

Article

Energy Aware Virtual Machine Scheduling in Data Centers

Yeliang Qiu ^{1,2}, Congfeng Jiang ^{1,2,*}, Yumei Wang ^{1,2}, Dongyang Ou ^{1,2}, Youhuizi Li ^{1,2} and Jian Wan ^{1,3,*}

¹ Key Laboratory of Complex Systems Modeling and Simulation, Ministry of Education, Hangzhou 310018, China; qiuyeliang@hdu.edu.cn (Y.Q.); wyumei1@gmail.com (Y.W.); oudongyang1@gmail.com (D.O.); huizi@hdu.edu.cn (Y.L.)

² School of Computer Science and Technology, Hangzhou Dianzi University, Hangzhou 310018, China

³ School of Information and Electronic Engineering, Zhejiang University of Science and Technology, Hangzhou 310023, China

* Correspondence: cjiang@hdu.edu.cn (C.J.); wanjian@zust.edu.cn (J.W.);
Tel.: +86-5718-6919-113 (C.J.); +86-5718-6919-111 (J.W.)

Received: 27 December 2018; Accepted: 13 February 2019; Published: 17 February 2019



Abstract: Power consumption is a primary concern in modern servers and data centers. Due to varying in workload types and intensities, different servers may have a different energy efficiency (EE) and energy proportionality (EP) even while having the same hardware configuration (i.e., central processing unit (CPU) generation and memory installation). For example, CPU frequency scaling and memory modules voltage scaling can significantly affect the server's energy efficiency. In conventional virtualized data centers, the virtual machine (VM) scheduler packs VMs to servers until they saturate, without considering their energy efficiency and EP differences. In this paper we propose EASE, the Energy efficiency and proportionality Aware VM SchEduling framework containing data collection and scheduling algorithms. In the EASE framework, each server's energy efficiency and EP characteristics are first identified by executing customized computing intensive, memory intensive, and hybrid benchmarks. Servers will be labelled and categorized with their affinity for different incoming requests according to their EP and EE characteristics. Then for each VM, EASE will undergo workload characterization procedure by tracing and monitoring their resource usage including CPU, memory, disk, and network and determine whether it is computing intensive, memory intensive, or a hybrid workload. Finally, EASE schedules VMs to servers by matching the VM's workload type and the server's EP and EE preference. The rationale of EASE is to schedule VMs to servers to keep them working around their peak energy efficiency point, i.e., the near optimal working range. When workload fluctuates, EASE re-schedules or migrates VMs to other servers to make sure that all the servers are running as near their optimal working range as they possibly can. The experimental results on real clusters show that EASE can save servers' power consumption as much as 37.07%–49.98% in both homogeneous and heterogeneous clusters, while the average completion time of the computing intensive VMs increases only 0.31%–8.49%. In the heterogeneous nodes, the power consumption of the computing intensive VMs can be reduced by 44.22%. The job completion time can be saved by 53.80%.

Keywords: data center; energy efficiency; energy proportionality (EP); scheduling; virtual machine (VM)

1. Motivation

While server energy efficiency has improved much since last decade due to technical breakthroughs, the explosion of cloud computing, online services, social media networks, and Internet

traffic has been the catalyst for an approximately 4.5% increase in worldwide server shipment and data center deployment annually from 2011 to 2015 on average [1]. Such an increase in server installation also contributes to the increase of energy consumption of data centers if the server energy efficiency remains. Nowadays, energy consumption of servers is becoming the major concern for data center operation and management [2]. Figure 1 plots prediction on electricity consumption of data centers in U.S. from some data sources.

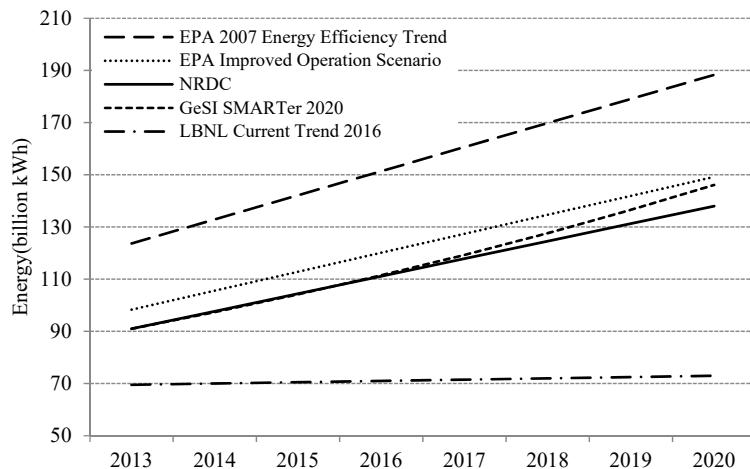


Figure 1. Prediction on electricity consumption of data centers in U.S. Data extrapolated on [1,3–5]. GeSI SMARTer 2020 data is extrapolated by 7% annual increase rate of Global data centers emission (GtCO₂e) with the baseline of NRDC 2011 data.

Although today's servers are more energy efficient than ever before, their power consumption under a low or idle workload accounts for more than 20% of the power at 100% workload for typical servers. Current data centers (DCs) of giant Internet companies like Google, Facebook, Microsoft, and Alibaba are highly energy efficient and have reached the power usage effectiveness (PUE) of 1.1 [6,7]. However, for small- and medium-sized data centers, they are not always in a high-load operating state and thus the servers still consume a lot of energy when they are idle. These small- and medium-sized data centers account for more than 95% percentage of total installed servers [3,8]. In such smaller scale data centers, server utilization is much lower than large scale data centers and their power use effectiveness (PUE) is also lower. Moreover, incoming requests are intermittent and workloads fluctuate on a daily or weekly basis. Therefore, energy proportional computing has emerged in both industry and academia. Energy proportionality (EP) is a proposed metric to measure the relationship of server power consumption and its utilization [9–12]. Ideally, the power consumption of the server with EP = 1.0 should be fully proportionate to its utilization when the server is working. For example, the power at idle or low utilization ideally should also be almost zero or proportional compared to power at 100% workload. That is, the power consumption of the server with utilization of 10% should be one-tenth of its power consumption with utilization of 100%. In a typical server system, the processor is more energy proportional than memory and other components [13,14]. Different hardware configuration may also affect the energy proportionality of the whole server system. Therefore, good knowledge of a server's energy proportionality and efficiency can help workload placement and job scheduling to achieve power and/or energy minimization in data centers.

Figure 2 gives the EP values and the SPECpower score of 509 commercial servers from the public valid SPECpower_ssj results ordered by hardware availability year (until 2017Q4 on 7 December 2017) [15]. The authors will use SPECpower and SPECpower_ssj interchangeably in the remainder of this article. We can observe that the energy efficiency and EP of commercial servers have improved in the last decade from the beginning of the first SPECpower_ssj benchmark results released in 2007.

For example, the average EP improved significantly from 0.32 to 0.91 (2007 to 2017); the average server energy efficiency improved from 489 to 11742 (ssj_ops per watt).

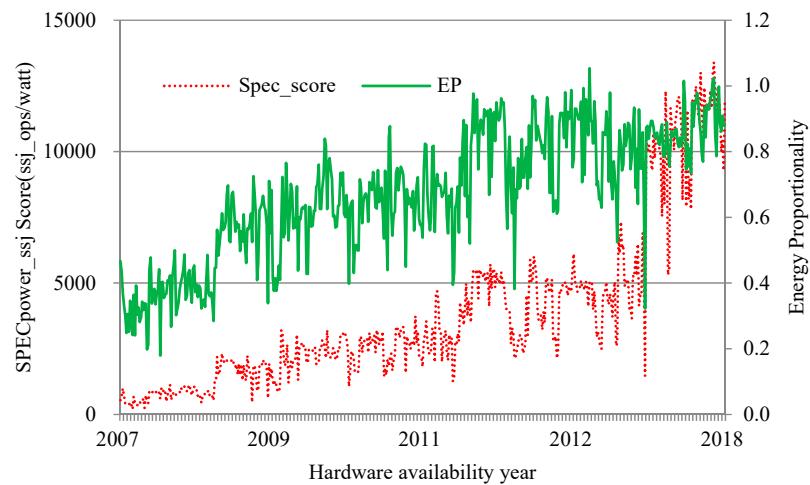


Figure 2. Energy Proportionality (EP) and energy efficiency (EE) trend of commercial servers.

With improvements in energy efficiency and EP, recent commercial servers can achieve their peak energy efficiency at a lower utilization than 100%. Table 1 is a sample SPECpower_ssj result published in 2017Q4 [15]. The server achieves peak energy efficiency 13,845 ssj_ops per watt at 60% utilization. The energy efficiency is 12,424 ssj_ops per watt at 100% utilization.

Table 1. Sample result of SPECpower_ssj2008.

