Cloud Energy Consumption Measurement and Reduction: an Overview of Methods

Martin Warnaar
Student number: 2067072
Vrije Universiteit Amsterdam

February 1, 2016

Abstract

As cloud computing is widely being used and still growing in terms of customers and providers, energy consumption and carbon dioxide production of data centers that are hosting cloud infrastructures has become a topic of importance over the last few years. Cloud consumers are, at the time of writing, only billed for used resources and time, not the consumed energy. In this article, an overview is given of methods in the literature that aim to determine the energy efficiency of data centers at different levels, as well as possible ways of improving their efficiency.

1 Introduction

Cloud computing has emerged as an infrastructure that allows its customers to purchase a specific set of resources at the moment they need it, instead of renting a fixed amount of physical server racks, providing support for *on-demand* computing. At any moment, they can decide to increase or decrease the resource capacity, allowing them to scale their applications to the demand using a *pay-as-you-go* model [23]. This is often done by allowing consumers to buy resources through different types of well-known services: Infrastructure as a Service (IaaS), Platform as a Service (PaaS), Software as a Service (Saas), and the more recent Container as a Service (CaaS) [18]. These forms of services are mostly realised through virtualisation technology, allowing consumers and providers to flexibly configure the resources to their needs. Virtualisation of hardware provides two techniques that make cloud computing to be considered energy efficient compared to renting bare-bone server racks: server consolidation and live migration [22]. These techniques allow migrating running virtual machines to different servers, so that higher server utilisation can be achieved and idle nodes can be shutdown or put in energy-preserving modes.

However, because cloud computing is becoming so large, more fine-grained energy efficiency optimisations within datacenters can still vastly decrease operating costs for cloud providers. The overall energy consumption of cloud computing is expected to be 1963 billion kWh by 2020, compared to 632 billion kWh in 2007 [4].

In this paper, an overview is given of existing methods in the literature that aim to measure and improve the energy efficiency of datacenters and cloud hosts. In section 2, the description of the conducted method for this literature study is given. Section 3 describes methods to determine and predict the energy consumption of data centers at different levels. Different methods to achieve higher energy efficiency are shown in section 4. Finally, we conclude in section 5.

2 Literature study method

This literature study was conducted by initially specifying inclusion and exclusion criteria that an article has to fulfill. If a paper matches all inclusion criteria and none of the exclusion criteria, we can use it as a resource for the study. By formulating a sound search query such as "(intitle:"virtual" OR intitle:"vm" OR intitle: "cloud") power electricity energy consumption measurement method tool" and using it on Google Scholar [8], we get a list of possible results. The problem with this is that this syntax is specific for Google Scholar. If we want to perform a search in a similar way on IEEE Xplore [11] for example, the syntax of the query and the resulting set of articles are different. To retrieve a similar result set to that of Google Scholar, the query that was used was "("Document Title":virtual OR "Document Title":vm OR "Document Title":cloud AND "Document Title":survey OR "Document Title":tutorial) AND (VM virtual machine power consumption electricity method tool)."

By filtering the results manually using the inclusion and exclusion criteria, we systematically get a list of papers that we are interested in. Examples of inclusion criteria are: the main subject of the paper is about virtualised environments, the main focus is about measuring or reducing power consumption, and the paper is written in English. Examples of used exclusion criteria are: The described virtualised environment is not a language-specific environment (such as the JVM), or the article does not apply to today's hardware.

As some papers provided more information about these specific subjects than others, such as [12] and [16], the references of these articles provided some additional resources that were not included in the initial result sets. The articles matching the criteria could therefore be used in this literature study as well.

Most of the articles about energy reducing methods for virtualised cloud environments also describe methods on how energy consumption of virtualised environments is determined in the first place. The distinction made between energy measurement and reduction in this article is mainly introduced for clarity, as there are multiple ways for doing both. Each detailed subsection has the article(s) named from which the information is originated.

