

**Karlsruhe Reports in Informatics 2014,4**

Edited by Karlsruhe Institute of Technology,  
Faculty of Informatics  
ISSN 2190-4782

**FRESCO: A Framework for the Energy  
Estimation of Computers**

Extended Version

Pavel Efros, Erik Buchmann and Klemens Böhm

2014

KIT – University of the State of Baden-Wuerttemberg and National  
Research Center of the Helmholtz Association



Fakultät für **Informatik**

**Please note:**

This Report has been published on the Internet under the following  
Creative Commons License:

<http://creativecommons.org/licenses/by-nc-nd/3.0/de>.

# FRESCO: A Framework for the Energy Estimation of Computers

Extended Version

Pavel Efros, Erik Buchmann and Klemens Böhm  
Karlsruhe Institute of Technology (KIT)  
76131 Karlsruhe, Germany

**Abstract**—Many application areas, e.g., demand response, energy accounting or energy-aware scheduling, require estimates of the energy consumption of computer systems. However, existing estimation approaches often make restrictive assumptions regarding the effort at setup time or run time that is acceptable, they are tailored for specific hardware or software, or they cannot provide accuracy guarantees for the estimates.

In this paper, we introduce *FRESCO*, a Framework for the Energy eStimation of COmputers. *FRESCO* is a flexible framework for the estimation of the energy consumption of a wide range of computer systems. In particular, *FRESCO* considers technical information of the hardware manufacturer, system specifications like the average load, information that have been sampled at run time, e.g., the time the CPU spent in a specific state, and energy consumption profiles that might have been learned at setup time. Based on accuracy requirements and information available, *FRESCO* deploys and executes appropriate estimators. We have evaluated *FRESCO* with three real-world use cases. Our evaluation shows that *FRESCO* produces meaningful estimates for a wide range of analytical scenarios.

## I. INTRODUCTION

[1] has estimated the share of energy consumed by computers to be 7.15% of the total electricity consumption, and it is estimated to grow further (approx. 14.6% by 2020). Thus, the energy consumption of IT systems is an important cost driver for any large enterprise, and quantifying this type of consumption reliably is a cornerstone for many current business models. For example, the energy consumption has a significant impact on the total costs of ownership of a data center. It must be considered for total absorption accounting. Furthermore, in the context of the smart grid, that data gives way to new business models, e.g., by scheduling data centers according to an oversupply of renewable energies.

One way to quantify this type of energy consumption and integrate it into the *Smart Grid* is to deploy a smart meter for each computer system [2], [3]. Since it is expensive to equip existing hardware components with such meters, this is doable only in rare cases. Recent research has provided methods to *estimate* the energy consumption of computer hardware, e.g., based on nameplate information [4], sophisticated hardware models [5], or by profiling system and component power usage [6], [7]. However, all of these approaches have different characteristics in terms of setup effort, estimation effort, estimation accuracy and hardware requirements. It is difficult to decide when to use which approach and how to

determine meaningful estimation parameters, as well as the accuracy requirements of a given application.

Our goal is to devise a flexible framework for the estimation of the energy consumption of computer systems that considers many different accuracy and effort measures. This is challenging, because today's computer systems come with a wide range of different usage parameters and technical specifications, and we have to consider use cases that differ very much in the accuracy required and the effort acceptable for the operator.

In this paper we introduce *FRESCO*, A Framework for the Energy eStimation of COmputers. *FRESCO* consists of a highly configurable set of estimators, and a workflow to set up and run an instance of an estimator. In particular, *FRESCO* is able to (a) suggest a set of appropriate estimators for computer energy consumption according to the effort the operator is willing to invest and to the requirements of a certain application, and (b) to execute an instance of the selected estimator with settings that are appropriate for the application. Depending on the kind of estimator, *FRESCO* can estimate the energy consumption of a computer from various parameters. Such parameters include (1) hardware characteristics, e.g., the energy consumption of a hard disk as specified by its vendor, (2) usage information like CPU load and network activity, and (3) calibration data, e.g., an energy consumption profile that has been recorded by an energy meter for a specific hardware configuration.

*FRESCO* explicitly models the trade-off between the accuracy of the estimation and the effort of obtaining technical specifications, building energy profiles or measuring CPU usage information. For example, *FRESCO* can estimate the energy consumption of a PC with a high precision in hourly intervals, but also with a lower precision in intervals of a few seconds. Furthermore, *FRESCO* can provide upper and lower bounds for the estimation, and it considers heterogeneous hardware components and heterogeneous loads.

In this paper, we make the following contributions:

- We introduce three flexible power-estimation models that generalize state-of-the-art approaches to cover a wide range of accuracy requirements and computer systems.
- We describe *FRESCO*, which integrates these models into an estimation workflow that allows to choose, configure and run an estimator depending on the use case.
- We evaluate our approach with three use cases, namely

energy-aware data center management, demand response and energy accounting.

Our evaluation confirms that FRESKO can estimate the energy consumption of IT systems by considering a wide range of parameters and with an accuracy of up to 95%, and that it supports many different use cases.

**Paper structure:** Section II describes three application scenarios for energy estimation. Section III explains the requirements and classes of effort our framework must take into account. Section IV introduces FRESKO, which Section V evaluates. Finally, Section VI reviews related work, and Section VII concludes.

## II. APPLICATION SCENARIOS

In this section, we describe three different use cases that cover the spectrum of energy-aware applications for FRESKO.

### A. Energy-Aware Management of Data Centers

Increasing the performance per watt is a key performance optimization for data centers [8], [9]. For this purpose, it is important to obtain the energy consumption of a complex IT system as early as the design time of the data center or the allocation time of the various computing workloads. Recent approaches, e.g., in the area of energy-aware cloud data centers [10] or energy management for warehouse-sized computing centers [11], distinguish (1) the (static) energy consumption at idle state, and (2) the (dynamic) energy consumption depending on the workload of the target system. This is important to design the power distribution infrastructure, to decide about computing hardware acquisitions or to find out if a scheduled workload exceeds the cooling capacity.

Thus, two different accuracy requirements exist: It must be possible (a) to provide estimates for the typical case that are sufficiently accurate to make educated decisions for hardware acquisitions, and (b) to provide upper bounds for the energy consumption in extreme cases. Both requirements must be fulfilled at design time or at allocation time, i.e., before the operator can measure the workload or the energy consumption. Furthermore, an estimator must consider that some in-depth hardware specifications might be unavailable at design time.