Target Load	Performance		Average Active Power (W)	Performance to Power Ratio
	Actual Load	ssj_ops		
100%	99.70%	11,725,627	944	12,424
90%	90.00%	10,580,169	851	12,427
80%	80.00%	9,411,437	716	13,151
70%	70.10%	8,241,170	598	13,779
60%	60.00%	7,056,523	510	13,845
50%	50.00%	5,876,594	431	13,621
40%	40.10%	4,709,344	373	12,614
30%	30.00%	3,527,435	324	10,895
20%	20.00%	2,352,157	277	8498
10%	10.00%	1,173,811	228	5154
Active Idle		0	82.9	0
$\sum \text{ssj_ops} / \sum \text{power} =$				12,120

Figure 3 illustrates the percentage of occurrences at peak energy efficiency per year. From 2004 to 2009, the servers achieved their peak energy efficiency at 100% utilization. The utilization spots for peak energy efficiency have varied since 2010. Most importantly, the latest manufactured servers tend to achieve peak energy efficiency at non-100% utilization. Among 30 published SPECpower_ssj results in 2017 (through 7 December 2017), 27 servers had hardware availability in 2017. Fourteen of the servers achieved peak energy efficiency at 70% utilization. Using VM scheduling and migration to keep the cluster server running at peak efficiency, we can improve the overall energy efficiency of the cluster and the whole data center.

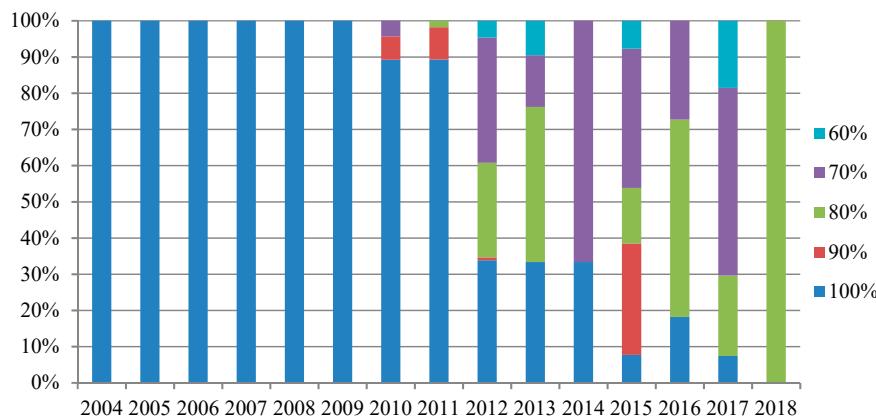


Figure 3. Distribution of server utilization spot with peak energy efficiency.

In data centers, servers are often provisioned and configured for peak workload and peak performance. Underutilized resources waste energy when working in lower utilization for non-energy proportional servers. Virtualization consolidation has been widely deployed in current data centers for service consolidation and energy reduction through resource multiplexing. In a multitenant cloud computing environment, cloud service providers pack as many VMs as possible to keep the physical server running at higher resource utilization regardless of the underlying server's EP. The packing policy masks the energy efficiency potentials of recent manufactured servers. In these cases, EP-blind scheduling wastes energy as it deteriorates system performance and service quality.

In this paper, we propose EASE, the energy efficiency and proportionality aware VM scheduling framework. Before scheduling, EASE identifies server's energy efficiency and EP by executing customized computing intensive, memory intensive, and hybrid benchmarks on the server. The rationale of EASE is to schedule VMs to servers to keep them working around their peak energy efficiency point, i.e., the near optimal working range. When workload fluctuates, EASE re-schedules or migrates VMs to other servers to make sure that all the servers are running as near their optimal working range as they possibly can. EASE can also calculate the global energy efficiency and energy savings of the data center through energy consumption data via an intelligent platform management interface (IPMI) enabled power data aggregation. Extensive experiments on a cluster were conducted, including computing intensive, memory intensive, and mixed workload VMs. The experimental results show that power consumption can save around 37.07%–49.98% in the homogeneous node cluster. The average completion time of the computing intensive workload increases 0.31%–8.49%. In the heterogeneous nodes, the power consumption of the computing intensive jobs can be reduced by 44.22%. In addition, the job completion time can be saved by 53.80%. The results presented in this article provide useful insight for the power and energy management of virtualized data centers, especially for heterogeneous servers. At the same time, EASE can serve as a reference for green data center operations based on server energy efficiency, EP, and workload characteristics.

The rest of the paper is organized as follows. In Section 2 we describe energy proportionality and the PEEP metric. In Section 3, we describe EASE, including its methodology and components for VM scheduling and migration. We provide experiment results, observations, and insights for different platforms with various configurations in Section 4. In Section 5 we review some related work on server energy proportionality and identify our unique approach and results. We conclude the paper with some future research directions in Section 6.

2. Energy Proportionality and the PEEP Metric

2.1. SPECpower Benchmark

Currently, the most authoritative industry standard organizations for performance measurements are the Transaction Processing Performance Council (TPC), the Standard Performance Evaluation Corporation (SPEC), and the Storage Performance Council (SPC). SPEC released the first industry standard benchmark, SPECpower_ssj, which measures power consumption related to performance for servers. Workload of SPECpower evaluates energy efficiency and performance running server-side Java applications for small- and medium-sized servers at graduated utilization levels. SPECpower reports power consumption for servers at different utilization levels (from 100% utilization to idle in 10% intervals) over a set time. The server's overall energy efficiency score is calculated by division of the sum of 11 average active power values (including idle) and the sum of 10 performances to power ratio values. Table 1, which is a sample result of SPECpower_ssj2008 published in 2017Q4, represents a server with hardware availability of Aug-2017. In a SPECpower result, the utilization is not the CPU utilization. The performance to power ratio at each target load level is often called the server's energy efficiency at that load level with unit of ssj_ops per watt. For the last column, performance to power ratio in Table 1, 12,424 is the energy efficiency at peak utilization (100%), 13,845 is the peak energy efficiency, and the server's overall energy efficiency score is 12,120.

2.2. Energy Proportionality Calculation And Its Implications

This article uses the EP metric proposed by Ryckbosch et al. in [10] as shown in Equation (1).

$$EP = 1 - \frac{Area_{real} - Area_{ideal}}{Area_{ideal}}. \quad (1)$$

where: $Area_{real}$ is the area covered by a real server's power curve (green zone in Figure 4). $Area_{ideal}$ is the area covered by an ideal server (black solid line zone in Figure 4).

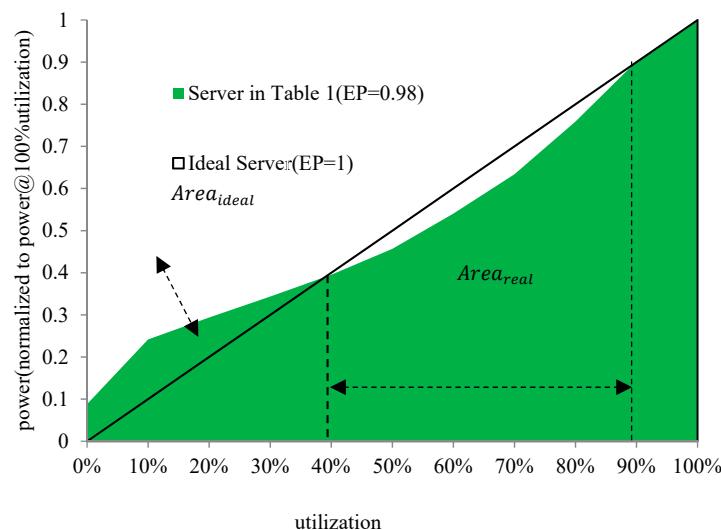


Figure 4. Comparison of EP curve (or power curve) of the server in Table 1 and an ideally energy proportional server (power normalized to power at 100% utilization).

We can approximate the EP value of the server as in Table 1 by calculating 10 subareas (10 utilization level from 10% to 100%) in its power-utilization curve in Figure 4. Without loss of generality, we can estimate any server's EP value given its power curve. From Equation (1) and Figure 4 we can see that higher EP means that: (1) the server has lower idle power; (2) it consumes less power than an ideal server at specific utilization level; or (3) both.

In our previous work [4], we identified the relationship between energy efficiency, EP, and processor architecture, hardware configuration, and server performance. Moreover, for servers with higher EP ($EP > 0.8$), the non-linearity of energy consumption corresponding to its performance is becoming more dominant. At the same time, the energy efficiency curve of the server is nonlinear.

Figure 3 shows that the peak energy efficiency spot from the latest server has shifted from 100% utilization to lower utilization. Table 2 lists the statistics of peak energy efficiency occurrence of all 509 servers with SPECpower_ssj results since 2007. Although 334 servers achieve their peak energy efficiency at 100% utilization, 95.2% of these servers (i.e., 318 servers) were manufactured before 2013. Only 4.8% (i.e., 16 servers) were manufactured between 2013 and 2018.

