Energy Efficiency Comparison of Hypervisors

Congfeng Jiang[1,2], Dongyang Ou[1], Yumei Wang[1],
Xindong You[1], Jilin Zhang[1], Jian Wan[1]

[1]Key Laboratory of Complex Systems Modeling and Simulation

Hangzhou Dianzi University

Hangzhou, Zhejiang, China

{jiangcongfeng, oudongyang1, wyumei1}gmail.com,

{youxindong, jilin.zhang,wanjian}@hdu.edu.cn

Bing Luo[2], Weisong Shi[2]

[2]Mobile and Internet Systems Laboratory

Department of Computer Science

Wayne State University

Detroit, MI, USA

{bing,weisong}wayne.edu

Abstract*—Current cloud data centers are fully virtualized for service consolidations and power/energy reduction. Although virtualization could reduce real time power and overall energy consumption, the energy characteristics of hypervisors hosting different workloads are not well profiled and understood. In this paper, we investigate the power and energy characteristics of mainstream hypervisors and container engine, i.e., VMware ESXi, Microsoft Hyper-V, KVM, XenServer and Docker, on five different platforms (two mainstream 2U rack servers, one emerging ARM64 server, one desktop server, and one laptop) with hundreds of hours' power measures. We use both computing intensive and mixed web server-database workloads to explore the power and energy characteristics of different hypervisors. Extensive experiment results of four workload levels (very light, light, fair, and very heavy workload) demonstrate that hypervisors expose different power and energy characteristics. We found that: (1) Hypervisors expose different power and energy consumption on the same hardware running same workloads. (2) Although mainstream hypervisors have different energy efficiencies aligned with different workload types and workload levels, there is no one hypervisor that outperforms all other hypervisors on all platforms in terms of power or energy consumptions. (3) Although container virtualization is considered as light-weight virtualization in terms of implementation and maintenance, it is not essentially more power efficient than traditional virtualization technology. (4) ARM64 server does have lower power consumption, but they finish computing jobs with longer execution time and sometimes consume more energy. And ARM64 servers has medium energy consumption per database operations for mixed workloads. The results presented in this paper provide useful insights to system designers, as well as data center operators for power-aware workload placement and virtual machine scheduling.

Keywords—energy efficiency; hypervisor;virtual machine; container virtualization

I. MOTIVATION

The explosion of cloud computing, big data analytics, e-commerce, and Internet traffic make the growing of power demand faster in data centers [1, 2]. Various hardware related approaches have been proposed to increase data center energy efficiency at different levels [3, 4], including circuits and chips [5, 6, 7], memory [8, 9], disk [10, 11], and network traffic routing [12, 13]. At the meantime, power and energy aware approaches are also proposed to deeply explore energy efficiency possibilities, such as performance tuning [14, 15], application centric power optimization [16, 17], resource scheduling and allocation [18, 19], and thermal aware power capping [20, 21]. Resource multiplexing in data centers provides data center wide power management opportunities [22, 23]. To further reduce the carbon footprints, renewable energy and liquid cooled systems are also introduced into modern data centers [24, 25]. Recently, ARM64 servers are also considered competitors in server market due to its lower power consumption compared to servers packaged with traditional x86 processors.

In cloud data centers server virtualization and consolidation are widely deployed for power and energy savings [26]. Usually over-commitment or over-subscription is used to further reduce energy costs more aggressively. For a virtualized platform, the hypervisor or virtual machine monitor (VMM) acts as the equivalent operating system and is responsible for resource scheduling and guest operating system hosting. Different from classic heavy weight virtualization, container virtualization (or containerization) is emerging as a new paradigm for cloud data centers for its light-weight simplicity and scalability of hosting more application instances than traditional virtual machines. Container technology makes it easier to deploy multiple copies of same applications/services than virtual machines because containers only need operating system, supporting programs and libraries, and system resources to run a specific program without any hardware abstraction. In containerization environment, the container engine acts as the hypervisor in traditional virtualized environment while it leverages the underlying operating system kernel for core resource management and allocation. In this paper we treat the container engine as hypervisor for simplification and use them interchangeably. We also use virtual machine and container interchangeably unless specified in the context.

However, due to hardware abstraction and semantic gap between virtual machine (or container) and underlying

978-1-5090-5117-5/16/$31.00 ©2016 IEEE

hardware, virtual machine operating system can't invoke real power aware management as in non-virtualized environment. Moreover, hypervisors may result in different energy efficiency that affects across different hardware platforms and guest virtual machines in many folds. Power management and energy accounting in such virtualized environment requires accurate knowledge and characteristics on energy efficiency of hypervisors [27]. From the data center wide point of view, differentiating energy efficiency of various hypervisors can help hypervisor selection, system design, and system operation. This motivates us to investigate and compare the energy efficiency of different hypervisors.

In this paper we do not intend to compare power and performance of different servers like previous work in [28, 29, 30] (see related work in Section IV). Instead we compare the energy efficiency of hypervisors on the same server. It's worth noting that currently most of the attention is focused on the publicly known large scale cloud data centers which consume less than 5 percent of U.S. data center electricity, many small, medium, corporate and multi-tenant data centers are still operating in lower energy efficiency ways [31]. Therefore in this paper ,we investigate not only the energy efficiency of rack servers used in typical large cloud data centers, but also desktop server, ARM64 server and laptop used in many other data centers. We select four mainstream hypervisors, i.e., Microsoft Hyper-V, VMware ESXi, KVM/QEMU, XenServer, and the Docker container engine because these hypervisors are widely deployed in nowadays virtualized data centers. We conducted extensive experiments in three orthogonal dimensions, i.e., hardware, hypervisor, and workload. Specifically, we compare energy efficiency of five hypervisors (including Docker,) on five hardware platforms running two types of workload at four workload levels. We collect the real time power of each experiment and calculate the energy consumption of each experiment and micro-operation.