### B. Demand-Response

Demand Response (DR) influences energy consumption patterns. For example, DR might be used to shift energy-intensive computing tasks to times of an energy surplus [12]. DR can be divided into (a) incentive-based DR and (b) time-based rates DR [12]. Incentive-based DR measures shift the energy consumption by providing, say, tariffs that reward to shift energy consumption into off-peak hours. In contrast, time-based rates DR makes use of static schedules.

Since a data center is a large, adjustable energy sink, it is particularly well suited to perform demand response measures [13]. To realize DR in a data center, an estimator must deliver continuous estimates of the energy consumption of the various IT components at run time. In particular, the estimates must be adequate to identify system states that produce energy-consumption peaks. Furthermore, the personal and computational effort of the estimation must not exceed

potential savings from DR. Finally, the estimator must cope with technical parameters on different levels of detail. For example, a coarse estimate could measure the average CPU load only, while a fine-grained approach might also consider the voltage and frequency of the CPU and the states of other hardware components.

### C. Computer Energy Accounting and Billing

Energy accounting and billing of the IT infrastructure becomes more and more important. For example, in the context of total absorption accounting, an enterprise might wish to assign each benefactor (a good or a service) the energy costs required for its production [14].

Typically, computer Energy Accounting requires estimates of the consumption with a frequency of 15 minutes to one hour. Furthermore, the estimator must provide stochastic accuracy guarantees (e.g., an accuracy of  $\pm 10\%$ ), that allows to assign the energy consumption of an IT system to a department or a product line. As for the previous scenarios, the estimator has to be applicable to a large variety of computer systems, with an effort that is adaptable.

## III. ESTIMATION REQUIREMENTS AND EFFORT

In this section, we compile requirements and effort classes for the energy estimation resulting from the application scenarios described in Section II.

### A. Requirements for the Estimation

We have compiled the following requirements for our framework from the application scenarios just described:

- R1: Generalizability** To be applicable to a wide spectrum of current and future IT systems, the estimation framework should enable the integration of a wide range of different approaches for energy estimation.
- R2: Adaptivity** Our framework must be able to provide (1) estimates and (2) accuracy guarantees on these estimates depending on the information available.
- R3: Accuracy and Effort Constraints** FRESKO should allow to trade off accuracy and effort. In particular, depending on the application scenario, FRESKO should allow the operator to rule out estimation approaches that are dominated by others in terms of accuracy and efforts.

### B. The Classes of Effort

We have identified two classes of effort, which our framework must take into account:

- E1: Setup** The setup effort is the one that is necessary to set the estimator up and running. This includes collecting technical specifications of the energy consumption of certain hardware components, e.g., the energy consumption of a CPU in activity states like idle or sleeping. Furthermore, it contains the effort of installing a monitoring application to measure run-time parameters of the hardware usage, e.g., disk activity. Finally, the setup effort includes the calibration of an energy consumption profile for a given hardware, e.g., measuring the energy consumption while executing a benchmark application.

**E2: Run-time** The run-time effort includes the network overhead and the computational overhead of the estimation process, and the overhead of a monitoring application collecting hardware parameters like CPU frequency or rotation speed of the hard disks, if required by the estimator. The more parameters the estimator samples, the higher is the data volume transferred, the computational overhead and the complexity of the estimation.

We expect that the two effort classes will be traded for each other. For example, the same accuracy requirement can be met either (a) by measuring an energy consumption profile at setup time or (b) by using detailed specifications and usage parameters at run time.

#### IV. FRESKO

In this section we describe the workflow and the estimators of *FRESKO*, our *FR*amework for the *E*nergy *e*Stimation of *C*omputers.

##### A. The *FRESKO* Workflow

With “*Target System*” we refer to the computer system whose energy consumption *FRESKO* must estimate. “*The Operator*” is responsible for installing and maintaining the estimator on the target system. *FRESKO* consists of three subsequent stages “*Setup*”, “*Configuration*” and “*Estimation*”, as shown in Figure 1, which we explain in the following.

*a) Setup:* At the first stage, the operator quantifies the trade-off between effort and estimation accuracy for the target system. In particular, the operator specifies the categories of information obtainable from the target system. This includes three kinds of information:

- The nameplate information available, e.g., if the energy consumption of the network card can be obtained.
- The parameters measurable on the target system, e.g., CPU frequency, CPU voltage or hard disk activity.
- If it is possible to measure a consumption profile for the target system, and with which accuracy.

Among the requirements the operator can have are the type of estimates and their type of error guarantees. These requirements influence the choice of estimator *FRESKO* suggests to use.

At the end of the setup stage, *FRESKO* either indicates the operator that, given his input, estimation is impossible or lets the operator choose one or a combination of estimators. In the latter case, *FRESKO* provides information on the estimation accuracy possible and the effort required for each estimator.

*b) Configuration:* At this stage, *FRESKO* helps the operator to configure the estimators selected, according to the following aspects:

- The usage parameters that must be measured to meet the accuracy specified in the setup stage.
- The frequency at which the parameters must be measured.
- The energy consumption profile, if necessary.

If the operator has chosen a calibration-based estimator, *FRESKO* provides a set of benchmarks and guides the operator through the process of measuring the energy consumption.