Table 2. Peak energy efficiency occurrence by utilization spot.

Utilization Spot Where Peak Energy Efficiency Occurs	Count	Count	Total
100%	334	65.49%	
90%	16	3.14%	
80%	65	12.74%	510
70%	81	15.88%	
60%	14	2.75%	

Note: One server achieves peak energy efficiency at both 80% and 90% utilization. Therefore, 509 servers have 510 peak energy efficiency spots.

The EP curves from 509 servers show that:

- (1) For all 334 servers that achieve peak energy efficiency at 100% utilization, their EP curves are always above the EP curve of the ideal server.
- (2) For the remaining 175 servers that achieve peak energy efficiency at non-100% utilization, six servers are always above the ideal server, 113 servers intersect with the ideal curve once, and 56 servers intersect with the ideal curve twice. Their energy efficiency and EP are listed in Table 3.

Table 3. Energy efficiency and proportionality of servers with non-100% peak energy efficiency.

Metrics	Above the Ideal Curve	Intersect Once	Intersect Twice
count	6	113	56
avg. EP	0.76	0.886	0.860
med. EP	0.76	0.891	0.855
avg. EE	5264	6781	6756
Med.EE	3849	5146	5318
avg. PEEP	1.003	1.080	1.050
med. PEEP	1.003	1.073	1.043
total		175	

We give some representative EP curves in Figure 5. From Figures 4 and 5 we derive two implications as follows:

- **Implication #1:** For a server whose power curve intersects with the ideal power curve before 100% utilization, there exists at least one work range where its energy efficiency is larger than the ideal server. This is defined as an optimal working range (OWR).

When in the utilization interval (40%, 90%), the server in Table 1 has an energy efficiency which is larger than an ideal server (see Figure 4). Table 4 lists its power and percentage compared to an ideal server. For example, at 50% utilization, its power consumption normalized to 100% utilization is 45.7%. For an ideal server, its power consumption at 50% utilization should also be 50%. The working range in Table 4 identified by the EP curve is the range that the server should work in.

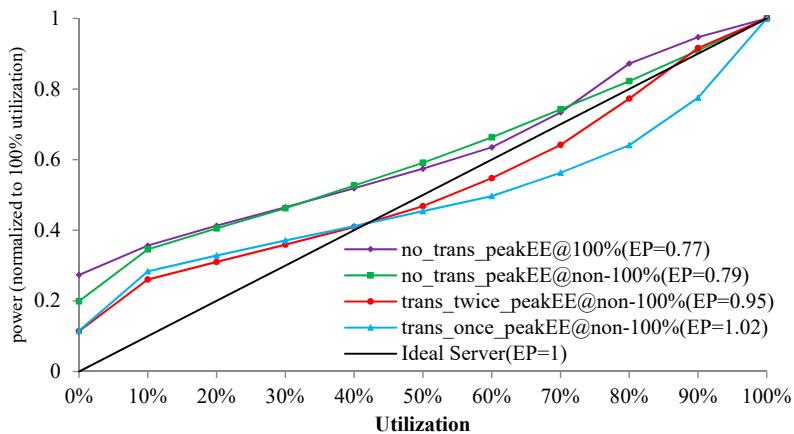


Figure 5. Typical EP curves of five real servers with peak energy efficiency at 100% and non-100% utilization.

Table 4. More energy efficient working range of the server in Table 1.

Utilization	Power (Normalized to Its 100% Utilization)	Power (Normalized to Ideal Server)	Peak Energy Efficiency Spot
40%	39.5%	98.8%	
50%	45.7%	91.3%	
60%	54.0%	90.0%	
70%	63.3%	90.5%	Yes
80%	75.8%	94.8%	

- **Implication #2:** The server achieves its peak energy efficiency at the utilization spot where its power is the lowest normalized to ideal server.

Table 4 shows that the server at 60% utilization consumes only 90% power compared to an ideal server, which is the lowest during all working range. Therefore, the server achieves its peak energy efficiency at 60% utilization (see Table 1).

2.3. The PEEP Metric

The PEEP metric is the ratio of peak energy efficiency over energy efficiency at peak utilization [16]. For example, the PEEP value of the server with SPECpower results in Table 1 is $13,845 / 12,424 = 1.114$. The PEEP value indicates the offset of a server's peak energy efficiency with respect to its energy efficiency at 100% utilization. PEEP values equal to 1.0 mean that the server performs peak energy efficiency at peak utilization (100% utilization). A PEEP value larger than 1.0 means that the server performs peak energy efficiency other than 100% utilization. Table 5 lists the counts of servers with a different utilization spot where peak energy efficiency occurs. Among 175 servers with $\text{PEEP} > 1.0$, 81 servers achieve their peak energy efficiency at 70% utilization.

Table 5. Count of utilization where peak energy efficiency occurs.

Utilization Where Peak Energy Efficiency Occurs	Count	Total
60%	14	
70%	81	
80%	65	510
90%	16	
100%	334	

Figure 6 draws the average and median values of PEEP, EP, and the SPECpower score of the 509 servers. These are grouped by their utilization spot where peak energy efficiency occurs. Servers

achieving their peak energy efficiency at lower utilization have higher PEEP and EP values than those at higher utilization.

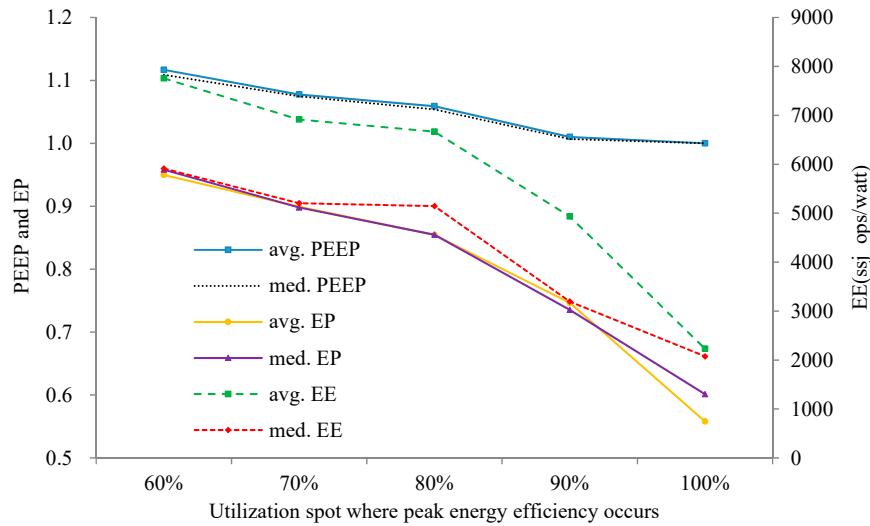


Figure 6. PEEP, EP, and EE with respect to utilization where peak energy efficiency occurs.

Figure 7 presents the EP and PEEP value for the 175 servers whose PEEP values are larger than 1.0. Figure 7 can be quantified as 0.85 based on Equation (2):

$$\text{PEEP} \approx 2.77 \cdot (\text{EP})^2 - 4.18 \cdot \text{EP} + 2.59 \quad (2)$$

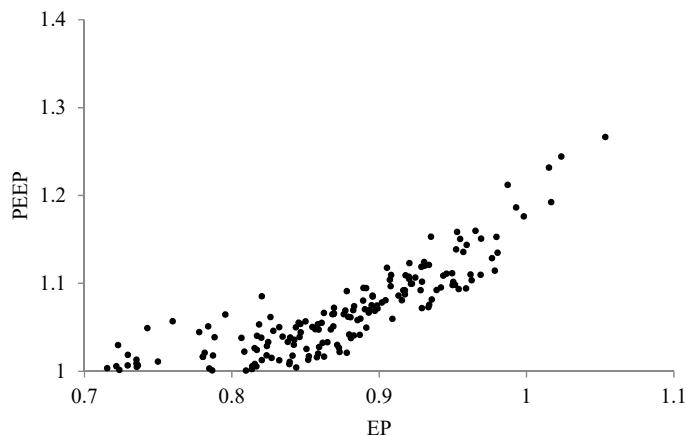


Figure 7. Scattering of PEEP and EP.

Figure 7 also suggests that servers whose peak energy efficiencies do not coincide at 100% utilization will have higher EP values, higher PEEP values, and vice versa.

Equations (4) and (5) show that power saving may correlate with a server's EP.