The remainder of this paper is organized as follows. In Section II we describe the details of the experiment methodology and system setup. In Section III we provide experiment results, observations and insights of energy efficiency of hypervisors. We summarize related work in Section IV and conclude the paper and make remarks on future work in Section V.

II. Experiment Methodology And Setup

A. Experiment Methodology

In this paper, we conduct experiments in three orthogonal dimensions, i.e., hardware, hypervisor, and workload. For example, for a specific hardware, we run different hypervisors on it and each hypervisor hosts a number of virtual machines running computing intensive and mixed workloads. In contrast to previous work, we conduct fine grained energy efficiency comparison of hypervisors under a series of workload intensities to get a profile of the energy proportionality of hardware and hypervisors. Previous work on energy efficiency comparison focus on fixed workload level for each workload type, which is not complete because each hardware platform has different energy proportionality. Instead, we classify the workload intensity into 4 workload levels, i.e., very light, light, fair, and very heavy workload. For computing intensive workload, the fair workload can stress the system at more than 95% utilization steadily and the very heavy workload saturates the system. Our methodology provides opportunity to investigate energy proportionality not only of the hardware itself but also the hypervisor and our experiments show that power and energy consumption change significantly on the same hardware with varying workload levels when running different hypervisors.

B. Experiment Setup

In order to extensively compare the energy efficiency of hypervisors, we select five typical hardware platforms in our lab as our testbeds. The testbed configuration is listed in Table 1 and each platform is sorted by CPU release date (Intel Q8300 is the oldest). Nowadays, some processors and motherboards dedicated for laptops are used in data centers for their better power performance. We select the Lenovo W541 laptop as a representative of such low power platform. When we run hypervisor on different platforms, we run Windows Server 2012 R2 for Hyper-V (data center version with GUI), CentOS Linux 7 (v1503 server with GUI) for KVM/QEMU 1.5.3, ESXi 6.0/vSphere 6.0 and XenServer 6.5.0, respectively. In order to provide an apples-to-apples comparison, the virtual machines are running the same operating systems with same software configurations including kernel version and software stack. For example, all virtual machines run the same x64 version CentOS 7 with Linux kernel 3.10. All the power data are measured by a WattsUP .Net power meter.

Table 1. Platform configurations

Platform	Type	CPU	Cores/threads	Memory	Storage
HP s5280t	Desktop	Intel Q8300	4/4	8 GB	1TB HDD
HP DL380 G6	2U server	Intel Xeon E5530 x2/wVT-x &VT-d	8/16	80GB	2TB SSD RAID10
Intel S2600GZ	2U server	Intel E5-2680x2/w VT-x &VT-d	16/32	64GB	2TB HDD RAID1
Lenovo W541	Laptop	Intel i7- 4710MQ with VT-x	4/8	32 GB	2TB SSD
APM X-C1	ARM64 server	APM X-Gene 2	8/8	16 GB	2TB SSD

C. Experiment Workload and Workload Level Classification

On top of hypervisors, we run a number of virtual machines with two workloads: i.e., one computing intensive workload and one mixed web server and database workload.

Computing intensive workload: We use a prime number computation program written in C language, namely, PrimeSearch, as the computing intensive workload. In one execution of the PrimeSearch, it will calculate and search the prime numbers in ten intervals, i.e., (1,1000000), (1,2000000), (1,3000000), (1,4000000), (1,5000000), (1,6000000), (1,7000000), (1,8000000), (1,9000000), (1,10000000), respectively. These ten intervals are ten sub searching tasks. And the completion time of searching for each interval will be calculated and the sum of the completion time of ten sub tasks is calculated as the task completion time of one Primesearch execution.

Mixed workload: We use the two-tier web server and database synthetic environment as mixed workload, namely, LAMP (Linux, Apache, MySQL, and PHP). In the mixed LAMP workload, we use PHP scripts in web page to insert data records into MySQL database via apache web server. Each execution will try to insert 1000000 rows of data. The Primesearch and LAMP workloads are static workload and during each experiment all virtual machines run the same code and therefore each machine are equivalently contributing to the workload to the system under test.

For convenience, we refer the above two workloads as computing workload and mixed workload respectively in the rest of the paper.

Workload level: Here we define four workload levels, *1/4*, *1/2*, *1/1*, and *2/1,* which stands for very light workload, light workload, fair workload, and very heavy workload. We use the workload level as workload intensity indicator. These workload levels mean the number of virtual machine running concurrently within a physical server is quarter of (1/4), half of (1/2), equal to (1/1), and twice (2/1) of the physical processor cores, respectively. For example, when we run very heavy computing workload on the HP DL380 G6 server, we run ***16*** virtual machines while we only run ***8*** virtual machines on the HP s5280t server for very heavy workload because the HP s5280t server only has four physical processor cores in total (see Table 1 for processor configurations). Although it's possible to compare energy efficiency of different servers with actual number of virtual machines of different hypervisors, it's fairer to compare them with the ratio of virtual machines number over its hardware configuration (specifically the number of processor cores) to investigate the energy efficiency and proportionality on different platforms.

III. Experimental Observations And Insights

A. *Computing Intensive Workloads*

We first run the PrimeSearch workload on each platform. The power and energy results are provided in Fig.1. The completion times are provided in Fig.2. The power and energy variation (ratio of the highest to the lowest) of different hypervisors running different workloads on different platforms are given in Table 2 and Table 3, respectively. We list the observations in the following.

