A New Perspective on Energy Accounting in Multi-Tenant Data Centers[*]

Mohammad A. Islam
University of California, Riverside

Shaolei Ren
University of California, Riverside

Abstract

Energy accounting plays a crucial role in optimizing data center energy efficiency. Nonetheless, in a multi-tenant data center, it is challenging to *fairly* account for non-IT energy on an individual tenant basis, because each non-IT system (e.g., power supply and cooling) is shared by multiple tenants and only the system-level non-IT energy consumption can be measured. Existing policies, e.g., proportionally or equally attributing non-IT energy to tenants, may attribute different energy to two "equivalent" tenants and hence are not fair. In this paper, we propose QSEA, a quick Shapley value-based energy accounting policy for multi-tenant data centers. QSEA is provably fair and also easy to implement with little to zero overhead. We run trace-based simulations and demonstrate that, compared to the exact Shapley value approach that has an exponential complexity, QSEA yields almost the same energy accounting result while having a negligible computation time.

1 Introduction

Data center energy consumption is soaring quickly and has undeniably become a central issue under scrutiny. Naturally, to tame the growing energy demand, the first and also crucial step is to measure it (referred to as *energy accounting* in this paper). In fact, accounting for and reporting energy usage has been increasingly mandated by tightening government regulations, pressured by environmental groups like Greenpeace [1], and/or required by green certifications that bring tax/zoning benefits and are being pursued by many large data centers.

In addition to those that operate their own data centers (like Google), energy accounting is also crucial for companies that house servers in third-party data centers. Such data centers are called multi-tenant data centers (or colocations) and located worldwide, where multiple companies/organizations (called tenants) manage their own physical servers in a shared space and the data center operator controls the non-IT assets. Almost all industry sectors, including finance, major websites (e.g., Wikipedia), medium-scale cloud providers and even gigantic IT companies (e.g., Microsoft), lease spaces in multi-tenant data centers to house servers, either as their primary IT deployment method or to complement their own infrastructure. Further, multi-tenant data centers are increasingly serving as *edge* data centers to bring computation closer to users, reducing latency and bandwidth costs. In the U.S., multi-tenant data centers collectively consumed as five times energy as Google-type data centers combined altogether in 2013 [2].

Although tenants do not operate the entire data center facility, their energy footprint has still been receiving attention and under a closer scrutiny than ever. For example, Greenpeace has begun to include in its 2014 sustainability report the energy usage by both multi-tenant data center operators and tenants (e.g., Akamai) [1]. Large IT companies are being increasingly pressured to disclose their energy portfolio in leased data centers as part of the transparency fulfilment. Moreover, energy accounting is also critical for the data center operator to price tenants' energy usage in an accurate and fair manner.

While power meters are readily in place for measuring tenants' IT server energy (even on a per-server basis), it is non-trivial and also challenging to *fairly* account for tenants' non-IT energy, because the non-IT systems, like uninterruptible power supply (UPS), power distribution unit (PDU) and cooling, are each shared by multiple tenants and inherently non-divisible. Multi-tenant data centers are often located in metropolitan and even downtown areas, limiting the applicability of "free" cooling (e.g., outside air economizer) due to location and airflow constraints. Thus, in an average multi-tenant data center where centralized UPS and chillers are commonly used, the non-IT part can take up 30% or even 50% of the total energy [3, 4]. For example, an industry-leading data

---

[*] This work was supported in part by the NSF under grants CNS-1551661 and CNS-1565474.

center in London has a power usage effectiveness (PUE) of 1.35 [5], i.e., 0.35 kWh of non-IT energy is consumed for each kWh of IT energy.

Currently, data center operator typically attributes the non-IT energy consumption to different tenants in proportion to their metered IT energy usage or their subscribed power capacities. These energy accounting policies, albeit simple, are not *fair*, in the sense that they do not accurately reflect the marginal contribution of a tenant in the total non-IT energy consumption [6]. This is because the non-IT energy grows *non-linearly* with the IT energy load (especially at low utilization levels due to infrastructure redundancy) and also has a *static* energy when active [7,8]. Section 3.2 provide more details.

