

Energy Measurement, Modeling, and Prediction for Processors with Frequency Scaling

Thomas Rauber¹ Gudula Rünger² Michael Schwind² Simon Melzner¹
Haibin Xu²

1: Computer Science Department, University Bayreuth, Germany

2: Computer Science Department, Chemnitz University of Technology, Germany

to appear in: The Journal of Supercomputing. Springer ISSN 0920-8542, December 2014,
Volume 70, Issue 3, pp 14511476

Abstract

The energy consumption is an important aspect of today's processors and a large variety of research approaches deal with reducing the energy consumption for specific application codes on different platforms under certain constraints. These research approaches are based on energy information acquired by very different means, such as hardware settings with power-meters, software methods with hardware counters available for more recent CPUs, or simulations based on theoretical models. In this article, all of these energy acquisition methods are investigated and compared. As application programs, we consider the SPEC CPU2006 integer and floating-point benchmark collections, which represent a large variety of applications from different areas. The investigations are done for single multicore CPUs with the goal to get more insight into their energy consumption behavior. An experimental evaluation is performed on three recent processor types with dynamic voltage frequency scaling. The article compares the measured energy and the energy provided by hardware counters with the energy predicted by simulation models. The comparison shows that the simulation models are able to capture the energy consumption quite accurately.

1 Introduction

Energy-aware computing and the efficient use of compute resources are now accepted to be as important for application codes as performance aspects. Energy efficiency has already been a critical concern in digital circuit design for the last two decades. This has led to several power-aware system features, including multicore-on-a-chip processors, core-independent functional units, dynamic frequency and voltage scaling, or clock gating, which are mainly aimed at a reduction of the energy consumption of the processors. To fully exploit the hardware capabilities for decreasing the energy consumption of application software running on that hardware, information about the energy consumption has to be available. Reliable energy models and measures are needed to plan and evaluate the software energy behavior on specific hardware.

One of the most effective hardware energy minimization techniques is dynamic voltage and frequency scaling involving a dynamic adjustment of the clock frequency and the corresponding supply voltage. Dynamic voltage-frequency scaling (DVFS) reduces the dynamic power consumption which is that part of the power needed to make a processor work. The static power is the power of the processor when it is not active and mainly includes the leakage power. There is also a considerable amount of power consumed by the entire hardware system and all its devices. The energy consumption of application code can be reduced by either decreasing the execution time, e.g., by algorithmic methods, which may result in a smaller energy consumption, or by scaling the frequency to decrease the power consumption, which may however increase the execution time. Since a low execution time is still an important goal, for instance when simulating large applications or when a predefined response time is required, energy-awareness and performance awareness have to be treated together. Thus, in addition to performance data, the availability of reliable energy data is a prerequisite to tune application code towards energy efficiency.

Measuring, evaluating or modeling energy consumption for either single specific applications on selected hardware platforms or an entire range of applications emphasizing more on the hardware facility is an active research area. The goal is to acquire information about the energy consumption with the aim to understand, to reduce or even to minimize the energy consumption under specific constraints. This requires a very solid foundation of energy data which are reliable in the sense that they reflect the reality qualitatively and quantitatively. Thus, the challenge is how to get good quality energy data. The methods used for gathering energy information of application codes running on modern hardware platforms are quite different and three main categories have evolved, which are hardware measurements with specific power-meters, software measurements with hardware counters available in more recent processors, and simulations based on theoretical energy models.

In this article, we consider all three methods for gathering energy data with the goal to investigate the correspondence or difference of the energy data provided by the methods. A contribution of the article is a detailed comparison of the energy measurement using power-meters and using the RAPL (Running Average Power Limit) interface of recent Intel processors. As application we have chosen the SPEC CPU2006 benchmarks, which are well-known to the community and reflect properties of typical application codes. The measurements with both of measurement techniques show a good correspondence, which allows us to use the energy information gathered by the faster and easier-to-use hardware counters. In contrast to earlier work, such as presented in [12, 29], the comparison focuses on the frequency scaling feature of the processors and is performed for a large set of complete application programs (the SPEC CPU2006 integer and floating-point programs). The correlation of the data acquired by these two measurement methods is crucial for getting experimental data on which a power prediction method can be based. A second contribution of the article is the comparison of the measured energy values with the energy values provided by different DVFS energy models. In particular, a physical and a heuristic energy model are investigated. These energy models can be used for an a priori energy estimation and as a basis for simulating the energy consumption of application programs.

The rest of the article is structured as follows. Section 2 describes measurement techniques for the energy consumption of processors and compares these techniques and their resulting energy values for the SPEC CPU2006 benchmarks. Section 3 presents energy measurements for the SPEC benchmarks on different modern processors. Section 4 describes different models for capturing the energy consumption of DVFS processors. Section 5 investigates how well these energy models are suited for predicting the energy consumption of processors. Section 6 discusses related work. Section 7 concludes the paper.

2 Measurement techniques for the energy consumption of DVFS processors

In this section, we describe the techniques with which we have measured the power and energy consumption of DVFS processors and present the resulting energy data. Hardware measurement techniques using specialized power-meters are described in Section 2.1. A software-based measurement technique based on accessing the hardware counters of the processors is described in Section 2.2. A comparison of the measured energy data is done in Section 2.3.

Modern microprocessors, such as the Intel Core i7 processors, incorporate a power management technology which supports different power management states: performance states (P-states), throttle states (T-states), idle states (C-states) and sleep states (S-states) [17]. P-states are predefined sets of paired frequency and voltage combinations at which an active core can operate [3]; the various P-states supported are implemented by using a combination of dynamic frequency scaling (DFS) and dynamic voltage scaling (DVS). A C-state is an idle state in which parts of the processor are powered down to save energy. Various C-states are supported by Intel processors, and a higher numbered C-state indicates more power savings.

For the experiments, three different Intel Core i7 processors have been used: (i) an Intel Core i7-2600 processor with the Sandy Bridge architecture; (ii) an Intel Xeon CPU E3-1225 V2 processor with the Ivy Bridge architecture; (iii) an Intel Core i7-4770 with the Haswell architecture. Table 1 describes some details of the processor architectures.

	Core i7-2600	Xeon E3-1225V2	Core i7 4770
architecture	Sandy Bridge	Ivy Bridge	Haswell
minimum frequency	1.6 GHz	1.6 GHz	0.8 GHz
maximum frequency	3.2 GHz	3.2 GHz	3.4 GHz
TDP	95 W	77 W	84 W
step size frequency	100 MHz	100 MHz	200 MHz
physical cores	4	4	4
hyperthreading	yes	no	yes
virtual cores	8	4	8
L1 data cache	32 KByte	32 KByte	32 KByte
L2 cache	256 KByte	256 KByte	256 KByte
L3 shared cache	8 MByte	8 Mbytes	8 MByte
RAM size	8 GByte	8 GByte	8 GByte

Table 1: Characteristics of the processors used for the experimental evaluation.




Figure 1: Illustration of the power measurement using the NI9205 device for measuring the supply voltage V_{cc} for the different wires.

2.1 Power measurement with power-meter

A precise method to capture the power consumption of computer systems is the direct measurement using a power-meter. We have implemented this approach using the data acquisition system National Instruments NI9205, which has been integrated into an NI CDAQ9181 chassis to enable the transfer of the measured power values via Ethernet to a separate system for data analysis, see Fig. 1 for an illustration of the overall configuration and Fig. 2 for a photo of the experimental setup. The power measurement technique is similar to the technique used by the PowerPack framework [13].

The NI9205 enables a fine-grain power measurement of different components of a computer system. In particular, it is possible to isolate the power consumption of the CPU, the main memory, disks, fans etc. by accessing the corresponding connectors used for the supply voltage of an individual component and measuring the power consumption at this connector. For example, the CPU is powered through a +12VDC pin whereas the memory is powered through a +3.3VDC pin. The overall power consumption of the computer system is the sum of the power consumption of the individual components.