Observation #1: Hypervisors expose different power consumption, completion time, and energy consumption on the same hardware running same workload.

As can be seen in Fig.1 and Fig.2, different hypervisors have different power, energy consumption, and completion time on same hardware. Here the difference has two-fold meanings. *Firstly*, one hypervisor has different power and energy consumption at different workload level on the same hardware. *Secondly*, different hypervisor has different power and energy consumption on the same hardware at same workload level.

Table 2. Power variation of computing intensive workloads

Platform	1/4 workload	1/2 workload	1/1 workload	2/1 workload
HP s5280t	**59.87%**	35.78%	3.52%	3.53%
HP DL380 G6	12.38%	12.98%	16.67%	**19.79%**
Intel S2600GZ	5.44%	3.03%	6.71%	**8.40%**
Lenovo W541	**25.03%**	22.71%	15.87%	10.81%
APM X-C1	**10.11%**	0.05%	1.18%	8.18%

Table 3. Energy variation of computing intensive workloads

Platform	1/4 workload	1/2 workload	1/1 workload	2/1 workload
HP s5280t	61.86%	36.74%	**72.46%**	14.88%
HP DL380 G6	15.70%	14.05%	28.05%	**41.90%**
Intel S2600GZ	13.56%	11.06%	**25.06%**	9.26%
Lenovo W541	**35.34%**	25.90%	16.62%	12.77%
APM X-C1	10.04%	0.05%	2.48%	**22.60%**

From Table 2 we observe that, the power variation on HP DL380 G6 is 19.79% for very heavy computing workload by comparing the highest power (from Hyper-V) to the lowest power (from KVM). But on another typical rack server, i.e., Intel S2600GZ, the power variation of the highest (from Docker) to the lowest (from Hyper-V) is 8.40% for very heavy computing workloads. We also observe in Table 2 that usually the heavier the workload is, the higher the power variation for two rack servers. But it is reverse for the desktop server and laptop. Since we only run KVM and Docker on APM X-C1, it does not expose any significant trend of power variation with respect to workload level. It's also interesting that on rack servers the maximum power variation occurs when the workload is the heaviest, while on desktop server, laptop and ARM64 servers, it occurs when the workload is the lightest. As can be seen in Table 3, the energy variations are larger than those of power variation in Table 2. This is because the energy is the product of power and task completion time. Due to space limitation, we only provide the variations of completion time (ratio of the highest over the lowest) in Table 4. Moreover, the energy variations of hypervisors on different platforms are more diversified and scattered on all workload levels except 1/2 workload level, in contrast to those of power variations.

Table 4. Variation of completion time for computing intensive workloads

Platform	1/4 workload	1/2 workload	1/1 workload	2/1 workload
HP s5280t	14.06%	20.41%	**70.48%**	17.32%
HP DL380 G6	6.23%	4.71%	13.03%	**18.95%**
Intel S2600GZ	12.13%	11.85%	**24.69%**	5.01%
Lenovo W541	11.18%	6.69%	11.08%	**14.76%**
APM X-C1	0.06%	0.00%	1.28%	**13.33%**

Observation #2: Although hypervisors have different energy efficiencies aligned with different workload types and workload levels, no one hypervisor outperforms all other hypervisors in terms of power, energy consumption, or completion time for all workload levels on all platforms.

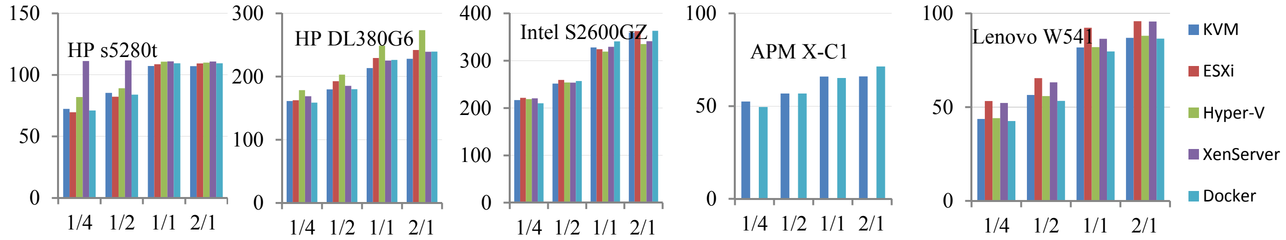

(a) Power consumption (x axis: workload level, y axis: power (watts))

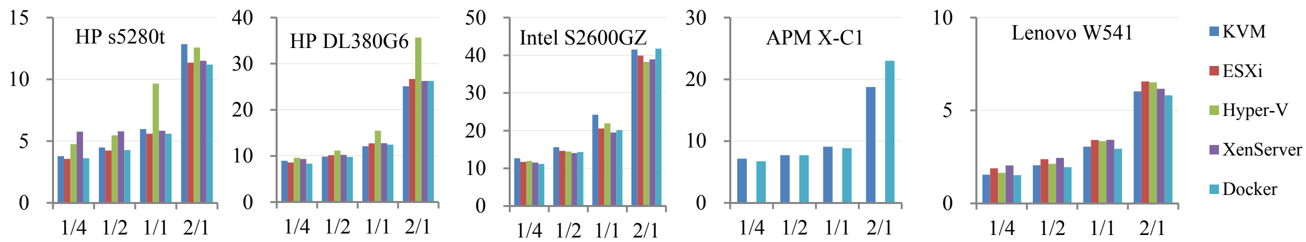

(b) Energy consumption (x axis: workload level, y axis: energy (10^6 joules))