To achieve fairness, we account for individual tenants' non-IT energy consumption by drawing upon the tool of Shapley value from game theory [6]. Shapley value has been proven to be the only *fair* method for cost/payoff sharing (in the sense of satisfying a set of desired axiomatic principles) [9]. Nonetheless, applying it in our context has two major challenges. *First*, calculation of Shapley value requires the knowledge of non-IT energy consumption had only a subset of tenants been present. In practice, however, we can only measure the energy consumption for each non-IT unit as a whole. *Second*, and more critically, deriving Shapley value has an exponential complexity that quickly becomes computationally prohibitive even for 10+ tenants, but a practical multi-tenant data center often has several tens or even hundreds of tenants.

In this paper, we address the challenges of using Shapley value for fair energy accounting in multi-tenant data centers. Specifically, we approximate the non-IT energy consumption as a quadratic function of the tenants' IT energy, which allows us to account for tenants' non-IT energy in real time and also has a high accuracy (compared to the original Shapley value calculation). We call our method QSEA (Quick Shapley value-based Energy Accounting). QSEA is not only provably fair, but also easy to implement with an interesting insight: it attributes dynamic energy of non-IT systems to tenants in proportion to their IT energy usage, and equally splits the static energy of non-IT systems among all active tenants. We also run extensive trace-based simulations to demonstrate the efficiency of QSEA and contrasting it with other energy accounting methods.

2 Model and Problem Statement

In a typical multi-tenant data center as illustrated in Fig. 1, utility power first enters into a centralized UPS through an automatic transfer switch (ATS). Then, the UPS feeds protected power (also called critical power) into multiple PDUs, each supporting a few tens of racks that are connected to tenants' servers.

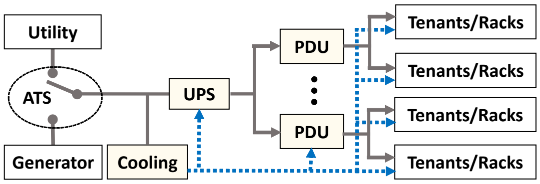

Figure 1: Data center infrastructure.

Model. In our notations, we omit the time index for presentation clarity, while noting that energy accounting needs to be performed periodically (e.g., every minute).

There are N tenants and M different non-IT units.[1] The entire set of tenants is denoted by $\mathcal{N}$. The non-IT units under consideration include PDU, UPS and cooling system, as highlighted in shaded boxes in Fig. 1. Each non-IT unit serves multiple tenants, while each tenant affects energy consumption of multiple non-IT units.

We denote $\mathcal{M}_i$ as the set of non-IT units whose energy is affected by tenant i. For example, if tenant i's IT energy affects the energy consumption/loss by one UPS, one PDU and one air economizer, then these three units will be included in the set $\mathcal{M}_i$. Letting $\mathcal{N}_j \subseteq \mathcal{N}$ be the subset of tenants that affect non-IT unit S_j, energy consumption of S_j can be written as $P_{S_j} = F_j(\sum_{i\in\mathcal{N}_j} P_{T_i})$, where P_{T_i} is energy consumption by tenant i and $F_j(\cdot)$ is referred to as the *energy function* that relates tenants' energy consumption to that of non-IT unit S_j. Denote $\Phi_{ij} \geq 0$ as tenant i's share of non-IT energy consumed by unit j, and thus we have $P_{S_j} = \sum_{i\in\mathcal{N}_j} \Phi_{ij}$.

