

Article

Analytical Energy Model Parametrized by Workload, Clock Frequency and Number of Active Cores for Share-Memory High-Performance Computing Applications

Vitor Ramos Gomes da Silva ^{1,*}, Carlos Valderrama ¹, Pierre Manneback ¹ and Samuel Xavier-de-Souza ²

¹ Department of Electronics and Microelectronics (SEMi), University of Mons, 7000 Mons, Belgium; carlosalberto.valderramasakuyama@umons.ac.be (C.V.); pierre.manneback@umons.ac.be (P.M.)

² Department of Computer Engineering and Automation, Universidade Federal do Rio Grande do Norte, Natal 59078-970, Brazil; samuel@dca.ufrn.br

* Correspondence: vitor.ramosgomesdasilva@umons.ac.be

Abstract: Energy consumption is crucial in high-performance computing (HPC), especially to enable the next exascale generation. Hence, modern systems implement various hardware and software features for power management. Nonetheless, due to numerous different implementations, we can always push the limits of software to achieve the most efficient use of our hardware. To be energy efficient, the software relies on dynamic voltage and frequency scaling (DVFS), as well as dynamic power management (DPM). Yet, none have privileged information on the hardware architecture and application behavior, which may lead to energy-inefficient software operation. This study proposes analytical modeling for architecture and application behavior that can be used to estimate energy-optimal software configurations and provide knowledgeable hints to improve DVFS and DPM techniques for single-node HPC applications. Additionally, model parameters, such as the level of parallelism and dynamic power, provide insights into how the modeled application consumes energy, which can be helpful for energy-efficient software development and operation. This novel analytical model takes the number of active cores, the operating frequencies, and the input size as inputs to provide energy consumption estimation. We present the modeling of 13 parallel applications employed to determine energy-optimal configurations for several different input sizes. The results show that up to 70% of energy could be saved in the best scenario compared to the default Linux choice and 14% on average. We also compare the proposed model with standard machine-learning modeling concerning training overhead and accuracy. The results show that our approach generates about 10 times less energy overhead for the same level of accuracy.

Keywords: energy model; dynamic frequency and voltage scaling; dynamic power management; high performance computing



Citation: Silva, V.R.G.d.; Valderrama, C.; Manneback, P.; Xavier-de-Souza, S. Analytical Energy Model Parametrized by Workload, Clock Frequency and Number of Active Cores for Share-Memory High-Performance Computing Applications. *Energies* **2022**, *15*, 1213. <https://doi.org/10.3390/en15031213>

Academic Editors: Andrea Bonfiglio and Andrea Mazza

Received: 15 November 2021

Accepted: 31 January 2022

Published: 7 February 2022

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2022 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Data center energy efficiency has become of crucial importance in recent years due to its high economic, environmental, and performance impact. For example, the leading petaflop supercomputers consume a range of 1–18 MW of electrical power, with 1.5 MW on average, which can be easily translated into millions of dollars per year in electricity bills [1]. Data center energy consumption was estimated to be between 1.1% and 1.5% of worldwide electricity usage in 2010 [2,3], generating as much pollution as a nation such as Argentina [4]. In some cases, the power costs exceed the cost of purchasing hardware [5]. Furthermore, the energy costs of powering a typical data center doubles every five years [6]. Therefore, with such a steep increase in power use, electricity bills have become a significant expense for today's data centers [7,8]. For these reasons, data center energy efficiency is now considered a primary concern for data center operators, often ahead of the traditional considerations of availability and security.

There are several approaches for green computing, from electrical materials to circuit design, systems integration, and software. These techniques may differ, but they share the same goal—to substantially reduce overall system energy consumption without a corresponding negative impact on delivered performance. The processor and main memory are the components that usually dominate power consumption, as shown in Figure 1. The processor can consume as much as 50% of the total energy [9–11]. For that reason, modern processors incorporate several features for power management [12–15], such as dynamic power management (DPM) and dynamic voltage and frequency scaling (DVFS). DPM encompasses a set of techniques for obtaining energy-efficient computing by deactivating or reducing the system components' performance when they are idle or partially utilized [16,17]. DVFS allows the frequency and voltage to be adjusted in run-time depending on current needs.




Figure 1. Power breakdown of a typical node of an HPC cluster at full use. The system used in this study (a) was built in 2016 and equipped with two Intel Xeon E5-2698, 128 GB of DDR4 memory and SSD as storage, while (b) the case study in [11] was built in 2012 and equipped with two Xeon E5507, 32GB of DDR3 memory and HDD as storage.

DVFS is motivated by the well-known fact that frequency and power have a near-cubic relationship [1,2]; this implies that running the CPU at a lower frequency causes a linear reduction in performance and a near-cubic reduction in power, which could lead to a near-square reduction in CPU energy. Because of this, it is possible to achieve dramatic energy savings just with frequency control, depending on the system and its architecture. Although very promising, the system software has yet to determine when and what voltage and frequency to use when running applications. Otherwise, not only will performance deteriorate, but, in the worst case, energy consumption would also increase [1]. Indeed, reducing the frequency results in a longer execution time, which increases the energy consumption of other system components, such as memory and disks. There is also an overhead of time and energy associated with a voltage and frequency switch that needs to be considered. Thus, finding the most appropriate voltage and frequency to use in all circumstances is not easy. Therefore, since its introduction in 1994 [1], there has been a tremendous amount of research on DVFS algorithms.

The DPM technique can achieve substantial energy savings on systems where the static power is high, or the system remains inactive for a long time. In that case, the problem is to determine when and which components to turn on/off. With DPM, energy savings of 70% have been reported [16,17].

However, at the same time, while these power-saving techniques reduce system energy, they can compromise performance leading to a complex trade-off that needs to be carefully exploited to produce more energy-efficient algorithms. Indeed, this study investigates whether the construction of an energy consumption model of an application can lead to significant energy savings.

We propose an analytical energy model for a given application in the function of the two control variables present in most HPC systems: CPU operating frequency and number of active cores. The model is composed of three application-dependent parameters and three parameters relating to the architecture of the system. The application parameters incorporate characteristics of the percentage of parallelism and the input size. The system

architecture parameters include power-related and technology-dependent components, such as dynamic, static, and leakage power.

The main contributions of the proposed model are:

- Simple model: faster to fit and compute, good for DVFS and DPM optimization.
- Parameters with logical meaning: helps to understand the contribution of each specific term.
- Analytical analysis: several analyses can be derived from the equation.
- Controllable variables: the equation is in the function of parameters that we can control directly.

We have organized the rest of this paper in the following way. In Sections 2 and 3, we present a general review of existing models showing the differences between each approach and their applications. In Section 4, we propose our model and derive its parameters alongside its constraints. In Section 5, we validate the model with the PARSEC benchmark applications. Further on, in Section 5.8, we present use cases of the model, as well as how we applied it in DVFS algorithms. Finally, we conclude with a discussion in Section 6.

2. Related Work

Merkel et al. [18] developed an energy model for processors based on events. Their model assumes a fixed energy consumption α_i for each activity, and by counting the number of occurrences c_i of every activity, they estimate the total energy as:

$$E = \sum_{i=1}^n \alpha_i c_i. \quad (1)$$

Another event-based model, introduced by Roy et al. [19], described the computational energy consumed by a CPU for an algorithm A as the Equation (2)

$$E(A) = P_{clk}T(A) + P_w W(A), \quad (2)$$

where P_{clk} is a processor clock leakage power, $T(A)$ is the total execution time, $W(A)$ is the total time taken by non-I/O operations, and P_w is used to capture the power consumption per operation performed by the CPU. $T(A)$ and $W(A)$ are estimated using performance features.

Models based on events present some drawbacks, they are highly dependent on the operating system and its architecture, making them problematic to port for other platforms. There are also limitations regarding the number of simultaneous events that can coexist without adding a non-negligible overhead. Additionally, there are cases where events need multiplexing, for example, when using more hardware events than the CPU can provide. There are also some well know problems regarding the precision of some events, as shown in many studies [20–25]. Some events that should be exact and deterministic (such as the number of executed instructions) show run-to-run variations and over-count on various architectures, even when running in strictly controlled environments. Because of that, our proposed model is not dependent on events and, therefore, not vulnerable to those drawbacks.

An instruction-level energy model was also proposed in [26] by Yakun et. al. Where they proposed an energy per instruction (EPI) characterization made on Xeon Phi. Their model is expressed as:

$$E(f) = \frac{(p_1 - p_0)(c_1 - c_0)/f}{N}, \quad (3)$$

where N is the total number of dynamic instructions, p_0 is the initial idle power, p_1 is the average dynamic power, and $(c_1 - c_0)$ refers to the cumulative number of cycles the micro-benchmark performs. This model is suitable for estimating the energy after the application finishes executing when it is possible to count the total cycles. However, it is challenging to use for optimization or forecasting since it does not have an application

model to predict the cycles. Our model integrates the behavior of the application, taking into account the execution time.