Figure 8 shows the 175 servers whose PEEP values are larger than 1.0. This can be quantified as 0.84 by Equation (3):

$$\text{PS} \approx 2.07 \cdot (\text{EP})^2 - 3.06 \cdot \text{EP} + 1.15 \quad (3)$$

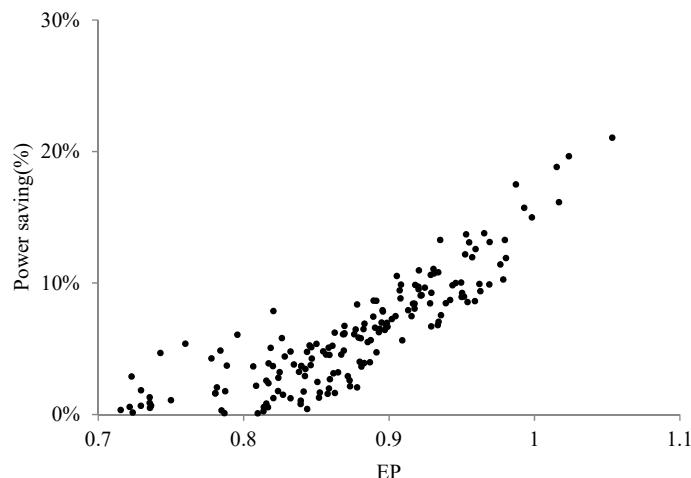


Figure 8. Scattering of power saving and EP.

This article shows that servers should perform in an optimal working range. However, for a data center with thousands of servers, it is more important to select which server to run multiple candidates, especially when the candidate servers have optimal working ranges. Equations (2)–(4) show that choosing a server with higher EP and PEEP values can save more power. This is the basic rationale of workload placement and scheduling in data centers with heterogeneous servers whose EP is also heterogeneous.

3. Energy Efficiency and Proportionality Aware Virtual Machine Scheduling

We provide the energy aware virtual machine scheduling flow chart in EASE in Figure 9.

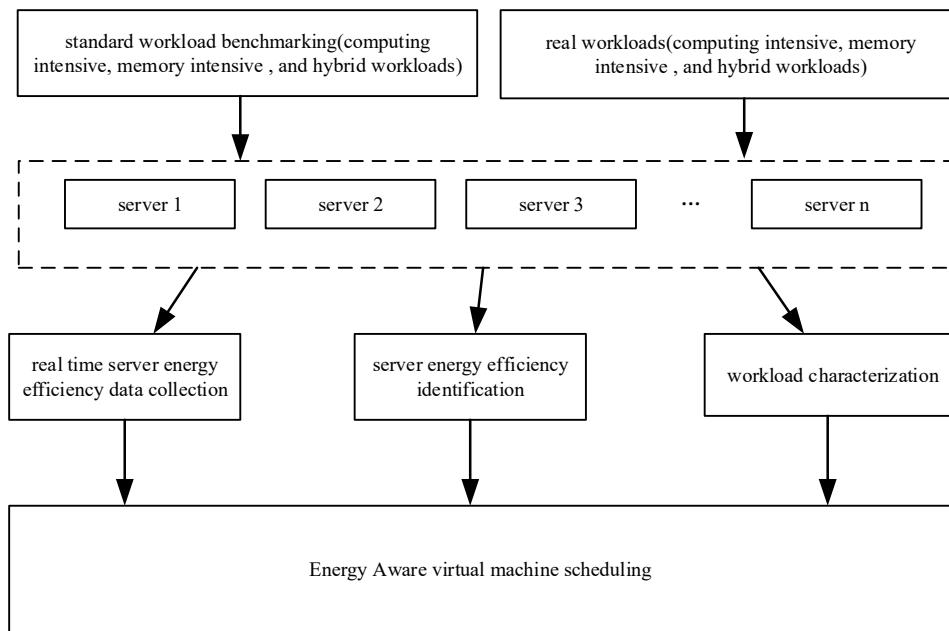


Figure 9. Flow chart of energy aware virtual machine scheduling.

3.1. Server Energy Efficiency Identification

Hardware configuration of different servers in the data center is usually heterogeneous. The operating system and workloads on each server also differ. The server's energy efficiency differs when running workloads under different configurations, including processor DVFS, memory bank powering mode, NIC transmitting settings, etc. [17]. To implement the energy efficiency and

proportionality aware VM scheduling, we must understand the energy efficiency and proportionality of each server in data centers. Although SPECpower_ssj score is an important reference source, not all servers in a data center have their own SPECpower_ssj score. Moreover, data center workload is very diverse and dynamic. This is not represented comprehensively by the SPECpower_ssj benchmark. Therefore, the authors have developed a customized benchmark for server energy efficiency identification. It has the following merits:

- (1) **Representative of Data Center Workloads:** We use SPECpower_ssj, PrimeSearch [17], STREAM [18], and their mixture. The workload is also proportional to the server's hardware configuration. It varies with multiple workload intensities to simulate a real multitenant virtualized cloud computing environment.
- (2) **Easy Implementation:** Recent commercial servers are equipped with IPMI support through which we get the real time power consumption from embedded sensors in a server mainboard. This includes system power, processor power, and memory bank power. With simulation workload- and IPMI-based power consumption, we can get the server's energy efficiency and proportionality under different workload types and workload intensity levels. In the authors' implementation, IPMI-based power data acquisition has negligible overhead on the physical server's system utilization and extra power consumption (< 0.25 watt).

For each server, we find the energy efficiency and proportionality by running computing intensive, memory intensive, and hybrid workload, respectively.

- (1) **Computing Intensive Workload:** The authors use a prime number computation program, namely, PrimeSearch, as the computing intensive workload. One execution of PrimeSearch will calculate and search prime numbers in 10 intervals (1,200000), (1,400000), (1,600000), (1,800000), (1,1000000), (1,1100000), (1,1200000), (1,1300000), (1,1400000), (1,1500000), respectively. These 10 intervals are 10 subsearching tasks. The completion time of the interval search is calculated and the sum of the completion time of 10 subtasks is calculated as the task completion time of one Primesearch execution.
- (2) **Memory Intensive Workload:** The authors use STREAM as the memory intensive workload. It is a single synthetic benchmark program to measure the amount of memory and corresponding computational rate of the simple vector kernel. The CPU of the computer system runs much faster than memory. As this progresses, more program performance will be limited to memory bandwidth rather than CPU computing performance.
- (3) **Hybrid Workload:** The authors use SPECpower_ssj2008 as the hybrid workload to stress the server's components, including CPU and memory synthetically.

3.2. VM Workload Characterization

Previous sections discussed how each server may have different energy efficiency and proportionality. In addition, each one must work in its own optimal working range. Therefore, it is necessary to determine each VM's workload type before scheduling it to the target physical server. In doing so, we can ensure that it is running in the optimal working range.

It is necessary to identify the workload characteristics of the VM. It is impossible for a VM to consume only CPU resources without memory resources in a real cloud computing environment. Therefore, an empirical setting for VM type identification is used in Table 6. This setting is based on empirical observations and experiences in testing environments and real data sets [19]. They can be updated by customized parameters. In the future version of EASE, the scheduler could learn and adjust it according to the actual workload trace data in data centers. The workload characterization procedure is done by the system profiling techniques like *sysstat* tool to collect the CPU and memory usage which indicates the workload type and intensity. For each VM, EASE will undergo workload characterization procedure by tracing and monitoring their resource usage including CPU, memory, disk, and network and determine whether it is computing intensive, memory intensive, or hybrid workload.

Table 6. Virtual machine (VM) workload type setting.

Type	CPU Utilization	Memory Utilization
Computing Intensive Workload	>70%	<30%
Memory Intensive Workload	>20% and <50%	>60%
Hybrid Workload	>30% and <60%	>30% and <60%

3.3. Virtual Machine Scheduling and Migration

In traditional data centers, VMs are packaged into many servers. The underutilized servers are powered off. As described, the latest manufactured servers achieve their peak energy efficiency at non-100% utilization. The effectiveness of reducing power consumption by increasing resource utilization will become increasingly unsatisfactory. Moreover, servers with low EP can be grouped together to run at their optimal working range to achieve better EP in cluster wide. Therefore, in EASE, we schedule VMs to servers with the best optimal working range.

EASE contains the following components:

1. **Server Monitoring and Power Data Collection:** CPU and memory contribute to most of the dynamic power consumption in a server. Therefore, the monitoring module of EASE is to collect real-time power consumption, CPU utilization, memory utilization, and other system access and activities statistics. The physical server provides the physical machine CPU, memory, input/output (I/O), and other hardware resource utilization using the `sysstat` utility.
2. **Benchmark Execution:** EASE runs benchmark tests on dedicated servers under multiple software and hardware configurations. It also collects performance data.
3. **Energy Efficiency and Proportionality Calculation:** Based on system monitoring data and benchmark execution results, the energy efficiency and proportionality are calculated in real time.
4. **VM Scheduling:** When EASE gets each server's current running status and optimal working range, it determines if the server is overloaded or underloaded. For new VMs, EASE looks for the appropriate target physical server to keep them working in an optimal working range. For existing VMs, EASE first characterizes the machine's workload type (i.e., computing intensive, memory intensive, or hybrid) and attempts to migrate the VM to an appropriate target physical server. Servers with different hardware and software configurations have different energy efficiency and proportionality. Therefore, if there are multiple servers with different resources (i.e., resource type, capacity, and configuration), EASE first sorts the servers with available resources by resource type in descending order of energy efficiency, EP, and PEEP values. Next, it selects a target host according to the priority that the target host can reach the peak energy efficiency or energy efficiency will be higher in its optimal working range.