3 Measuring Energy Consumption

Measuring the energy consumption of virtualised hardware is far from trivial. Resource contention of multiple applications trying to use hardware can lead to cause significant overheads. On the other hand, resource sharing can also lead to a decrease in energy consumption, as is the case in sharing a multicore CPU for example. Accessing shared memory CPU caches by multiple cores often leads to a performance increase, while less energy is being consumed.

In this section, we make a distinction between utilisation-based energy prediction models and performance monitoring counter-based (PMC) energy prediction models. In utilisation-based models, linear regression is often applied to a specific set of hardware utilisation statistics and the measured energy consumption. In PMC-based models, logged (virtualised) hardware event counters are used to form a prediction model. In general, it is easier to obtain OS-provided inputs such as CPU, disk, and memory utilisation than to obtain hardware event counters [12, 13, 17, 19, 20]. However, these models can often not be generalised to all hardware setups, as efficiency of hardware components differ greatly. Therefore, using relative PMC (rPMC) as is done in [21] can offer a more generic and less error-prone model for estimating the energy consumption for workloads.

3.1 Utilisation

OS performance counters such as the CPU utilisation rate, disk read/write speeds, memory and network interface controller (NIC) usages are often used to determine to what extent applications are consuming hardware resources.

Using these measurements together with the measured consumed energy, energy consumption prediction models can be formed. In this section, we consider how these estimation models can be used in cloud platforms, and the problems that arise when doing so. Application profiling is also considered, in which key temporal resource requirements of application workflows are determined, which can then be used for consolidation of VMs as described in section 4.1.

3.1.1 Estimation Models

The power consumption of computers is not proportional to the work they accomplish [16]. The amount of idle nodes should be kept to a minimum in data centers, because they consume a considerate amount of energy while not performing any operations. It can be the case that a server does need to be running, but only low performance is required. Dynamic voltage and frequency scaling (DVFS) can then help in reducing the energy cost, allowing it to run in a lower power mode by downscaling the CPU voltage and frequencies [14].

In energy consumption estimation models, the energy consumption is often determined as the sum of two components: the idle and dynamic energy consumptions. In order to do this, many articles in the literature describe using linear regression to create estimation models based on a set of OS performance counters [2, 12, 15, 16, 20]. The CPU is almost always the biggest energy consuming component of the system, which is why it is sometimes used as only input to form a linear model to predict the energy consumption. This is done to simplify the model and is adequate for some cases, depending on the context [20]. However, when dealing with strict service level agreements (SLAs) as is the case for cloud environments, such a simple model does often not suffice.

A possible problem of these linear regression models, and this also holds for PMC-based models, is that they heavily depend on the hardware that is used to generate them [16].

3.1.2 Application Profiling

Instead of measuring energy consumption at the (virtualised) hardware level, one can also form application profiles that describe the used resources and utilisation patterns of runnable tasks over time, as is done in [23]. This works specifically well for tasks for which the execution patterns are always nearly the same, as it is the case in weather predictions for example.

After forming such resource profiles, tasks that have different resource requirements can be consolidated to maximise server utilisations and save, in this example case, save up as much as 56% in resource and energy costs, while having a performance degradation of only 15%. This is done by calculating the distance between tasks in a multi-dimensional space, clustering the ones that have a distance smaller than a set threshold, and finally mapping them to VMs.

3.2 PMC

Instead of taking OS-provided inputs, hardware event counters can also be used to create linear regression models for estimating the resource and energy consumption of VMs and runnable tasks. This requires more model tuning, as it is much more fine grained than OS utilisation values.

3.2.1 Vmeter

Modelling the energy consumption for virtualised clouds using hardware event counters is described in [1]. Here, event counters of the virtualised CPU, cache, DRAM, and disk are used to form regression models using a hybrid approach. One part of the hybrid system is a set of hardware performance counters, while the second part is a disk monitoring utility program.

Using principal component analysis (PCA) on the input data set, a high correlation was found between the (CPU, cache) pair and energy consumption, and the (disk, DRAM) pair and energy consumption. Therefore, I/O-bound and CPU-bound processes were considered as the two major application categories, in which CPU-bound processes tend to have a higher energy consumption.