Lewis et al. [27] described the overall system energy consumption using the following equation:

$$E = A_0(E_{proc} + E_{mem}) + A_1E_{em} + A_2E_{board} + A_3E_{hdd}, \quad (4)$$

where, A_0 , A_1 , A_2 , and A_3 are unknown constants that are calculated via linear regression analysis and those remain constant for a specific server architecture. This model, as the previous one, relies on knowledge of energy spent on each component, being a suitable option for estimation after the application has already run, but not for optimization of the run itself, which is the aim of our model.

In another energy consumption model based on system utilization, Mills et al. [28] modeled the energy consumed by a compute node with CPU (single) executing at speed σ as Equation (5)

$$E(\sigma, [t_1, t_2]) = \int_{t_1}^{t_2} \sigma^3 + \rho\sigma_{max}^3 dt, \quad (5)$$

where ρ stands for the overhead power consumed regardless of the processor speed, t_1 and t_2 are the application's initial and final execution times. The overhead includes power consumption by all other system components, such as memory, network, and more. For this reason, although the authors mentioned the energy consumption of a socket, their power model is generalized to the entire server. This model lacks a closed-form, i.e., it depends on the definition of $\alpha(t)$ to be complete. Our model has a closed-form which facilitates analyses.

Although much work has been done on DVFS, the focus is still on the consumer electronics and laptop markets. For HPC, the notion of energy perception is relatively new [29]. Moreover, the operational characteristics of non-HPC and HPC systems are significantly different. First, the workload on non-HPC systems is very interactive with the end-user, but the workload on the HPC platform is not. Second, activities conducted on a non-HPC platform tend to share more machine resources. In contrast, in HPC, each job often runs with dedicated resources. Third, an HPC system is usually much larger than a non-HPC system, making it more challenging to gather information, organize, and execute global decisions. Therefore, it is worthwhile to investigate whether a DVFS scheduling algorithm, which works well for conventional computing, remains effective for HPC.

Our paper proposes a full-system energy model based on the CPU frequency and the number of cores. The model aims to understand and optimize the energy behavior of parallel applications in HPC systems according to application parameters, such as the degree of parallelism and CPU parameters related to dynamic and static power. The proposed model differs from existing ones, including the frequency and number of cores in the same equation for estimating the energy for a specific application in a given configuration. This model can serve as a base for considering DVFS and DPM optimization problems, including frequency and active cores. It can also be used to analyze the contribution of each parameter (ex: level of parallelism) to energy consumption. Furthermore, the number of cores is essential in HPC since applications are designed to run on multiple cores.

The proposed energy model is the product of an application-agnostic power model and an architecture-specific application performance model. The power model is based on the CMOS logic gates power draw as a function of the frequency [30,31] augmented to include the number of cores. The performance model is based on Amdahl's law [32–34], which can be used to estimate runtime in multi-core systems. In addition, this model has been extended to include execution frequency and input size, characterizing the application on the target architecture.

Table 1 summarizes the models comparing the system dependencies and the controllable variables.

Table 1. Related work summary.

Model	System Dependency	Variable	Controllable Variables
Merkel et al. [18]	performance counters	number of activities	-
Roy et al. [19]	performance counters	io operations, total time	-
Yakun et al. [26]	number of instructions	frequency	frequency
Lewis et al. [27]	energy of subcomponents	energy of subcomponents	-
Mills et al. [28]	power of subcomponents	total time, frequency	frequency
Our model	-	frequency, cores, input size	frequency, cores, input size

3. Theoretical Background

A model is a formal representation of a natural system. The representation of computer system models includes equations, graphical models, rules, decision trees, representative collections of examples, and neural networks. The choice of representation affects the model's accuracy, as well as its interpretability by people [35–37]. Accurate energy and power consumption models are essential for many energy efficiency schemes employed in computing equipment [5], and they can have multiple uses, including the design, forecasting, and optimization of data center systems. This study focuses on analytical models that could aid energy optimization and analyses of crucial factors in the total energy draw.

The desirable properties of a full-system model of energy consumption include accuracy, speed, generality and portability, inexpensiveness, and simplicity [38]. However, modeling an HPC system's exact energy consumption behavior is not straightforward, either at the whole-system level or at the level of individual components. Data centers' energy consumption patterns depend on multiple factors, such as hardware specifications, workload, cooling requirements, or the type of the applications. Some of these factors cannot be measured easily. Furthermore, it is impractical to perform detailed measurements of the energy consumption of lower-level components without additional overhead.

Several proposed models have already been classified concerning their input parameters, as shown by Dayarathna et al. [2], who analyzed more than 200 models according to their characteristics and limitations and classified them into categories where the model is more suited to its objectives:

- System utilization or workload
- Frequency
- Other system states, such as cache miss, branch prediction, number of instructions executed, and more

Often, energy models are described as a combination of two main parts, the power model of the system and the performance model of the application. This is because the concept of energy (E) is the total amount of work performed by a system over a period of time (T), while power (P) is the rate at which the system performs the work. The relation between these three amounts can be expressed as:

$$E = \int_0^T P(t)dt. \quad (6)$$

3.1. Power Models

The modeling of system parameters is becoming popular nowadays with the advantage of performance counters provided by the CPU or the operating system. These counters can measure micro-architectural events, such as instructions executed, cache hits, miss-predicted branches, and more; thus, providing a base for many different estimations of power usage. This makes this type of model very suitable for power estimation because it can use information about several internal states of the computer.

Frequency-based models are the most common kind of model. They serve as a base for many power models [30,31,39]. These models utilize the fact that every digital circuit (including modern processors) is composed of transistors. Thus, modeling one transistor's interaction and scaling this to the chip can give a reasonable estimate of the entire system's

energy. One of the most common frequency-based model approximations is defined as follows:

$$P = \alpha + \beta f^3, \quad (7)$$

where α and β are model parameters, and f is the operating frequency (details of this equation are covered in Section 4). This type of model is suitable for optimization problems since these are a function of the operating frequency, which can be easily controlled.

3.2. Performance Models

The most common way to model the application performance is using the workload. The workload is an abstract representation of the amount of work done for a given time and speed. The workload (W) can be defined in many different ways. One common way, used in many studies, such as Paolillo et al. [40], Francis et al. [1], and Kim et al. [41], is the following:

$$W = \int_0^\tau s(t)dt = s\tau, \quad (8)$$

where τ is total active time, and s is the execution speed in instructions/second.

Utilization models [1,42] are also found in the literature, defined as the ratio between the time that the system is active and the total time (idle and active). These models are present in many DVFS algorithms present in Linux. They can be viewed as a good alternative to the workload since it is impossible to measure workload in real-time. Equation (9) defines workload in terms of CPU utilization (u):

$$u = \frac{\tau}{T} = \frac{W/s}{T}, \quad (9)$$

where T is the total execution time (idle and active), and τ is the active time, meaning when the processor was executing instructions. Models based on CPU utilization are the basis for DVFS algorithms. Even though this is not a controllable parameter, it is straightforward to measure system utilization with almost no overhead, and it is also very portable in terms of operating systems and architectures.

4. Modeling Energy with Performance and Power

This section describes the models proposed for power, performance, and energy.

4.1. Power Model

The developed power model is based on the developed frequency models [43–46]. In this approach, the idea is to reduce the complexity of the processor dynamics by looking only at the main element that it is composed of, the transistor. Thus, modeling the power consumption can be resumed to model the logic gates and multiply this by the total number of gates, reducing the complexity of the modeling process.

FINFET and MOSFET comprise the main techniques to manufacture transistors. However, FINFET is the more recent, and has gradually replaced the mature technology MOSFET. Despite having different characteristics, they have aspects in common that can be modeled [43–46]. These are static power P_{static} , dynamic power $P_{dynamic}$, and leakage power P_{leak} , which, in combination, comprise and approximate the total power draw.

The dynamic power and leakage power behavior can be approximated by the following equations, respectively, as shown by Sarwar et al. [30] and Butzen et al. [31].

$$P_{dynamic} = CV^2f, \quad (10)$$

$$P_{leak} \propto V, \quad (11)$$

where C is the load capacitance, V is the voltage applied to the circuit, and f is the switching frequency.

Another common approximation is to assume a linear relationship between the voltage and the applied frequency [39], such that:

$$f \propto V, \quad (12)$$

These approximations have been demonstrated to be very precise. In the work of Silva et al., the mean percentage error was calculated to be 0.75% [47].