For the actual power acquisition and profiling, the software tool LabView [23] has been used for which we have configured several modules operating in a client-server fashion so that the collection of the power data is automated. The client program initiates the LabView server program to acquire power data and then starts the application program for which the power profile is requested. After the termination of the application program, the LabView module saves the collected power consumption data to disk and initiates the post-processing of the data. LabView modules were also needed for the synchronization of the power profiling process with the running application program. The software modules developed are also able to capture the power consumption of specific parts of the application program to be analyzed and to determine which part of a computer system has which power consumption at which point of program execution. As example, Fig. 3 shows the power consumption of the SPEC CPU2006




Figure 2: Photo of the power measurement experimental setup using the NI9205 device.

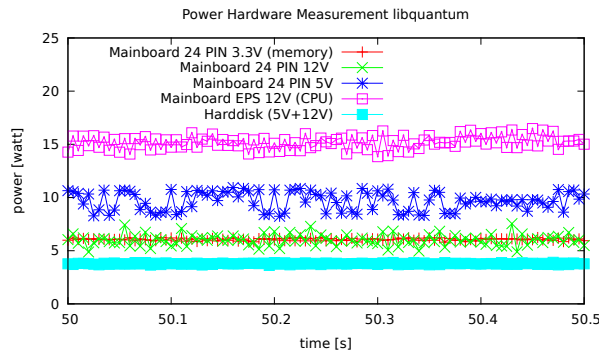


Figure 3: Detailed results of the power measurement for a specific time interval of the SPEC benchmark libquantum using the NI9205 device on an Intel Ivy Bridge processor.

benchmark libquantum running on a Core i7 Ivy Bridge architecture measured with the NI9205 power-meter. The time interval between 50 and 50.5 sec is shown, which has been selected arbitrarily. Other intervals show a similar behavior. The total runtime of the benchmark is 292 sec in the turbo mode at 3.6 GHz. The following five connectors have been measured: (i) the 24 pin ATX 3.3V main-board connector supplying the memory modules; (ii) the 24 pin ATX 12V main-board connector supplying peripheral devices and the CPU; (iii) the 24 pin ATX 5V main-board connector also supplying peripheral devices and the CPU; (iv) the main-board 12V EPS connector which is mainly used by the CPU; (v) the 12V and 5V supply lines for the hard disk. Figure 3 shows that the power consumption of the memory modules is quite small compared to the power consumption of the CPU, and the power consumption due to disk accesses is even smaller.

Figure 4 shows the power consumption for the different connectors for the SPEC CPU2006 benchmark libquantum depending on the frequencies. Especially the power consumption of the 12V EPS connector, which is the main power supply of the CPU, is significantly increasing with the frequency. The power consumption of the ATX 5V main-board connector, which is also partly used by the CPU, also increases slightly. The power consumption of the remaining connectors remains nearly constant over the range of available frequencies, since they are not affected by the frequency change.

The power-meter approach can be extended to capture the power consumption of parallel systems, such as clusters, by measuring the power consumption of the nodes of the cluster system separately and adding up the measured power consumption of the nodes. However, this would require a significant amount of measuring devices, especially for larger parallel systems. An alternative approach is to use software techniques with hardware counters as described in the next Subsection 2.2 requiring much less overhead. The question arises how accurate the measurement with hardware counters is in comparison with a measurement with a power-meter. We address this question in Subsection 2.3.

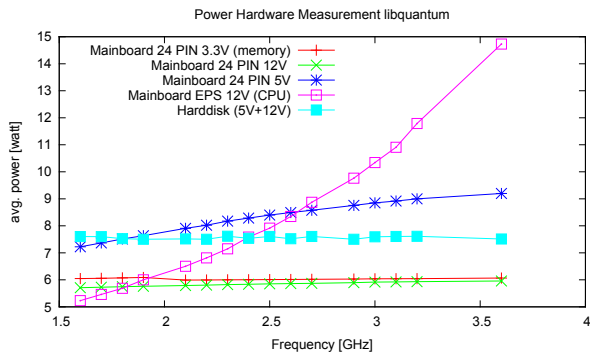


Figure 4: Detailed power measurement for the SPEC benchmark libquantum executed at different frequencies; measurement using the NI9205 device on an Intel Ivy Bridge processor.

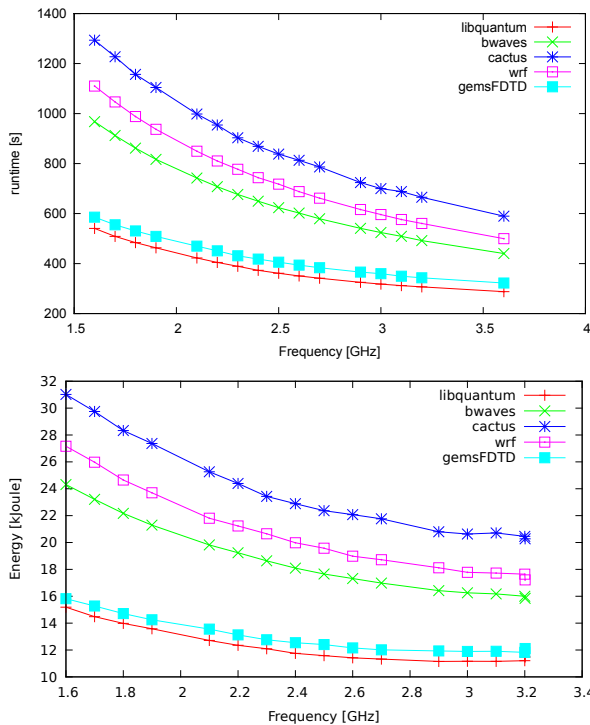


Figure 5: Runtime and energy consumption of selected SPEC CPU206 benchmarks with varying frequencies. For the energy, the total energy consumption is shown, collected with the NI9205 device on an Intel Ivy Bridge processor.

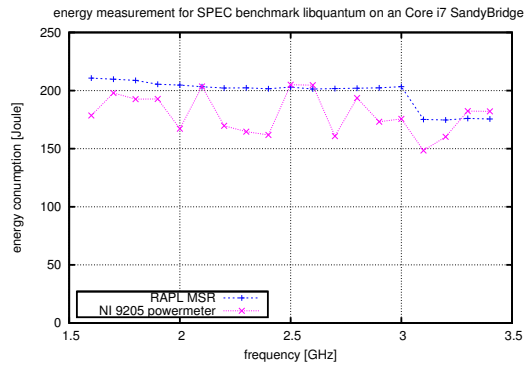


Figure 6: Comparison of the results of the energy measurement with power-meters and the energy measurement using RAPL for the SPEC benchmark libquantum on an Intel Core i7 Sandy Bridge processor.

2.2 Power measurement with RAPL sensors

Many recent processors provide hardware counters which capture the energy consumption. Starting with the Sandy Bridge architecture, Intel has introduced the Running Average Power Limit (RAPL) feature, which provides sensors that allow the measurement of the power consumption of CPU components [29, 17]. The Core i7 processor family has two power planes on chip, which are PP0 containing the processor cores and all caches, and PP1, also referred to as Uncore, containing additional devices, such as graphics devices or the Package Control Unit (PCU) [29]. The corresponding RAPL sensors are RAPL_PP0 and RAPL_PP1 measuring the power consumption of the processor core and the processor uncore. In addition, there are the RAPL sensors RAPL_DRAM and RAPL_PKG for measuring the power consumption of the memory controller and the whole CPU package, respectively.

RAPL sensors can be accessed by control registers, known as Model Specific Registers (MSRs), which are updated in intervals of about 1 *ms* [17]. The MSR can be accessed by a pair of instructions *rdmsr* and *wrmsr*. These instructions are privileged (privilege level 0) and are to be executed in kernel mode. The MSRs provide information about the energy status of the PP0 and PP1 power planes using the specific registers MSR_PP0_ENERGY_STATUS and MSR_PP1_ENERGY_STATUS. The MSR MSR_PKG_ENERGY_STATUS captures the energy status of the whole CPU package. The corresponding energy status unit is obtained via the MSR_RAPL_POWER_UNIT MSR; the default value is 15.3 micro-Joule.