3.3 Evaluating Model-Based Power Characterization

In [15], the effectiveness of a combination of PMC and OS-input models is evaluated. Forming these types of regression models becomes increasingly complex due to multicore CPU optimisations, hidden device states and other dynamic power components. Moreover, linear regression models as we have described often work well for workloads that are trivial or homogenous, but in practice cloud applications tend to be dynamic and resource usage patterns can change rapidly over short periods of time.

Therefore, automatic feature selection is introduced; a method in which system components are selected that are statistically correlated to the energy consumption. Using this method, the energy model is formed using multiple linear regression models based on both OS-provided inputs as well as hardware event counters. This results in a 2-6% mean relative error for predicting the energy consumption based on multicore CPU scenarios.

4 Reducing Energy Consumption

Using the linear regression models and application usage patterns, the energy consumption of cloud hosting hardware can be reduced with different scheduling techniques and consolidation methods.

At lower levels, energy credit scheduling methods can help reduce the energy consumption, while also helping gain insights into which users or applications use which resources.

4.1 Virtual Machine Consolidation

Consolidation, see figure 1, is a technique in which (running) virtual machines are migrated to a different physical host, allowing for higher resource usages while idle nodes can be put in power-saving modes, or shutdown. "Ideally,

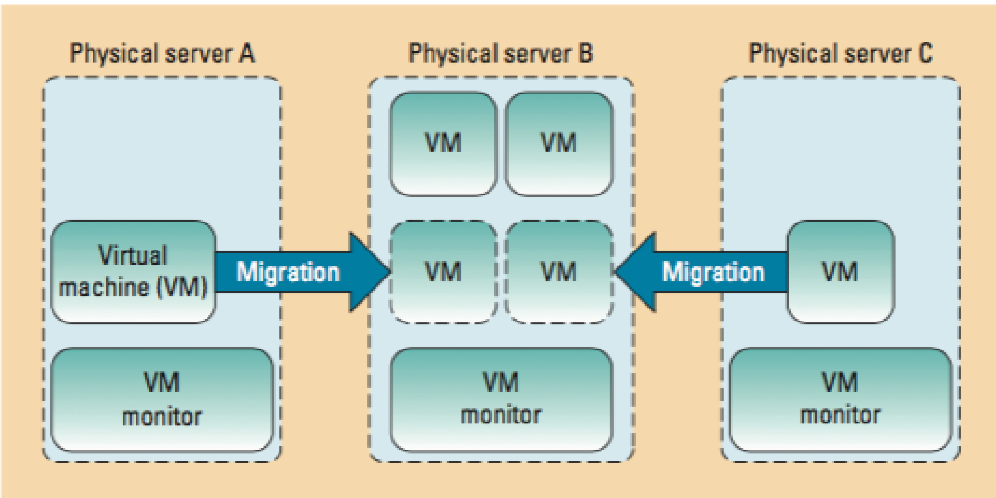

Figure 1: Consolidation

Source: [10]

due to the static energy consumed by server (especially the CPUs), running servers at the maximum utilization level is more energy efficient. This is typically referred to as the lack of energy-proportionality in modern server hardware." [10]

Again, this is beneficial to the energy consumption as long as resource contention does not occur, as it introduces overhead and thus energy efficiency will start decreasing. Therefore, using application resource usage patterns of VMs combined with hardware profiles to keep resource contention to a minimum, is a possible solution to improve the energy efficiency of cloud hosts.

4.1.1 Container as a Service

VMs can often be purchased with varying resource specifications, such as the memory capacity, CPU, and disk size. Container as a Service (CaaS) considers clustering sets of tasks that need to be executed, and running these clustered "containers" on VMs with suited specifications for each type of container. In [18], the public Google Cloud tracelog [9] is used to perform this clustering optimisation task using X-means, an extension of K-means.

A cluster consists of a group of tasks with usage similarities, for which virtual machine types can then be identified to run each cluster on. Other task specifications than the resource usage that are taken into account are the task length, scheduling class (production vs. development priority), and the submission rate.