Thus, the proposed model for one processing core of a multi-core processor is derived by using Equations (10)–(12) to write Equation (13).

$$P(f) = c_1 f^3 + c_2 f + c_3, \quad (13)$$

where c_1 , c_2 , and c_3 are the model's parameters associated with the dynamic, leakage and static power aspects, respectively. Including the number of active cores p , the proposed estimation of the power consumption of the whole processor becomes Equation (14)

$$P(f, p) = p(c_1 f^3 + c_2 f) + c_3, \quad (14)$$

4.2. Performance Model

We consider a program as a set of instructions executed on a mean frequency f with c_k instructions per cycle to model the application execution time. The time T_f that this program will take to complete at a given frequency is devised as follows:

$$T_f = \frac{I}{c_k f}, \quad (15)$$

where I is the total number of instructions and c_k the ratio of instructions per unit of time.

The next step is to include the number of cores in the equation. Amdahl's law [32], gives the theoretical background for that. It describes the speedup in latency of the execution of a task at a fixed workload.

$$S = \frac{T_s}{T_p} = \frac{1}{1 - w + \frac{w}{p}}, \quad (16)$$

where T_s is the serial time, T_p the parallel time, S is the theoretical speedup of the execution of the whole task, w is the proportion of the execution time that benefits from improving system resources, and p is the speedup part of the task that benefits from improved system resources. Combining this with Equation (15), the parallel time at frequency f can be written as:

$$T_p = \frac{T_s}{S} = \frac{T_f}{\frac{1}{1 - w + \frac{w}{p}}}, \quad (17)$$

We can then write the equation of the program execution time as a function of frequency, the number of cores, and parallelism as Equation (18) and subsequently derive Equation (19):

$$T(f, p) = \frac{I}{\frac{c_k f}{1 - w + \frac{w}{p}}}, \quad (18)$$

$$T(f, p) = \frac{d_1(p - wp + w)}{fp}, \quad (19)$$

where d_1 is a constant.

Finally, to fully characterize the application, a parameter representing the application's workload, called input size N , is introduced, representing the number of basic operations needed to complete a problem [48]. In Oliveira et al. [49], they showed that this parameter could generally be described as exponential. Therefore the proposed performance model is

presented in Equation (20). This resulting equation describes the behavior of the execution time of a program for an input N , frequency f , and active cores p :

$$T(f, p, N) = \frac{d_1 N^{d_2} (p - wp + w)}{fp}, \quad (20)$$

where d_1 , d_2 and w are constants that depend on the application.

4.3. Energy Model

Combining the power model output described in Section 4.1 and the characterization of the application performance described in Section 4.2, the total energy can be modeled as:

$$E(f, p, N) = P(f, p) \times T(f, p, N), \quad (21)$$

where $P(f, p)$ is the total power modeled by Equation (14), $T(f, p, N)$ is the execution time estimated by the Equation (20), f is the frequency, p is the number of active cores, and N is the input size. The final equation can be written as:

$$E(f, p, N) = \frac{d_1 N^{d_2} (p - wp + w) (p(c_1 f^3 + c_2 f) + c_3)}{fp}. \quad (22)$$

5. Experimental Validation

In this section, the models presented in Sections 4.1 and 4.2 were validated with a benchmark specific for multi-core architectures. Additionally, in order to assess the modeling overhead and accuracy, our proposal was then compared to machine learning approaches. We compared against support vector regression (SVR) [50], decision tree [51], k-nearest neighbors [52], multilayer perceptron [53], and some new methods, such as Gao et al. [54]. However, SVR was chosen as the most representative because it performed best in our tests without aggressive fine-tuning, as shown in Figure 2.




Figure 2. Average of the mean squared error for all applications of our study case Section 5.2.

5.1. Case Study Architecture

The experiments were executed in one computer node equipped with two Intel Xeon E5-2698 v3 processors with sixteen cores each and two hardware threads for each core. The

maximum non-turbo frequency was 2.3 GHz, and the total physical memory of the node was 128 GB (8×16 GB). Turbo frequency and hardware multi-threading were disabled during all experiments. The operating system used was Linux CentOS 6.5, kernel 4.16.

The Linux kernel has many different policies for power management, depending on the driver. In the default driver, the `acpi-cpufreq`, the options are Powersave, Performance, Ondemand, Conservative, and Userspace. Each governor has a policy on how the frequency is selected. In this investigation, the frequency control was performed using the Userspace governor, which allows the user or any userspace program to set the CPU to a specific frequency. The core control was accomplished by modifying the appropriate system files with the default CPU-hotplug driver.

The architecture was equipped with the intelligent platform management interface (IPMI), a set of interfaces allowing out-of-band management of computer systems and platform-status monitoring via the local network [55]. It can monitor variables and resources, such as the system's temperature, voltage, fans, and power supplies, with independent sensors attached to the hardware.

5.2. Case Study Applications

The applications `blackscholes`, `bodytrack`, `canneal`, `dedup`, `fluidanimate`, `freqmine`, `raytrace`, `swaptions`, `vips` and `x264` from the PARSEC <https://parsec.cs.princeton.edu/download.htm> (accessed on 20 February 2020). parallel benchmark suite, version 3.0 [56], OpenMC [57] and LINPACK (HPL) [58], were chosen as case studies. The PARSEC benchmark focused on emerging workloads and was designed to represent the next-generation shared-memory programs for chip-multiprocessors. It covers an ample range of areas, such as financial analysis, computer vision, engineering, enterprise storage, animation, similarity search, data mining, machine learning, and media processing. The OpenMC and the LINPACK are two classic HPC programs.

5.3. Verifying Hypothesis

In this section, we validate whether the assumptions of our model are valid for the system used.

5.3.1. Frequency and Voltage Relation

One of the assumptions was that the frequency and the voltage have a linear relationship, as indicated by Equation (12). To verify that, we build an experiment that sets the frequency to a specific value while sampling the voltage using the APERF and MPERF registers that provide feedback on the current CPU frequency. The average result of the sampling voltages is shown in Figure 3, where we can observe a near-perfect linear relation. This is because manufacturers implement this curve in the processors, using tables that relate ranges of frequencies to voltages so that they can precisely define any curve that will better suit their design.




Figure 3. Frequency voltage relation.

5.3.2. Input Size and Instructions

We ran the applications with different inputs assuming linear growth in the amount of work for one input to the other when building our model. However, measuring and controlling the amount of work would require much instrumentation and tuning to find an input corresponding to a certain amount of work. Therefore, to build our models, we use the time to reference the amount of work, assuming that the work is proportional to the executing time. Figure 4 corresponds to the verification of this supposition.

Table 2 shows that the assumption was reasonable since the average correlation was 0.96 for all applications, indicating that growth in the number of instructions will follow the time. This was the case for all applications that we ran in our benchmark and should hold for any data parallelism type of application.

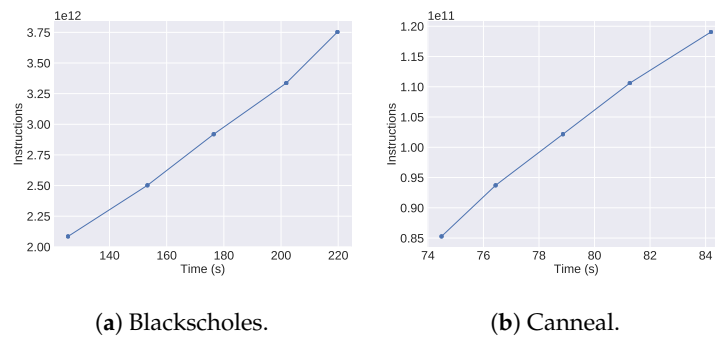


Figure 4. Relation between time and instructions for each input size.

Table 2. Correlation of time and instructions for all applications.

Application	Blackscholes	Bodytrack	Canneal	Dedup	Ferret	Fluidanimate	Freqmine	Openmc	Raytrace	Swaptions	Vips	x264	HPL
Correlation	0.99	0.99	0.99	0.99	0.96	0.99	0.99	0.94	0.99	0.99	0.98	0.99	0.79

The next assumption was that the application’s behavior was the same when varying the workload. This condition is necessary for using the model with an unknown input size because, if the behavior is the same, we can interpolate the known inputs. One way to verify this is to measure the rate of instructions per second normalized by the frequency, as shown in Figure 5.

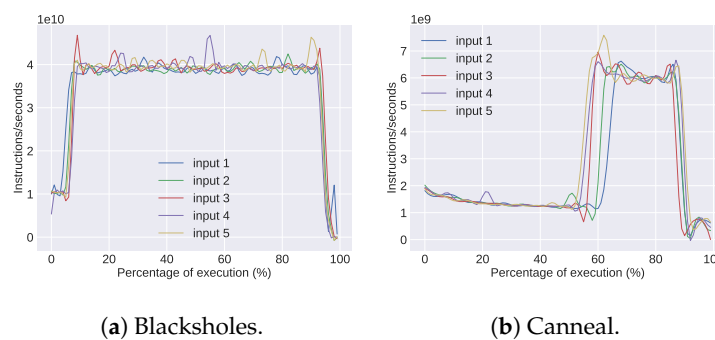


Figure 5. Rate of instructions per second varying the input size normalized by the frequency.

