

Article

Performance and Energy Trade-Offs for Parallel Applications on Heterogeneous Multi-Processing Systems

Demetrios A. M. Coutinho ^{1,*}, Daniele De Sensi ², Arthur Francisco Lorenzon ³,
Kyriakos Georgiou ⁴, Jose Nunez-Yanez ⁴, Kerstin Eder ⁴ and Samuel Xavier-de-Souza ⁵

¹ Instituto Federal do Rio Grande do Norte, Pau dos Ferros 59900-000, Brazil

² Computer Science Department, Università di Pisa, 56127 Pisa, Italy; desensi@di.unipi.it

³ Computer Science Department, Universidade Federal de Pampa, Alegrete 97546-550, Brazil; aflorenzon@unipampa.edu.br

⁴ Department of Computer Science, University of Bristol, Bristol BS8 1UB, UK; kyriakos.georgiou@bristol.ac.uk (K.G.); j.l.nunez-yanez@bristol.ac.uk (J.N.-Y.); kerstin.eder@bristol.ac.uk (K.E.)

⁵ Department of Computer Engineering and Automation, Universidade Federal do Rio Grande do Norte, Natal 59078-970, Brazil; samuel@dca.ufrn.br

* Correspondence: demetrios.coutinho@ifrn.edu.br

Received: 17 March 2020; Accepted: 21 April 2020; Published: 11 May 2020

Abstract: This work proposes a methodology to find performance and energy trade-offs for parallel applications running on Heterogeneous Multi-Processing systems with a single instruction-set architecture. These offer flexibility in the form of different core types and voltage and frequency pairings, defining a vast design space to explore. Therefore, for a given application, choosing a configuration that optimizes the performance and energy consumption is not straightforward. Our method proposes novel analytical models for performance and power consumption whose parameters can be fitted using only a few strategically sampled offline measurements. These models are then used to estimate an application's performance and energy consumption for the whole configuration space. In turn, these offline predictions define the choice of estimated Pareto-optimal configurations of the model, which are used to inform the selection of the configuration that the application should be executed on. The methodology was validated on an ODROID-XU3 board for eight programs from the PARSEC Benchmark, Phoronix Test Suite and Rodinia applications. The generated Pareto-optimal configuration space represented a 99% reduction of the universe of all available configurations. Energy savings of up to 59.77%, 61.38% and 17.7% were observed when compared to the performance, ondemand and powersave Linux governors, respectively, with higher or similar performance.

Keywords: heterogeneous multi-processing; energy efficiency; power model; Pareto frontier

1. Introduction

Low energy consumption is a key requirement in the design of modern embedded systems, affecting the size, cost, user experience, and the capability to integrate more high-level features. Single Instruction-Set Architecture (Single-ISA) Heterogeneous Multi-Processors (HMP) are known for delivering a significantly higher performance-power ratio than their counterparts. There exist many commercially available designs in the embedded and mobile world [1–3]. Nevertheless, in this type of architecture it is increasingly more complicated to find the number of cores, operating frequency, and voltage that optimize performance and energy consumption to meet the requirements of a given application [4–6]. This complexity increases when application characteristics need to be extracted

at runtime to increase performance or save energy at different phases during the execution of an application [7].

One of the commercial widely used heterogeneous architectures is the ARM big.LITTLE [2]. It consists of two clusters with two types of processing cores, each one containing one or more cores. Cluster type *big* is composed of higher performance larger cores that are also more power-hungry. Cluster type *LITTLE* consists of smaller cores that are slower and more energy-efficient. The increased sophistication of this type of architecture delivers a challenging task to develop better energy management solutions. Running a multi-threaded application only on big cores may not justify the performance gain related to energy consumption. Also, using only small cores may not be the best choice for reducing energy consumption, as the application may suffer a significant increase in execution time. Besides, it is possible to scale up or scale down the core frequencies to improve the performance or save energy, respectively. Each arrangement, or configuration, of operating frequency and the number of available cores of each cluster type to execute a given application may produce a different performance-energy trade-off.

As simple approach to obtaining one or more configurations in HMPs that provide most beneficial performance-energy trade-offs is to collect the performance and energy consumption of a given application in all available configurations. However, in current heterogeneous systems, this would require a high cost to be performed on-line. For example, an embedded system with two clusters each with four cores and 16 clock frequency levels yields 4096 configurations to explore. Such a large number of configurations could be unfeasible for on-line approaches since, in many cases, the search process for the best configuration can outlast the execution itself. Even if executed offline, although it could be possible to evaluate all configuration, the expected overheads will make the method unfeasible.