Problem statement. Our problem can be formally stated as: *How to fairly determine tenant i's non-IT energy* $\Phi_i = \sum_{j\in\mathcal{M}_i} \Phi_{ij}$*?* It is non-trivial to *fairly* decompose $F_j(\cdot)$ into multiple shares, each for one tenant, because one can only measure the energy consumption P_{S_j} of a shared non-IT unit as a whole. Further, non-IT unit j's energy function $F_j(\cdot)$ is non-linear in tenants' IT energy (e.g., UPS energy loss grows quadratically in its load [7] and outside air economizer energy grows cubically [3]). Note that, when a non-IT unit operates at a high utilization (e.g., 90%), its static energy is less significant and dynamic energy becomes almost linear [7]. Nonetheless, non-IT units typically run at a fairly low utilization (often less than 40%) because of infrastructure redundancy in practice. For example, for a commercial data center with "2N" redundancy, every non-IT unit is duplicated and equally shares the IT load, resulting in a maximum utilization of 50% given a full IT load. Thus, static energy and non-linearity of dynamic energy of non-IT units cannot be ignored for energy accounting.

---

[1] Each *tenant* represents a minimum set of servers (typically one rack) that are served by the operator's non-IT infrastructure. Hence, in practice, a single tenant with a large cluster corresponds to multiple "tenants" in our model.

3 QSEA: Quick Shapley Value-Based Energy Accounting

In this section, we present QSEA, a quick energy accounting policy based on Shapley value. QSEA is provably fair and also easy to implement: it attributes dynamic energy of non-IT systems to tenants in proportion to their IT energy usage, and equally splits the static energy of non-IT systems to active tenants.

3.1 Axiomatic Principles for Energy Accounting

The non-IT energy accounting problem can be essentially viewed as *attributing a shared cost/payoff to individual players*, which is a classic problem in economics with *fairness* as a key consideration [6, 9]. While there is no uniform definition for fairness, prior studies have commonly used a set of four axiomatic principles, and an allocation policy (i.e., energy accounting policy) satisfying all of them is said to be *fair*. Below, we introduce these four axioms and explain them in our context.

Efficiency. The sum of accounted non-IT energy by individual tenants is equal to the total non-IT energy, i.e., $\sum_{i\in\mathcal{N}_j}\Phi_{ij} = P_{S_j}, \forall j \in \mathcal{M}$.

Symmetry. If two tenants are interchangeable and indistinguishable for their contribution to non-IT energy increase, they should account for the same non-IT energy, i.e., if $F_j(\sum_{l\in\mathcal{X}\cup\{i\}} P_{T_l}) = F_j(\sum_{l\in\mathcal{X}\cup\{k\}} P_{T_l})$ for any $\mathcal{X} \subseteq \mathcal{N}_j \setminus \{i,k\}$, then $\Phi_{i,j} = \Phi_{k,j}$.

Null player. If the non-IT energy does not change when we add or remove a tenant, zero non-IT energy should be attributed to this tenant (also called null player in a game), i.e., if $F_j(\sum_{l\in\mathcal{X}\cup\{i\}} P_{T_l}) = F_j(\sum_{l\in\mathcal{X}} P_{T_l})$ for any $\mathcal{X} \subset \mathcal{N}_j$, then $\Phi_{i,j} = 0$.

Additivity. The total non-IT energy attributed to a tenant over time should be the sum of energy consumption by each individual non-IT unit attributed to this tenant at each time instance (or sub-accounting period).

3.2 Existing Energy Accounting Policies

We now discuss the existing energy accounting policies.

Policy #1: *Based on Average IT Energy.* The total non-IT energy is attributed in proportion to each tenant's average IT energy over a predefined billing/accounting period (e.g., a month). This policy is commonly used for charging tenants' energy consumption.

Policy #2: *Based on Short-term IT Energy.* The billing period is divided into multiple time slots (e.g., hourly), during each of which tenants are accountable for non-IT energy in proportion to their short-term IT energy.

Policy #3: *Based on Power Subscription.* Tenants subscribe and pay for a power capacity in multi-tenant data centers. This policy accounts for tenants' non-IT energy in proportion to their subscribed capacities, regardless of their actual IT energy consumption.

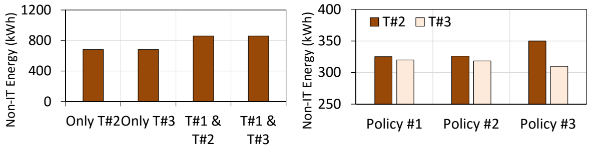

(a) (b)

Figure 2: **(a)** Symmetry condition for tenants #2 and #3. **(b)** Non-IT energy accounting.