Figure 5 shows that the applications have roughly the same curve when normalized; this also happens for all other applications in our benchmark.

The final assumption is that the workload should also not vary depending on the number of cores or frequency. To verify, we measure the total number of executed instructions while varying the cores from 1 to 32. Table 3 shows the results.

Table 3. Variation of the number of instructions when changing the number of cores for the same input.

Application	Average Number of Instructions	Standard Deviation	Standard Deviation (%)
Vip	7.97×10^{11}	7.16×10^6	0.00
Openmc	8.17×10^7	1.65×10^4	0.02
Rtview	9.91×10^{12}	1.55×10^9	0.02
X264	4.52×10^{11}	5.81×10^7	0.01
Bodytrack	1.86×10^{12}	3.95×10^{10}	2.13
Fluidanimate	2.09×10^{12}	8.44×10^{10}	4.04
HPL	1.14×10^8	1.24×10^5	0.11
Blackschole	3.75×10^{12}	1.40×10^9	0.04
Dedup	1.02×10^{11}	5.74×10^7	0.06
Swapti	2.43×10^{12}	8.87×10^8	0.04
Canneal	1.19×10^{11}	4.46×10^7	0.04
Freqmine	1.27×10^{12}	4.78×10^8	0.04
Ferret	4.76×10^{11}	7.04×10^7	0.01

Table 3 shows the standard deviation and what that corresponds to in terms of the total number of instructions as a percentage.

The same test was performed for the frequency, varying from 1.2 to 2.2 GHz with 100 MHz steps. The results are shown in Table 4.

These results show that all the assumptions were reasonable, and we can safely move to the validation of the model's prediction.

Table 4. Variation of the number of instructions when changing the frequency for the same input.

Application	Average Number of Instructions	Standard Deviation	Standard Deviation (%)
Vip	7.97×10^{11}	1.16×10^6	0.00
Openmc	8.17×10^7	4.52×10^3	0.01
Rtview	9.91×10^{12}	6.64×10^5	0.00
X264	4.52×10^{11}	1.54×10^5	0.00
Bodytrack	1.84×10^{12}	2.54×10^5	0.00
Fluidanimate	2.38×10^{12}	1.70×10^9	0.07
HPL	1.14×10^8	5.95×10^3	0.01
Blackschole	3.75×10^{12}	4.36×10^5	0.00
Dedup	1.02×10^{11}	8.32×10^7	0.08
Swapti	2.43×10^{12}	1.48×10^5	0.00
Canneal	1.19×10^{11}	3.01×10^5	0.00
Freqmine	1.27×10^{12}	3.70×10^8	0.03
Ferret	4.76×10^{11}	5.63×10^7	0.01

5.4. Fitting the Models

To find the parameters of Equation (22), 10 uniformly random configurations of frequencies (f), cores (p) and inputs (N) were chosen from the range $1 \leq p \leq 32$, $1.2 \leq f \leq 2.2$ and $1 \leq N \leq 5$, respectively. The application was executed for each

chosen configuration, and the measured energy and time values were collected. For the input size, if we assume that all CPU instructions take approximately the same time to execute, the number of basic operations will be directly correlated with the time. Thus, we can estimate the input size by looking at the execution time, allowing us to divide a large input size into several smaller ones, knowing their relationship, as performed in the work of Oliveira [49]. The unity can also vary depending on the definition. For simplicity, we assign numbers from 1 to 10, increasing the problem linearly, so it is also possible to interpolate any input in between these values.

For each configuration, samples of the power were collected using IPMI every 1 second. This sampling rate was chosen based on the magnitude of the mean run time of the applications, which is in the order of minutes. Therefore, this rate provides enough samples to measure average power. Additionally, timestamps and the total run time were collected. The total energy spent on each configuration is estimated by first interpolating the power samples using the first-order method and then integrating this function in the time.

The model's parameters are calculated by solving an optimization problem of finding the values that minimize the squared error of the prediction to the measured values using the non-linear least-squares method.

The Python library Scikit-Learn was used to build the SVR model [59]. The SVR was trained using the same data used for parameter estimation of Equation (22) with a grid search used to find the best kernel function and the best values for the hyper-parameters penalty for the wrong (C) and (γ). For this data, the best function was the radial base function (RBF), and the hyper-parameters were $C = 10^4$ and $\gamma = 0.5$.

5.5. Measured versus Modeled Energy

To validate the model, we ran all possible configurations in the tested machine, varying the cores in a range of $1 \leq p \leq 32$, the frequency in $1.2 \leq f \leq 2.2$, and the input in $1 \leq N \leq 5$. The total number of configurations varies from 400 to over 1000 depending on the application, as some applications have restrictions on the number of cores that they can run. Once the data was collected, we computed the mean percentage error (MPE) according to the following equation:

$$MPE = \frac{1}{N} \sum_i \frac{|y_{\text{estimated}} - y_{\text{measured}}|}{y_{\text{measured}}}. \quad (23)$$

5.5.1. Frequency \times Cores

Figure 6 plots the measured and modeled energy consumption for some of the applications modeled. In addition, some of the possible shapes that the model can take while varying the number of active cores, and operating frequency, are shown.

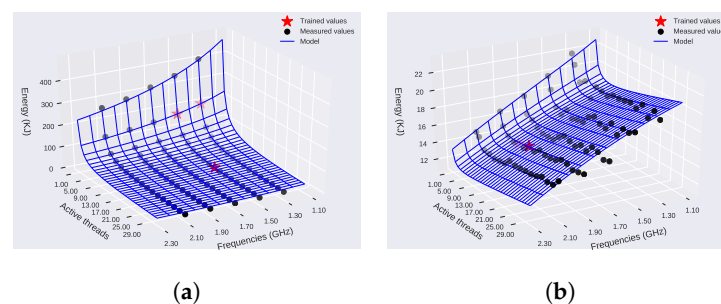


Figure 6. Example fit for a specific input size: Blackscholes (a) and Canneal (b). “measured values” are the sensor data, and “minimum energy” is the minimum energy model prediction.

5.5.2. Frequency × Input

Figure 7 plots the measured and modeled energy consumption for some of the applications modeled. The diagrams show some of the possible shapes that the model can take while varying the operating frequency, and input size.

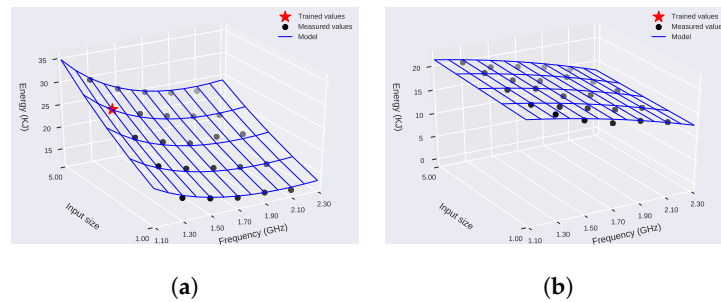


Figure 7. Example fit for a specific input size: Blackscholes (a) and Canneal (b). “measured values” are the sensor data and “minimum energy” is the minimum energy model prediction.

5.5.3. Cores × Input

Figure 8 plots the measured and modeled energy consumption for some of the applications modeled. The diagrams show some of the possible shapes that the model can take while varying the number of active cores, and input size.

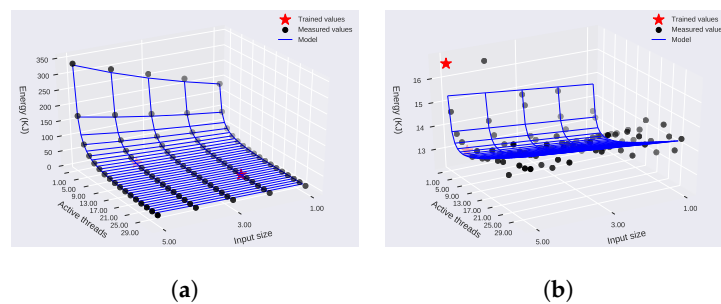


Figure 8. Example fit for a specific input size: Blackscholes (a) and Canneal (b). “measured values” are the sensor data and “minimum energy” is the minimum energy model prediction.

5.5.4. Validation

The average results for each application were calculated using a model trained with only 10 configurations, and the comparison is displayed Figure 9.