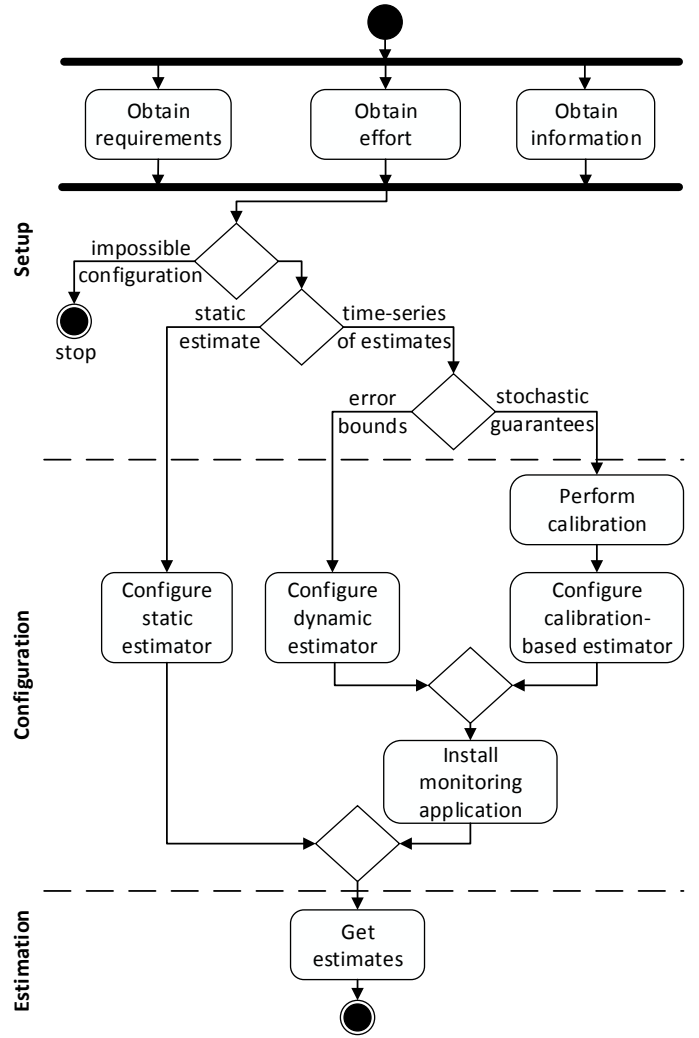


Fig. 1: *FRESKO* Workflow

The result of the configuration stage is the combination of configured estimators.

*c) Estimation:* Finally, *FRESKO* runs instances of the chosen estimators with the configuration parameters just fixed on the target system, estimates its energy consumption and sends the estimates to the operator.

##### B. The *FRESKO* Estimators

*FRESKO* considers three estimators which differ regarding the effort required, as sketched in Figure 2: The *Static estimator* makes use of static information on the computational load and hardware specifications. The *Dynamic estimator* predicts the energy consumption from hardware specifications and run-time parameters measured. The *Calibration-based estimator* uses an energy consumption profile that has been calibrated at the configuration stage. In the following, we describe each estimator, and we discuss its effort and accuracy.

*1) Static Estimator:* Our static estimator has been inspired by [4]. It is tailored for servers running a single application, e.g., similarly to the TPC benchmark suite, at peak load. Thus, this estimator aggregates the peak power consumption of all




Fig. 2: Effort Required by Our Estimators




Fig. 3: Classes of Information Used by the Static Estimator

hardware components. Since we are interested in estimates for a wide range of target systems operating at different loads, we have extended this approach in two ways:

**Generalization** FRESCO models a wide range of computer architectures and hardware components, e.g., laptop screens or motherboards of standard PCs.

**Detail Levels** FRESCO considers three levels of detail to model the computer architecture, namely the PSU level, the component level and the component state level, as shown in Figure 3.

The *PSU level* considers only the specification of the Power Supply Unit (PSU) powering the target system. In this case, the estimator calculates the total energy consumption  $E$  by using the maximum power  $PSU_{max}$  the PSU is able to supply and the run time  $\Delta T$  of the target system:

$$E = \Delta T \cdot PSU_{max} \quad (1)$$

If a higher accuracy is required, if more information is available, and if more effort at setup time is acceptable, FRESCO considers information at the *component level*. In this case, the consumption  $E$  is the sum of the power consumptions  $P_i$  of each component  $i \in C$ ,  $C = \{CPU, RAM, Hard\ Disk, \dots\}$ , multiplied with the run time  $\Delta T$ :

$$E = \Delta T \cdot \sum_{i \in C} P_i \quad (2)$$

Note that we represent the CPU of a multicore system as a set of components. One component constitutes the core-independent consumption of the CPU, e.g., due to caches and data transmission facilities shared by all cores. The other components represent the individual CPU cores.

The *component state level* is the highest level of detail of our static estimator. It includes information on (1) the power consumption of the different states of the components of the target system and (2) the time the components typically spent in certain states, given the typical computational load. In this case, FRESCO obtains the total consumption  $E$  by summing up the energy consumption  $P_{ij}$  of each component  $i \in C$  in a particular state  $j \in S_i$  (the set of states of component  $i$ ), multiplied with the time  $\Delta T_{ij}$  each component typically is in this state:

$$E = \sum_{i \in C} \sum_{j \in S_i} \Delta T_{ij} \cdot P_{ij} \quad (3)$$

In the case of a virtualized environment, the static estimator can incorporate virtual machines at the component state level. In this case, the operator can map components to virtual machines and use these together with their predicted states instead of the hardware components.

Component	Model	Power Consumption
CPU	Intel i5-3320M	Idle – 2.9 W Minimum active – 7.5 W TDP – 35 W Maximum active – 80.56 W
Memory	Micron Technology 2x4 GB DDR3L SDRAM 800 MHz	Minimum – 0.3 W Typical – 1.48 W Maximum – 1.68 W
Hard Disk	Hitachi HTS725050 500 GB at 7200rpm	Sleep – 0.1 W Standby – 0.2 W Low power idle – 0.7 W Active idle – 1.0 W Performance idle – 1.7 W Read/write – 1.8 W Seek – 2.0 W Startup – 5.5 W

TABLE I: Laptop Components

The accuracy of a static estimator depends on the detail level used and on confidence intervals on the input parameters. In the following, we will briefly discuss two extreme cases:

*Minimal information:* By using Equation 1, the operator obtains a very coarse upper bound of the energy consumption of the target system. This is because the PSU intake, as described on its nameplate, is usually overestimated for safety reasons [11]. However, if the manufacturer has provided tolerance bounds for the PSU intake in typical settings, the operator might be able to narrow down this upper bound.

*Full information:* Each hardware manufacturer provides detailed data sheets containing the minimal, typical and maximal energy consumption of any hardware component. If the operator specifies parameters on the component state level (cf. Equation 3), this information can be used to obtain hard upper and lower bounds for the energy consumption.

*Example 1:* Consider a laptop as described in Table I. The lower bound on the consumption is the sum of the minimum

values of each of the components, i.e.,  $2.9 \text{ W} + 0.3 \text{ W} + 0.1 \text{ W} = 3.3 \text{ W}$ . Likewise, the upper bound is the sum of the maximum values:  $80.56 \text{ W} + 5.5 \text{ W} + 1.68 \text{ W} = 87.74 \text{ W}$ .