In addition, some multi-tenant data centers also use *equal distribution* as their energy accounting policy, by equally dividing the non-IT energy among all the tenants. This method is clearly not fair, as it disregards the differences in actual IT energy usage by different tenants (violating the null player axiom). Thus, we do not discuss it in this paper for brevity.

Example of axiom violation. Now, we show that the existing policies violate the symmetry axiom. As an example, we consider the measurement data of a real double-conversion UPS reported by [7]. This type of UPS incurs energy losses during AC/DC and DC/AC conversions while feeding power into PDUs. In particular, the UPS energy loss (as a percentage of its full load) can be expressed as $F(x) = 0.03455 \cdot x^2 + 0.00959 \cdot x + 0.03234$, where x is the UPS load [7]. We see that there exists a quadratic term and a static term in the UPS energy loss: the quadratic term is due to the UPS's circuit heat that increases quadratically with current (which is roughly linear with the load), while the static term represents idle energy to keep UPS active [7].

To illustrate the symmetry violation by the existing energy accounting policies, we consider an example of three tenants, with a power subscription of 150KW, 130KW and 115KW, respectively. We consider 192 time slots (each representing 15 minutes), and sum up the non-IT energy by each tenant accounted over time. We generate tenants' power traces based on collected power usage data in a real data center. We then re-scale the data such that the three tenants have an average power capacity utilization of 72%, 65% and 72%, respectively, which results in the satisfaction of symmetry condition for tenants #2 and #3 (shown in Fig. 2(a)). Specifically, the symmetry condition requires $\sum_t F^t(P_{T_2}) = \sum_t F^t(P_{T_3})$ and $\sum_t F^t(P_{T_2} + P_{T_1}) = \sum_t F^t(P_{T_3} + P_{T_1})$, where t represents time index. Then, according to the symmetry axiom, a *fair* energy accounting policy should attribute equal energy to both tenants #2 and #3. As shown in Fig. 2(b), however, all the policies produce different energy shares for these two tenants. This is because static energy should be equally split among tenants (shown in Section 3.3), whereas the existing policies incorrectly split the static energy in proportion to tenants' IT energy (or power subscription). Further, Policy #3 disregards tenants' actual IT energy when attributing dynamic energy.

Therefore, the existing energy accounting policies are not *fair*, in the sense that they violate the symmetry axiom and do not reflect tenants' true contribution to non-IT energy [6, 9].

3.3 Shapley Value-Based Energy Accounting

Now, we develop our new energy accounting policy, called QSEA, which builds upon Shapley value [6, 9].

Shapley value. Fairly sharing the total cost/payoff among multiple players has been studied extensively in the economics literature. The key result is that Shapley value is the **only** allocation rule that satisfies the four axiomatic principles listed in Section 3.1 [6, 9]. Applying Shapley value in our context, the energy share of non-IT unit j attributed to tenant i is calculated as

$$\begin{aligned}\Phi_{ij} = &\sum_{X \subseteq \mathcal{N}_j \setminus \{i\}} \frac{|X|!(|\mathcal{N}_j|-|X|-1)!}{|\mathcal{N}_j|!} \\ &\cdot [F_j(P_X+P_{T_i})-F_j(P_X)]\end{aligned} \tag{1}$$

where X is a subset of tenants (excluding tenant i) in $\mathcal{N}_j$ supported by non-IT unit j, and $P_X = \sum_{k \in X} P_{T_k}$. The complexity of calculating Φ_{ij} is $\mathcal{O}(2^N)$, resulting in a total complexity of $\mathcal{O}(M \cdot N \cdot 2^N)$ for energy accounting.