Fig.1 Power and energy consumption of varying computing workloads

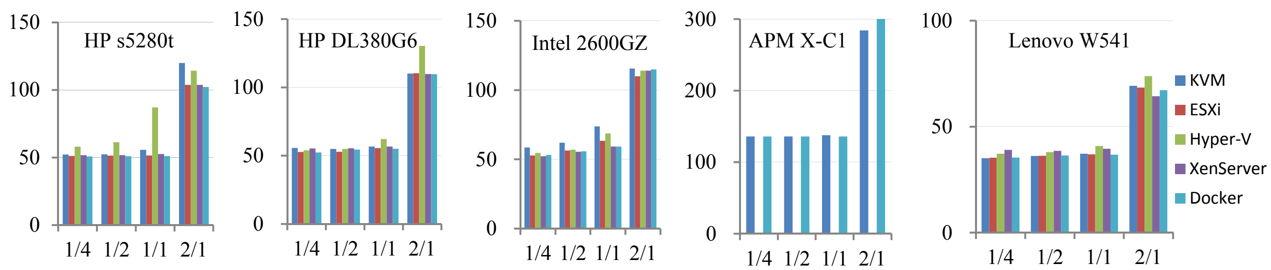

Fig.2 Completion time of varying computing workloads (x axis: workload level, y axis: completion time (10^3 seconds))

We list the highest and lowest power consumption at all workload levels on all platforms in Table 5.

Table 5 Hypervisor with the highest and the lowest power consumption for computing intensive workload

Platform	Power	Workload level			
		1/4	1/2	1/1	2/1
HP s5280t	highest	XenServer	XenServer	XenServer	XenServer
	lowest	ESXi	ESXi	KVM	KVM
HP DL380 G6	highest	Hyper-V	Hyper-V	Hyper-V	Hyper-V
	lowest	Docker	KVM	KVM	KVM
Intel S2600GZ	highest	ESXi	ESXi	Docker	Docker
	lowest	Docker	KVM	Hyper-V	Hyper-V
Lenovo W541	highest	ESXi	ESXi	ESXi	ESXi
	lowest	Docker	Docker	Docker	Docker
APM X-C1	highest	KVM	Docker	KVM	Docker
	lowest	Docker	KVM	Docker	KVM

From Table 5 we can see that there is no one hypervisor that always has the highest or lowest power consumption for all workload levels on all platforms. Instead, the distribution of hypervisors with the highest or lowest power consumption is neither platform dependent nor workload level dependent. One hypervisor may have the highest power consumption for some or all workload levels on one platform, but it may have the lowest power consumption at some or all workload levels on another platform. In other words, there is no hypervisor-affinity across different hardware and workload levels. For example, Hyper-V has the highest power and energy consumption on HP DL380 G6 server running computing intensive application at all workload levels, while it has the lowest power consumption for fair and very heavy computing workloads on Intel S2600GZ server.

More importantly, although the ARM64 server APM X-C1 has the lowest power consumption, it also has the highest energy consumption because it has the longest completion time. Moreover, XenServer has the highest power consumption only on HP s5280t for all workload levels and Hyper-V has the highest power consumption only on HP DL380 G6 at all workload levels. Here we take ESXi as another example. ESXi has the highest power consumption on Lenovo W541 at all workload levels and on Intel S2600GZ for 1/4 and 1/2 workload levels. However, ESXi has the lowest power consumption on HP s5280t for 1/4 and 1/2 workload levels. It’s worth noting that

XenServer only exposes the highest power consumption while KVM only exposes the lowest power consumption if we exclude the APM X-C1 platform because we do not run as many hypervisors on it as other platforms (only KVM and Docker on it). In fact, the APM X-C1 ARM64 server has the lowest power variation for all workloads no matter which hypervisor is running on it. Similarly, the distribution of hypervisors with the highest or lowest energy consumption is more diversified than power consumption in Table 5 since the energy consumption is affected by power and task completion time jointly.

Observation #3: Although container virtualization is considered as light-weight virtualization in terms of implementation simplicity and easy maintenance, it is not essentially more power efficient than traditional virtualization technology for computing workloads.

We compare the power and energy consumption of Docker with the highest, lowest, and average power and energy consumption of all hypervisors (including Docker itself) in Table 6 and Table 7.

Table 6 Comparison of Docker and the hypervisor with the highest, average, and lowest power

Platform	Power	Workload level			
		1/4	1/2	1/1	2/1
HP s5280t	Highest	**-36.17%**	-24.82%	-1.39%	-1.35%
	Average	**-12.59%**	-7.20%	0.03%	0.03%
	Lowest	2.04%	2.08%	2.08%	**2.13%**
HP DL380 G6	Highest	-11.01%	-11.41%	-9.10%	**-12.39%**
	Average	-4.31%	**-4.37%**	-1.03%	-2.02%
	Lowest	0.00%	0.09%	**6.05%**	4.95%
Intel S2600GZ	Highest	**-5.16%**	-0.86%	0.00%	0.00%
	Average	-3.44%	0.78%	**3.96%**	3.29%
	Lowest	0.00%	2.14%	6.71%	**8.40%**
Lenovo W541	Highest	**-20.02%**	-18.51%	-13.69%	-9.75%
	Average	**-9.76%**	-9.44%	-5.63%	-4.54%
	Lowest	0.00%	0.00%	0.00%	0.00%
APM X-C1	Highest	**-9.18%**	0.00%	-1.17%	0.00%
	Average	**-4.81%**	0.03%	-0.59%	3.93%
	Lowest	0.00%	0.05%	0.00%	**8.18%**
Average	Highest	**-16.31%**	-11.12%	-5.07%	-4.70%
	Average	**-6.98%**	-4.04%	-0.65%	0.14%
	Lowest	0.41%	0.87%	2.97%	**4.73%**

Note: Positive value means Docker consumes more power.