*Summary: A static estimator might be sufficient for any application that does not need time series of estimates. It requires a small effort at setup time for obtaining the hardware specifications, and no effort at run time. The accuracy of this estimator depends on the detail level of its input values and the availability of tolerance bounds. In particular, the static estimator can provide bounds on the energy consumption.*

2) *Dynamic Estimator:* Our dynamic estimator models the energy consumption similarly to the static estimator, but installs a monitoring application on the target system to periodically measure detailed load information in real-time, e.g., CPU load or sleep times of the hard disk. Thus, our dynamic estimator generates time series of energy consumption data. Our dynamic estimator uses a monitoring application to record at run time in which state  $j \in S_i$  the component  $i \in C$  operates at time  $t$ . The energy consumption  $E_t$  at time  $t$  is the sum of the consumptions  $P_{it}$  of the components  $i$ :

$$E_t = \sum_{i \in C} P_{it} \quad (4)$$

The consumption  $E$  for a time interval  $[t_p; t_q]$  is the sum of the consumption at each point in time, multiplied with the period of time  $\Delta t$  between taking two consecutive samples:

$$E = \sum_{i=t_p}^{t_q} E_i \cdot \Delta t \quad (5)$$

Note that the energy consumption  $P$  of the components can be modeled in different ways. For example, consider the consumption  $P_t^{CPU}$  of the CPU. Suppose that the monitoring application measures the state information ‘‘CPU load’’  $l_t^{CPU}$ , and the operator knows the minimum and maximum power  $P_{min}^{CPU}$  and  $P_{max}^{CPU}$  the CPU can consume. In this case,  $P_t^{CPU}$  is:

$$P_t^{CPU} = P_{min}^{CPU} + l_t^{CPU} \cdot (P_{max}^{CPU} - P_{min}^{CPU}) \quad (6)$$

We have used this approach in our evaluation. An alternative approach [5] computes the consumption of the CPU from its operating frequency  $f_t$  and voltage  $V_t$  measured at time  $t$ , and the CPU capacitance  $c$ :

$$P_t^{CPU} = c \cdot f_t \cdot V_t^2 \quad (7)$$

It is also possible to integrate specific models for multi-core systems [15] and to model virtual machines as components of the target system. Orthogonally to this, FRESCO can model the energy consumption on different levels, similarly to the static estimator. Another option is to combine the static and the dynamic estimator, e.g., the energy consumption of the CPU can be estimated dynamically, while the consumption of all other components is a static estimation.

While the accuracy of the static estimator depends on the knowledge about the typical load of the target system, our dynamic estimator samples such parameters. Thus, the accuracy of our dynamic estimator depends on the sampling

frequency of the monitoring application. The reason is as follows: If a state changes between taking two consecutive samples, the estimator does not know to which extent the states were active. However, FRESCO provides upper and lower bounds on the energy consumption by assuming that a state change has taken place immediately before or after taking a sample.




Fig. 4: Example Sampling of Hard-Disk Power Consumption

*Example 2:* We illustrate this with the hard disk of a laptop (cf. Table I): The hard disk has three states, as shown on the y-axis in Figure 4. The x-axis shows the time intervals the samples were taken (dotted lines). Assume the disk has been observed in standby at time  $t_3$ , and in the idle state at  $t_4$ . Thus, the lower bound for the energy consumption in interval  $[t_3 : t_4]$  is  $0.2W \cdot \Delta t$ , and the upper bound is  $1.0W \cdot \Delta t$ .

Formally, consider a component with the sequence of states  $S = (s_1, s_2, \dots, s_n)$ , ordered by the power  $P_i$  the component consumes in state  $i$ ,  $i = 1, \dots, n$ . Let  $s_t = (s_{t_1}, s_{t_2}, \dots)$  be the time series of the states of the component sampled at times  $t_1, t_2, \dots$ . The upper bound  $E_{\Delta t}^u$  on the energy consumption during time interval  $\Delta t$  between consecutive samples  $t_j$  and  $t_{j+1}$  is:

$$E_{\Delta t}^u = \begin{cases} P_1 \cdot \Delta t & \text{if } s_{t_{j+1}} = s_1 \wedge s_{t_j} = s_1 \\ \dots \\ P_i \cdot \Delta t & \text{if } s_{t_{j+1}} = s_i \wedge s_{t_j} \leq s_i \vee \\ & s_{t_j} = s_i \wedge s_{t_{j+1}} \leq s_i \\ \dots \\ P_n \cdot \Delta t & \text{if } s_{t_{j+1}} = s_n \wedge s_{t_j} \leq s_n \vee \\ & s_{t_j} = s_n \wedge s_{t_{j+1}} \leq s_n \end{cases}$$

We calculate the lower bound  $E_{\Delta t}^l$  likewise.

*Summary: Our dynamic estimator is suitable for applications that require time series of the energy consumption of the target system at run time. Since this estimator also needs technical specifications, it requires a similar effort at setup time as a static estimator. The effort at run time depends on the number of parameters that the estimator samples, and on the sampling frequency. The accuracy of our dynamic estimator depends on the tolerances of the technical specifications and the sampling frequency. The estimator can compute bounds on the energy consumption.*

3) *Calibration-Based Estimator*: This estimator borrows from the Mantis approach [16], which estimates the power consumption of a system by correlating AC power measurements from a calibration phase with performance counters of the CPU. Our calibration-based estimator executes a detailed benchmark at setup time, which gradually stresses each system component in isolation. At the same time, a digital power meter records the actual energy consumption, and our monitoring application measures load information such as CPU frequency and voltage, hard disk usage, etc. FRESKO then performs a regression analysis to build a consumption profile that relates any load information to the total energy consumption of the target system.