We explain the implication of Shapley value in (1) as follows. Suppose that tenants join the non-IT unit sequentially, and consider a certain subset X of tenants that have already joined the non-IT unit j before tenant i. Then, $F_j(P_X+P_{T_i})-F_j(P_X)$ is the marginal contribution of tenant i to the non-IT unit j's energy increase. Note that the subset X of tenants can join the system in $|X|!$ ways due to all different permutations, while the tenants that join the non-IT unit after tenant i can happen in $(|\mathcal{N}_j|-|X|-1)!$ ways. The term $|\mathcal{N}_j|!$ in the denominator is to take the average of all the possible permutations of tenants joining the non-IT unit. Thus, by taking the average, the marginal contribution of tenant i is obtained.

While Shapley value satisfies all the axioms, applying it for our problem has two major challenges.

Challenge 1: We see from (1) that energy accounting based on Shapley value requires the value of $F_j(\sum_{k \in X} P_{T_k})$, i.e., the non-IT unit j's energy consumption when only a subset of tenants are connected. In practice, however, we can only measure non-IT unit j's total energy consumption $P_{S_j} = F_j(\sum_{i \in \mathcal{N}_j} P_{T_i})$ as a whole.

Challenge 2: Shapley value requires an exponential number of calculations to get the non-IT energy share for only one tenant. This is because Shapley value averages over all the possible combinations of tenants in the system, resulting in an intolerable computational complexity (e.g., over 20 minutes for only 18 tenants, whereas a real multi-tenant data center may have several tens or even hundreds of tenants).

Our solution. We propose a novel quick and fair energy accounting method, called QSEA. More concretely, QSEA leverages a quadratic function to *approximate* energy usage of each non-IT unit as follows:

$$F_j(x) = \begin{cases} 0, & \text{when } x \leq 0 \\ a_j \cdot x^2 + b_j \cdot x + c_j, & \text{otherwise} \end{cases} \tag{2}$$

where x is the total IT energy by tenants served by non-IT unit j, and a_j, b_j, and c_j are modeling parameters that we learn and calibrate online as we measure the non-IT unit j's energy. Note that the quadratic approximation comes from real-world measurements [4, 7] and allows a quick calculation of Shapley value. Even though certain type of non-IT units (e.g., outside air economizer) do not follow a quadratic energy function, we show through simulations that QSEA is still fairly accurate compared to the exact Shapley value approach.

When tenant i has a zero IT energy, its non-IT energy is clearly also zero, according to the null player axiom. Now, we consider the case when tenant i has a non-zero IT energy during an accounting period (e.g., every minute). By applying the quadratic function $F_j(x)$ into (1) and letting $|X| = r_X$ and $|\mathcal{N}_j'| = n_j$, where $\mathcal{N}_j' \subseteq \mathcal{N}_j$ is the set of tenants that have non-zero IT energy, we derive tenant i's energy share of non-IT unit j as

$$\begin{aligned}\Phi_{ij} = &\frac{2a_jP_{T_i}}{n_j!} \sum_{X \neq \emptyset, X \subseteq \mathcal{N}_j' \setminus \{i\}} [r_X!(n_j-r_X-1)!P_X] \\ &+ a_jP_{T_i}^2 + b_jP_{T_i} + \frac{c_j}{n_j},\end{aligned} \tag{3}$$

where $P_X = \sum_{k \in X} P(T_k)$. Note that over all nonempty subsets $X \subseteq \mathcal{N}_j' \setminus \{i\}$ that have the same size/cardinality of u, each tenant i's IT energy P_{T_i} appears $\binom{n_j-2}{u-1} = \frac{(n_j-2)!}{(u-1)!(n_j-u-1)!}$ times. Using this, we have

$$\begin{aligned}&\sum_{X \neq \emptyset, X \subseteq \mathcal{N}_j' \setminus \{i\}} [r_X!(n_j-r_X-1)!P_X] \\ =&\sum_{u=1}^{n_j-1} \sum_{X, s.t., |X|=u} [u!(n_j-u-1)!P_X] \\ =&\sum_{u=1}^{n_j-1} u(n_j-2)! \sum_{k \in \mathcal{N}_j' \setminus \{i\}} P_{T_k} = \frac{n_j!}{2} \sum_{k \in \mathcal{N}_j' \setminus \{i\}} P_{T_k}.\end{aligned} \tag{4}$$