Considering the vast diversity of performance-energy trade-offs that the large configuration space of HMP systems yields for a given application, the understanding of these trade-offs becomes critical to energy-efficient software development and operation [8]. We present a novel methodology to find the best performance and energy trade-off configurations of parallel applications running on (HMP) systems. It is intended to be used by the operating system to make scheduling decisions at runtime according to performance and energy consumption requirements of the given application and according to the system it is running on. For that, we propose analytical models for performance and power that are fitted offline using a few measurements of the application's execution.

The combination of the outcomes of the performance and power models yields the energy model that can be used to estimate the energy consumption of all configurations for a given application on a specific architecture. The performance and energy models are used to assess the whole configuration space, commonly in the order of thousands of points, and the outcomes are used to select the configurations that lay on the performance and energy Pareto frontier, whose number is often in the order of tens of points. The Pareto frontier is the set of all optimal solutions from which it is not possible to improve a criterion (performance or energy) without degrading another. Hence, we propose a methodology that can estimate the best relationship between performance and energy consumption, without requiring an expensive exploration of the full search space.

We employed the methodology in an ODROID-XU3 board on eight multi-thread applications. We present the fitting of the proposed models and the trade-off Pareto frontier configurations for each application. The results show that the proposed methodology and models can successfully estimate the Pareto trade-off consistently with measured data. On average, we obtained gains in performance and energy consumption in the order of 33% and 40% respectively, when compared to the performance, ondemand and powersave standard Linux DVFS governors. The main contributions of this paper are:

- Simple analytical models for HMP performance and power consumption that only require measurements of execution time and power, avoiding performance hardware counters, which may not be accessible on all architectures.

- A practical methodology to sample the configuration space of a given application efficiently and select points that are estimated to be Pareto-optimal regarding trade-offs of performance and energy consumption.
- Results show that the specific knowledge of the application performance and the system power embedded into analytical models can deliver significantly better results than conventional DVFS governors.
- The technique is easy to use on new applications running on the same HMP system; the power model of the HMP system can be reused, and the performance model can be re-fitted based solely on the execution time of the new application.

The rest of this paper is arranged as follows. Section 2 presents our motivation and places this paper in the context of existing research. Section 3 introduces our modeling techniques and our methodology to find the trade-off configurations. Section 4 describes the experimental setup used to evaluate our methodology, and compares our results to those obtained using the standard Linux DVFS governors. Our conclusions and a discussion of future work follow in Section 5.

2. Related Work

We attempt to distinguish the main characteristics that differentiate our work from others with respect to existing approaches for finding applications' configurations to improve energy-efficiency. The most relevant works are characterised in Table 1.

Table 1. Related work categorized by various aspects.

| Reference | Scenarios | | | Techniques | | | Decision guided by | | | Decision Strategy | Arch. | | |
|-----------------------|----------------|------------|----------|------------|-------------|----------|--------------------|-----------|------------|-------------------|-------------------|----------|---------------|
| | Multi-threaded | Multi-apps | Power | Energy | Performance | DVFS | DPM | Placement | Monitoring | Profiling | Prediction | Pareto | Heterogeneous |
| Endrei et. al [9] | X | - | - | X | X | X | - | - | - | - | Analytical | X | - |
| De Sensi et al. [10] | X | - | X | - | X | X | - | X | X | - | ML | - | - |
| Gupta et al. [11] | X | X | - | X | X | X | X | - | X | X | - | X | X |
| Manumachu et al. [12] | X | - | - | X | X | X | - | - | - | - | Analytical | X | - |
| Tzilis et al. [13] | - | X | - | - | X | X | X | X | X | X | Analytical | - | X |
| Loghin et al. [14] | X | - | - | - | X | - | - | - | - | - | Analytical | - | X |
| Vasilakis et al. [15] | X | - | - | X | X | X | - | X | - | - | Analytical | - | X |
| De Sensi [16] | X | - | X | - | X | X | - | - | - | - | Analytical | - | - |
| Aalsaud et al. [17] | - | X | X | - | - | X | X | X | - | X | - | - | X |
| Proposed Work | X | - | - | X | X | X | - | X | - | X | Analytical | X | X |

Research can be categorized concerning the application scenario focus. Some works concentrate on concurrent execution of multiple applications [11,13,17] and others exclusively on single-application scenarios [9,10,12]. Tzilis et al. [13] propose a runtime manager that estimates the performance and power of applications choosing the most efficient configuration by making use of heuristics to select candidate solutions. They do not consider multi-threaded versions, but instead, they allow multiple instances of the same application to run simultaneously. Indeed, energy efficient management of concurrent applications is harder to accomplish due to dynamically changing situations. However, guarantee an energy consumption requirement with a minimum performance level of a single application affects the energy efficiency of a whole cluster-based system.