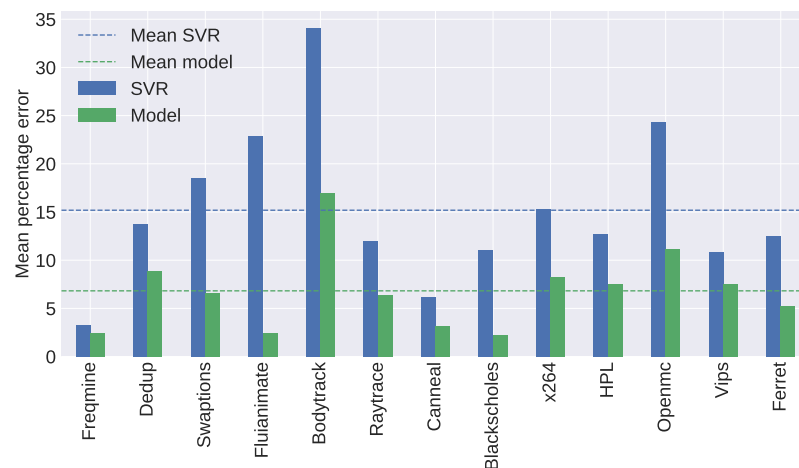


Figure 9. Comparison of the mean percentage error between the proposed model and SVR. “Model mean” and “SVR mean” are the average of all MPE values for all applications.

Figure 9 shows that the proposed model always performed better, with a lower MPE than SVR, when we were limited to 10 training points. This result is further explored in the next Section 5.6, where we undertake a comparison with different training sizes.

5.6. Overheads on Training

It is known that machine learning is data-driven; in that sense, the SVR model obtained using only 10 configurations could be improved, but what about the analytical model? To answer that question, the proposed model and the SVR were also trained with a varying number of configurations. We then compared the MPE and the amount of energy spent to create each model. This accuracy-energy trade-off is crucial since building models' energy overhead defeats the primary goal of saving power when running applications.

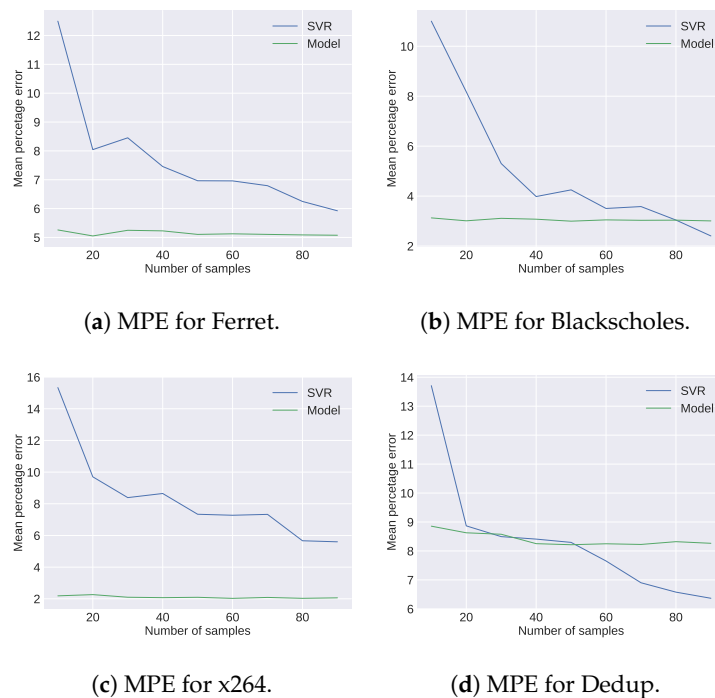


Figure 10. MPE of the case studies versus training size, comparing how many training points is necessary to reach an acceptable result.

Figure 10 shows the comparisons of MPE and energy spent to create each model for two selected applications. According to the results, the analytical model is very stable, not changing much as more data is added, while the SVR keeps reshaping to adapt to the data. The error of the analytical model is almost constant but that of the SVR, initially very high, drops as more data is used in the training process.

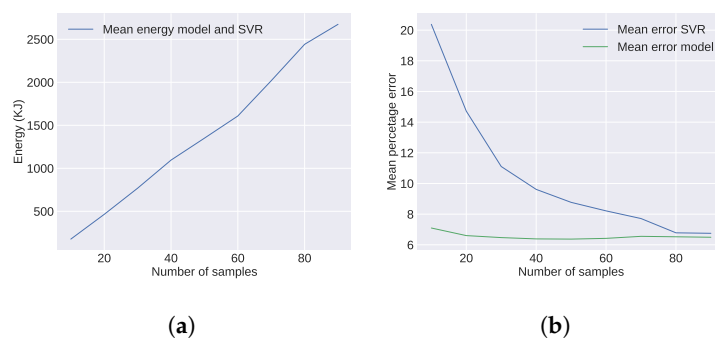


Figure 11. Overall results for energy and MPE for each training size. (a) Average energy spent on all applications during model creation. The two curves are identical because the same data were used to adjust the SVR and the model. (b) MPE of all applications: SVR needs 10 times more data to have an MPE lower than the proposed model.

Figure 11 presents the overall results, with the mean energy overhead and MPE for all applications. The meeting point of the MPE for the SVR and the proposed model can be extracted from Figure 11b. It shows that, in around 90 configurations, the SVR starts to have a smaller error. The cost of that is the linear increase in energy spent on training. The increase in energy, about 10 times more, can be observed in Figure 11a.

5.7. Analysis

One of the most significant advantages of using an analytical model is the understanding of the problem that an equation provides, making many different kinds of analysis possible that are otherwise impossible with a machine learning model. In this section, we discuss one of the possible analyses. In the following figures, we try to understand the contribution of each parameter of the equation to the total energy consumption.

For this analysis, we took the model of one of the applications and, varying one parameter of the equation, we display the energy versus performance (time) for all configurations. After that, we computed the Pareto frontier, a set of all Pareto efficient allocations, i.e., all the configurations where resources cannot be reallocated to make one individual better off without making at least one individual worse off. This gives us all the configurations where we have an optimal trade-off of performance and energy to choose from.

Figure 12 shows the Pareto frontier for several values for the static power parameter (c_3 in Equation (22)) with configurations of frequency ranging from 1.2 to 5 GHz and cores from 1 to 64, so that we can also have an idea of what is the tendency when we increase the frequency and number of cores.




Figure 12. Pareto frontier for several values of static power parameter. The arrows with blue heads indicate the maximum energy, while the arrows with a red head the minimal energy for each corresponding curve.

From this figure, we can see that when increasing the value of the static power parameter, the total energy consumption increases as expected. We can also observe that

the values that minimize the total energy consumption tend to be high frequency and multiple cores. This is one of the consequences of increasing the static power factor. As the dynamic factor proportionally decreases, its variables tend to have less impact on total consumption, enabling configurations with high frequency and several cores. This also enables chip-level optimization for choosing components that change the ratio between static and dynamic power.

Figure 13 shows the Pareto frontier in the same ranges described before but for the parameter corresponding to the level of parallelism of the application (w in Equation (22)).




Figure 13. Pareto frontier for several values of static power parameter. The arrows with blue heads indicate the maximum energy, while the arrows with a red head, the minimal, for each corresponding curve.

In Figure 13, we observe that, as the parallelism level increases the total energy decreases. The number of cores tends to be higher with a higher level of parallelism as expected, and the frequency shows an inverse relation.

5.8. DVFS and DPM Optimization

The effectiveness of the proposed approach during optimization was evaluated with a simple algorithm that finds the optimal frequency and number of active cores from the proposed equation. The results were then compared to the Linux default choices for power management.

With Equation (22), it is possible to calculate energy consumption estimates for each possible configuration since there is a finite range of possible values for the frequency and number of cores. It is also possible to apply constraints on the execution time, frequency, and the number of active cores. Then, the configuration that minimizes energy consumption for a given input can be selected. The complete workflow is shown in Figure 14. We can see that any optimization problem can be structured with our model and the system's constraints. In the following examples, the optimization problem that we build is to minimize the

energy equation given the constraints of possible frequencies and the number of cores that our system can run. The algorithm selected to minimize was the newton-CG [60].

Current HPC managers leave to the user the choice of how many cores to use. On this basis, three situations were analyzed in relation to the number of cores:

1. Worst choice: number of cores that maximize the total energy consumed;
2. Random choice: energy consumed for a random choice of the number of cores;
3. Best choice: number of cores that minimize the total energy consumed (oracle).