4.1.2 Ant Colony Systems

VM consolidation can also be considered a multi-objective optimisation problem, taking the minimisation of power consumption and resource wastage as the objectives, as described in [5] and [6]. Resource wastage is seen as the sum of remaining unused resources of a physical host.

The goal is to obtain a non-dominated set of solutions, the Pareto set, using Ant Colony Optimization heuristics. "Ant Colony Optimization (ACO) is a metaheuristic inspired by the observation of real ant colonies and based upon their collective foraging behavior." [6] In order to do this, linear regression models for resource wastage and power consumption are formed, of which the results are used as the two objectives of the optimisation algorithm. The algorithm takes a requested set of VMs and physical hosts as input and calculates the desirability and probability of moving a VM to a specific host, while making sure that each host does not surpass its resource usage thresholds. When the Pareto optimal set has been calculated, VM consolidation can take place using the resulting VMs to hosts mapping.

4.2 Migration

Migration, the part of consolidation that is responsible for placing VMs to different hosts, does not come for free. Migrating a virtual machine to a different host imposes overheads on both servers, as data from the VMs has to be sent from one server to another. This possibly causes performance degradation not only for the migrating VM, but also for other VMs running on the same server because of resource contention [10].

There exist three variants of live VM migration. In *pre-copy* live migration, the destination VM is created while the source VM continues to operate. The source memory is copied to the destination host at first. Then, memory pages that have changed in-between are copied iteratively until all memory has been copied to the destination host. In the second variant, *post-copy*, the destination VM takes control directly and requests memory from the source when it cannot access it locally. A third, hybrid variant, has a pre-copy phase in which most of the memory is transferred to the destination, followed by a post-copy phase in which any modified memory is transferred [3]. These variants all have advantages and disadvantages. In *pre-copy*, the source VM remains untouched until the destination is ready to take over. However, as dirty pages are being copied until the memory states are converged at the destination, this could go on indefinitely.

In *post-copy*, there is a risk that the VM state can be lost entirely if networking errors occur, but the amount of memory that needs to be transferred will often be lower and cannot get into an infinite loop of memory convergence.

4.2.1 Sustainable Cloud Federation

So far, only VM placement and energy consumption within datacenters of the same provider have been considered. Instead, cloud federations allow for migrating VMs between multiple datacenters of different providers at multiple geographical locations.

A cloud federation consists of often smaller cooperating cloud providers that aim for sustainability, reducing carbon dioxide emissions, and lowering power costs [7].

As small and medium cloud providers can often not compete with the big cloud providers such as Amazon, Google, and Rackspace, they form a federation that lets each participant borrow external resources from each other, while minimising the energy costs and carbon-dioxide emissions of their datacenters. Power consumption can be minimised by taking the local power prices into consideration, while consolidation can take place on a larger scale than these smaller cloud providers were able to than before.

4.3 VM Scheduling

Instead of looking at virtual machine placement algorithms over multiple physical hosts, optimisations can also be made on the physical host level. Here, we consider VM scheduling that makes use of power budgeting: a limitation of power that a VM is allowed to consume within a certain time interval.

4.3.1 Energy-Credit Scheduling

In [13], a modification to the default Credit-scheduler of the Xen-hypervisor is proposed. By default, the scheduler uses a fair-share policy of resources for each VM. The modification, energy-credit scheduler (ECS), uses energy budgets per VM instead.

Each VM is assigned an energy budget per fiscal time interval. When this budget is exceeded, the system idles until the new fiscal time interval begins. In each fiscal time interval, the budgets for all VMs in the next fiscal time interval are calculated using an estimation of resource usage for the VMs using linear regression.

4.3.2 VPM Tokens

In [17], the VirtualPower framework is introduced. This framework uses so-called VPM tokens as an abstraction of resource utilisation, used for specifying power budget policies. These power budget policies are then implemented as custom VM scheduling schemes in the Xen hypervisor, allowing for a utility increase of up to 43% in the experimental results. The power budgeting policies are calculated with respect to the SLA that each VM is bound to.