For example, Docker has the highest power consumption running computing intensive workloads on Intel S2600GZ for fair and very heavy workloads. Docker consumes 8.40% more power than Hyper-V (the lowest), and 3.29% more power than average power, for very heavy computing workload on Intel S2600GZ. Docker also consumes 6.71% more power than Hyper-V (the lowest), and 3.96% more power than average power, for fair computing workload on Intel S2600GZ. From Table 6 we can see that in average Docker consumes less power than the highest and average power consumption of all hypervisors on all platforms running all levels of workloads except that it consumes 0.14% more power than average of all hypervisors. Although Docker consumes less power, it does not consume the lowest power in all experiments. The heavier the workload level is, the higher Docker's power consumption is (from 0.41% to 4.73%). In average, Docker's power consumption is very close to the lowest power consumption for 1/4 and 1/2 workload level on all platforms (only slightly 0.41% and 0.87% more than lowest ones). However, for 1/1 and 2/1 workload levels, Docker consumes 2.97% and 4.73% more power than the lowest power of other hypervisors.

On ARM64 server, Docker also has the highest power consumption running computing intensive workloads for light and very heavy workloads. On Lenovo W541 Docker has the lowest power and energy consumption consistently for all workload levels. Only on such platform, Docker outperforms all other hypervisors in terms of both power and energy.

Table 7 Comparison of Docker and the hypervisor with the highest, average, and lowest energy consumption

Platform	Energy	Workload level			
		1/4	1/2	1/1	2/1
HP s5280t	Highest	-37.34%	-26.12%	**-42.02%**	-12.95%
	Average	**-15.91%**	-11.78%	-14.32%	-5.93%
	Lowest	**1.41%**	1.02%	0.00%	0.00%
HP DL380 G6	Highest	-13.57%	-12.32%	-19.58%	**-26.35%**
	Average	**-7.25%**	-4.50%	-4.97%	-6.20%
	Lowest	0.00%	0.00%	2.98%	**4.51%**
Intel S2600GZ	Highest	-11.94%	-8.15%	**-16.69%**	0.00%
	Average	**-5.35%**	-1.93%	-5.07%	4.40%
	Lowest	0.00%	2.02%	4.19%	**9.26%**
Lenovo W541	Highest	**-26.11%**	-20.57%	-14.25%	-11.33%
	Average	**-12.42%**	-11.20%	-9.28%	-6.41%
	Lowest	0.00%	0.00%	0.00%	0.00%
APM X-C1	Highest	**-9.13%**	0.00%	-2.42%	0.00%
	Average	-4.78%	0.03%	-1.22%	**10.15%**
	Lowest	0.00%	0.05%	0.00%	**22.60%**
Average	Highest	**-19.62%**	-13.43%	-18.99%	-10.13%
	Average	**-9.14%**	-5.88%	-6.97%	-0.80%
	Lowest	0.28%	0.62%	1.43%	**7.27%**

Note: Positive value means Docker consumes more energy.

B. Mixed Workloads

The power, energy, and completion time for mixed workloads (LAMP) are provided in Fig.3 and Fig.4. Since the LAMP workload is web server and database centric, it uses fewer CPU cycles than Primesearch. Although the power consumption of hard disk and solid state disks are much less than processors configured in the same server, the real time power fluctuates at higher magnitude than the computing intensive workload. This means that although the average power of mixed workload is less than computing workload on almost all the servers for all 4 workload levels, the power variation during LAMP execution is greater than that of Primesearch workload. We also list the power and energy variation (ratio of the highest over the lowest) of different hypervisors running different level of workloads on all platforms in Table 8 and Table 9, respectively.

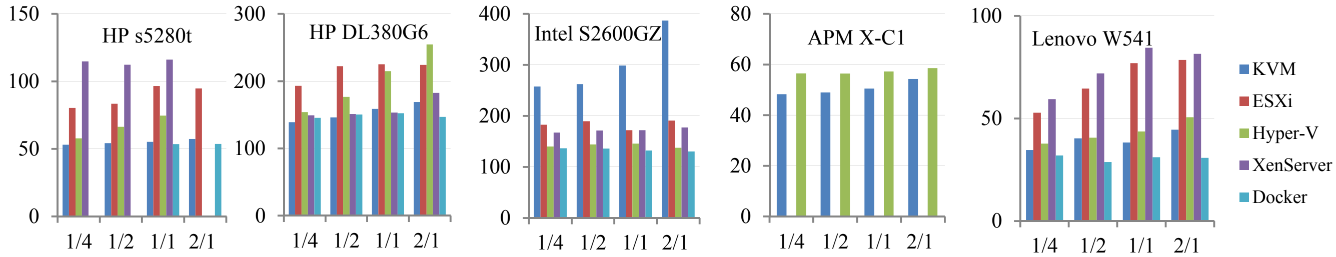

(a) Power consumption (x axis: workload level, y axis: power (watts))