Next, by plugging (4) into (3), the share of non-IT unit j's energy that QSEA attributes to tenant i is derived as

$$\Phi_{ij} = \begin{cases} 0, & \text{if } P_{T_i} = 0 \\ P_{T_i} \cdot \left[a_j \sum_{k \in \mathcal{N}_j'} P_{T_k} + b_j\right] + \frac{c_j}{|\mathcal{N}_j'|}, & \text{otherwise.} \end{cases}$$

Now, we discuss two important properties of QSEA. **First**, QSEA is fair, as it is derived based on Shapley value and hence satisfies all the four axioms in an approximate sense (due to the usage of approximate quadratic energy functions). If the non-IT energy function is indeed quadratic, then QSEA strictly satisfies the four axioms. **Second**, QSEA offers a closed-form expression of energy accounting with an interesting insight: the static energy of a non-IT unit is equally split among all the served tenants with non-zero energy, while the dynamic energy is attributed in proportion to tenant's IT energy usage (since the term "$a_j \sum_{k \in \mathcal{N}_j'} P_{T_k} + b_j$" is the same for all tenants served by non-IT unit j). In other words, QSEA is very easy to implement with little to zero overhead, by combining two existing non-IT energy accounting policies (i.e., proportional for dynamic energy and equal for static energy). Therefore, QSEA can be applied at runtime, as opposed to the original Shapley value approach in (1) that is even infeasible for 10+ tenants due to its exponential complexity (see Fig. 5(a)).

4 Performance Evaluation

In this section, we conduct a trace-based simulation to evaluate QSEA and show its differences from the existing energy accounting policies. We also demonstrate that QSEA is reasonably accurate compared to the exact Shapley value, while having a negligible complexity.

4.1 Setup

As illustrated in Fig. 1, we simulate the full hierarchy of data center power distribution with "2N" redundancy (i.e., duplicating each unit), including UPS and PDU. The cooling load includes both IT energy and energy losses by UPS and PDUs. We consider two different cooling systems: chiller-based cooling and outside air economizers, whose energy consumptions are quadratic and cubic in cooling load, respectively, as reported by [3,4]. In our evaluation, the centralized UPS supports four PDUs, each serving three tenants. The UPS and PDU capacities are set as 1MW and 250kW, respectively. Operated in $2N$ redundancy, each UPS/PDU equally shares the load. The 12 tenants served by the four PDUs have power subscription capacities in kW as follows: $\{25, 75, 150\}, \{75, 75, 100\}, \{125, 100, 25\}, \{25, 75, 150\}$, where each bracket represents one group of tenants served by one PDU. QSEA also applies to larger data centers, and we omit the results for brevity.

In the simulator, we create a data center energy module that simulates the energy consumption of each non-IT unit (including PDU, UPS and cooling system) given the tenants' IT energy usage as inputs. The input power trace for tenants in our simulation is collected from rack-level measurement every 15 minutes in a real data center, and it is scaled to our system configuration with an average power capacity utilization of 70% for each tenant. We also use the produced non-IT energy consumption to learn and calibrate our approximated quadratic energy functions in (2) based on minimum mean square errors. Note that the non-IT energy consumption produced by our simulator is not necessarily a quadratic function of the IT energy; instead, it is based on real energy models reported by [3,4,7], e.g., the UPS energy loss is quadratic in its own load, which includes both IT energy and PDU energy loss. In the sensitivity study, we will further add noise to the energy model to reflect runtime system disturbances and evaluate the accuracy of QSEA.

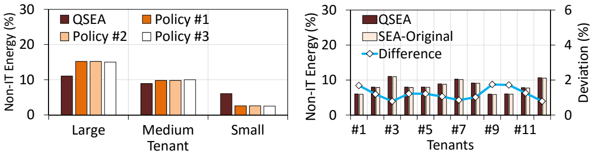

(a) (b)

Figure 3: Energy accounting with chiller-based cooling.