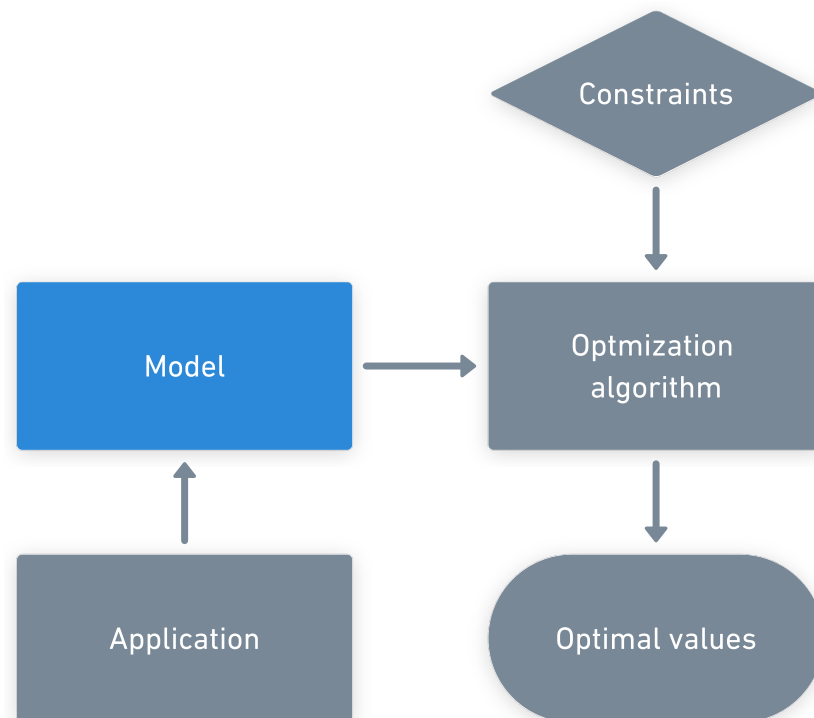


Figure 14. Optimization workflow showing how DVFS and DPM optimization could be implemented from our model.

The default option for the Linux governor is Ondemand, and, by default, it has no DPM control for the number of active cores. As Ondemand only performs DVFS, for comparison, each application was executed with all available cores in the system, from 1 to 32.

Figures 15–17, show the energy savings with respect to Ondemand, i.e., $\frac{\text{Ondemand} - \text{Model}_{\min}}{\text{Ondemand}}$ for the three cases described above. The savings and losses for each case are:

1. Worst choice: save 69.88% on average;
2. Random choice: save 12.04% on average;
3. Best choice: lost 14.06% on average.

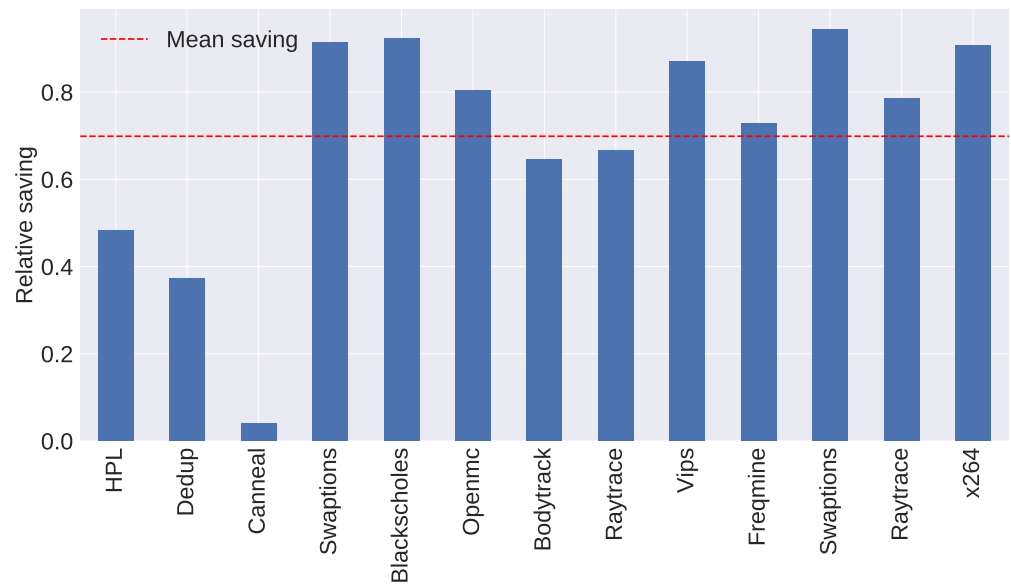


Figure 15. Energy savings comparisons between the proposed model and the Worst case.

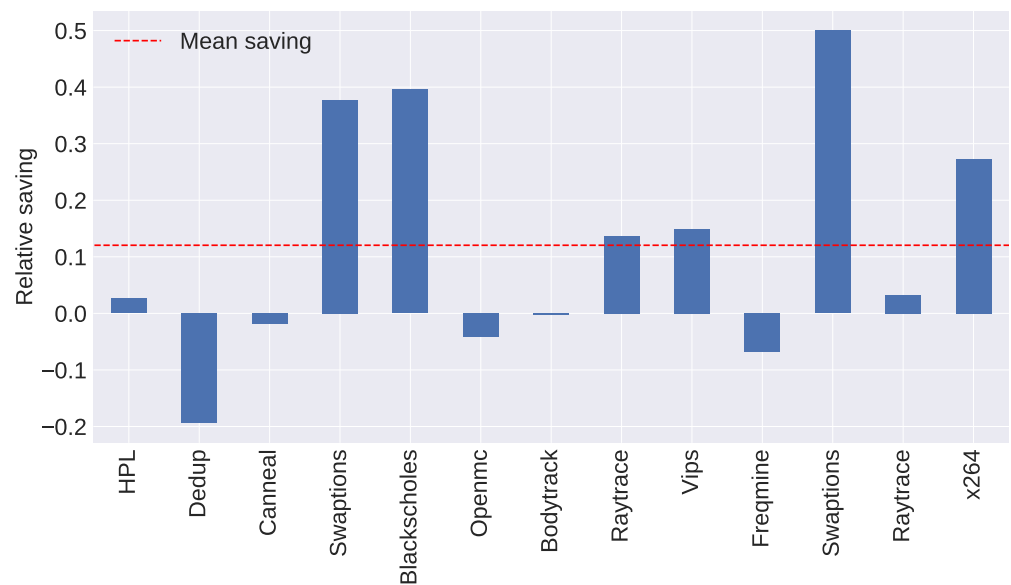


Figure 16. Energy savings comparisons between the proposed model and the Random case.

By default, operating systems do not implement DPM at the core level, and, in HPC, the user usually explicitly chooses the number of cores to run their job. To give a better idea of the impact on the energy consumption of DPM at the core level, we analyzed the choices of the number of cores over a period of one year in the HPC center at UFRN. The result is plotted in Figure 18.

It is of note that the most common choice of many regular users is a single core requested per job, matching the worst-case choice for all applications analyzed in this investigation. The best choice was quite often 32 cores, which is the third most popular choice among users, but it is 72 times less frequent than 1 core. This led us to envision how much energy could be saved and encouraged us towards future research using the proposed model for DPM or more advanced optimization algorithms.

In practice, this approach can be implemented by allowing the resource manager to perform these changes for the user using pre-scripts and post-scripts for high energy consumption job submissions.

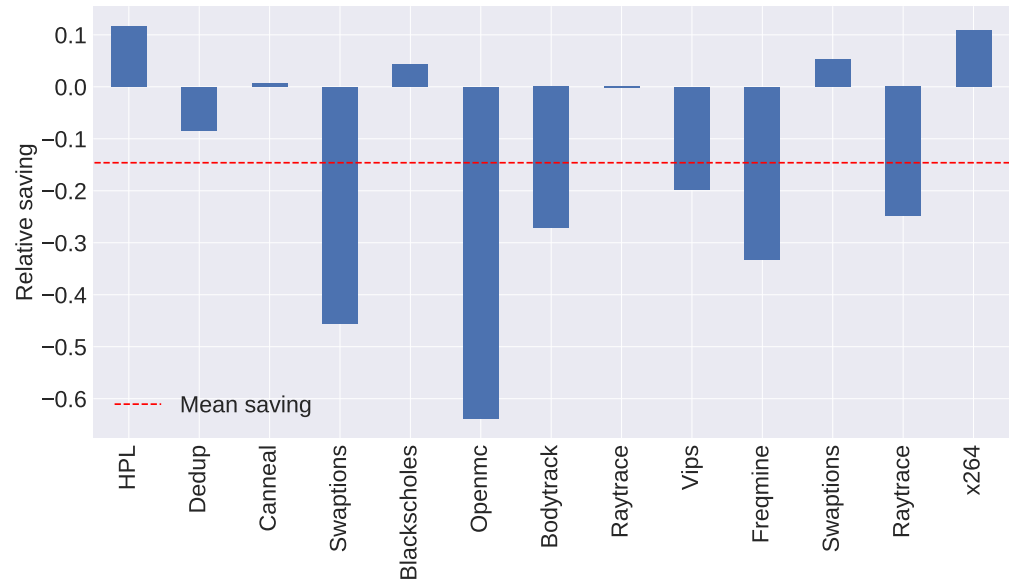


Figure 17. Energy savings comparisons between the proposed model and the Best case.




Figure 18. Number of CPU requests during one year in HPC cluster, sorted by the number of cores requested per job.