Our investigations include power and energy measurements exploiting RAPL sensors for the SPEC CPU2006 benchmarks. For the RAPL-based experiments, we have used the likwid tool-set (Version 3.0) [33], which provides access to the MSRs introduced above. In the following subsection, the measurement methods with power-meters and the measurement method with RAPL sensors are compared and evaluated. The goal is to clarify the influence of the measurement overhead on the measurement results.

2.3 Comparison of the measurement techniques

Hardware counters, such as the MSR registers, can capture only the power consumption of the CPU. On the other hand, measurements with specialized devices, such as the NI9205 system, enable a more detailed measurement that can also take the memory system, fans, and peripheral devices into account. However, a significant hardware and software overhead is required to perform accurate power-meter measurements with such a device. One might argue that the main contribution to power consumption comes from the CPU and that the measurement with hardware counters is good enough to capture the relevant effects. In this context, it is interesting to investigate how accurate the power measurement with hardware counters is compared with detailed measurements using special power-meter devices. We address this issue in this section and present a comparison for selected application programs.

To set the frequencies of the cores to a fixed value, we have used the *cpufreq_set* tool; the minimum and the maximum core frequencies have been set to the same value. During program execution, the frequency of each core can be observed with the *i7z* tool. The runtime experiments have been performed with no other application

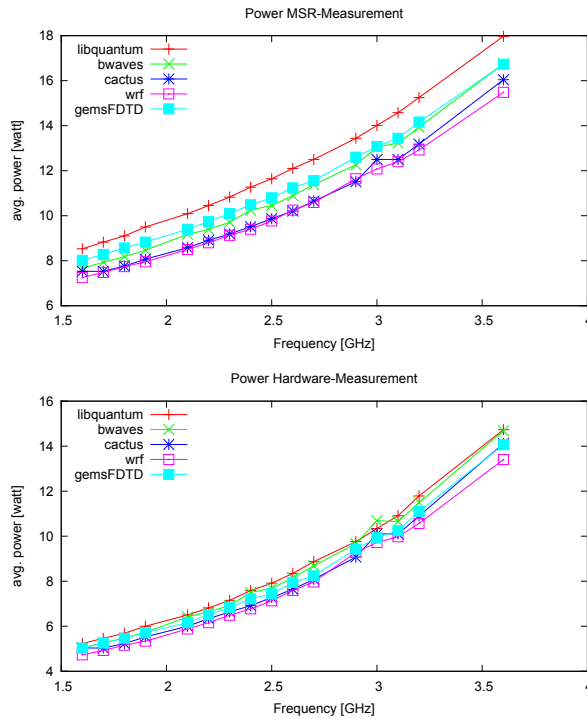


Figure 7: Comparison of power measurement with MSR registers (top) and NI9205 device (bottom) on an Intel Ivy Bridge processor.

program running on the machine, i.e., interferences can only come from jobs of the operating system. To reduce the effect of such interferences, the runtime experiments have been performed several times.

For the Sandy Bridge architecture, the operational frequency can be set between 1.6 GHz and 3.4 GHz in 100 MHz steps. The same holds for the Ivy Bridge architecture with the exception that the maximal frequency is 3.2 GHz. For the Haswell architecture, the frequency can be set between 0.8 GHz and 3.4 GHz; the stepsize is usually 200 MHz with two exceptions (1.4/1.5 GHz and 2.7/2.8 GHz).

Figure 6 compares the two energy measurement techniques described above. The measurement has been performed for the SPEC benchmark libquantum on an Intel Core i7 Sandy Bridge processor. The benchmark runs have been repeated 20 times with the "test" input. The power measurement shown in Figure 6 using the NI 9205 acquisition system takes only the main-board 12V EPS connector into consideration, which is the main power supply of the CPU. The 12V ATX and the 5V ATX connector, which also partially supply the CPU, but also other peripheral devices, are not taken into consideration, as it cannot be determined which part of the power actually goes to the CPU. On the other hand, the measurement with the RAPL interface captures the complete power consumption of the CPU. The diagram shows a good match of the energy values obtained with both techniques for some of the frequencies. For other frequencies, the energy values obtained with RAPL are up to about 20% larger than the energy values obtained with the NI 9205 device. A reason for these deviations could be that there is an additional power flow to the CPU via the 12V and the 5V ATX connectors, which is not captured. Our investigations have shown that these deviations are smaller for most of the other SPEC benchmarks on the Sandy Bridge architecture. The differences in the deviations for different frequencies are not present for the Ivy Bridge processor as described in the following.

Figure 7 shows the power consumption obtained with the two measurement methods for the Core i7 Ivy Bridge architecture using five selected SPEC benchmarks. The figure shows that the power consumption increases more than linearly when the frequency is increased. Both measurement techniques can capture this behavior. The figure also shows that the two measurement methods lead to similar values for the power consumption. In particular, the relative order of the power consumption of the different benchmarks is the same for both measurement techniques. However, the power values obtained with RAPL are systematically larger than the power values obtained with the

NI 9205 device for the same reason as mentioned for Figure 6. In particular, the power measured with the RAPL interface is typically about 3 Watt larger than the power measured with the NI 9205 device using only the 12V EPS connector. Adding the power consumption of the 12V ATX and 5V ATX connectors to the power measured with the NI 9205 device would lead to larger values as measured with the RAPL interface. In summary, it can be concluded that the power measurement via RAPL is accurate enough to replace the usage of power-meters.

3 Energy consumption of the SPEC CPU benchmark programs

In this section, we present results of power and energy measurements for the entire set of applications in the SPEC CPU2006 benchmark suite, see Figs. 8 - 10. Based on the findings of the previous section, the RAPL interface has been used to obtain the power and energy values. The likwid toolset [33] has been used to access the MSR registers of the different architectures.

The SPEC CPU2006 benchmark suite has been developed with the goal to capture the performance of desktop systems and single-processor servers. The benchmark suite consists of integer and floating-point benchmarks, which are real programs covering different application areas for computer systems. The benchmarks are sequential C, C++, or Fortran programs. The integer programs include, for example, a compression program (bzip2), a C compiler (gcc), a video compression program, a chess game, and an XML parser. The floating point programs include, for example, several simulation programs from physics, a speech recognition program, a ray-tracing program (povray), as well as programs from numerical analysis and a linear programming algorithm (soplex), see, e.g., [15] and www.spec.org for more details. In the experimental evaluation, the benchmarks have been compiled with gcc 4.7.2 using the compiler option `-ftree-parallelize-loops=4`, enabling an automatic parallelization at loop level.

Figure 8 shows the runtimes (top) in seconds, the energy consumption (middle) in Joule, and the power consumption (bottom) in Watt as functions of the frequency for the SPEC CPU2006 integer benchmarks on an Intel Core i7 Sandy Bridge architecture. All available frequencies are shown. The figure shows that the runtime decreases nearly linearly with increasing frequencies for all benchmarks. Only for smaller frequencies, the increase of the execution time with a reduced frequency is more than linear. The diagram shows that for most of the SPEC CPU2006 integer benchmarks, there is only a slight variation of the energy consumption for the entire range of the frequencies possible. The power consumption values P have been calculated from the energy values E obtained from the MSR registers and the runtime t measured also with hardware registers using $E[\text{Joule}] = P \times t[\text{Watt} \cdot \text{s}]$. Looking at the power consumption, it can be observed that different applications lead to slightly different power consumptions, i.e., there is a small dependence of the power consumption from the characteristics of the application. Two benchmarks (gcc and libquantum) lead to a larger power consumption than the rest of the integer benchmarks. This shows that in the case of the SPEC CPU2006 integer benchmarks, the power consumption mainly depends on characteristics of the hardware and that the different energy consumptions for the different benchmarks are mainly caused by the different execution times.