The VM scheduling algorithm is listed in Figure 10. We assume that the scheduler can always find a target server receiving VM with enough free CPUs, memory, and disk to run the migrated VM.

To ensure that the data center is running in optimal working range, EASE monitors and evaluates all the servers. As it checks that it is running in its optimal working range, it may migrate VMs to other servers.

Live migration is perfectly implemented in state-of-the-art VM monitoring. Therefore, the authors did not consider VM migration overheads. Moreover, the migration is only conducted inside a data center network. The quality of service-related issues (i.e., service interruption and data dependence redirection) are negligible.

The server is *overloaded* if it is not running in its optimal working range and nears the upper bound of its optimal working range. The VM on the *overloaded* server will migrate until the server returns to its optimal working range. Similarly, the server is *underloaded* if it is not running in its optimal working range and nears the lower bound of its optimal working range. The newly created VMs or to-be-migrated VMs will be scheduled to the *underloaded* server until the server is running in its optimal working range.

Algorithm 1: Virtual machine scheduling algorithm based on energy efficiency and proportionality in EASE

Input: virtual machines and available server set

Output: scheduling between VMs and servers

- Energy efficiency and proportionality calculation for each workload type (computing intensive, memory intensive or hybrid) and hardware software configurations
 - for each server in available server set
 - calculate its peak energy efficiency and optimal working range
 - sort the server by energy efficiency, EP and PEEP
 - end for
 - for each VM
 - workload characterization by server tracing and monitoring
 - schedule it to available server with highest PEEP
 - end for
 - update VM set
 - update available server set
-

Figure 10. VM scheduling in EASE.

For power minimization, the first goal is to migrate out VMs on overloaded servers. This ensures that there is no overloaded server. The second goal is to keep servers running in optimal working range. If there are multiple *underloaded* servers, we merge VMs on *underloaded* servers to use as few as possible servers.

Figure 11 lists the VM migration algorithm. The rationale to migrate the VM with the largest resource requirement on an overloaded server is to get it back to optimal working range as soon as possible with minimal migration operations. For example, if we migrate a VM with smallest resource requirement, the server may still be in overloaded status. After migration, EASE updates the original server status information. It will continue the above process until there is no overloaded server. The following section conducted experiments on real clusters with EASE-enabled VM scheduling.

Algorithm 2: virtual machines migration

Input: server set S with VM running set V

Output: server set S' with VM running set V'

1. max = optimal working range threshold //set the upper threshold of server load
 2. for each server
 3. get the load status W_i
 4. if $W_i > \text{max}$
 5. group server to SO //overloaded server set
 6. end for
 7. for S_i in SO
 8. for each vm in S_i
 9. sort each vm by resource needed descending
 10. While $W_i > \text{max}$
 11. select the vm with largest resource
 12. requirement to migrate
 13. update S_i status
 14. end for
 15. update SO
 16. end for
-

Figure 11. VM migration in EASE.

4. Experiment Results and Analysis

4.1. Experimental Platform

In order to evaluate EASE's feasibility and performance, this article conducts experiments on the following platforms (see Table 7). The physical platform and the VM run CentOS 7.2; the hypervisor is KVM/QEMU. Each VM is configured with two vCPUs, 4GB memory, and 20GB hard disk space. We developed a simulator to evaluate our algorithm's scalability in our work in progress. Due to space limitations, we only provide experiments on three machines in our lab. However, EASE can scale well in large data centers with hundreds or thousands of servers for two reasons. (1) EASE can identify server's energy efficiency and energy proportionality offline and in advance. This process only runs once and can be done automatically after server deployment before service in production. (2) The workload characterization process in scheduling algorithm can scale with flexible window size for performance tradeoff in computing time and workload categorization accuracy in large scale data centers.

Table 7. Server types and experimental platform configuration.

No.	Platform	Year of Manufacture	CPU	Total CPU Core	CPU TDP (Watt)	Memory (GB)	Hard Disk
1	Sugon A620r-G	2012	2*AMD Opteron 6272	32	115	64(8G*8) DDR3 1600MHz	4*SAS 300GB 10K rpm (RAID10)
2	ThinkServer RD640	2014	2*Intel Xeon E5-2620 v2	12	80	160(16G*10) DDR4 2133MHz	1*SSD 480GB
3	ThinkServer RD450	2015	2*Intel Xeon E5-2620 v3	12	85	192(16G*12) DDR4 2133MHz	2*HDD 4TB

The calculation interval of PrimeSearch is (0,200000), (0,400000), (0,600000), (0,800000), (0,1000000), (0,1100000), (0,1200000), (0,1300000), (0,1400000), (0,1500000). The matrix size in STREAM is 170000000, which needs 3.8GB memory for matrix manipulations.

4.2. Experimental Results

Homogeneous Cluster. Since the hardware configurations of the homogeneous servers are the same, the server's energy efficiency and proportionality are the same for the same workload. For example, Figure 12 shows the power consumption when running different workloads on server #3.

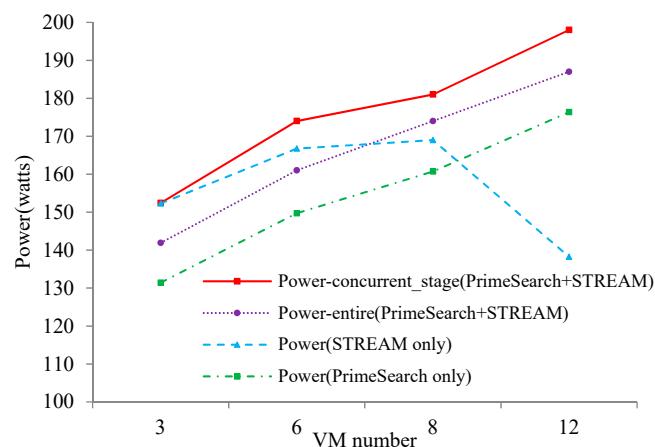


Figure 12. Power consumption of server #3.

It is important to note that when 12 VMs are running STREAM workload concurrently on server #3, the power consumption (138 watts) is lower than when eight VMs are running (each VM continues running STREAM load). This is mainly due to overheads introduced when the kernel virtual machine (KVM) is in memory virtualization, which results in a sharp decline in the efficiency of memory virtualization. A large number of STREAM computing tasks are blocked, resulting in multiple VMs in an idle and waiting state. This eventually leads to reduced server power consumption. This can be observed from the completion time of the STREAM in Figures 13–15.

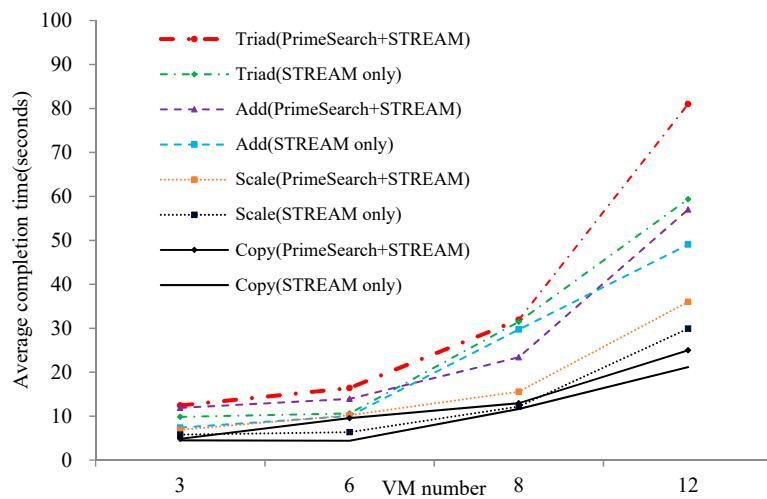


Figure 13. Average completion time of STREAM (server #3).