More specifically, let  $M^{CPU}(l, f)$ ,  $M^{Disk}(l)$ ,  $M^{RAM}(l)$  be the regression models obtained through calibration for the CPU, Hard Disk and RAM, and let  $l$  be the load of the component and  $f$  the frequency of the CPU. Given the load information  $l_t^{CPU}$ ,  $l_t^{Disk}$ ,  $l_t^{RAM}$  and  $f_t$  at time  $t$ , our calibration-based estimator computes the energy consumption  $E$  at time interval  $[t_1, t_n]$  as follows:

$$E = \sum_{i=t_1}^{t_n} (M^{CPU}(l_i^{CPU}, f_i) + M^{RAM}(l_i^{RAM}) + M^{Disk}(l_i^{Disk})) \quad (8)$$

The calibration-based estimator can derive stochastic accuracy guarantees from the regression model used, by considering the maximal and minimal energy consumption that has been recorded for each distinct benchmark load. For example, think of a regression model which uses only CPU load  $l$  and CPU frequency  $f$  to estimate the total energy consumption. Let  $(f, l)$  be the pair of values for the current load and frequency of the CPU and  $E_{l,f} = \{e_1, e_2, \dots, e_m\}$  be the set of values for the real energy consumption that the calibration phase had measured for the pair  $(f, l)$ . The upper bound  $E_{\Delta t}^u$  for the energy consumption during the time  $\Delta t$  when the CPU operates at frequency  $f$  and at load  $l$  is

$$E_{\Delta t}^u = \max E_{l,f} \cdot \Delta t \quad (9)$$

while the lower bound  $E_{\Delta t}^l$  is:

$$E_{\Delta t}^l = \min E_{l,f} \cdot \Delta t \quad (10)$$

We calculate upper and lower bounds for other pairs of values of frequency and load the same way.

*Example 3*: Assume that the benchmark has resulted in a CPU load of 50% and in a frequency of 2.4 GHz for some time interval, and the energy consumption measured has been  $\{97.5 W, 99 W, 100 W, 102 W, 102.5 W, 103 W\}$ . Thus, for this load and frequency FRESKO would stochastically guarantee a maximal (minimal) energy consumption of  $103W \cdot \Delta t$  ( $97.5W \cdot \Delta t$ ). Since the dynamic estimator and the calibration-based estimator use the same monitoring application, it is possible to obtain upper and lower bounds for the energy consumption from our dynamic estimator in tandem.

*Summary*: *The calibration-based estimator is well-suited for applications that require a calibrated zero point and stochastic guarantees on the estimation quality, such as billing or accounting. Due to the extensive calibration, this estimator*

*comes with a very high effort at setup time. At run time, the effort of this estimator depends on the number of parameters that must be sampled and on the sampling frequency, similarly to the dynamic estimator.*

Summing up everything, FRESKO can model a wide range of computer architectures and environments, including multi-core and virtualized systems. It can generate static ex-ante estimates solely on hardware specifications as well as time series of estimates at run time with a configurable estimation frequency and different accuracy guarantees. FRESKO considers different kinds of effort at setup time and run time. In the next section, we show that FRESKO covers a wide spectrum of use cases.

## V. EVALUATION

Our framework operates as intended if it provides estimates that are appropriate for a wide range of applications. That is, FRESKO must let the operator decide on a tradeoff between accuracy and effort, according to the requirements of the application. Thus, we evaluate FRESKO by means of three use cases, and we measure the accuracy that can be obtained with a certain effort, in terms of estimation frequency and technical specifications provided.

### A. Measures

To evaluate how well FRESKO can estimate the real data measured by our digital multimeter, we have computed two metrics.

The *Mean Absolute Percentage Error (MAPE)* measures the average of the percentual deviation between the actual values and the estimates:

$$MAPE = \frac{1}{n} \sum_{i=1}^n \frac{|y_i - y'_i|}{y_i}$$

where  $y_i$  are the actual values and  $y'_i$  are the estimates.  $n$  is the number of records in the dataset. MAPE of zero means that the estimated values perfectly match the ones measured.

Correspondingly, the *Maximum Absolute Percentage Error (MaxAPE)* measures the maximum percentual deviation between the actual values and the estimates:

$$MaxAPE = \max_{i=1}^n \left( \frac{|y_i - y'_i|}{y_i} \right)$$

### B. Evaluation Setup

We have tested three datasets that have been obtained from a server, a desktop PC, and a laptop. The architectures of these target systems range from a multicore machine with redundant components to a mobile architecture that has been optimized to save energy. Furthermore, system usage differs very much, as described next.

*a) Server Dataset*: This dataset is about a mail server constantly executing SpamAssassin [17]. Its workload is a daily pattern with a low usage during the night and a high usage in the morning and afternoon hours. Load peaks occur when the server checks bulks of e-mails sent to large mailing lists. Table II shows the hardware components of this system. To obtain our *server dataset*, we have used a digital multimeter



Wattsup PRO [18] (accuracy: 1.5%) to measure the energy consumption at every minute as a reference. Furthermore, our monitoring application has logged CPU usage, CPU frequency and hard disk drive usage with a sampling frequency of one second. Our measurements cover a period of three weeks.

Component	Model	Power Consumption
CPU	2 x AMD Opteron 275	Maximum – 95.2 W P-State #1 – 90.3 W P-State #2 – 75.9 W Minimum P-State – 36.1 W Halt Mode – 16.6 W
Memory	Micron Technology 4x1 GB DDR400 PC3200	Minimum – 9.9 W Typical – 36.4 W Maximum – 87.48 W
Hard Disk	2 x Seagate ST937401 2x74 GB at 10000 rpm	Maximum – 10.2 W Idle – 5.07 W Minimum – 4.69 W

TABLE II: Server Dataset

*b) Desktop Dataset:* The *desktop dataset* contains measurements of three weeks of energy consumption, CPU usage and CPU frequency measured on an office computer with a sampling frequency of one second. This target system is equipped as shown in Table III. Its workload is the result of typical secretarial tasks, e.g., MS Office, Internet Explorer and a number of custom-made administrative applications. Thus, the workload rarely reaches the maximal computing capacity, and the computer is active only during office hours.

Component	Model	Power Consumption
CPU	Intel Pentium Dual Core E5300	Deeper Sleep – 4 W Extended Halt – 8 W TDP – 65 W Maximum – 92.9 W
Memory	Crucial Memory - 2x2 GB DDR2 SDRAM 800 MHz	Minimum – 3.65 W Typical – 5.1 W Maximum – 10.4 W
Hard Disk	Western Digital WD2500AAJS 250 GB 7200 rpm	Standby – 0.73 W Sleep – 0.73 W Idle – 4.92 W Read/Write – 5.36 W

TABLE III: Desktop Dataset

*c) Laptop Dataset:* The *laptop dataset* consists of two weeks of energy consumption, CPU usage, -frequency and -voltage, measured with a time resolution of one second on a laptop as shown in Table I. The laptop has been used for research purposes, i.e., the system load does not follow any regular pattern and shows idle periods as well as maximum load conditions.