Figure 9 shows the runtimes (top), energy consumption (middle), and power consumption (bottom) for the SPEC CPU2006 floating-point benchmarks for different frequencies on an Intel Core i7 Sandy Bridge architecture. The SPEC CPU2006 floating-point benchmarks typically have a slightly larger execution time than the integer benchmarks. Accordingly, a larger energy consumption results. The energy consumption shows a similar behavior as for the integer benchmarks and there is no great variation of the energy consumption in dependence of the frequency. Compared to the integer benchmarks, a slightly larger variation of the power consumption can be observed for the different SPEC floating-point benchmarks. This suggests that there is a larger dependence of the power consumption from the usage of processor resources by the specific application.

Figures 10 and 11 show the runtimes (top), energy consumption (middle), and power consumption (bottom) of the SPEC CPU2006 integer and floating-point benchmarks for different frequencies on an Intel Core i7 Haswell architecture. Similar to the Sandy Bridge architecture, the figures show a decrease of the runtime of the different benchmarks which is nearly linear to the operational frequency. For smaller frequencies, the more-than-linear increase of the execution time with reduced frequency is stronger than on the Sandy Bridge architecture, since the Haswell architecture allows smaller frequencies. Especially for larger frequencies, most of the benchmarks are faster on the Haswell architecture than on the Sandy Bridge architecture, presumably due to the more recent architecture with a larger L3 cache. Especially for small frequency values, the energy consumption increases more

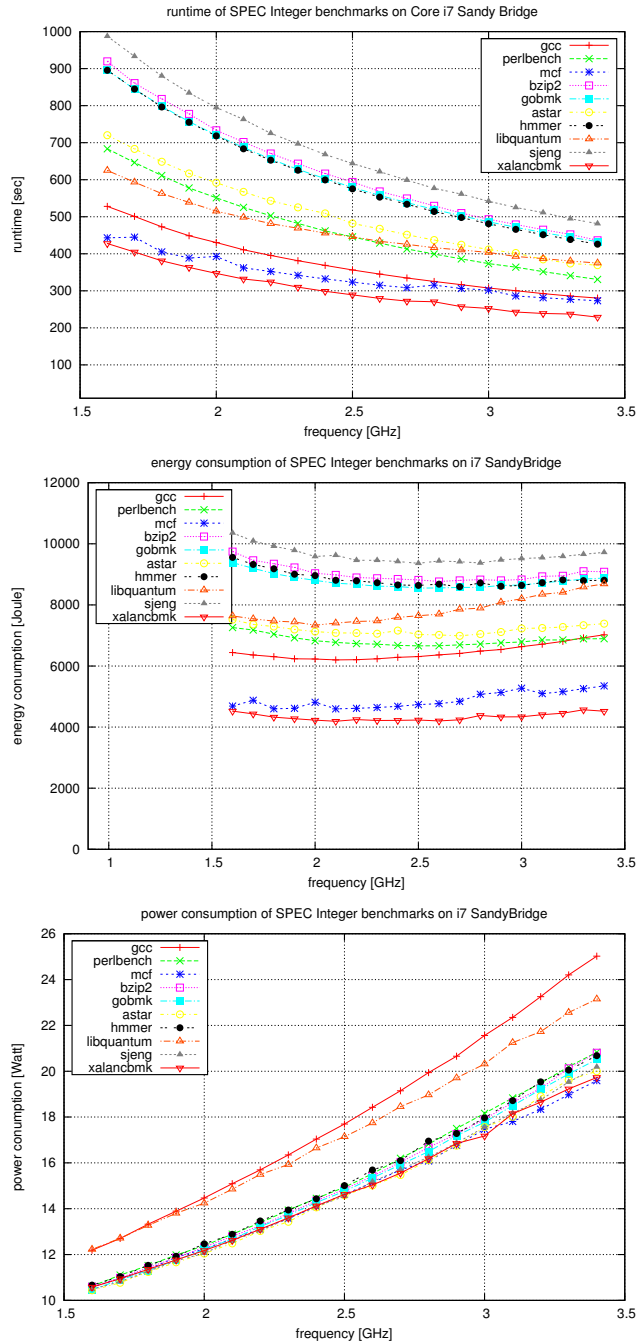


Figure 8: SPEC CPU2006 integer benchmarks on an Intel Core i7 Sandy Bridge processor: runtime (top), energy consumption (middle), and power consumption (bottom) for varying frequencies.




Figure 9: SPEC CPU2006 floating-point benchmarks on an Intel Core i7 Sandy Bridge processor: runtime (top), energy consumption (middle), and power consumption (bottom) for varying frequencies.

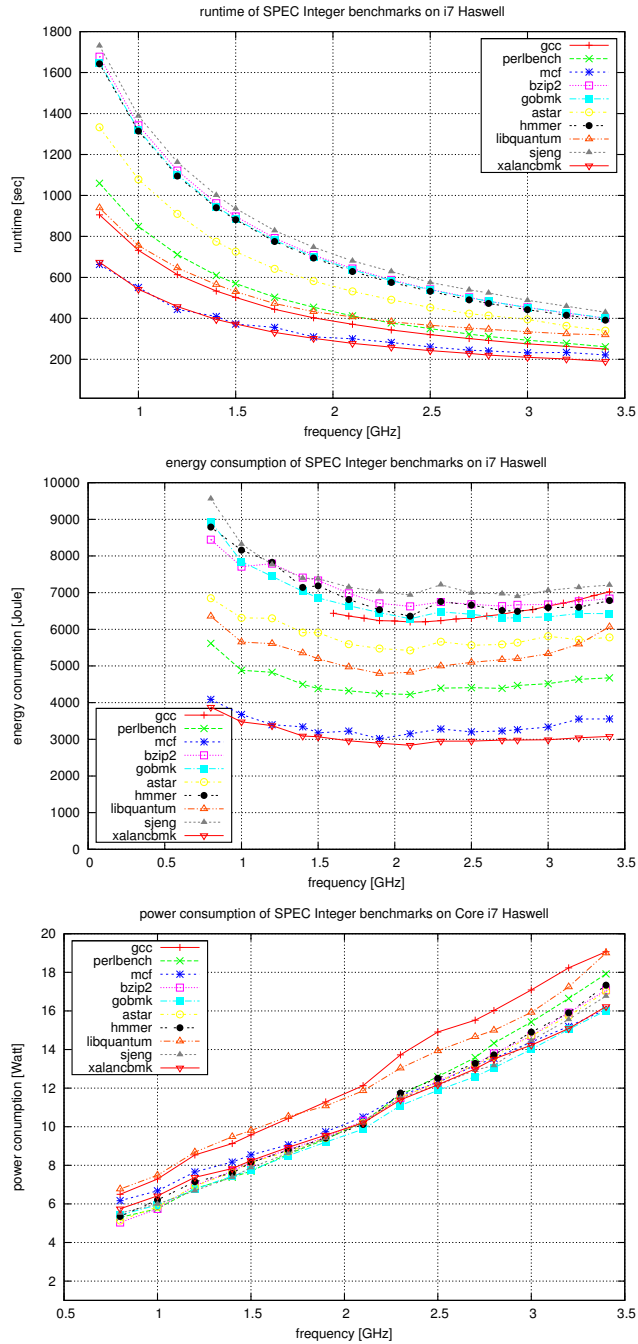


Figure 10: SPEC CPU2006 integer benchmarks on an Intel Core i7 Haswell processor: Runtime (top), energy consumption (middle), and power consumption (bottom) for varying frequencies.

visibly than for the Sandy Bridge architecture, caused by the increase in execution time. The power consumption on the Haswell architecture is typically smaller than on the Sandy Bridge architecture for the same frequency. The same holds for the energy consumption, i.e., the Haswell architecture usually needs less energy for the same benchmarks than the Sandy Bridge architecture using the same frequency.