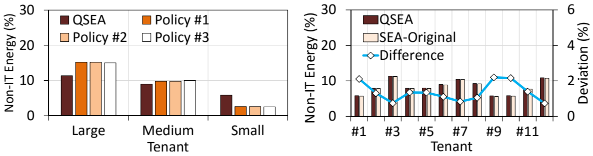

(a) (b)

Figure 4: Energy accounting with outside air cooling.

4.2 Results

We run the simulation for one month and compare QSEA with three existing energy accounting policies listed in Section 3.2. For the clarity of figures, we only show the average non-IT energy accounted for three tenants, labeled as "Large", "Medium" and "Small" corresponding to a power capacity subscription of 150kW, 100kW and 25kW (served by the first two PDUs), respectively.

Non-IT energy accounting. We first show the energy accounting result for data center with a chiller-based cooling system illustrated in Fig. 3(a). The y-axis shows the percentage of non-IT energy accounted for a tenant compared to all the non-IT energy consumption. We see a significant difference between QSEA and the existing energy accounting policies. In particular, we observe that QSEA favors large tenants. This is partly because non-IT energy has a static power when the non-IT unit is active, and QSEA equally distributes the static energy to tenants rather than in proportion to tenants' sizes (in power/energy consumption). Next, we compare QSEA against a baseline policy (called SEA-Original) that uses exact Shapley values and real energy models. Fig. 3(b) shows that QSEA and SEA-Original yield almost the same energy accounting result with a less than 2% difference, validating our choice of approximation based on quadratic forms.

Now, we consider outside air-based cooling, whose

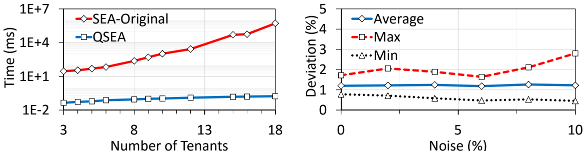

(a) (b)

Figure 5: **(a)** Computation time. **(b)** Robustness of QSEA.

energy increases cubically with the cooling load [3]. As shown in Fig. 4, the results are similar to those with chiller-based cooling. In particular, although QSEA uses quadratic form to approximate a cubic cooling energy model, it still produces an energy accounting outcome that is fairly close to that by SEA-Original.

Efficiency and Robustness. We first show the computation time required by QSEA and SEA-Original, with different number of tenants. We execute the two policies on a Dell desktop with Intel Core i7 and 16GB memory. Fig. 5(a) shows the computation time for 10 runs. We see that SEA-Original becomes computationally prohibitive for practical energy accounting when there are more than 10 tenants, while QSEA takes very little time. Next, we examine in Fig. 5(b) the robustness of QSEA against random system disturbances by showing the deviation of QSEA from SEA-Original (averaged over all tenants). In practice, given the same IT energy usage, non-IT energy can be disturbed by various factors such as humidity and thermal noise. To account for these factors, we add random noises to the non-IT energy produced by the simulator, and shows the energy accounting differences of QSEA compared to SEA-Original. We see that even when 10% noise is added to the non-IT energy, our proposed quadratic-based approximation still yields almost the same energy accounting result as SEA-Original.

5 Related Work

Energy accounting/profiling has received much attention in recent years (see [10] for a survey). For example, [11, 12] study model-based power metering for virtualized systems, with the goal of better utilizing the expensive power infrastructure. Further, [3] develops IT and non-IT energy models to minimize data center energy usage, [4] proposes measurement-based power models for different sub-components (e.g., UPS and cooling) in data centers, [13] profiles power usage by different applications on hyper-threaded processors. In contrast, we focus on non-IT energy accounting in multi-tenant data centers, which has not been well investigated.