### C. Use Cases

We now evaluate FRESKO with the use cases described in Section II, namely energy-aware data center management, demand response and energy accounting.

*1) Energy-Aware Data Center Management:* This scenario requires ex-ante estimates of the energy consumption depending on a predefined workload. The estimates must be sufficiently accurate for informed management decisions, e.g.,

it must be possible to find out if one target system requires significantly more energy for a certain workload than another one. Furthermore, it must be possible to find out if a certain workload might exceed the cooling capacity in the worst case. Thus, FRESKO proposes the static estimator model. To evaluate this scenario, we let FRESKO estimate upper and lower bounds on the energy consumption, and the average energy consumption for a typical workload.

*a) Minimal and Maximal Consumption:* We let FRESKO exemplarily estimate upper and lower bounds on the consumption of a server (cf. Table II). With our use case, the operator specifies manufacturer information on the component level for CPU, RAM and disk. The CPU consumes the least possible energy in HALT mode, which corresponds to a consumption of 16.6 W. The maximum power consumption of the CPU is the maximal current intake multiplied with the highest voltage allowed by the manufacturer, which is 95.2 W. The minimum power the hard-disk consumes is its power consumption in sleep mode: 4.69 W. The maximum power it consumes is equal to the power consumption in read/write mode at the highest rate of I/Os per second: 10.2 W. Concerning the RAM, it consumes at least 9.9 W and at most 87.48 W. Thus, for our example system, the lower bound for the total power consumption is  $(2 \cdot 16.6 + 2 \cdot 4.69 + 9.9) \text{ W} = 52.48 \text{ W}$ . The upper bound for the energy consumption is  $(2 \cdot 95.2 + 2 \cdot 10.2 + 87.48) \text{ W} = 298.28 \text{ W}$ . Our measurements confirm that these bounds apply for the server dataset. The bounds can be narrowed if the operator is able to specify a limit for the time each component is in a specific state (cf. Equation 3).




Fig. 5: MAPE Depending on the Length of Interval - Static Estimator




Fig. 6: MaxAPE Depending on the Length of Interval - Static Estimator

b) *Average Consumption*: Now we assume that the operator wants FRESKO to estimate the energy consumption for an average CPU load of 50%, in order to assess the typical cooling requirement. To evaluate the accuracy of the estimates, we aggregate the energy consumption, which we have measured with a frequency of up to one second, to time intervals from one second to 16 minutes. Furthermore, we let FRESKO use the static estimator to provide estimates for the same time intervals. Figure 5 shows the MAPE on the y-axis and the length of the time interval on the x-axis, for each of our three datasets. A value of 30% at the interval length of 1 minute for the laptop dataset means that, on average, the energy consumption estimates summed up for intervals of 1 minute deviate by 30% from the corresponding consumption measured. Figure 6 shows the MaxAPE for all datasets. The figures indicate that for longer time intervals, FRESKO can provide more accurate estimates. In particular, for the server dataset, the maximum error goes from around 23% for an aggregation level of one second to around 11% for a higher aggregation level of 16 minutes, corresponding to a two-fold decrease in value. Smaller decreases (69–62% and 31–30%) occur for the other two datasets. This is because longer interval lengths mitigate the effect of short-term deviations in the workload. Evidently, the accuracy of the estimation can be improved if the operator provides more accurate information on the workload of the target system.

*Summary: FRESKO has provided upper bounds for the energy consumption. Furthermore, it is able to provide reasonable estimates for the average workload by requiring only little data from the operator. Thus, we conclude that FRESKO is able to deal with the requirements of this use case.*

2) *Demand Response*: Our second case study is a Demand Shifting scenario as described in Section II-B. Demand Shifting requires time series of estimates to identify periods of time with high energy consumptions (peaks), together with upper and lower bounds. As the operator is willing to invest only a small effort, FRESKO suggests our dynamic estimator model.

To evaluate this scenario, we let FRESKO estimate the consumption based on the CPU load and on information on the maximal and minimal energy consumptions of our three target systems, cf. Equation 6. We use these estimates to identify points in time when the energy consumption is above a given threshold. In particular, we evaluate two dynamic thresholds that consider the difference between the largest and smallest possible values of a time series  $T$ :

$$\theta_1 = 0.8 \cdot \left( \max_{i=1}^{|T|} (T_i) - \min_{i=1}^{|T|} (T_i) \right) \quad (11)$$

$$\theta_2 = 0.95 \cdot \left( \max_{i=1}^{|T|} (T_i) - \min_{i=1}^{|T|} (T_i) \right) \quad (12)$$

Given a time series of energy consumption  $T$ , we use a filter  $\phi$  to compute a time series of peak consumption  $T^{peak}$  that does not include values smaller than  $\theta$ .

$$\phi(v, \theta) = \begin{cases} v & \text{if } v \geq \theta \\ \perp & \text{otherwise} \end{cases} \quad (13)$$

We compute such time series of peak consumption from our measured values as well as for the time series FRESKO has estimated. We use an estimation frequency of one second. If

our estimates are accurate, FRESKO can identify periods with high energy consumption and can thus enable operators to perform Demand Shifting.

Figure 7 illustrates the cumulative distribution function (CDF) of the real energy consumption during specific intervals for the desktop dataset. The first set of intervals is when FRESKO estimated the consumption to be greater than  $\theta_1$  (continuous line). The second set is when FRESKO estimated the consumption to be greater than  $\theta_2$  (dashed line). We observe that, if the estimator predicts a value greater than  $\theta_1$ , then the real energy consumption is greater or close to  $\theta_1$ . Thus, in around 88% of all cases, a value predicted to be greater than  $\theta_1$ , is also greater than  $\theta_1$ . Concerning estimates predicted to be greater than  $\theta_2$ , these are also close to or greater than  $\theta_2$ . In 80% of the predicted cases, the real energy consumption was greater than 90% of  $\theta_2$ .