For all SPEC CPU2006 benchmarks and both architectures, the diagrams of the runtime, the power and the energy consumption show a similar qualitative behavior as a function of the frequency. The runtime is a decreasing, slightly convex function with the highest runtime for the smallest available frequency and the smallest runtime for the highest available frequency. The power consumption is a slightly convex, increasing function depending on the frequency. As a result, the energy consumption, being the power integrated over the runtime, is a convex function that is slightly descending for smaller frequencies and slightly ascending for higher frequencies. This effect is more distinct for the Haswell processor. These energy functions have a minimum in the frequency range between 2 GHz and 2.7 GHz, varying for the different benchmarks and architectures. Since the SPEC CPU2006 benchmarks cover a wide range of different applications, one may conclude that it is a typical behavior to have an energy minimum, and to achieve the minimum, a suitable frequency should be used. However, the minimum of the runtime, being at the largest frequency, and the minimum of the energy consumption do not coincide. Thus, the user has to decide whether to minimize the runtime or the energy consumption before selecting a frequency.

4 Energy models with frequency scaling

Many different energy models have been proposed and used for simulating the energy consumption of application programs on CPUs, see [8, 20] for an overview. These models usually concentrate on the dynamic power consumption, which used to be the most dominant part in earlier processors, and more recently also include the static power consumption, memory power consumption or other significant power aspects of a specific CPU or computer system. The dynamic power consumption is related to the supply voltage and the switching activity during the computing activity of the processor. The static power consumption captures the leakage power consumption as well as the power consumption of peripheral devices. The total power consumption includes both power components.

For DVFS processors, the power consumption depends on the operational frequency f , see, e.g., Fig. 8 for an experimental evaluation of this dependence. In this section, we investigate two models simulating the energy consumption, a physically based model which has been proposed in the literature and a heuristic model that is derived from the experimental results in Section 3.

4.1 Physical energy model with frequency scaling

Analytical energy models describe the power consumption of a processor by power models developed for digital circuits used for the construction of the processors. The energy consumption is calculated by a multiplication with the runtime of an application program. This class of energy models is typically derived from physical considerations of the expected power behavior, and the models are used within optimization or scheduling methods which are aimed at the minimization of the energy consumption. In the following, we exploit a physical energy model that has been proposed in the literature, see [8, 20]. In this article, we use this model and determine parameters for the SPEC benchmarks.

The specific energy model that we consider approximates the dynamic power consumption by $P_{dyn} = \alpha \cdot C_L \cdot V^2 \cdot f$, where α is the switching probability, C_L is the load capacitance, V is the supply voltage, and f is the operational frequency. The static power consumption is intended to capture the leakage power consumption which consists of several components, including sub-threshold leakage, reverse-biased-junction leakage, gate-induced-drain leakage, gate-oxide leakage, gate-current leakage, and punch-through leakage [20]. The exact power values for these components are varying and depend on the specific architecture considered; however, only approximations are needed. Such an approximation has been proposed by [7], modeling the static power consumption due to leakage power as $P_{stat} = V \cdot N \cdot k_{design} \cdot I_{leak}$, where V is the supply voltage, N is the number of transistors, k_{design} is a design dependent parameter, and I_{leak} is a technology-dependent parameter.

For DVFS processors, we are interested in investigating the dependence of the power consumption on the operational frequency, which can be scaled within a predefined interval $[f_{min}, f_{max}]$. The scaling can be expressed

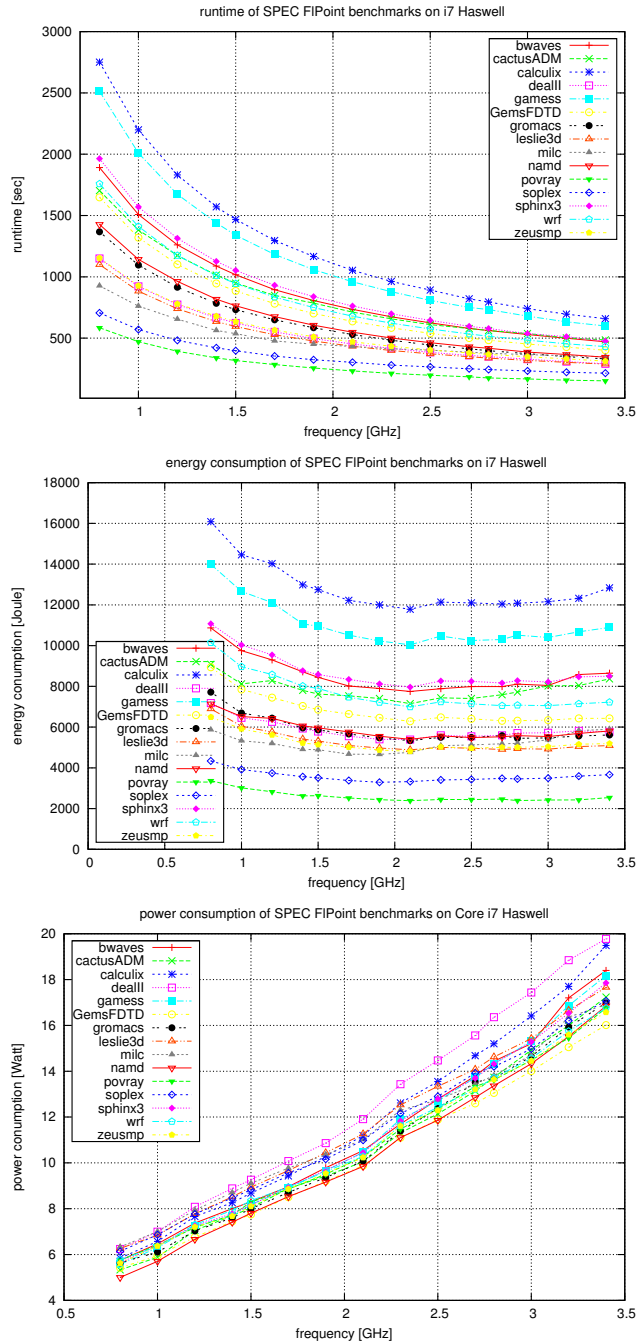


Figure 11: SPEC CPU2006 floating-point benchmarks on an Intel Core i7 Haswell processor: runtime (top), energy consumption (middle), and power consumption (bottom) for varying frequencies.

by a dimensionless scaling factor $s \geq 1$ which describes a smaller frequency $\tilde{f} < f_{max}$ as $\tilde{f} = f_{max}/s$. The following functional dependencies have to be considered: The frequency f depends linearly on the supply voltage V , i.e., $V = \beta \cdot f$ with some appropriate constant β . Thus, the dependence of the dynamic power consumption on the frequency f can be expressed as $P_{dyn}(f) = \gamma \cdot f^3$ with $\gamma = \alpha \cdot C_L \cdot \beta^2$ or with the corresponding scaling factor s as $P_{dyn}(s) = s^{-3} \cdot P_{dyn}(1)$ where $P_{dyn}(1)$ is the dynamic power consumption in the un-scaled case. This means that the dynamic power increases cubically with the operational frequency, which can be used to study the change of the dynamic power consumption with respect to varying frequency values. Using $V = \beta \cdot f$ for the static power consumption P_{stat} leads to a linear dependence of the static power on f , i.e., $P_{stat}(f) = \delta \cdot f$ with $\delta = N \cdot k_{design} \cdot I_{leak} \cdot \beta$ or $P_{stat}(s) = s \cdot P_{stat}(1)$ where $P_{stat}(1)$ is the static power consumption in the un-scaled case.

Reducing the operational frequency of a processor by a scaling factor of s usually decreases the power consumption, however, it also increases the execution time $T_P(1)$ of a program P by the same factor compared to an un-scaled execution, i.e., $T_P(s) = s \cdot T_P(1)$. This has to be taken into consideration for the time integration to compute the energy consumption $E_P(s) = (P_{dyn}(s) + P_{stat}(s)) \cdot s \cdot T_P(1) = (s^{-3} \cdot P_{dyn}(1) + s \cdot P_{stat}(1)) \cdot s \cdot T_P(1)$. Runtime experiments on an Intel Sandy Bridge show that the linear dependence of the runtime on the frequency scaling factor can more or less be observed. However, for smaller frequencies the runtime increases more than linearly, see Figs. 8 and 9. This effect is even stronger for the Core i7 Haswell architecture, see Figs. 10 and 11, since this architecture supports smaller frequencies.