Shapley value has been used in other contexts such as energy accounting on mobile systems [14] and peak demand cost splitting across users in cloud data centers [15]. Our work differs from [14, 15] in that we study an orthogonal problem and propose a novel low-complexity algorithm with little to zero implementation overhead. Our method exploits the unique characteristics of data center non-IT energy model, and also differs from the generic random sampling-based fast Shapley value calculation that may yield large errors [16]. Finally, note that *fair* multi-resource allocation in computer systems [17] has a fundamentally different goal than our work: it aims at fairly improving system utilization by encouraging users' resource sharing, whereas we re-attribute non-IT energy to different tenants and "fairness" in our context, as supported by Shapley value theorem [6,9], means satisfying the axioms in Section 3.1.

6 Conclusion

In this paper, we study non-IT energy accounting in multi-tenant data centers and propose QSEA, a quick Shapley value-based energy accounting policy. We show that QSEA is fair (in the sense of approximately satisfying all the four desired axioms) and offers a different perspective on non-IT energy accounting than existing policies. QSEA is easy to implement: it equally splits static energy of a non-IT unit among tenants with non-zero IT energy and attributes dynamic energy in proportion to tenants' IT energy. Further, compared to the exact Shapley value approach, QSEA yields almost the same accounting result while having a negligible complexity.

References

[1] Greenpeace, "Clicking clean: How companies are creating the green internet," 2014.

[2] NRDC, "Scaling up energy efficiency across the data center industry: Evaluating key drivers and barriers," *Issue Paper*, Aug. 2014.

[3] Z. Liu, Y. Chen, C. Bash, A. Wierman, D. Gmach, Z. Wang, M. Marwah, and C. Hyser, "Renewable and cooling aware workload management for sustainable data centers," in *SIGMETRICS*, 2012.

[4] S. Pelley, D. Meisner, T. F. Wenisch, and J. W. VanGilder, "Understanding and abstracting total data center power," *Workshop on Energy-Efficient Design*, 2009.

[5] Equinix Telecity, "Powergate, London," `http://www.telecitygroup.com/uk-collateral/powergate_brochure.pdf`.

[6] L. S. Shapley, "A value for *n*-person games," *Contributions to the Theory of Games (by H.W. Kuhn and A.W. Tucker) — Annals of Mathematical Studies*, vol. II, no. 28, pp. 307–317, 1953.

[7] R. L. Sawyer, "Making large UPS systems more efficient (Rev 3)," 2012.

[8] Q. Zhang and W. Shi, "UPS-aware workload placement in enterprise data centers," *IEEE Computer Magazine*, May 2015.

[9] R. Gopalakrishnan, J. R. Marden, and A. Wierman, "Potential games are necessary to ensure pure nash equilibria in cost sharing games," in *EC*, 2013.

[10] A. Noureddine, R. Rouvoy, and L. Seinturier, "A review of energy measurement approaches," *SIGOPS Oper. Syst. Rev.*, vol. 47, pp. 42–49, Nov. 2013.

[11] A. Kansal, F. Zhao, J. Liu, N. Kothari, and A. A. Bhattacharya, "Virtual machine power metering and provisioning," in *SoCC*, 2010.

[12] H. Lim, A. Kansal, and J. Liu, "Power budgeting for virtualized data centers," in *USENIX ATC*, 2011.

[13] Y. Zhai, X. Zhang, S. Eranian, L. Tang, and J. Mars, "Happy: Hyperthread-aware power profiling dynamically," in *USENIX ATC*, 2014.

[14] M. Dong, T. Lan, and L. Zhong, "Rethink energy accounting with cooperative game theory," in *MobiCom*, 2014.

[15] N. Nasiriani, C. Wang, G. Kesidis, B. Urgaonkar, L. Y. Chen, and R. Birke, "On fair attribution of costs under peak-based pricing to cloud tenants," in *MASCOTS*, 2015.

[16] J. Castro, D. Gómez, and J. Tejada, "Polynomial calculation of the shapley value based on sampling," *Computers & Operations Research*, vol. 36, no. 5, pp. 1726–1730, 2009.

[17] A. Ghodsi, M. Zaharia, B. Hindman, A. Konwinski, S. Shenker, and I. Stoica, "Dominant resource fairness: Fair allocation of multiple resource types," in *NSDI*, 2011.