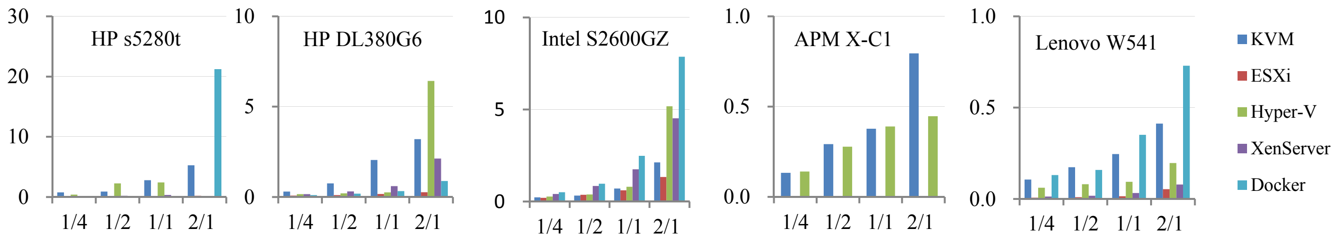

(b) Energy consumption (x axis: workload level, y axis: energy (10^6 joules))

Fig.3 Power and energy consumption of mixed workloads

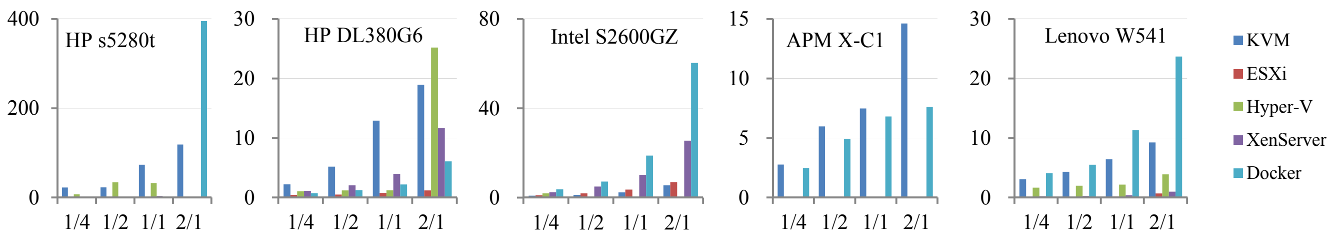

Fig.4 Completion time of mixed workloads (x axis: workload level, y axis: completion time (10^3 seconds))

Table 8. Power variation of mixed workloads

Platform	1/4 workload	1/2 workload	1/1 workload	2/1 workload
HP s5280t	**116.51%**	107.00%	110.24%	65.35%
HP DL380 G6	38.95%	52.22%	47.74%	**73.44%**
Intel S2600GZ	88.62%	92.58%	126.10%	**196.85%**
Lenovo W541	85.86%	118.48%	**171.63%**	164.87%
APM X-C1	**17.11%**	15.21%	13.45%	7.98%

Table 9. Energy variation of mixed workloads

Platform	1/4 workload	1/2 workload	1/1 workload	2/1 workload
HP s5280t	779.35%	**2702.96%**	2060.54%	2423.03%
HP DL380 G6	279.90%	592.76%	1093.76%	**2303.57%**
Intel S2600GZ	152.20%	200.32%	305.40%	**490.85%**
Lenovo W541	1699.02%	1814.92%	**2288.77%**	1256.78%
APM X-C1	5.17%	5.09%	3.25%	**77.83%**

We already see in Table 2, that the power variation on typical rack server is 8.4%~19.79% for very heavy computing workloads comparing to the lowest power consumption. However, from Table 8 and Table 9, we can see that the power and energy variations of mixed workloads are much greater than those of computing intensive workloads in Table 2 and Table 3. Except for the much greater power and energy variations among all hypervisors on the same platform running mixed workloads, the power variations during the execution of the mixed workload are also greater than those of computing intensive workloads.

C. *Insights on Energy Efficiency of Hypervisors*

From the above observations, we derive some insights for hypervisor selection in data centers in terms of power and energy consumption.

Insight #1: The hypervisor, hardware, and workload type are coupled with each other and such complication reminds the system designers of virtualized infrastructure and cloud data centers to carefully selection hypervisors.

Insight #2: The power and energy efficiency changes with the workload levels. One hypervisor consumes the lowest power or energy for lightly loaded workload may consume the highest power or energy if the workload level increases dramatically and vice versa.

Insight #3: ESXi should be deployed in a non-power sensitive environment for aggressive computing performance while KVM should be deployed in power sensitive environment to squeeze as many as possible virtual machines and achieve a reasonable computing performance. Moreover, ESXi uses power more aggressively to achieve high performance and high throughput, especially in highly contending conditions (for very heavy workload). ESXi and XenServer consume power more aggressively than Hyper-V and KVM on laptop (tailored mobile servers).

Insight #4: Typical 2U servers have higher idle power percentage and should be always running in highly heavy virtual machine workload since it its idle power percentage decreases when workload increases.

Insight #5: Although ARM64 server has lower max power than laptop, it has higher idle power percentage and less dynamic range. ARM64 server should be deployed in low-contending environment where virtual machines have light computing workload and power supply is very precious and there is no enough surplus power capacity for typical 2U servers. In other words, ARM64 server should also be deployed in a steady power usage server room, while laptop (mobile tailored server) should be deployed for highly dynamic workloads to leverage and complement the power fluctuation due to workload variations. For example, the ARM64 server configured with the KVM virtualization environment consumes 71.1% less power and 25.4% less energy than the HP DL380 G6 rack server (both packaged with 8 processor cores) for same computing workloads while the ARM64 server has 158% longer execution time.

IV. Related Work

To the best of our knowledge, we are the first to compare the energy efficiency of four mainstream hypervisors plus container engine across multiple platforms with different workload types and workload levels in the literature. We briefly review some related work in this section.