Thus, considering this aspect important, our methodology manages energy and performance requirements of individual applications as approached by other works [9,11,12,15]. Some works are only concerned about the performance constraints [13,14]. Furthermore, for parallel or multi-threaded

applications, some works like ours extend Gustafson's [18] and Amdahl's law [19] in order to characterize the application [14,16,20]. Others attempt to identify the phases of each application thread through monitoring performance counters [7,11]. As shown in Loghin et al. [14], the applications' workloads with large sequential fractions present small energy savings regardless of the heterogeneous processing system. Therefore, it is clearly vital to exploit the energy efficiency in HPM architectures as well as the parallelism of single multi-threaded applications.

A notable number of power-management approaches target a reduction of power dissipation and a performance increase by combining three techniques. First, Dynamic Voltage and Frequency Scaling (DVFS) [21], which selects the optimal operating frequency. Second, Dynamic Power Management (DPM) [22], which turns off system components that are idle. Finally, application placement (allocation) [23], which determines the number of cores or the cluster type an application executes on. Some strategies decide the application placement without taking control of the DVFS aspect [14,24–26]. Gupta et al. [11] and Tzilis et al. [13] have different approaches, yet combine DVFS, DPM, and application placement by setting the operating frequency/voltage and the type and number of active cores simultaneously. Also, there are works where the only concern is finding the best cluster frequency performances [9,12,14,16]. In this work, we find the Pareto-optimal core and frequency configurations in heterogeneous architecture that deliver the optimal performance-energy trade-off. Consequently, our method merges DVFS and DPM by choosing the operating frequency/voltage and number of active cores of each cluster type. Thus, our approach covers two of the main techniques for saving energy.

Next, we describe the points that are essential to guide the configuration choice. Generally, we distinguish:

- i offline application profiling;
- ii runtime performance or power monitoring;
- iii predictions performed using a model.

Gupta et al. [11] characterize the application by collecting power consumption, processing time and six performance counters. They use the data to train classifiers that map performance counters to optimal configurations. Then, at runtime, these classifiers and performance counters are used to select the optimal configuration concerning a specific metric, e.g. energy consumption. The major problem of this approach is the growth of the number of cores and cluster types, leading to this strategy losing reliability. Tzilis et al. [13] use a strategy that requires a total number of runs that is linear concerning the number of cores. They use a similar number of performance counters as [11] to profile the application. When the profiled application spawns, they match the online measurement to the closest profiled value and use it as a starting point to predict its performance in the current situation. De Sensi et al. [10] focus on a single application and also monitor data to refine power consumption and throughput prediction models. They compared the predicted outcomes from their models against the data monitored in the current configuration. In case the prediction error is lower than a defined threshold, the calibration phase completes. The problems of online monitoring are the overhead incurred in refining models or making decisions. In contrast, our proposal characterises the application performance and power consumption of a given HMP system using only a few strategically sampled offline measurements. Further, they are used for predicting the optimal performance and energy consumption trade-off configurations.

A sort of performance and power modelling is commonly used to guide the configuration selection. Many works [10,11,13,15,17] are based on analytical models that rely upon runtime measurements retrieved from traces or instrumentation to collect performance counters, which is often not a trivial process. Usually, performance and energy events should be recorded in different runs to prevent counter multiplexing. Otherwise, it may cause application disturbance and decrease measurement accuracy [27]. As in other works [9,12,16], our approach only relies on a minimum set of parameters, such as the number of cores of each cluster and the cluster frequency. Thus, our approach is simple

to automate because it does not require any external instrumentation tools, and its utilisation across different architectures is less restricted.

Loghin and Teo [14] derived equations to evaluate the speedup and energy savings of modern shared-memory homogeneous and heterogeneous computer systems. They introduced two parameters: the inter-core speedup (ICS), describing the execution time relation among diverse types of cores in a heterogeneous system for a given workload; and the active power fraction (APF), representing the ratio between the average active power of one core and the idle power of the whole system. They validated their work against measurements on different types of heterogeneous and homogeneous modern multicore systems. Their results show that energy savings are limited by the system's APF, particularly on performance cores and by the workload's sequential fraction when running on more efficient cores. We used a similar approach when devising our performance and power equations. Our performance model has a parameter that represents the median of the larger (big) cores speedup compared with the smaller (LITTLE) cores. The value of this parameter can be determined independently of any target application, i.e., it is application-agnostic. Our power model describes the dynamic and the static power consumption of the active cores of each cluster type. The aim is to capture the properties of the hardware architecture so that we do not need to fit the power model and the performance speedup parameter for each application. Ultimately, our energy model is a combination of a performance model and an application-agnostic power model.