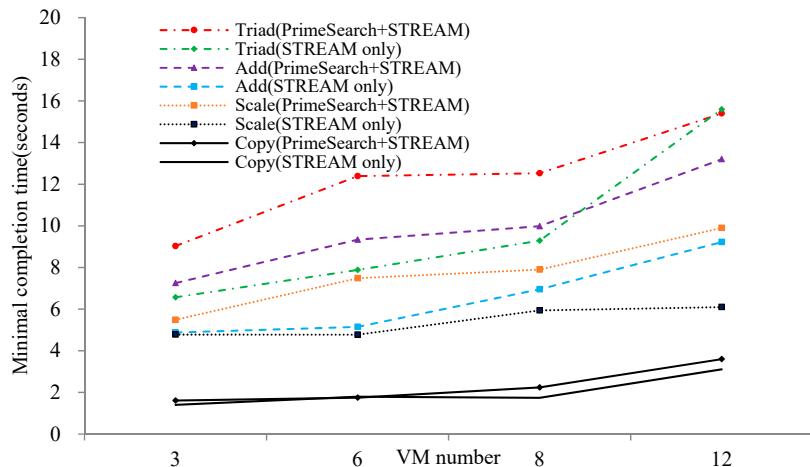


Figure 14. Minimum completion time of STREAM (server #3).

The shorter STREAM calculation time (calculation time of single Copy, Add, Scale, or Triad are generally not longer than one minute and the total run time is generally 10 min) is significantly less than the PrimeSearch completion time (3–12 VMs is about 1620 sec, 24 VMs is 2980 sec). Even in the concurrent running of VMs with STREAM and PrimeSearch, the server power consumption remains significantly less than the sum of the power consumption of two servers running separately.

Figure 16 shows the maximum memory access bandwidth for STREAM running inside each VM while running STREAM separately and running STREAM and PrimeSearch concurrently.

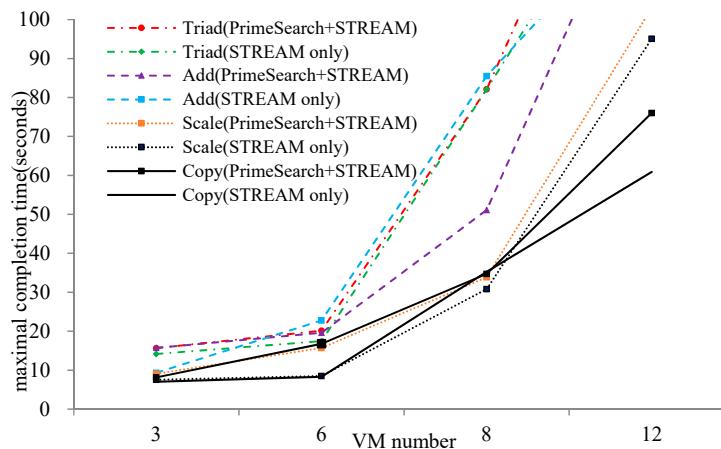


Figure 15. Maximum completion time of STREAM (server #3).

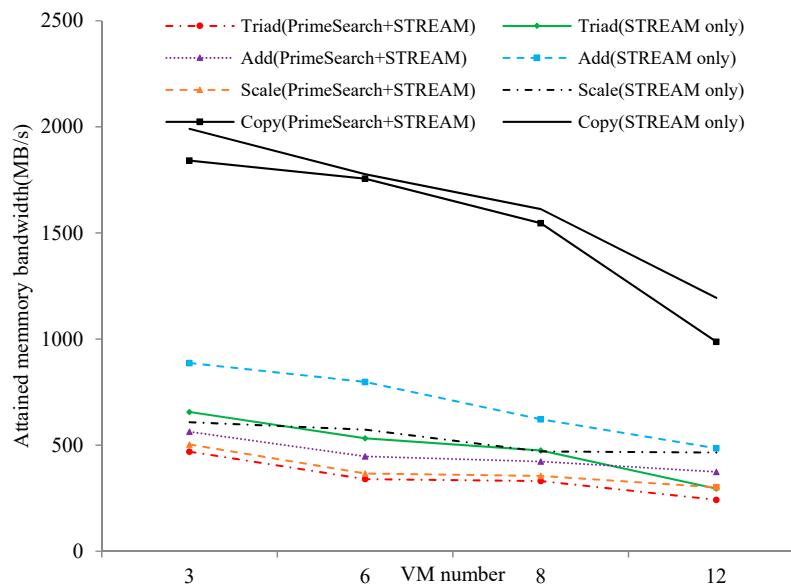


Figure 16. Maximum memory access bandwidth (server #3).

Table 8 shows the homogeneous node cluster's power consumption and performance comparison using EASE. The experimental results on server #3 show that power consumption can achieve 37.07%–49.98% savings in both the computing intensive and memory intensive workload.

Table 8. Power consumption and performance comparison of different scheduling algorithms.

Number of VMs	Power Consumption		Prime Search Complete Time	STREAM			
	Whole Running	Concurrent Phase		Bandwidth	Average Time	Min Time	Max Time
3	−49.98%	−46.28%	0.31%	−22.45	28.44%	29.00%	28.84%
6	−49.13%	−45.02%	7.53%	−27.46%	67.80%	48.15%	47.10%
8	−47.23%	−45.11%	7.52%	−22.34%	4.80%	35.15%	−7.96%
12	−40.56%	−37.07%	8.49%	−19.72%	16.75%	42.54%	56.48%

The average completion time of the computing intensive workload only increases from 0.31% to 8.49%. Although the completion time of the memory intensive workload is deteriorated (increased 4.8% to 67.8%) in most cases (except the case of eight PrimeSearch VMs and eight STREAM VMs concurrently), the maximum completion time reduced by 7.96% compared to STREAM alone. The

reason may be that when all VMs run the same memory intensive workloads, the hypervisor's memory virtualization and page replacement performance deteriorate due to a different ratio between VM numbers and physical processor count.

Heterogeneous Cluster. Figures 17 and 18 demonstrate how runtime configuration affects the server's energy efficiency using server #1 and server #3 under different processor frequency. As can be seen from these figures, the EP, peak power consumption, and peak EP of the three heterogeneous servers are different (i.e., server #3 has higher EP and EE than the others). It should be noted that Figures 17 and 18 do not suggest that only the CPU impacts a server's energy efficiency and proportionality. Instead, these show that runtime configuration, like frequency scaling, can affect server's energy efficiency and proportionality.

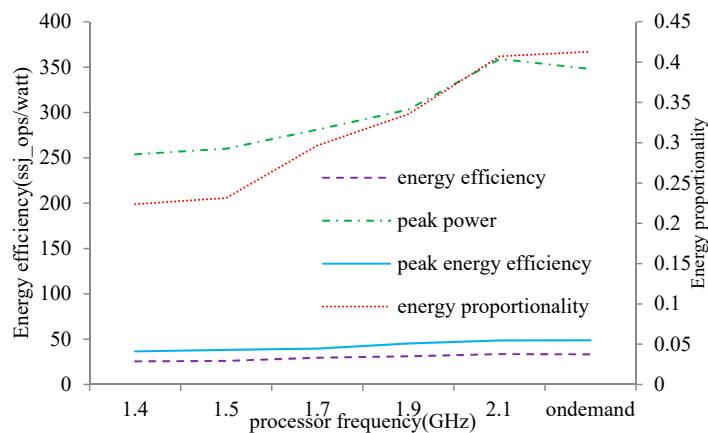


Figure 17. Server #1 energy efficiency and proportionality under different frequency scaling configurations.

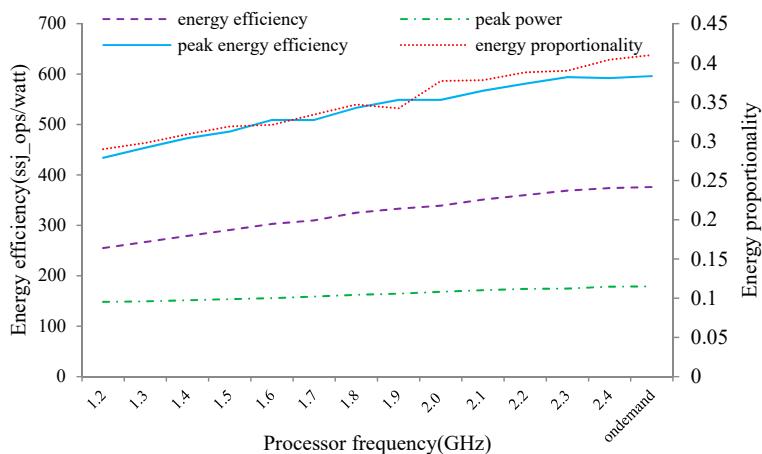


Figure 18. Server #3 energy efficiency and proportionality under different frequency scaling configurations.

Tables 9 and 10 show the power consumption of different scheduling algorithms. They also show the average completion time of a PrimeSearch workload on the heterogeneous virtualized cluster consisting of the three types of servers. The power consumption can be saved by 44.22%; the task completion time can be saved by 53.80%. After using EASE, the power consumption saving is 48.27%; task completion time saving is 30.49%.

Table 9. Power consumption and performance comparison of heterogeneous cluster.