4.2 Heuristic energy model

The second energy model considered in this article is derived by an approximation approach based on the discrete energy data measured. This model considers the entire power consumption of the CPU processing one of the benchmarks and uses data fitting to derive a closed formula. Thus, the derivation of this energy model applies the opposite approach compared to the model in Section 4.1: it starts with the discrete data from the experimental setting and uses curve fitting methods with the purpose of getting a continuous approximation of the discrete data. The heuristic model described in the following is new and has not been used before.

Due to the power data that we have collected in our experimental setup, e.g. given in Figs. 10 and 11, we assume an almost linear dependence of the power from the frequency f with unknown parameters, i.e. we assume the power to be approximated by $P_{total}(f) = a + bf^{1+\epsilon}$ with parameters a and b to be determined by curve fitting. Although these parameters do not correspond to physical constants, but are the result of a mathematical calculation, the parameter a can be interpreted as a static part of the power consumption that does not change with the frequency, whereas b captures a dynamic part of the power consumption that increases with the operational frequency of the CPU. For the parameter ϵ , several fixed values can be used so that the parameters a and b can be determined by applying the least squares method. The values of a and b reflect the characteristics of the specific CPU. However, different benchmarks may have different values for the parameters a and b due to their specific computational and memory access behavior, leading to a different usage of the processor resources and different power requirements. The results of the curve fitting are included in the next section reporting the comparison of the predicted and measured energy values.

5 Evaluation of the energy models

In this section, we investigate to which extent the energy models presented in the last section are able to capture the power and energy consumption of standard benchmark programs on recent processors. The emphasis lies on a modeling of the energy consumption as a function of the frequency in the frequency range provided by the processor. For those frequencies, we quantitatively and qualitatively compare the measured energy values as reported in Section 3 with energy values provided by the models from Section 4. Subsection 5.1 considers an application-specific energy modeling and Subsection 5.2 addresses an application-independent modeling and discusses the differences.




Figure 12: Intel Core i7 Sandy Bridge processor: Comparison of measured and predicted energy consumption of the SPEC CPU2006 integer and floating-point benchmarks for the frequencies $f = 1.6$ GHz (top), $f = 2.5$ GHz (middle), and $f = 3.4$ GHz (bottom).

5.1 Application-specific energy modeling

The physical energy model based on frequency scaling from Section 4.1 contains the parameters γ in the equation for the dynamic power and δ in the equation for the static power consumption. The values for γ and δ have been determined by curve fitting with the least-squares method using the measured power values for the different frequencies for each individual application of the SPEC CPU2006 benchmark suite. The resulting parameters γ and δ are application-specific and can be used for a modeling of the energy consumption of the specific application as described in Section 4.1.

For the heuristic model from Section 4.2, $\epsilon = 0.2$ has been used. This value has been determined by experiments which have shown that this value is well suited for all applications investigated. The model again contains parameters a and b , and as just described, we have applied curve fitting to determine these parameter values for the individual application of the SPEC CPU2006 benchmark suite. Based on these application-specific parameters, Figure 12 compares the measured and predicted energy consumption for the different SPEC CPU2006 benchmarks for the Sandy Bridge architecture using selected frequencies. The smallest and the largest frequency as well as a medium frequency have been selected for the diagrams. The entry *predicted 1* corresponds to the physical energy model from Section 4.1, and the entry *predicted 2* corresponds to the heuristic energy model from Section 4.2. Figure 13 shows the same comparison for the Haswell architecture and also includes a comparison with an application-independent modeling to be discussed in the next subsection.

From Figures 12 and 13 it can be observed that for the application-specific modeling both energy models are well suited to describe the energy consumption of most benchmark programs and both architectures. The correspondence between modeled and measured energy consumption is especially good for medium and large frequencies, as shown for frequencies $f = 2.5$ GHz and $f = 3.4$ GHz. Similar results are obtained for other frequencies between these frequency values. The deviations usually lie below 10%.

The comparison of the modeling quality using the two different models shows that for the Sandy Bridge processor, see Figure 12, there is no significant difference between the energy values provided by the two models. On the Haswell architecture, larger deviations between the measured and predicted values can be observed for smaller frequencies, when the physical energy model from Section 4.1 is used, see the results for frequency $f = 0.8$ GHz in Figure 13 (top) as example. These deviations are much smaller for the heuristic energy model. For small frequencies, the physical energy model systematically underestimates the energy consumption for both the Sandy Bridge and the Haswell processor. The reason for this underestimation is assumed to be an underestimation of the static power consumption when using $P_{stat}(f) = \delta \cdot f$. Other energy models [24, 38] propose to use $P_{stat}(f) = constant$ for all frequencies, but this would lead to a significant overestimation of the total energy consumption for small frequencies (not shown in a figure), leading to even bigger deviations from the measured energy consumption. In contrast, the heuristic energy model captures the dynamic and the static energy consumption together as one physical entity in one formula, and this seems to be better suited for the whole range of frequencies. Thus, the heuristic model provides energy values that are better suited for a prediction than those provided by the physical energy model from Section 4.1.

Averaging over all frequencies and benchmarks, the average percentage difference between the measured energy values and the predicted values using the heuristic model with application-specific parameters is 1.8% with a minimum difference of 0.9% and a maximum difference of 6.0%. Using the physical model, the average percentage difference is 5.6% with a minimum difference of 4.4% and a maximum difference of 10.7%.

5.2 Application-independent energy modeling

One goal for modeling the energy consumption values is to provide a prediction model for the energy consumption of application programs. Such a model suitable as prediction model has to be easy to use and has to be application-independent. Both the models introduced in Section 4 are represented by energy functions that are quite easy to evaluate. So both models meet the goal of simplicity. However, so far we have performed an application-specific modeling due to the parameters γ , δ , a , and b , respectively. An application-independent model requires parameters which are independent from a specific benchmark application but may depend on the processor architecture. In this subsection, we propose to use the average of the parameter values γ and δ or a and b over all benchmarks executed on the same processor.

For the different programs of the SPEC CPU2006 benchmark suite, the values of the parameters γ and δ resulting from the physical energy model are quite close to each other for most of the benchmarks on the same architecture; their difference is typically below 10 %. Thus, using the average values of the parameters also leads to a good correspondence between measured and modeled values. This represents an energy modeling that is independent from the specific application but is still architecture-dependent. For the different architectures, different values for the parameters γ and δ result.

The modeling with the average parameter values is additionally included in Figure 13. The entry "predicted 1 av" corresponds to the modeling with the physical energy model from Section 4.1 using the average values for




Figure 13: Intel Core i7 Haswell processor: Comparison of measured and predicted energy consumption of the SPEC CPU2006 integer and floating-point benchmarks for the frequencies $f = 0.8$ GHz (top), $f = 2.5$ GHz (middle), and $f = 3.4$ GHz (bottom).

the parameters γ and δ taken over all SPEC CPU2006 applications on the Haswell processor. Similarly, the entry "predicted 2 av" corresponds to the modeling with the heuristic model using the average values for the parameters a and b . Using the average parameter values for the modeling leads to energy values that are slightly less close to the measured energy values than for the application-specific modeling, but the deviations are still acceptable. However, in some cases, especially for smaller frequencies, using the average parameter values for the modeling even leads to better predictions than the application-specific modeling. This is due to the strongly convex behavior of the measured energy values for frequencies between 0.8 GHz and 1.6 GHz, see Figures 10 and 11, which indicate a high overhead of the runtime even though the power function remains linear. The architectural effects are difficult to capture in one continuous functional pattern. However, the modeling with a piecewise continuous function would result in a too complex prediction model. For some benchmarks, the averaging of the parameters seems to level off extreme values. The heuristic energy model using the average parameter values for the modeling leads to slightly better predictions than the physical energy model with the average parameter values for the modeling.