The main concern with using heuristics, models refinement or any other strategy to make decisions at runtime is the added overhead. Thus, the most significant contribution of our work is to avoid a full exploration of the search space to find the best configurations to execute an application at runtime. A similar goal, but with a different approach, has been pursued by De Sensi [16]. However, the architecture is not HMP, and it offers only 312 possible configurations. Moreover, Endrei et al. [9] and Manumachu et al. [12] use Pareto Frontier as a strategy to minimize energy consumption without degrading application performance. Still, they do not work on ISA heterogeneous architecture as Gupta et al. [11] and our approach. Furthermore, our approach has the advantage of not depending on performance counters, resulting in less restricted employment.

Using the Pareto frontier method, we can find optimal configurations, since it represents the optimal trade-off between energy consumption and performance for the target system [28]. Thus, users can take advantage of configurations at the Pareto frontier to guide development toward energy-optimal computing [29]. Also, these configurations can provide a power-constrained performance optimization to identify the optimized performance under a power budget [30]. Moreover, the Pareto-optimal fronts offer a trade-off zone that can be used to produce the optimal performance and energy efficiency prediction of a model [9,31].

In summary, we propose a method to find the optimal energy-performance trade-offs on a two-cluster heterogeneous architecture of a single multi-threaded application. The energy estimation is obtained by combining the offline performance and power models. The offline performance application characterization is straightforward; it does not use performance counters. The power model involves only architecture-specific parameters, so it is sufficient to fit only once for a specific platform. The individual applications' performance and power requirements can be achieved by the energy-performance trade-off zone provided by the Pareto frontier.

3. Estimating Pareto-Optimal HMP Configurations

This section describes the methodology proposed to estimate the Pareto-optimal configurations for HMP systems. A configuration is defined as a point in the vast parametric space that defines parallel software operation, including the number of processing cores used in each cluster and the operating frequency of those cores. We assume that the clusters have their own voltage and frequency domain, which is shared among the cores of the cluster. Design attributes such as the number of clusters and the number of cores and frequency levels of each cluster are known parameters. The model matches a large subset of actual HMP systems, to which the proposed methodology can be applied. Five steps

compose this methodology: sampling of configurations, runtime measurements, fitting of models, model validation, and the Pareto frontier selection. Figure 1 presents an overview of these steps.

In the first step—sampling of configurations, we need to carefully choose the hardware configurations that will be used to obtain the required data to fit the parameters of the models. To obtain parameters that generate higher accuracy, data measurements should be as diverse as possible, i.e., two or more measurement points should not have a similar number of cores and operating frequency, since those points will most likely not add significant information to the modeling. For this reason, we used a Halton low discrepancy number generator to choose configurations that are not too similar to each other. Morokoff et al. [32] describe how the low discrepancy number generator ensures that it covers the domain more evenly compared to pseudo-random number generators. Also, De Sensi [16] shows that by picking equidistributed points, it is possible to achieve higher accuracy compared to pseudo-random selection of the points.

In the second step—runtime measurements, we evaluate the sampled configurations at runtime, measuring performance and power data to fit the models. First, for each configuration, we execute a generic stress-test to stress the processor and capture power consumption. Second, we execute the target parallel application to measure execution time for another subset of sampled configurations using the Halton number generator.

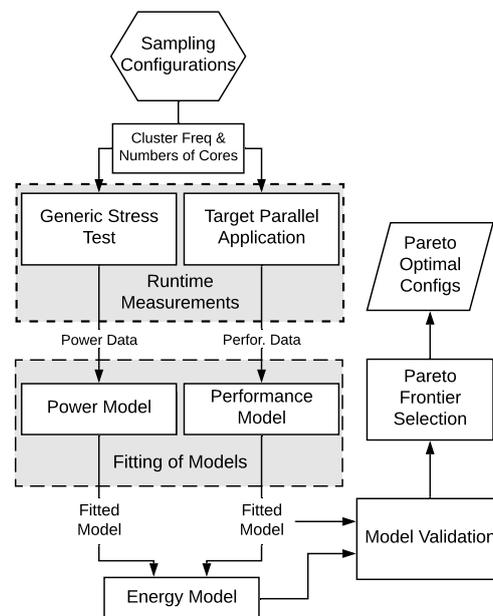


Figure 1. Flowchart of the methodology applied for predicting the optimal performance and energy efficiency trade-offs of a parallel application.