Fig. 7: Distribution of Real Energy Consumption Values – Desktop Dataset

Figure 8 illustrates the cumulative distribution function (CDF) of the predicted energy consumption during specific intervals for the desktop dataset. The first set of intervals is when the real consumption was greater than  $\theta_1$  (continuous line). The second set is when the real consumption was greater than  $\theta_2$  (dashed line). Thus, the estimator predicted a value of at least 75% of  $\theta_1$  for values which were greater than  $\theta_1$  in 90% of the cases. The estimator predicted a value of at least 80% of  $\theta_2$  for all values which were greater than  $\theta_2$  in all cases.




Fig. 8: Distribution of Predicted Energy Consumption Values – Desktop Dataset

For the laptop dataset (Figure 9), if our dynamic estimator predicts that the energy consumption during an interval is greater than  $\theta_1$ , then the actual consumption is greater than 75% of  $\theta_1$  in 92% of the cases. Furthermore, a value greater

than  $\theta_2$  of the estimated energy consumption is greater than 75% of  $\theta_2$  of the real consumption in around 92% of the cases. On the other hand (Figure 10), our estimator predicted a value of at least 75% of  $\theta_1$  for intervals with a consumption greater than  $\theta_1$  in 95% of the cases. Our estimator predicted values greater than 80% of  $\theta_2$  for all intervals with a consumption greater than  $\theta_2$ .




Fig. 9: Distribution of Real Energy Consumption Values – Laptop Dataset




Fig. 10: Distribution of Predicted Energy Consumption Values – Laptop Dataset

Concerning the server dataset, for which we use minute-by-minute measurements and estimations, the accuracy of our dynamic estimator is better. In 90% of the cases where an estimate is greater than  $\theta_1$ , the real value also is greater than  $\theta_1$ . Additionally, the estimator correctly predicted all values greater than  $\theta_2$ . On the other hand, our estimator predicted a value of at least 75% of  $\theta_1$  for intervals with a consumption greater than  $\theta_1$  in 90% of the cases. Our estimator predicted values greater than 80% of  $\theta_2$  for all intervals with an actual consumption greater than  $\theta_2$ .

*Summary: Our dynamic estimator can identify periods of time with peak energy consumptions with a reasonable accuracy with a low estimation effort. That is, it uses only static information on the minimal and maximal consumption of the target system, and it samples only the CPU load. Moreover, the estimator is flexible, i.e., it can sample more usage parameters in order to improve its accuracy. We conclude that FRESKO fulfills the requirements of this use case.*

3) *Energy Accounting:* Our third use case is the energy accounting scenario as described in Section II-C. This use case requires estimates with stochastic guarantees. Thus, FRESKO suggests a calibration-based estimator, which provides estimates based on calibrated energy profiles.

Again, we evaluate this use case with our three target systems, as described in Section V-B. In particular, we let FRESKO calibrate energy profiles for each of our target systems at the setup time of the estimator. At run time, our monitoring application samples the CPU load with different sampling frequencies. In order to compare the estimates with the real values, we have calculated MAPE and MaxAPE for all datasets. Figure 11 shows the MAPE on the y-axis and the length of the time interval on the x-axis. The figure indicates that the estimation accuracy is better for longer time intervals. For all three datasets, the mean error decreases by around a fourth (14–37% decrease) when the estimator aggregates estimates for intervals of 16 minutes instead of one second. Similarly, the MaxAPE decreases significantly for all datasets with longer estimation intervals, as shown in Figure 12. In particular, for the server dataset, the maximum error is around 6.5 times smaller, decreasing from about 35% to 5.3%. For the other two datasets, the decrease in maximum error is significant as well (2.4 times for the laptop dataset and 22 times for the desktop dataset, respectively).

*Summary: With estimation intervals that make sense in energy accounting, FRESKO is able to provide estimates of a high accuracy. While the effort at run time is similar to the one of the dynamic estimator, the effort at setup time is very high. However, many energy accounting scenarios make use of numerous target systems with identical hardware, e.g., for typical office tasks. In such scenarios, the calibration effort at setup time takes place only once. Thus, we conclude that FRESKO can perform well in an energy accounting scenario.*




Fig. 11: MAPE Depending on the Length of Interval - Calibration-Based Estimator




Fig. 12: MaxAPE Depending on the Length of Interval - Calibration-Based Estimator

## VI. RELATED WORK

The energy consumption of a computing system can be monitored directly. This is done by measuring the energy consumption using common digital meters [2], custom-designed devices [3] or integrated hardware power sensors [19]. In the case of large and heterogeneous computing centers, installing digital meters or power sensors at every subsystem (server, PC, etc.) is costly.

A related domain of significant interest in recent research is computer power characterization at the system and subsystem level. Part of recently developed power characterization models are based on collecting microarchitectural events using hardware registers. These models consider both subsystem [20], [21], [22] and system [23] levels, as well as virtualized environments [24].

A drawback of power characterization models which use hardware registers is that these models are tailored to specific hardware. This makes such models less portable and general. Thus, in the case of large heterogeneous deployments of computing systems, this lack of genericity would require a significant effort for power consumption modeling and estimation of the entire deployment. Another issue is that the number of hardware performance events tends to be large [25]. Moreover, only a small part of them can be measured at the same time [25], due to the limited number of hardware registers. [26] proposes a solution, namely time-multiplexing different sets of events on the hardware registers. While this approach allows for a greater number of performance events to be monitored, it increases the overhead and reduces the accuracy.

Using high-level statistical information provided by the operating system avoids the need for specific detailed (low-level) hardware knowledge when designing power estimation models. Recent work has proposed power consumption models based solely on high-level performance or usage metrics provided by the operating system to maximize energy efficiency using various optimizations. Thus, [11] uses CPU utilization in order to estimate the power consumption of large numbers of servers, reaching a mean error of 1% when considering groups of several hundreds of servers. The power consumption model proposed in [27] uses the expected load on a server cluster in order to estimate its power consumption. *JouleMeter* [28] is a solution for virtual machine power metering which infers the power consumption from resource usage at runtime. [29] proposes a model of the power consumption of idle servers.

The advantages of high-level black-box models are the low overhead, simplicity and relatively good accuracy. However, these models estimate full-system power consumption and do not allow for a more fine-grained repartition of the power consumption, such as per process or per application power consumption. Moreover, the majority of the models has been developed and tested on computer systems with a big share of static energy, such as servers [11] or clusters of virtual machines [28]. In order to model such computer systems, these black-box models are easily integrable into FRESKO.