From these observations it can be concluded that both energy models considered, although different, are able to capture the energy consumption and its dependence on the operational frequency with reasonable accuracy for most situations. Thus, both energy models can be used for an a priori prediction of the energy consumption as it is needed when solving problems such as optimization or scheduling. Compared to the physical model, the heuristic model leads to better predictions in most situations and should therefore be preferred. This is due to a better modeling of the leakage power, which may not be best represented by the linear term in the physical model. Other models proposed in the literature use constant, quadratic, or even exponential terms for the leakage power consumption. However, these models are usually not validated for a wide range of application programs.

6 Related Work

Today, power-management mechanisms are integrated in computer systems of almost every size and class, from handheld devices to large servers [31]. Correspondingly, the analysis, measurement, or simulation of power and energy consumption is an active research area. An important feature is the DVFS technique, which trades off performance for power consumption by lowering the operating voltage and frequency if this is possible, see [3, 34, 38] for an overview. Frequency scaling due to the DVFS technique has been applied in practical programming as well as in theoretical investigations, optimizing or minimizing the energy consumption under certain constraints.

Approaches to determine the frequency scaling factor that minimizes the total CPU energy consumption by taking both the dynamic power and the leakage power into consideration have been discussed in [18, 19, 38] for sequential programs. Voltage scaling has also been considered in [21].

Energy simulation models are often used in approaches to minimize the energy consumption by scheduling algorithms or other heuristics. Algorithmic research on speed scaling processors and related scheduling algorithms using the total energy consumption as objective function has been initiated by the article [36]. Many different scenarios and algorithms have been investigated in the literature, see [1, 5, 10] for an overview. For example, theoretical foundations of scheduling algorithms in a setting with dynamic speed scaling processors are investigated in [2], considering the scheduling of n jobs on m identical variable speed processors working in parallel, where each job is specified by a release date, a deadline, and a processing volume. Different scenarios concerning the job size, release dates and deadlines are considered and approximation algorithms for the resulting NP-hard scheduling problems are presented. In most of the articles in this research line, the emphasis lies on a theoretical investigation of the approximation algorithm derived and no simulations or measurements on real hardware systems are provided.

A comparison of energy measurements using PowerPack and RAPL (accessed via PAPI) for dense linear algebra algorithms is given in [12], showing that the RAPL measurements are a good alternative to physical power-meters. However, frequency scaling has not been taken into consideration in [12]. The power management architecture of the Sandy Bridge processor is described in [29] and includes a comparison of the actual measured power and the so-called architectural power-meter which predicts the active power consumption. The dependence of the power and the frequency is not directly investigated nor is the modeling with a theoretical power model included. An approach using RAPL for memory power estimation is discussed in [11]. The correspondence with simulation methods and prediction possibilities has also not been considered in the publication. A detailed comparison of power measurement techniques including RAPL has been given in [14]. However, the DVFS features with a

variation of the frequency and the resulting impact on the power are not investigated. The performance API PAPI now also covers the report of energy and power values based on RAPL assumed that the code is previously instrumented, see [35]. The PAPI approach corresponds to the software approach accessing hardware counters in our work. [35] also lists state-of-the-art measurement techniques based on power-meters and hardware counters, but does not provide a comparison.

Performance prediction for DVFS processors is addressed in [30] with an emphasis on green supercomputing. The energy consumption of parallel algorithms for shared memory architectures based on the parallel external memory (PEM) model [4] has been discussed in [22]. [32] proposes a system-level iso-energy-efficiency model to analyze, evaluate and predict energy-performance of data intensive parallel applications running on cluster systems. Based on measurements from smaller configurations, the power performance for larger systems with increasing number of nodes is predicted and analyzed. As example applications, the FT, EP, and CG benchmarks of the NAS Parallel Benchmarks are used. A scalability analysis shows a good correspondence of the predictions of the model. The interaction between the parallel execution time and the energy consumption is considered in [9] by partitioning a parallel algorithm into sequential and parallel regions and computing optimal frequencies for these regions. Approaches for an energy complexity metric are discussed in [6]. The application of a physical energy model and its usage for the scheduling of sequential tasks in fork-join patterns with energy as objective function has been addressed in [27]; the energy-based scheduling of parallel tasks has been considered in [28].

In the domain of real-time scheduling, many techniques for utilizing available waiting times based on DVFS have been considered, see, e.g., [16, 26, 37]. These approaches are usually based on heuristics and are not based on an analytical model as presented in this work. The effects of dynamic concurrency throttling (DCT) and DVFS in the context of a hybrid MPI/OpenMP programming model are considered in [25]. In particular, frequency selection is formulated as a variant of the 0-1 knapsack problem and dynamic programming is used to compute an approximation.

7 Conclusions

In this article, we have investigated and evaluated measurement methods providing energy consumption values for recent DVFS processors. The article reports the experimental setting for the hardware as well as for the software measurement approach and presents the resulting measurement data. The CPUs investigated are the Intel Core i7 Sandy Bridge, Ivy Bridge and Haswell architectures. The application codes measured are the SPEC CPU2006 benchmarks. The analysis of the energy data gathered with the power-meter based and the RAPL based energy measurements show a good correspondence across all frequency values, CPUs and benchmarks, leading to the insight that both measurement methods are suitable and interchangeable. We consider this result to be especially valuable for research projects concerned with bigger applications as well as larger machines based on these recent CPUs. A second result of our investigations is that physical and heuristic energy models are well suited for an early prediction of the energy consumption of application programs. Thus, these energy models can be used as a basis for simulation approaches of the energy consumption of application programs.

References

- [1] Albers, S.: Energy-efficient algorithms. *Commun. ACM* **53**(5), 86–96 (2010). DOI 10.1145/1735223.1735245. URL <http://doi.acm.org/10.1145/1735223.1735245>
- [2] Albers, S., Müller, F., Schmelzer, S.: Speed scaling on parallel processors. In: *Proc. of the 19th Annual ACM Symp. on Parallel Algorithms and Architectures, SPAA '07*, pp. 289–298. ACM, New York, NY, USA (2007). DOI 10.1145/1248377.1248424. URL <http://doi.acm.org/10.1145/1248377.1248424>
- [3] Anshumali, K., Chappell, T., Gomes, W., Miller, J., Kurd, N., Kumar, R.: Circuit And Process Innovations to Enable High-Performance, and Power and Area Efficiency on the Nehalem and Westmere Family of Intel processors. *Intel Technology Journal* **14**, 104–127 (2010)
- [4] Arge, L., Goodrich, M., Nelson, M., Sitchinava, N.: Fundamental parallel algorithms for private-cache chip multiprocessors. In: *SPAA '08: Proc. of the 20th Ann. Symp. on Parallelism in Algorithms and Architectures*, pp. 197–206. ACM (2008). DOI <http://doi.acm.org/10.1145/1378533.1378573>