The primary goal here is to improve the overall utilisation of the physical host. As an improvement of Energy-Credit scheduling, applications that don't use the entirety of their budget allow other applications running on that host to use the remaining budget by redistributing it. The budget manager takes care of converting token values to power limitation values, designating the normalised maximum system performance levels within the budget.

4.3.3 Hypervisor-based Energy Management

In [19], a framework for energy management on the operating system level is introduced. Using an energy-aware host OS allows for machine-wide energy constraints, while energy-aware guest OSs allow for using application-specific constraints.

Recursive power consumption of hardware components is taken into account in order to have detailed information about energy consuming events that took place in the system. In this model, the CPU, disk, host-level OS, and recursive power consumption are used to form a system-wide energy consumption model. Then, using the event logs within a time-based interval of hardware components such as the CPU and disk, the energy consumption for that interval is calculated.

By throttling the CPU allocation at runtime, CPU energy consumption of individual machines can be regulated. By employing a stride scheduling algorithm in the hypervisor, proportional CPU shares are allocated over the virtual processors. On the host-level, a simple policy for the energy management enforces device power limits per VM.

As noted in [12] on using hardware event and counter logs for enforcing power limits in this way; availability of detailed power models for each hardware component are not available in practice, and are in fact difficult to provide. So while enforcing power limits, assuming that detailed event counter information is available for all hardware of the (virtualised) system, is very precise in theory, it will be very hard if not impossible to implement in practice.

5 Conclusion

We have seen multiple methods on different levels and scales for measuring and reducing the energy consumption for data centers that are hosting cloud infrastructures. It depends on the availability and quantity of (virtualised) hardware information how energy and resource consumption models can be formed.

Different migration and consolidation algorithms can allow for higher server utilisations, allowing to put idle nodes in energy-preserving modes or shutting them down.

Using different VM-scheduling policies can also help in reducing energy consumption by either limiting the resources that a VM or physical host can use, or by increasing the overall utilisation and maximising the energy efficiency of physical hosts.

Models that seem promising in theory, can sometimes offer only limited use in practice because hardware information and event counter information can be scarce. Also, generalisability of linear regression models for estimating resource and energy consumption is not guaranteed, as hardware specifications change rapidly over time.

Furthermore, it depends on the application type how much the energy efficiency can be improved. If an application is performance-critical, SLAs are often stricter and consolidation becomes harder as it might be unacceptable that downtime occurs. When the usage pattern of an application is known upfront, this can be taken into account in the scheduling policy. The dynamic usage patterns of applications of typical cloud consumers make consolidation and resource estimations to be harder compared to applications that have a repeating usage pattern.

References

[1] A. Bohra and V. Chaudhary. Vmeter: Power modelling for virtualized clouds. In *2010 IEEE International Symposium on Parallel & Distributed Processing, Workshops and Phd Forum (IPDPSW)*, pages 1–8, 2010.

[2] Q. Chen, P. Grosso, K. van der Veldt, C. de Laat, R. Hofman, and H. Bal. Profiling energy consumption of vms for green cloud computing. In *2011 IEEE Ninth International Conference on Dependable, Autonomic and Secure Computing (DASC)*, pages 768–775, 2011.

[3] V. Cima, B. Grazioli, S. Murphy, and T. Bohnert. Adding energy efficiency to openstack. In *Sustainable Internet and ICT for Sustainability (SustainIT)*, pages 1–8, 2015.

[4] K. Djemame, D. Armstrong, R. Kavanagh, A. Ferrer, D. Perez, D. Antona, J.-C. Deprez, C. Ponsard, D. Ortiz, M. Macias, J. Guitart, F. Lordan, J. Ejarque, R. Sirvent, R. Badia, M. Kammer, O. Kao, E. Agiatzidou, A. Dimakis, C. Courcoubetis, and L. Blasi. Energy efficiency embedded service lifecycle: Towards an energy efficient cloud computing architecture. In *CEUR Workshop Proceedings. Joint Workshop Proceedings of the 2nd International Conference on ICT for Sustainability*, pages 1–6, 2014.