Server	VM Scheduling Output		
	Initial (no EASE)	Packing Scheduling	EASE
#1	8	0	0
#2	2	8	4
#3	6	8	12
power(watts)	554	317	309
average completion time (s)	3476	1611	1606

Table 10. Power consumption and performance comparison of heterogeneous cluster.

Server	VM Scheduling Output		
	Initial (no EASE)	Packing Scheduling	EASE
#1	16	6	0
#2	6	12	6
#3	8	12	24
Power	636	604	329
Average completion time	3767	2370	2618

As presented in Table 10, EASE schedules the VM to the server with high energy efficiency and high EP. Although it costs 10.48% more time than packing scheduling, the power consumption is reduced by 45.53%. Therefore, EASE can be applied in a virtualized data center with limited power consumption or power capping.

Table 11 shows the performance of EASE for mixed workload cases.

Table 11. Power consumption and performance comparison of a heterogeneous cluster before and after scheduling (mixed load).

Server	VM scheduling output		
	Initial (no EASE)	Packing Scheduling	EASE
#1	8prime+8stream	0	0
#2	3prime+3stream	8prime+8stream	4prime+8stream
#3	6prime+6stream	9prime+9stream	12prime+9stream
Power	612	352	325
Prime average completion time	3366	1610	1606
STREAM average maximum bandwidth	574	514	513
STREAM average completion time	16	22	25

EASE saves power by 46.90% and 7.67% compared to the initial (without EASE) and packing scheduling, respectively. It saves completion time by 52.29% and 0.27% compared to the initial (without EASE) algorithm and packing scheduling, respectively. For the memory-based workload, EASE reduces the maximum bandwidth by -10.70% and -0.34% compared to the initial (without EASE) and the packing scheduling of STREAM, respectively. As for average completion time, EASE is increased by 40.9% and 4.64% compared to the initial (without EASE) and intermediate scheduling of STREAM, respectively.

5. Related Work

Thanks to the rapid development and technological breakthroughs of hardware, server energy efficiency has significantly improved. However, in many scenarios, dynamic voltage frequency scaling (DVFS) and dynamic power management (DPM) are still used for single server's power

aware adaptation and performance tuning [20–23]. Those approaches dynamically adjust the running state for a single resource (i.e., CPU or memory) or focus on single workloads. They ignore the interdependence of different resources and their impact on server performance. In addition, they ignore power consumption when running different types of applications.

For data center power management, the main goal is to coordinate all resources for multiplexing and service provisioning. In a virtualized data center, it is more challenging to initiate DVFS because VM has no knowledge of underlying hardware chips although virtualization enables service consolidation and power reduction. The power and energy aware scheduler must balance the overall system energy consumption and performance. It must also improve resource utilization while trying to satisfy service level agreements (SLAs). Genetic algorithm, Ant Colony algorithm, linear programming, adaptive heuristics, and utility-based approaches are proposed for resource scheduling and VM migration [24–34]. For typical air-cooled data centers, researchers propose thermal-aware workload allocation strategy with respect to the chip temperature constraint [35] and use computational fluid dynamics (CFD) to model and validate Airflow in a Data Center [36]. However, state-of-the-art VM scheduling approaches do not consider the server's energy efficiency and EP.

Server's power consumption varies when its utilization level changes. Moreover, a server consumes large amount of power even when it's idle. Thus energy proportional computing is proposed. Energy proportionality typically refers to that the system's power consumption changes with the workload level. The research results in [37] show that it is difficult to achieve the energy proportional computing on a single server. The reason is that CPUs, RAMs, graphics cards and other hardware on a single server consume about 70% of peak energy consumption even in idle state [38]. Therefore, energy proportional computing is more achievable in data center scale where multiple servers are located together for service provisioning.

The latest commercial servers achieve their peak energy efficiency at non-100% utilization. Therefore, it is possible to save more energy when the server is underutilized compared to the 100% utilization case [38–41]. Moreover, efficiency in hypervisor itself can also impact virtualized server's energy efficiency, such as inter-domain communication performance [42]. However, EP requires more exploration for energy efficiency and proportionality-aware VM scheduling.

This paper proposes EASE, the energy efficiency and proportionality-aware VM scheduling, for virtualized cloud data centers. EASE first identifies server energy efficiency and EP by running customized and standard benchmarks. It schedules VMs to servers with the highest energy efficiency or the servers with the highest PEEP value to save more power and energy. Our work is based on the energy efficiency and energy proportionality characteristics of each server in data centers. And we schedule virtual machines in virtualized environment rather than jobs to physical servers. Our core contribution is that we integrate the energy efficiency and energy proportionality awareness in the scheduling framework. According the real benchmarking results, our algorithms outperform related work in energy minimization and performance optimization.

6. Conclusions and Future Work

A typical data center contains heterogeneous servers manufactured in different generations from different vendors. However, the energy efficiency and proportionality-aware scheduling approach mask heterogeneity, the complexity of hardware configuration, and performance tuning. The authors propose EASE to schedule VMs according to workload characteristics, energy efficiency, and proportionality of servers in a data center. EASE keeps the servers running in their optimal range in terms of energy efficiency through VM scheduling and migration. It can provide server consolidation and ensure the quality of service guarantee of multi-tenants. In addition, it can reduce energy consumption in virtualized data centers while not sacrificing significant application performance inside VMs. Experiments on real homogeneous clusters show that EASE achieves 37.07%–49.98% power consumption savings for computing and memory intensive workloads. The average completion time of the workload is only increased by 0.31% to 8.49%. Experiments on heterogeneous clusters

show that the power consumption of the computing intensive workload can be reduced by 44.22%; the task completion time can be saved by 53.80%.

The use of EASE on recent server platforms equipped with power management options (i.e., IPMI, intelligent power distribution cabinet with power capping PDM or PDU) shows that it can be easily integrated into existing power management systems to deploy energy efficiency-aware VM scheduling. For legacy servers without IPMI capability, EASE may function well if it can read power measurements from power meters in real time if the power meter has real time data acquisition function. We are developing an energy efficiency simulator for data centers to evaluate our algorithm's scalability [43]. Moreover, in the future, we also would like to adapt EASE to diverse workload types by self-learning workload characterization in data centers.

Author Contributions: C.J. conceived the idea and designed the algorithms. Y.Q., Y.W. and D.O. implemented the software modules. Y.L. validated the modeling. J.W. provided the formal analysis.

Funding: This research was funded by the National Natural Science Foundation of China (Grant No. 61472109, 61672200, and 61572163).

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Worldwide Server unit Shipments from 1st Quarter 2009 to 1st Quarter 2016. Available online: <http://www.statista.com/statistics/267390/global-server-shipments-by-vendor/> (accessed on 26 December 2018).
2. Fernández-Cerero, D.; Fernández-Montes, A.; Velasco, F. Productive Efficiency of Energy-Aware Data Centers. *Energies* **2018**, *11*, 2053. [[CrossRef](#)]
3. Report to Congress on Server and Data Center Energy Efficiency Public Law 109-431. Available online: <https://escholarship.org/uc/item/74g2r0vg> (accessed on 26 December 2018).
4. Natural Resources Defense Council, Data Center Efficiency Assessment. Available online: <https://www.nrdc.org/sites/default/files/data-center-efficiency-assessment-IP.pdf> (accessed on 26 December 2018).
5. GeSI SMARTer2020: The Role of ICT in Driving a Sustainable Future. Available online: <http://gesi.org/smarter2020> (accessed on 26 December 2018).
6. Barroso, L.A.; Clidaras, J.; Hölzle, U. *The Datacenter as a Computer: An Introduction to the Design of Warehouse-Scale Machines*, 2nd ed.; Morgan & Claypool Publishers: San Rafael, CA, USA, 2013.
7. Wu, Q.; Deng, Q.; Ganesh, L.; Hsu, C.; Jin, Y.; Kumar, S.; Li, B.; Meza, J.; Song, Y. Dynamo: Facebook’s Data Center-Wide Power Management System. In Proceedings of the 2016 ACM/IEEE 43rd Annual International Symposium on Computer Architecture (ISCA), Seoul, South Korea, 18–22 June 2016.
8. Zafar, S.; Chaudhry, S.A.; Kiran, S. Adaptive TrimTree: Green Data Center Networks through Resource Consolidation, Selective Connectedness and Energy Proportional Computing. *Energies* **2016**, *9*, 797. [[CrossRef](#)]
9. Barroso, L.A.; Hölzle, U. The case for energy-proportional computing. *Computer* **2007**, *12*, 33–37. [[CrossRef](#)]
10. Ryckbosch, F.; Polfliet, S. Trends in server energy proportionality. *Computer* **2011**, *9*, 69–72. [[CrossRef](#)]
11. Sen, R.; Wood, D. Energy-Proportional Computing: A New Definition. *Computer* **2017**, *8*, 26–33. [[CrossRef](#)]
12. Jiang, C.; Wang, Y.; Ou, D. Energy Proportional Servers: Where Are We in 2016? In Proceedings of the 2017 IEEE 37th International Conference on Distributed Computing Systems (ICDCS), Atlanta, GA, USA, 5–8 June 2017; pp. 1649–1660.
13. Malladi, K.T.; Lee, B.C.; Nothaft, F.A.; Kozyrakis, C.; Periyathambi, K.; Horowitz, M. Towards energy-proportional datacenter memory with mobile DRAM. In Proceedings of the 39th Annual International Symposium on Computer Architecture (ISCA), Portland, OR, USA, 9–13 June 2012.
14. Malladi, K.T.; Shaeffer, I.; Gopalakrishnan, L.; Lo, D.; Lee, B.C.; Horowitz, M. Rethinking DRAM Power Modes for Energy Proportionality. In Proceedings of the 2012 45th Annual IEEE/ACM International Symposium on Microarchitecture, Vancouver, BC, Canada, 1–5 December 2012; pp. 131–142.
15. SPECpower_ssj2008. Available online: https://www.spec.org/power_ssj2008/ (accessed on 26 December 2018).
16. Jiang, C.; Wang, Y.; Ou, D.; Qiu, Y.; Li, Y.; Wan, J.; Luo, B.; Shi, W.; Cerin, C. EASE: Energy Efficiency and Proportionality Aware Virtual Machine Scheduling. In Proceedings of the 30th International Symposium on Computer Architecture and High Performance Computing (SBAC-PAD2018), Paris, France, 25 September 2018.