- [5] Bansal, N., Kimbrel, T., Pruhs, K.: Speed scaling to manage energy and temperature. *J. ACM* **54**(1), 3:1–3:39 (2007). DOI 10.1145/1206035.1206038. URL <http://doi.acm.org/10.1145/1206035.1206038>
- [6] Bingham, B., Greenstreet, M.: Computation with Energy-Time Trade-Offs: Models, Algorithms and Lower-Bounds. In: ISPA '08: Proc. of the 2008 IEEE Int. Symp. on Parallel and Distributed Processing with Applications, pp. 143–152. IEEE Computer Society (2008). DOI <http://dx.doi.org/10.1109/ISPA.2008.127>
- [7] Butts, J., Sohi, G.: A static power model for architects. In: In Proc. of the 33rd Int. Symp. on Microarchitecture (MICRO-33) (2000)
- [8] Chen, H., Shi, W.: Power Measurement and Profiling. In: I. Ahmad, S. Ranka (eds.) *Handbook of Energy-Aware and Green Computing*, pp. 649–674. CRC Press (2012)
- [9] Cho, S., Melhem, R.: Corollaries to Amdahl's Law for Energy. *IEEE Comput. Archit. Lett.* **7**(1), 25–28 (2008). DOI <http://dx.doi.org/10.1109/L-CA.2007.18>
- [10] Chrobak, M.: Algorithmic Aspects of Energy-Efficient Computing. In: I. Ahmad, S. Ranka (eds.) *Handbook of Energy-Aware and Green Computing*, pp. 311–329. CRC Press (2012)
- [11] David, H., Gorbatoev, E., Hanebutte, U., Khanaa, R., Le, C.: RAPL: memory power estimation and capping. In: Proc. International Symposium on Low Power Electronics and Design (ISLPED), pp. 189–194. ACM (2010)
- [12] Dongarra, J., Ltaief, H., Luszczek, P., Weaver, V.: Energy Footprint of Advanced Dense Numerical Linear Algebra Using Tile Algorithms on Multicore Architectures. In: Proc. 2nd Int. Conf. on Cloud and Green Computing (CGC), pp. 274–281. IEEE (2012)
- [13] Ge, R., Feng, X., Song, S., Chang, H.C., Li, D., Cameron, K.: PowerPack: Energy Profiling and Analysis of High-Performance Systems and Applications. *IEEE Transactions on Parallel and Distributed Systems* **21**(5), 658–671 (2010)
- [14] Hackenberg, D., Ilsche, T., Schöne, R., Molka, D., Schmidt, M., Nagel, W.: Power measurement techniques on standard compute nodes: A quantitative comparison. In: 2013 IEEE International Symposium on Performance Analysis of Systems & Software, pp. 194–204 (2013)
- [15] Hennessy, J., Patterson, D.: *Computer Architecture - A Quantitative Approach* (5. ed.). Morgan Kaufmann (2012)
- [16] Horvath, T., Abdelzaher, T., Skadron, K., Liu, X.: Dynamic Voltage Scaling in Multitier Web Servers with End-to-End Delay Control. *IEEE Trans. Comput.* **56**(4), 444–458 (2007). DOI <http://dx.doi.org/10.1109/TC.2007.1003>
- [17] Intel: Intel 64 and IA-32 Architecture Software Developer's Manual, System Programming Guide (2011)
- [18] Irani, S., Shukla, S., Gupta, R.: Algorithms for power savings. *ACM Trans. Algorithms* **3**(4), 41 (2007). DOI <http://doi.acm.org/10.1145/1290672.1290678>
- [19] Jejurikar, R., Pereira, C., Gupta, R.: Leakage aware dynamic voltage scaling for real-time embedded systems. In: DAC '04: Proceedings of the 41st annual Design Automation Conference, pp. 275–280. ACM (2004)
- [20] Kaxiras, S., Martonosi, M.: *Computer Architecture Techniques for Power-Efficiency*. Morgan & Claypool Publishers (2008)
- [21] Kim, T.: Power Saving by Task-level dynamic Voltage Scaling: A Theoretical Aspect. In: I. Ahmad, S. Ranka (eds.) *Handbook of Energy-Aware and Green Computing*, pp. 361–383. CRC Press (2012)
- [22] Korthikanti, V., Agha, G.: Towards optimizing energy costs of algorithms for shared memory architectures. In: SPAA '10: Proceedings of the 22nd ACM symposium on Parallelism in algorithms and architectures, pp. 157–165. ACM, New York, NY, USA (2010). DOI <http://doi.acm.org/10.1145/1810479.1810510>
- [23] LabVIEW: LabVIEW Measurement data format, <http://www.ni.com/white-paper/4139/en>
- [24] Lee, Y., Zomaya, A.: Minimizing Energy Consumption for Precedence-Constrained Applications Using Dynamic Voltage Scaling. In: CCGRID '09: Proc. of the 2009 9th IEEE/ACM Int. Symp. on Cluster Computing and the Grid, pp. 92–99. IEEE Computer Society (2009)
- [25] Li, D., de Supinski, B., Schulz, M., Nikolopoulos, D., Cameron, K.: Strategies for Energy Efficient Resource Management of Hybrid Programming Models. *IEEE Transaction on Parallel and Distributed Systems* (2012)
- [26] Mishra, R., Rastogi, N., Zhu, D., Mossé, D., Melhem, R.: Energy Aware Scheduling for Distributed Real-Time Systems. In: IPDPS '03: Proc. of the 17th Int. Symp. on Parallel and Distributed Processing, p. 21.2. IEEE Computer Society (2003)
- [27] Rauber, T., Rüniger, G.: Energy-aware Execution of Fork-Join-based Task Parallelism. In: Proc. of the 20th Int. Symp. on Modeling, Analysis and Simulation of Computer and Telecommunication Systems (MASCOTS'12). IEEE (2012)
- [28] Rauber, T., Rüniger, G.: Towards an Energy Model for Modular Parallel Scientific Applications. In: IEEE International Conference on Green Computing and Communications (GreenCom 2012), pp. 523–532. IEEE (2012). DOI 10.1109/GreenCom.2012.79

- [29] Rotem, E., Naveh, A., Ananthakrishnan, A., Rajwan, D., Weissmann, E.: Power-Management Architecture of the Intel Microarchitecture Code-Named Sandy Bridge. *IEEE Micro* **32**(2), 20–27 (2012). DOI <http://doi.ieeecomputersociety.org/10.1109/MM.2012.12>
- [30] Rountree, B., Lowenthal, D.K., Schulz, M., de Supinski, B.R.: Practical performance prediction under Dynamic Voltage Frequency Scaling. In: *Proc. of the 2011 Int. Green Computing Conference and Workshops, IGCC '11*, pp. 1–8. IEEE Computer Society, Washington, DC, USA (2011). DOI 10.1109/IGCC.2011.6008553. URL <http://dx.doi.org/10.1109/IGCC.2011.6008553>
- [31] Saxe, E.: Power-efficient software. *Commun. ACM* **53**(2), 44–48 (2010). DOI <http://doi.acm.org/10.1145/1646353.1646370>
- [32] Song, S., Su, C.Y., Ge, R., Vishnu, A., Cameron, K.: Iso-energy-efficiency: An approach to power-constrained parallel computation. In: *Proc. of the 25th IEEE Int. Parallel and Distributed Processing Symp. (IPDPS 11)*. IEEE (2011)
- [33] Treibig, J., Hager, G., Wellein, G.: LIKWID: A Lightweight Performance-Oriented Tool Suite for x86 Multicore Environments. In: *39th International Conference on Parallel Processing Workshops, ICPP '10*, pp. 207–216. IEEE Computer Society (2010)
- [34] Usman, S., Khan, S., Khan, S.: A comparative study of voltage/frequency scaling in NoC. In: *2013 IEEE Int. Conf. on Electro/Information Technology (EIT)*, pp. 1–5 (2013). DOI 10.1109/EIT.2013.6632716
- [35] Weaver, V., Johnson, M., Kasichayanula, K., Ralph, J., Luszczek, P., Terpstra, D., Moore, S.: Measuring Energy and Power with PAPI. In: *Proc. of the ICPP Workshop on Power Profiling and Evaluation, Workshop Proc. of ICPP 2012*, pp. 262–268. IEEE Computer Society (2012)
- [36] Yao, F., Demers, A., Shenker, S.: A scheduling model for reduced CPU energy. In: *Proc. of the 36th Annual Symp. on Foundations of Computer Science, FOCS '95*, pp. 374–. IEEE Computer Society, Washington, DC, USA (1995). URL <http://dl.acm.org/citation.cfm?id=795662.796264>
- [37] Zhu, D., Melhem, R., Mossé, D.: Energy efficient redundant configurations for real-time parallel reliable servers. *Real-Time Syst.* **41**(3), 195–221 (2009). DOI <http://dx.doi.org/10.1007/s11241-009-9067-8>
- [38] Zhuo, J., Chakrabarti, C.: Energy-efficient dynamic task scheduling algorithms for DVS systems. *ACM Trans. Embed. Comput. Syst.* **7**(2), 1–25 (2008)