can cost the client loss in business revenues and can spike the operating costs of the provider. As a result, automatic scaling of VMs becomes essential and algorithms independent of application types and application specific performance metrics can instead monitor energy spikes in order to determine scaling needs and perform them accordingly.
4. In [20], they suggest a power management model for cluster-wide power powering on and off cluster nodes based on the cluster's overall load in order to reach energy efficiency. It will be an interesting direction to pair such mechanism with our ASEE ranking in order to use the most efficient nodes and power off the least

efficient ones. Thus, optimizing the overall energy efficiency of the data center.

VI. Related Work

WattApp [21] discusses the need for application-aware power meter for shared data centers where they took the application parameters (e.g. throughput) into consideration when building their power modeling framework. Therefore, they found a linear relationship between marginal power and marginal application throughput. Their model differs from ours because even though they dealt with heterogeneity in application, they did not consider the heterogeneity of servers as we did.

In [22], they discuss a VM-level power utilization metering and explore the feasibility and challenges in black box monitoring of the power usage of VM. They experimentally observed that there is a substantial rise in power consumption when increasing the cores. Though the paper deals with modeling the power for VM they did not consider the impact of energy efficiency.

One of the earlier works in this related field is discussed in [23] which provide energy distributed accounting on vertical structured OS with Virtual machines. They provide a frame work for managing energy in multilayered OS and accounts recursive energy consumption spent in virtualization layer of driver components.

[24] is one of the similar work which deals with evaluating energy efficient cloud on a multicore platform. Their consideration includes only the cores and evaluate only with CPU intensive benchmarks and the impacts of energy consumption during migration of VM's.

Many related works such as [25] are built upon the fact that the energy consumption scales linearly with the processor and did not consider the impact of memory associated with it. Cloud resources are not sole dependent upon various types of cores but also upon various ranges of memory. From [25] it is evident that energy consumption can be reduced when two or more tasks are consolidated as opposed to be solely assigned to one resource. But the performance hit of such task was not considered.

VII. Conclusion

In this paper, we defined the notion of ASEE, Application Specific Energy Efficiency and we performed experimental analysis while running three different application types within a heterogeneous data center. We concluded that the application type and its workload, node type, and VM size all have an impact on the energy efficiency of each node. Finally, using the results of our experiments can be used in order to improve the overall energy efficiency of a datacenter. Our future work focuses on updating Eucalyptus to have an automatic ASEE profiler to determine the ASEE ranking for each node upon its addition to the cluster and then updating the scheduler to assign VMs to the available nodes based on their ASEE ranking.

Acknowledgment

This work is in part supported by NSF Career Award CCF-0643521. The views expressed herein are those of the authors and do not necessarily reflect the views of NSF.

References

[1] L. A. Barroso, J. Dean, and U. Holzle, "Web search for a planet: The Google cluster architecture,"Micro IEEE, vol 23, 2003, pp.22-28

[2] P. Bohrer, D. Cohn, E. Elnozahy, T. Keller, M. Kistler, C. Lefurgy, R. Rajamony, F. Rawson and E. V. Hensbergen."Energy Conservation for Servers", IEEE Computer Society, vol 36, 2003, pp. 39-48.

[3] A.Greenberg ,J.Hamilton , A.David ,p.Patel, "The cost of a cloud: research problems in data center networks", ACM SIGCOMM Computer Communication, vol 39, 2009, pp. 68-73.

[4] Y. Lee, A. Zomaya, "Energy efficient utilization of resources in cloud computing systems" , J. of Supercomputing, vol 13, 2010, pp. 243-254.

[5] J. Koomey , "Estimating total power consumption by servers in the U.S. and the wold". Lawrence Berkely National Laboratory, http://enterprise.amd.com/Downloads/svrpwrusecompletefinal.pdf.

[6] TPC-W http://www.tpc.org/tpcw/

[7] P. Chen, S.J. Cokus, M.Pellegrini , "BS Seeker: precise mapping for bisulfite sequencing" Journal of BMC, vol 11,2010, pp. 123-133.

[8] Data Intensive Systems (DIS) Benchmark Performance Summary. http://www.dtic.mil/cgi-bin/GetTRDoc?AD=ADA418752

[9] B. Gaeke, P. Husbands, X. Li, L. Oliker, K. Yelick, R. Biswas. "Memory-Intensive Bencmarks: IRAM vs. Cache-Based Machines. International Paralleland Distributed Processing Symposium, 2002.

[10] D.Nurmi et al., "The Eucalyptus Open-Source CloudComputing System," Proc. IEEE Symp.on Cluster Computing and Grid, 2009, pp. 124-131.

[11] Watts up? Meters. www.wattsupmeters.com

[12] Oprofile http://oprofile.sourceforge.net/about/

[13] Java TPC-W Implementation http://pharm.ece.wisc.edu/tpcw.shtml

[14] J. Rao, C.Xu, "CoSL: a coordinated statistical learning approach to measuring the capacity of multi-tier websites", Proc. IEEE Symp.on Parallel and Distributed Processing, April 2008, pp 1–12.

[15] BS Seeker http://pellegrini.mcdb.ucla.edu/BS_Seeker/BS_Seeker.html

[16] B.Langmead et al,"Ultrafast and memory-efficient alignment of short DNA sequences to the human genome", The Journal of Genome Biology, vol 25, 2009 , pp. 423-433.

[17] D. Garcia and J. Garcia. TPC-W E-Commerce Benchmark Evaluation. *IEEE Computer*, pages 42–48, February 2003.

[18] J.S.Chase, C.Darrell, Anderson, P. N. Thakar and A.M. Vahdat. "Managing Energy and Server Resources in Hosting Centers," Proc. IEEE Symp.on Operating System Principles, 2001 , pp. 103-116.

[19] R.Bradshaw, P.Zbiegiel,"Experiences with eucalyptus: deploying an open source cloud". Proc. IEEE Symp.on large installation system administration, pp 1–16.

[20] C. Rusu, A. Ferreira, C. Scordino, A. Watson, R. Melhem andD. Mosse," Energy-Efficient Real-Time Heterogeneous Server Clusters", Proc. IEEE Symp.on Real-Time and Embedded Technology and Applications. 2006. pp. 418-428.

[21] R.Koller, A.Verma, A.Neogi, "WattApp: An Application Aware Power Meter for Shared data Centers". Proc. IEEE Symp. On Autonomic computing (ICAC '10) 2010 , pp. 31-40.

[22] B.krishnan, Hrishikesh and A. Karsten, "VM Power Metering: Feasibility and Challenges", ACM SIGMETRICS ,vol 38 , 2010, pp. 56-60.

[23] J. Stoess ,L. Bellosa ,"Energy Management for Hypervisor-based Virtual Machines", Proc. IEEE Symp.on USENIX Annual Technical Conference Jun 2007, pp. 28-37.

[24] L. Lefevre , A. Orgerie ,"Design and evaluating an energy efficient Cloud", Journal of Supercomputing ,Vol 51, Mar 2010, pp. 352-373.

[25] Y.Lee, A.Y.Zomaya ,"Energy efficient utilization of resources in cloud computing systems ", J. of Supercomputing ,Vol 51, 2010, pp. 374-387.