## VII. CONCLUSIONS

As the share of computer energy consumption increases, it becomes increasingly important to quantify it in a solid

manner. Many important enterprise applications, accounting procedures and business models require such data to allow informed management decisions, e.g., in the context of energy-aware management of IT resources, IT-energy accounting or demand response. However, most existing estimators are tailored to specific use cases, hardware architectures and usage profiles. Moreover, due to their different characteristics in terms of effort and accuracy, the choice of estimation method to use for a given application is far from obvious. In this article, we have proposed FRESKO – a general and flexible framework for the estimation of the energy consumption of computers. Depending on the effort the operator is willing to invest and on the requirements of the application, FRESKO can propose and run appropriate estimators with good parameters settings. It is able to give quality guarantees on the output estimates. FRESKO considers heterogeneous hardware components and loads, as well as the frequency of the estimation. Comprehensive experimental results based on three representative real-world datasets show that our framework is useful in many analytical scenarios.

## ACKNOWLEDGMENT

The authors would like to thank Olaf Hopp and Manfred Alef for their help with the energy and software measurements.

## REFERENCES

- [1] Vereecken W. et al., "Overall ICT Footprint and Green Communication Technologies," in *International Symposium on Communications, Control and Signal Processing (ISCCSP)*, 2010.
- [2] Ge R. et al., "PowerPack: Energy Profiling and Analysis of High-Performance Systems and Applications," *IEEE Transactions on Parallel and Distributed Systems*, vol. 21, 2010.
- [3] D. C. Snowdon, S. M. Petters, and G. Heiser, "Power Measurement as the Basis for Power Management," in *Workshop on Operating Systems Platforms for Embedded Real-Time applications*, 2005.
- [4] M. Poess and R. O. Nambiar, "Power Based Performance and Capacity Estimation Models for Enterprise Information Systems," *IEEE Data Engineering Bulletin*, vol. 34, 2011.
- [5] Nouredine A. et al., "Runtime Monitoring of Software Energy Hotspots," in *IEEE/ACM International Conference on Automated Software Engineering*, 2012.
- [6] S. Rivoire, P. Ranganathan, and C. Kozyrakis, "A Comparison of High-Level Full-System Power Models," in *HotPower*, 2008.
- [7] A. Kansal, F. Zhao, J. Liu, N. Kothari, and A. A. Bhattacharya, "Virtual Machine Power Metering and Provisioning," in *ACM Symposium on Cloud Computing*, 2010.
- [8] Rivoire S. et al., "JouleSort: A Balanced Energy-Efficiency Benchmark," in *ACM SIGMOD international conference on Management of data*, 2007.
- [9] J. Laudon, "Performance/Watt: The New Server Focus," *ACM SIGARCH Computer Architecture News*, vol. 33, 2005.
- [10] M. Milenkovic, E. Castro-Leon, and J. R. Blakley, "Power-Aware Management in Cloud Data Centers," in *International Conference on Cloud Computing (CloudCom)*, 2009.
- [11] X. Fan, W.-D. Weber, and L. A. Barroso, "Power Provisioning for a Warehouse-Sized Computer," in *Annual International Symposium on Computer Architecture*, 2007.
- [12] P. Palensky and D. Dietrich, "Demand Side Management: Demand Response, Intelligent Energy Systems, and Smart Loads," *IEEE Transactions on Industrial Informatics*, vol. 7, 2011.
- [13] A. Berl et al., "Modelling Power Adaption Flexibility of Data Centres for Demand-Response Management," in *Energy Efficiency in Large Scale Distributed Systems*, 2013.
- [14] Jimenez V. et al., "Energy-Aware Accounting and Billing in Large-Scale Computing Facilities," *IEEE Micro*, vol. 31, 2011.

- [15] R. Basmadjian and H. de Meer, "Evaluating and Modeling Power Consumption of Multi-core Processors," in *International Conference on Future Energy Systems (e-Energy)*, 2012.
- [16] D. Economou, S. Rivoire, and C. Kozyrakis, "Full-System Power Analysis and Modeling for Server Environments," in *Workshop on Modeling Benchmarking and Simulation (MOBS)*, 2006.
- [17] [Online]. Available: <http://spamassassin.apache.org/>
- [18] [Online]. Available: <https://www.wattsupmeters.com/secure/products.php?pn=0&wai=638&spec=8>
- [19] Intel, "Intelligent platform management interface." [Online]. Available: <http://www.intel.com/design/servers/ipmi/index.htm>
- [20] J. Janzen, "Calculating Memory System Power for DDR SDRAM," *Designline*, vol. 10, 2001.
- [21] A. Merkel and F. Bellosa, "Balancing Power Consumption in Multiprocessor Systems," in *ACM SIGOPS/EuroSys European Conference on Computer Systems*, 2006.
- [22] Bertran R. et al., "A Systematic Methodology to Generate Decomposable and Responsive Power Models for CMPs," *IEEE Transactions on Computers*, vol. PP, 2012.
- [23] W. L. Bircher and L. K. John, "Complete System Power Estimation Using Processor Performance Events," *IEEE Transactions on Computers*, vol. 61, 2012.
- [24] G. Dhiman, K. Mihic, and T. Rosing, "A System for Online Power Prediction in Virtualized Environments Using Gaussian Mixture Models," in *Design Automation Conference*, 2010.
- [25] S. M. Rivoire, "Models and Metrics for Energy-efficient Computer Systems," Ph.D. dissertation, Stanford University, 2008.
- [26] R. Azimi, M. Stumm, and R. W. Wisniewski, "Online Performance Analysis by Statistical Sampling of Microprocessor Performance Counters," in *Annual International Conference on Supercomputing*, 2005.
- [27] Heath T. et al., "Energy Conservation in Heterogeneous Server Clusters," in *ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, 2005.
- [28] Kansal A. et al., "Virtual Machine Power Metering and Provisioning," in *ACM Symposium on Cloud Computing*, 2010.
- [29] R. Basmadjian, F. Niedermeier, and H. De Meer, "Modelling and Analysing the Power Consumption of Idle Servers," in *Sustainable Internet and ICT for Sustainability (SustainIT)*, 2012.