6. Conclusions

This paper proposes an energy model based on the operating frequency and the number of cores for a shared memory system. This model serves as a reference for DVFS and DPM optimization problems.

Results from three different HPC benchmarks demonstrate the potential of the proposed model while consuming 10 times less energy than a machine learning approach, such as SVR, to characterize applications. Moreover, it can provide knowledge-based hints

to improve DVFS and DPM algorithms by enabling analysis of the contribution of each model parameter (e.g., level of parallelism) to the energy consumption. Indeed, as shown in Section 5.8, when no oracle is available to choose the frequency and the number of cores the application should use, the proposed model can save around 12% of energy for a random choice and up to 70% for the worse possible choice. Considering the job history of our own HPC center, which shows the prevalence of worse possible choices made by users, the potential energy savings are very significant and encourage further research.

Although the model is promising, it still has some limitations. The main one is related to the input size, which needs to be estimated to create the application model and optimize the application. Another limitation concerns the power model, which does not consider the load variation, so our model ends up using an average of the energy consumption, which is enough to obtain good results but limits its implementation in real-time optimization. Future research is intended to solve both problems, first adapting the model to use the ratio of executed instructions as input size, something which is more tangible and easy to measure in modern systems without much overhead, and adding new parameters to the power model to account for the load. This would allow us to develop more advanced DVFS models that could identify different phases of a target program with more subtle changes in frequency, and, perhaps, in the number of active cores to further improve the results presented here.

Author Contributions: Conceptualization, V.R.G.d.S., S.X.-d.-S.; methodology, V.R.G.d.S., S.X.-d.-S. and C.V.; software, V.R.G.d.S.; validation, V.R.G.d.S., S.X.-d.-S. and C.V.; formal analysis, V.R.G.d.S.; investigation, V.R.G.d.S.; resources, V.R.G.d.S., S.X.-d.-S., C.V. and P.M.; data curation, V.R.G.d.S.; writing—original draft preparation, V.R.G.d.S.; writing—review and editing, V.R.G.d.S., S.X.-d.-S., C.V.; visualization, V.R.G.d.S.; supervision, C.V., P.M. and S.X.-d.-S.; project administration, C.V., P.M. and S.X.-d.-S.; funding acquisition, C.V., P.M. and S.X.-d.-S. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Informed Consent Statement: Not applicable.

Data Availability Statement: The data presented in this study are openly available on Github at <https://github.com/VitorRamos/analytical-energy-model>.

Acknowledgments: The experiments performed in this investigation used the compute nodes of the High-Performance Computing Center (NPAD/UFRN).

Conflicts of Interest: The authors declare no conflict of interest.

Abbreviations

The following abbreviations are used in this manuscript:

DVFS	Dynamic voltage and frequency scaling
DPM	Dynamic power management
SVR	Support vector regression
RBF	Radial base function
HPC	High performance computing
IPMI	The intelligent platform management interface
RBF	Radial base function
MPE	Mean percentage error

References

1. Ishfaq, A.; Sanjay, R. *Handbook of Energy-Aware and Green Computing*; Chapman & Hall/CRC: London, England, 2012; Volume 1, pp. 702–713.
2. Dayarathna, M.; Wen, Y.; Fan, R. Data Center Energy Consumption Modeling: A Survey. *IEEE Commun. Surv. Tutor.* **2016**, *18*, 732–794. [[CrossRef](#)]
3. Corcoran, P.; Andrae, A. *Emerging Trends in Electricity Consumption for Consumer ICT*; National University of Ireland: Galway, Ireland, 2013; pp. 1–56.

4. Mathew, V.; Sitaraman, R. K.; Shenoy, P. Energy-aware load balancing in content delivery networks. In Proceedings of the 2012 Proceedings IEEE INFOCOM, Orlando, FL, USA, 25–30 March 2012; pp. 954–962.
5. Rivoire, S.; Shah, M.A.; Ranganathan, P.; Kozyrakis, C.; Meza, J. Models and Metrics to Enable Energy-Efficiency Optimizations. *Computer* **2007**, *40*, 39–48. [[CrossRef](#)]
6. Buyya, R.; Vecchiola, C.; Selvi, S.T. *Mastering Cloud Computing*; Morgan Kaufmann Publishers Inc.: San Francisco, CA, USA, 2013; pp. 3–27
7. Poess, M.; Nambiar, R.O. Energy cost, the key challenge of today's data centers. *Proc. VLDB Endow.* **2008**, *1*, 1229–1240. [[CrossRef](#)]
8. Gao, Y.; Guan, H.; Qi, Z.; Wang, B.; Liu, L. Quality of service aware power management for virtualized data centers. *J. Syst. Archit.* **2013**, *59*, 245–259. [[CrossRef](#)]
9. Fan, X.; Weber, W.D.; Barroso, L.A. Power provisioning for a warehouse-sized computer. *ACM SIGARCH Comput. Archit. News* **2007**, *35*, 13–23. [[CrossRef](#)]
10. Barroso, L.A.; Hözl, U. The Case for Energy-Proportional Computing. *Computer* **2007**, *40*, 33–37. [[CrossRef](#)]
11. Malladi, K.T.; Nothaft, F.A.; Periyathambi, K.; Lee, B.C.; Kozyrakis, C.; Horowitz, M. Towards energy-proportional datacenter memory with mobile DRAM. In Proceedings of the 2012 39th Annual International Symposium on Computer Architecture (ISCA), Portland, OR, USA, 9–13 June 2012; pp. 37–48.
12. Rotem, E.; Naveh, A.; Ananthakrishnan, A.; Weissmann, E.; Rajwan, D. Power-Management Architecture of the Intel Microarchitecture Code-Named Sandy Bridge. *IEEE Micro* **2012**, *32*, 20–27. [[CrossRef](#)]
13. Brown, L.; Moore, R.; Li, D.S.; Yu, L.; Keshavamurthy, A.; Pallipadi, V. ACPI in Linux. *Symposium* **2005**, *51*, 1–51.
14. Hackenberg, D.; Schone, R.; Ilsche, T.; Molka, D.; Schuchart, J.; Geyer, R. An Energy Efficiency Feature Survey of the Intel Haswell Processor. In Proceedings of the 2015 IEEE International Parallel and Distributed Processing Symposium Workshop, Hyderabad, India, 25–29 May 2015; pp. 896–904.
15. Intel. *12th Generation Intel® Core™ Processors*; Intel: Santa Clara, CA, USA, 2020; pp. 420–430
16. Shuja, J.; Madani, S.A.; Bilal, K.; Hayat, K.; Khan, S.U.; Sarwar, S. Energy-efficient data centers. *Computing* **2012**, *94*, 973–994. [[CrossRef](#)]
17. Benini, L.; Bogliolo, A.; De Micheli, G. A survey of design techniques for system-level dynamic power management. *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.* **2000**, *8*, 299–316. [[CrossRef](#)]
18. Merkel, A.; Bellosa, F. Balancing power consumption in multiprocessor systems. *ACM SIGOPS/EuroSys Eur. Conf. Comput. Syst.* **2006**, *40*, 403–4014.
19. Roy, S.; Rudra, A.; Verma, A. An energy complexity model for algorithms. In Proceedings of the 4th conference on Innovations in Theoretical Computer Science, New York, NY, USA, 9–12 January 2013.
20. Weaver, V.M.; McKee, S.A. Can hardware performance counters be trusted? In Proceedings of the 2008 IEEE International Symposium on Workload Characterization, Seattle, WA, USA, 14–16 September 2008; pp. 141–150.
21. Weaver, V.M.; Terpstra, D.; Moore, S. Non-determinism and overcount on modern hardware performance counter implementations. In Proceedings of the 2013 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS), Austin, TX, USA, 21–23 April 2013; pp. 215–224.
22. Das, S.; Werner, J.; Antonakakis, M.; Polychronakis, M.; Monroe, F. SoK: The Challenges, Pitfalls, and Perils of Using Hardware Performance Counters for Security. In Proceedings of the 2019 IEEE Symposium on Security and Privacy (SP), San Francisco, CA, USA, 19–23 May 2019; pp. 20–38.
23. Mc Guire, N.; Okech, P.; Schiesser, G. Analysis of Inherent Randomness of the Linux Kernel. In Proceedings of the Eleventh RealTime Linux Workshop, Dresden, Germany, 28–30 September 2009.
24. Ramos, V.; Valderrama, C.; Xavier de Souza, S.; Manneback, P. An Accurate Tool for Modeling, Fingerprinting, Comparison, and Clustering of Parallel Applications Based on Performance Counters. In Proceedings of the IEEE International Parallel and Distributed Processing, Rio de Janeiro, Brazil, 20–24 May 2019; pp. 797–804.
25. Silva-de Souza, W.; Iranfar, A.; Bráulio, A.; Zapater, M.; Xavier-de Souza, S.; Olcoz, K.; Atienza, D. Containergy—A Container-Based Energy and Performance Profiling Tool for Next Generation Workloads. *Energies* **2020**, *13*, 2162. [[CrossRef](#)]
26. Shao, Y.S.; Brooks, D. Energy characterization and instruction-level energy model of Intel's Xeon Phi processor. In Proceedings of the International Symposium on Low Power Electronics and Design (ISLPED), Beijing, China, 4–6 September 2013; pp. 389–394.
27. Lewis, A.; Ghosh, S.; Tzeng, N.F. Run-time energy consumption estimation based on workload in server systems. In Proceedings of the 2008 Conference on Power Aware Computing and Systems, San Diego, CA, USA, 8–10 December 2008; pp. 3–4.
28. Mills, B.; Znati, T.; Melhem, R.; Ferreira, K.B.; Grant, R.E. Energy Consumption of Resilience Mechanisms in Large Scale Systems. In Proceedings of the 2014 22nd Euromicro International Conference on Parallel, Distributed, and Network-Based Processing, Turin, Italy, 12–14 February 2014; pp. 528–535.
29. Feng, W.c. Making a Case for Efficient Supercomputing. *Queue* **2003**, *1*, 54–64. [[CrossRef](#)]
30. Sarwar, A. Cmos power consumption and cpd calculation. In *Proceeding: Design Considerations for Logic Products*; Texas Instruments: Dallas, TX, USA, 1997.
31. Butzen, P.; Ribas, R. *Leakage Current in Sub-Micrometer CMOS Gates*; Universidade Federal do Rio Grande do Sul: Porto Alegre, Brazil, 2007; pp. 1–30.
32. Amdahl, G.M. Validity of the single processor approach to achieving large scale computing capabilities. In Proceedings of the Spring Joint Computer Conference on—AFIPS '67 (Spring), New York, NY, USA, 18–20 April 1967.