[5] M. H. Ferdaus, M. Murshed, R. N. Calheiros, and R. Buyya. Virtual machine consolidation in cloud data centers using aco metaheuristic. *Lecture Notes in Computer Science*, 8632:306–317, 2014.

[6] Y. Gao, H. Guan, Z. Qi, Y. Houb, and L. Liuc. A multi-objective ant colony system algorithm for virtual machine placement in cloud computing. *Journal of Computer and System Sciences*, 79(8):1230–1242, 2013.

[7] M. Giacobbe, A. Celesti, M. Fazio, M. Villari, and A. Puliafito. An approach to reduce carbon dioxide emissions through virtual machine migrations in a sustainable cloud federation. In *Sustainable Internet and ICT for Sustainability (SustainIT)*, pages 1–4, 2015.

[8] Google. Google scholar. `https://scholar.google.com`. Last visited: February 2016.

[9] Google. Google cluster workload traces v2. `https://github.com/google/cluster-data`, 2011. Last visited: February 2016.

[10] P. Graubner, M. Schmidt, and B. Freisleben. Energy-efficient virtual machine consolidation. *IT Professional*, 15(2):28–34, 2013.

[11] IEEE. Ieee xplore digital library command search. `http://ieeexplore.ieee.org/search/advsearch.jsp?expression-builder`. Last visited: February 2016.

[12] A. Kansal, F. Zhao, J. Liu, N. Kothari, and A. A. Bhattacharya. Virtual machine power metering and provisioning. In *Proceedings of the 1st ACM symposium on Cloud computing*, pages 39–50, 2010.

[13] N. Kim, J. Cho, and E. Seo. Energy-based accounting and scheduling of virtual machines in a cloud system. In *2011 IEEE/ACM International Conference on Green Computing and Communications (GreenCom)*, pages 176–181, 2011.

[14] T. Mastelic, A. Oleksiak, H. Claussen, I. Brandic, J.-M. Pierson, and A. V. Vasilakos. Cloud computing: Survey on energy efficiency. *ACM Computing Surveys (CSUR)*, 47(2), 2015.

[15] J. C. McCullough, Y. Agarwal, J. Chandrashekar, S. Kuppuswamy, A. C. Snoeren, , and R. K. Gupta. Evaluating the effectiveness of model-based power characterization. In *Proceedings of the 2011 USENIX conference on USENIX annual technical conference*, 2011.

[16] C. Mobius, W. Dargie, and A. Schill. Power consumption estimation models for processors, virtual machines, and servers. *IEEE Transactions on Parallel and Distributed Systems*, 25(6):1600–1614, 2013.

[17] R. Nathuji, K. Schwan, A. Somani, and Y. Joshi. Vpm tokens: virtual machine-aware power budgeting in datacenters. *Cluster Computing*, 12(2):189–203, 2009.

[18] S. Piraghaj, A. Dastjerdi, R. Calheiros, and R. Buyya. Efficient virtual machine sizing for hosting containers as a service. In *2015 IEEE World Congress on Services (SERVICES)*, pages 31–38, 2015.

[19] J. Stoess, C. Lang, and F. Bellosa. Energy management for hypervisor-based virtual machines, 2007.

[20] D. Versick, I. Wamann, and D. Tavangarian. Power consumption estimation of cpu and peripheral components in virtual machines. *ACM SIGAPP Applied Computing Review*, 13(3):17–25, 2013.

[21] P. Xiaoa, Z. Hub, D. Liua, G. Yana, and X. Qua. Virtual machine power measuring technique with bounded error in cloud environments. *Journal of Network and Computer Applications*, 36(2):818–828, 2013.

[22] P. Xiaoa, Z. Hub, D. Liua, X. Zhanga, and X. Qua. Energy-efficiency enhanced virtual machine scheduling policy for mixed workloads in cloud environments. *Computers & Electrical Engineering*, 40(5):1650–1665, 2014.

[23] Q. Zhu, J. Zhu, and G. Agrawal. Power-aware consolidation of scientific workflows in virtualized environments. In *Proceedings of the 2010 ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis*, pages 1–12, 2010.