17. Jiang, C.; Wang, Y.; Ou, D.; Li, Y.; Zhang, J.; Wan, J.; Luo, B.; Shi, W. Energy efficiency comparison of hypervisors. *Sustain. Comput. Inform. Syst.* **2017**. [[CrossRef](#)]
18. STREAM. Available online: <https://www.cs.virginia.edu/stream/> (accessed on 26 December 2018).
19. Jiang, C.; Han, G.; Lin, J.; Jia, G.; Shi, W.; Wan, J. Characteristics of Co-allocated Online Services and Batch Jobs in Internet Data Centers: A Case Study from Alibaba Cloud. *IEEE Access* **2019**. [[CrossRef](#)]
20. Meisner, D.; Gold, B.T.; Wenisch, T.F. PowerNap: Eliminating server idle power. In Proceedings of the 14th international conference on Architectural support for programming languages and operating systems, Washington, DC, USA, 7–11 March 2009; pp. 205–216.
21. Isci, C.; McIntosh, S.; Kephart, J.; Das, R.; Hanson, J.; Piper, S.; Wolford, R.; Brey, T.; Kantner, R.; Ng, A.; et al. Agile, efficient virtualization power management with low-latency server power states. In Proceedings of the 40th Annual International Symposium on Computer Architecture, Tel-Aviv, Israel, 23–27 June 2013; pp. 96–107.
22. Liu, Y.; Draper, S.C.; Kim, N.S. SleepScale: Runtime joint speed scaling and sleep states management for power efficient data centers. In Proceedings of the 2014 ACM/IEEE 41st International Symposium on Computer Architecture, Minneapolis, Minnesota, 14–18 June 2014; pp. 313–324.
23. Chiaravaggio, L.; Cianfrani, A.; Listanti, M.; Liu, W.; Polverini, M. Lifetime-Aware Cloud Data Centers: Models and Performance Evaluation. *Energies* **2016**, *9*, 470. [[CrossRef](#)]
24. Ferreto, T.C.; Netto, M.A.; Calheiros, R.N.; De Rose, C.A. Server consolidation with migration control for virtualized data centers. *Future Gener. Comput. Syst.* **2011**, *27*, 1027–1034. [[CrossRef](#)]
25. Speitkamp, B.; Bichler, M. A mathematical programming approach for server consolidation problems in virtualized data centers. *IEEE Trans. Serv. Comput.* **2010**, *3*, 266–278. [[CrossRef](#)]
26. Beloglazov, A.; Buyya, R. Optimal online deterministic algorithms and adaptive heuristics for energy and performance efficient dynamic consolidation of virtual machines in cloud data centers. *Concurr. Comput. Pract. Exp.* **2012**, *24*, 1397–1420. [[CrossRef](#)]
27. Liu, H.; Jin, H.; Xu, C.; Liao, X. Performance and energy modeling for live migration of virtual machines. *Cluster Comput.* **2013**, *16*, 249–264. [[CrossRef](#)]
28. Xu, C.; Zhao, Z.; Wang, H.; Shea, R.; Liu, J. Energy efficiency of cloud virtual machines: From traffic pattern and CPU affinity perspectives. *IEEE Syst. J.* **2017**, *11*, 835–845. [[CrossRef](#)]
29. Lagen, S.; Pascual-Iserte, A.; Munoz, O.; Vidal, J. Energy efficiency in latency-constrained application offloading from mobile clients to multiple virtual machines. *IEEE Trans. Signal Process.* **2018**, *66*, 1065–1079. [[CrossRef](#)]
30. Jiang, C.; Duan, L.; Liu, C.; Wan, J.; Zhou, L. VRAA: Virtualized resource auction and allocation based on incentive and penalty. *Cluster Comput.* **2013**, *16*, 639–650. [[CrossRef](#)]
31. Belabed, D.; Secci, S.; Pujolle, G.; Medhi, D. Striking a balance between traffic engineering and energy efficiency in virtual machine placement. *IEEE Trans. Netw. Serv. Manag.* **2015**, *12*, 202–216. [[CrossRef](#)]
32. Yan, S.; Xiao, S.; Chen, Y.; Cui, Y.; Liu, J. GreenWay: Joint VM placement and topology adaption for green data center networking. In Proceedings of the 26th International Conference on Computer Communication and Networks, Vancouver, BC, Canada, 31 July–3 August 2017; pp. 1–9.
33. Lago, D.G.; Madeira, E.R.; Medhi, D. Energy-aware virtual machine scheduling on data centers with heterogeneous bandwidths. *IEEE Trans. Parall. Distr. Syst.* **2018**, *29*, 83–98. [[CrossRef](#)]
34. Liu, X.; Zhan, Z.; Zhang, J. An Energy Aware Unified Ant Colony System for Dynamic Virtual Machine Placement in Cloud Computing. *Energies* **2017**, *10*, 609. [[CrossRef](#)]
35. Bai, Y.; Gu, L.; Qi, X. Comparative Study of Energy Performance between Chip and Inlet Temperature-Aware Workload Allocation in Air-Cooled Data Center. *Energies* **2018**, *11*, 669.
36. Wibron, E.; Ljung, A.-L.; Lundström, T.S. Computational Fluid Dynamics Modeling and Validating Experiments of Airflow in a Data Center. *Energies* **2018**, *11*, 644. [[CrossRef](#)]
37. Tsirogiannis, D.; Harizopoulos, S.; Shah, M.A. Analyzing the energy efficiency of a database server. In Proceedings of the 2010 ACM SIGMOD International Conference on Management of data, Indianapolis, Indiana, 6–10 June 2010; pp. 231–242.
38. Lang, W.; Patel, J.M. Energy management for mapreduce clusters. *Proc. VLDB Endow.* **2010**, *3*, 129–139. [[CrossRef](#)]

39. Wong, D. Peak efficiency aware scheduling for highly energy proportional servers. In Proceedings of the 2016 ACM/IEEE 43rd Annual International Symposium on Computer Architecture, Seoul, South Korea, 18–22 June 2016; pp. 481–492.
40. Leverich, J.; Kozyrakis, C. On the energy (in) efficiency of Hadoop clusters. *ACM SIGOPS Oper. Syst. Rev.* **2010**, *44*, 61–65. [[CrossRef](#)]
41. Schall, D.; Hudlet, V. WattDB: An energy-proportional cluster of wimpy nodes. In Proceedings of the 2011 ACM SIGMOD International Conference on Management of data, Athens, Greece, 12–16 June 2011; pp. 1229–1232.
42. Jiang, C.; Fan, T.; Qiu, Y.; Wu, H.; Zhang, J.; Xiong, N.; Wan, J. Interdomain I/O Optimization in Virtualized Sensor Networks. *Sensors* **2018**, *18*, 4395. [[CrossRef](#)] [[PubMed](#)]
43. Jiang, C.; Qiu, Y.; Shi, W.; Cerin, C.; Xiong, N.; Wan, J. Escope: An Energy Efficiency Simulator For Data Centers. In Proceedings of the IEEE 39th International Conference on Distributed Computing Systems (ICDCS), Dallas, Texas, USA, 7–10 July 2019. submitted.



© 2019 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>).