33. Eyerman, S.; Eeckhout, L. Modeling critical sections in Amdahl's law and its implications for multicore design. In Proceedings of the 37th Annual International Symposium on Computer Architecture—ISCA '10, New York, NY, USA, 19–23 June 2010.
34. Gustafson, J.L. Reevaluating Amdahl's law. *Commun. ACM* **1988**, *31*, 532–533. [[CrossRef](#)]
35. Seel, N.M. *Encyclopedia of the Sciences of Learning*; Springer: Berlin/Heidelberg, Germany, 1988; pp. 223–242.
36. Roy, P.; Mahapatra, G.S.; Dey, K.N. Forecasting of software reliability using neighborhood fuzzy particle swarm optimization based novel neural network. *IEEE/CAA J. Autom. Sin.* **2019**, *6*, 1365–1383. [[CrossRef](#)]
37. Zhu, W.; Liu, X.; Xu, M.; Wu, H. Predicting the results of RNA molecular specific hybridization using machine learning. *IEEE/CAA J. Autom. Sin.* **2019**, *6*, 1384–1396. [[CrossRef](#)]
38. Rivoire, S.; Ranganathan, P.; Kozyrakis, C. A comparison of high-level full-system power models. In Proceedings of the 2008 Conference on Power Aware Computing and Systems, San Diego, CA, USA, 8–10 December 2008; pp. 1–5.
39. Usman, S.; Khan, S.U.; Khan, S. A comparative study of voltage/frequency scaling in NoC. In Proceedings of the IEEE International Conference on Electro-Information Technology, Rapid City, SD, USA, 9–11 May 2013; pp. 1–5.
40. Paolillo, A. Optimisation of Performance Metrics of Embedded Hard Real-Time Systems using Software/Hardware Parallelism, Ph.D. Thesis, Université libre de Bruxelles, Brussels, Belgium, 2018.
41. Kim, D.H.; Imes, C.; Hoffmann, H. Racing and Pacing to Idle: Theoretical and Empirical Analysis of Energy Optimization Heuristics. In Proceedings of the 2015 IEEE 3rd International Conference on Cyber-Physical Systems, Networks, and Applications, Hong Kong, China, 19–21 August 2015; pp. 78–85.
42. Fu, C.; Chau, V.; Li, M.; Xue, C.J. Race to idle or not: Balancing the memory sleep time with DVS for energy minimization. *J. Comb. Optim.* **2018**, *35*, 860–894. [[CrossRef](#)]
43. Rauber, T.; Rüniger, G.; Schwind, M.; Xu, H.; Melzner, S. Energy measurement, modeling, and prediction for processors with frequency scaling. *J. Supercomput.* **2014**, *70*, 1451–1476. [[CrossRef](#)]
44. Goel, B.; McKee, S.A. A Methodology for Modeling Dynamic and Static Power Consumption for Multicore Processors. In Proceedings of the IEEE International Parallel and Distributed Processing Symposium, Chicago, IL, USA, 23–27 May 2016; pp. 273–282.
45. Du, Z.; Ge, R.; Lee, V.W.; Vuduc, R.; Bader, D.A.; He, L. Modeling the Power Variability of Core Speed Scaling on Homogeneous Multicore Systems. *Sci. Program.* **2017**, *2017*, 1–13. [[CrossRef](#)]
46. Gonzalez, R.; Gordon, B.; Horowitz, M. Supply and threshold voltage scaling for low power CMOS. *IEEE J. Solid-State Circuits* **1997**, *32*, 1210–1216. [[CrossRef](#)]
47. Silva, V.R.G.; Furtunato, A.F.A.; Georgiou, K.; Sakuyama, C.A.V.; Eder, K.; Xavier-de Souza, S. Energy-Optimal Configurations for Single-Node HPC Applications. In Proceedings of the 2019 International Conference on High Performance Computing & Simulation (HPCS), Dublin, Ireland, 15–19 July 2019; pp. 448–454.
48. Kumar, V.; Gupta, A. Analyzing Scalability of Parallel Algorithms and Architectures. *J. Parallel Distrib. Comput.* **1994**, *22*, 379–391. [[CrossRef](#)]
49. Oliveira, V.H.F.; Furtunato, A.F.A.; Silveira, L.F.; Georgiou, K.; Eder, K.; Xavier-de Souza, S. Application Speedup Characterization. In Proceedings of the ACM/SPEC International Conference on Performance Engineering, Berlin, Germany, 9–13 April 2018; pp. 43–44.
50. Smola, A.J.; Schölkopf, B. A tutorial on support vector regression. *Stat. Comput.* **2004**, *14*, 199–222. [[CrossRef](#)]
51. Kitts, B. Regression Trees Lecture. *Data Min.* **2006**, 6–7.
52. Altman, N.S. An Introduction to Kernel and Nearest-Neighbor Nonparametric Regression. *Am. Stat.* **1992**, *46*, 175–185.
53. Murtagh, F. Multilayer perceptrons for classification and regression. *Neurocomputing* **1991**, *2*, 183–197. [[CrossRef](#)]
54. Gao, S.; Zhou, M.; Wang, Y.; Cheng, J.; Yachi, H.; Wang, J. Dendritic Neuron Model With Effective Learning Algorithms for Classification, Approximation, and Prediction. *IEEE Trans. Neural Netw. Learn. Syst.* **2019**, *30*, 601–614. [[CrossRef](#)]
55. Schwenkler, T.; Deutschland, S. Intelligent Platform Management Interface. In *Sicheres Netzwerkmanagement*; Springer: Berlin/Heidelberg, Germany, 2006; pp. 169–207.
56. Bienia, C.; Kumar, S.; Singh, J.P.; Li, K. The PARSEC benchmark suite. In Proceedings of the 17th international conference on Parallel architectures and compilation techniques—PACT '08, New York, NY, USA, 25–29 October 2008.
57. Romano, P.K.; Horelik, N.E.; Herman, B.R.; Nelson, A.G.; Forget, B.; Smith, K. OpenMC: A state-of-the-art Monte Carlo code for research and development. *Ann. Nucl. Energy* **2015**, *82*, 90–97. [[CrossRef](#)]
58. Dongarra, J.J. The LINPACK Benchmark: An explanation. In Proceedings of the 1st International Conference on Supercomputing, Athens, Greece, 8–12 June 1987.
59. Pedregosa, F.; Varoquaux, G.; Gramfort, A.; Michel, V.; Thirion, B.; Grisel, O.; Blondel, M.; Prettenhofer, P.; Weiss, R.; Dubourg, V.; et al. Scikit-learn: Machine Learning in {P}ython. *J. Mach. Learn. Res.* **2011**, *12*, 2825–2830.
60. Royer, C.W.; O'Neill, M.; Wright, S.J. A Newton-CG algorithm with complexity guarantees for smooth unconstrained optimization. *Math. Program.* **2020**, *180*, 451–488. [[CrossRef](#)]