There are quite a few works on performance comparison of hypervisors [28, 29, 30, 32, 33, 34, 35]. All these work compare the performance of different hypervisors with each other or native physical performance on one platform. The difference of our work, however, is that we compare the energy efficiency of different hypervisors and container engine, on different platforms. There are some works on power or energy efficiency comparison of hypervisors [36, 37, 38, 39]. In [36, 37], they compared the energy overhead of Xen and KVM when running 3 virtual machines in total on a small server with 2GB memory and 500GB hard disk. The results reflect the energy overheads of Xen and KVM compared to physical machine. In [38], the author compared the power consumption of Xen, KVM, Docker and LXC when running up to 8 virtual machines on a desktop server with 4 processor cores and 12GB memory. They focused on the network traffic benchmark and the workload level is fixed. In [39, 40, 41], they compared the energy efficiency of Xen, KVM, and VMware when running dedicated high performance computing workloads. In [42, 43], the authors compared the power consumption of KVM, Xen and OpenVZ when running network transactions with those of non-virtualized environment. What distinguishes our work from those is that we compare the energy efficiency of XenServer, KVM, VMware ESXi, Hyper-V, and Docker container engine on various platforms, including high end rack servers, desktop server, laptop, and emerging ARM64 server. We run computing intensive and mixed web server benchmarks with varying workload levels on these platforms to investigate the energy efficiency difference among these hypervisors. We tried to mimic real multi-tenant cloud computing environments with massive virtual machines at different workload levels.

In summary, our study is more comprehensive with fine-grained power samples from five platforms. We also provide the power variation for different levels of workload on different platforms.

V. Conclusions

Understanding energy efficiency of hypervisors on different servers under different workloads can help data center designers and system operators in many folds, including system capacity planning, power shifting, virtual machine placement, migrations, and resource scheduling. In this paper, we conducted extensive experiments and measurements to investigate the power and energy characteristics of different mainstream hypervisors on different kind of servers. Experiment results show that hypervisors expose different power and energy characteristics on same hardware with same workloads. Moreover, different hypervisor poses different attributes and aligns with different workload types and workload levels. And they may be deployed in different workload levels in different power situations. Our experiments also show that container virtualization is light-weight in terms of system implementation and maintenance, but not essentially more power efficient than traditional virtualization technology. Last but not least, ARM64 servers do have lower power consumption, but they finish computing jobs with much longer execution time and sometimes consume more energy. Laptop processors and motherboards are strong competitors to ARM64 server both in power and energy consumption.

As future work, we plan to characterize the instruction execution of different hypervisors that results in the power and energy differences presented in this paper.

Acknowledgment

This work is supported by Natural Science Foundation of China (No.61472109, 61003077, 61402140, 61572163), and Natural Science Fund of Zhejiang Province (LY14F020045, LY16F020018, Y1101092). This work is also supported in part by National Science Foundation (NSF) grant CNS-1205338 and CNS-1561216, and by the Introduction of Innovative R&D team program of Guangdong Province (No. 201001D0104726115).

References

[1] S. Subramanya, Z. Mustafa, D. Irwin, Prashant Shenoy. Beyond energy-efficiency: evaluating green datacenter applications for energy-agility. ACM/SPEC ICPE 2016, pp. 185-196.

[2] T. Mastelic, A. Oleksiak, H. Claussen, et al. Cloud computing: survey on energy efficiency. ACM Computing Surveys, 2014, 47(2): 1-36.

[3] P. Pande, F. Clermidy, D. Puschini, Imen Mansouri,Paul Bogdan,Radu Marculescu,Amlan Ganguly. Sustainability through massively integrated computing: are we ready to break the energy efficiency wall for single-chip platforms?. Proceedings of the Design, Automation and Test in Europe (DATE), 2011, pp.1656-1661.

[4] Z. Zhang and S. Fu. Characterizing power and energy usage in cloud computing systems. IEEE CloudCom 2011, pp.146-153.

[5] S. Deb, K. Chang, X. Yu, et al, Design of an energy efficient CMOS compatible NoC architecture with millimeter wave wireless interconnects, IEEE Transactions on Computers, 2012, 62(12):2382-2396.

[6] D. Fan, Z. Tang, H. Huang, Hailin Huang, G. R. Gao. An energy efficient TLB design methodology. Proceedings of the 2005 international symposium on Low power electronics and design (ISLPED), 2005, pp.351-356

[7] R. Rodrigues, A. Annamalai, I. Koren, et al Reducing energy per instruction via dynamic resource allocation and voltage and frequency adaptation in asymmetric multicores, IEEE ISVLSI 2014, pp.436-441.

[8] M. Bi, S. Chandrasekharan, and C. Gniady. IAMEM: interaction-aware memory energy management. USENIX ATC 2013, pp.267-278.

[9] Q. Deng, D. Meisner, A. Bhattacharjee, et al. CoScale: coordinating CPU and memory system DVFS in server systems. IEEE/ACM MICRO-45,2012, pp.143-154.

[10] D. Tiwari, S. Vazhkudai, Y. Kim, et al. Reducing data movement cost using energy-efficient active computation on SSD. USENIX FAST 2013, pp.119-132.

[11] Y. Zhang, J. Liu, M. Kandemir. Software-directed data access scheduling for reducing disk energy consumption. Proceedings of the 2012 IEEE 32nd International Conference on Distributed Computing Systems (ICDCS),2012,pp.596-605.

[12] R. Krishnaswamy, V. Nagarajan, K. Pruhs, et al. Cluster before you hallucinate: approximating node-capacitated network design and energy efficient routing. ACM STOC 2014, pp.734-743.

[13] Y. Ohsita, and M. Murata, Optical data center networks; architecture, performance, and energy-efficiency, Handbook on Data Centers, 2015.

[14] T. Banerjee and S. Ranka, A genetic algorithm based autotuning approach for performance and energy optimization, IEEE IGSC 2015, pp.1- 8.

[15] C. Wang, N. Nasiriani, G. Kesidis, et al. Recouping energy costs from cloud tenants: tenant demand response aware pricing design. ACM e-Energy 2015, pp.141-150

[16] Y. Guo, D. Zhu and H. Aydin, Reliability-aware power management for parallel real-time applications with precedence constraints, IEEE IGCC 2011, pp.1-8.

[17] D. Tiwari, S. Boboila, S. Vazhkudai, et al. Active flash: towards energy-efficient, in-situ data analytics on extreme-scale machines . USENIX FAST, 2013, pp.119-132

[18] H.-K. Chen, X.-M. Zhu, J.-H. Zhu, et al.Towards energy-efficient scheduling for real-time tasks under uncertain a cloud computing environment. Journal of Systems and Software, 2015, 99(1): 20-35.

[19] H. Sun, P. Stolf, JM. Pierson, et al. Energy-efficient and thermal-aware resource management for heterogeneous datacenters. Sustainable Computing: Informatics and Systems, 2014, 1-15.

[20] W. Huang, M. Allen-Ware, J.B. Carter, et al. TAPO: Thermal-aware power optimization techniques for servers and data centers. IEEE IGCC 2011, pp.1-8.

[21] S. Yeo, M. M. Hossain, J.-C. Huang,et al. ATAC: ambient temperature-aware capping for power efficient datacenters. ACM SOCC 2014, pp.1-14.

[22] J. Carter, K. Rajamani, Designing energy-efficient servers and data centers, Computer, 2010, 43(7): 76-78.

[23] A. Gandhi, V. Gupta, M. Harchol-Balter,et al. Optimality analysis of energy-performance trade-off for server farm management. Performance Evaluation, 2010, 67(11):1155-1171

[24] I. Goiri, W. Katsak, K. Le, et al. Designing and managing datacenters powered by renewable energy. IEEE Micro, 2014,34(3):8 -16

[25] L. Li, W. Zheng , X. Wang , et al. Data center power minimization with placement optimization of liquid-cooled servers and free air cooling.Sustainable Computing: Informatics and Systems, 2016.

[26] A. Khosravi, S. Kumar Garg, R. Buyya. Energy and carbon-efficient placement of virtual machines in distributed cloud data centers. LNCS,vol 8097, pp.317-328.

[27] M. A. Islam and S. Ren. A new perspective on energy accounting in multi-tenant data centers. USENIX CoolDC 2016.

[28] I. Ahmad, J.Anderson, A. Holler, A. An analysis of disk performance in VMware ESX server virtual machines. IEEE WWC 2003.

[29] T. Deshane, Z. Shepherd, J. Matthews,et al, Quantitative comparison of Xen and KVM. XenSummit 2008.

[30] R. Morabito, J. Kjallman, M. Komu. Hypervisors vs. lightweight virtualization: a performance comparison. IEEE IC2E 2015, pp. 386-393.

[31] Natural Resources Defense Council. Data center efficiency assessment: scaling up energy efficiency across the data center industry:evaluating key drivers and barriers, 2014.

[32] J. Hwang, S. Zeng, F. Wu, Timothy Wood, A component-based performance comparison of four hypervisors. IFIP/IEEE IM 2013, pp. 269 - 276.

[33] J. Li ,Q. Wang,D. Jayasinghe, et al. Performance overhead among three hypervisors: an experimental study using hadoop benchmarks. IEEE Big Data 2013, pp.9-16.

[34] VMware. A performance comparison of hypervisors. https://www.vmware.com/pdf/hypervisor_performance.pdf

[35] W. Felter, A. Ferreira, R. Rajamony, et al. An updated performance comparison of virtual machines and linux containers.IBM Research Report, RC25482, 2014.

[36] Y. Jin, Y. Wen, Q. Chen, et al. An empirical investigation of the Impact of server virtualization on energy efficiency for green data center.The Computer Journal, 2013:1-14

[37] Y. Jin,Y. Wen,Q. Chen. Energy efficiency and server virtualization in data centers: an empirical investigation. INFOCOM WKSHPS 2012, pp.133-138.

[38] R. Morabito. Power consumption of virtualization technologies: an empirical investigation. IEEE/ACM UCC 2015, pp. 522-527.

[39] S. Varrette ,M. Guzek, V. Plugaru, et al. HPC performance and energy-efficiency of Xen, KVM and VMware Hypervisors. SBAC-PAD 2013, pp.89-96.

[40] M. Guzek , S. Varrette, V. Plugaru, A holistic model of the performance and the energy-efficiency of hypervisors in an HPC environment, LNCS, vol8046, pp. 133-152.

[41] M. Guzek , S. Varrette, V. Plugaru, et al. A holistic model of the performance and the energy-efficiency of hypervisors in an HPC environment. Concurrency and Computation: Practice and Experience 2014; 26(15):2569–2590.

[42] R. Shea,H. Wang, J. Liu. Power consumption of virtual machines with network transactions: measurement and improvements. IEEE INFOCOM 2014, pp. 1051-1059.

[43] C. Xu,Z. Zhao, H. Wang, Energy efficiency of cloud virtual machines: from traffic pattern and CPU affinity perspectives. IEEE Systems Journal, 2015, PP(99) 1-11.