In the third step—fitting of models, we fit the measured data to the proposed architecture- and application-specific performance model and the architecture-specific, application-agnostic power model. The performance model intends to estimate the execution time of a parallel application for any point in the configuration space. We adjust the parameters of the power and performance models separately. The goal of the proposed power model is not to be accurate per se, but to capture the trend changes in power when a change of the operating frequency and number of active cores occurs, independent of the application running. The power consumption prediction accuracy is not strictly crucial if there are no power restrictions to apply on system budget. More importantly, we aim for decreasing energy consumption in HMP since it has more impact on batteries, which is essential on embedded devices. Moreover, being application-agnostic avoids the need for power measurements for every application while preserving the architecture’s signature. Thus, we do not need to fit the power model for each application. Besides, the process of acquiring these measurements also increases the

difficulty of automating these models as practical tools. For each model, we use a non-linear regression minimizing the Root Mean Square Error, given by

$$\text{RMSE} = \sqrt{\frac{\sum_{i=1}^n (A_i - E_i)^2}{n}}, \quad (1)$$

where A_i is the i th measured data of E_i , E_i is the i th estimated value given by the model, and n is the total number of configurations. Sections 3.1 and 3.2 detail the proposed performance and power consumption models, respectively. The proposed energy model is a combination of the power consumption model—previously fitted to a specific HMP system, and the performance model—fitted to a given parallel application and the same HMP system. Section 3.3 presents the proposed energy model.

In the fourth step—model validation, we evaluate the performance and energy consumption models for every point in the vast configuration space that defines parallel software operation.

Finally, in the fifth step—Pareto frontier selection, we select all modeled configurations whose models' outcomes lay on the Pareto frontier, i.e., those configurations that yield the optimal performance and energy consumption trade-offs.

3.1. Application Performance Modelling

In this section, we devise a performance model for an HMP platform with two processing clusters. We assume that a parallel application coherently runs in b big cores and L LITTLE cores. Moreover, we do not consider or require different frequencies for each core in a cluster, i.e., all the cores in a cluster run at the same frequency. Inspired by other works [33–35], we devised the following model that can be used to estimate the performance of a given parallel application running on a two-cluster HMP.

$$T_{\text{HMP}}(F, F_b, F_L, b, L) = T_L(F) \left(\frac{(1-f)F}{\text{perf} \cdot F_b} + \frac{fF}{b \cdot \text{perf} \cdot F_b + LF_L} \right), \text{ if } b > 0, \quad (2)$$

where T_{HMP} is the performance target, that is the total execution time goal for a given application. $T_L(F)$ is the application's execution time running on a single LITTLE core with operating frequency F , which is not necessarily the same frequency of F_L . The number of active big and LITTLE cores in the processor is denoted by b and L , respectively. F_b and F_L are the operating frequencies of the big and LITTLE cores, respectively. perf is the performance improvement when moving computation from a LITTLE to a big core, independent of an application. Note that, perf depends on the hardware design, so we assume it has a fixed value representing how fast a big core is when compared to any LITTLE core. Thus, we do not need to fit this performance speedup parameter for each application.

The value of f represents the parallel fraction of the application. It is the only parameter that characterizes the application. Indeed, by modelling complex parallel applications using their parallel fraction only, we are neglecting the effects that a frequency change has on the performance of the memory hierarchy, on the parallel overhead, and on the distribution of load across the heterogeneous cores [36]. Besides, as those features should limit the parallel speedup, it is expected that the sequential fraction includes them. It is our priority to keep our model straightforward to achieve low runtime overheads. We observed that, despite the many simplifications, the model still provides reasonably accurate energy consumption estimations.

We consider that the sequential fraction of the application is accelerated by using one of the big cores when such core is available. In this approach, the sequential portion expects to include the parallel overhead, i.e., communication or synchronization among the threads, which does not accelerate by making the big core run faster—as it would do for actual sequential code. Otherwise, we can extend Equation (2) as follows:

$$T_{\text{HMP}}(F, F_b, F_L, b, L) = T_L(F) \left(\frac{(1-f)F}{F_L} + \frac{fF}{LF_L} \right), \text{ if } b = 0. \quad (3)$$

In Equation (3), when there is no big core available or active, the application would not take advantage of the performance improvements from the big cores, therefore the parameter perf is removed. Also, the big core operating frequency is not used by the parallel application.

This performance model assumes that the modelled parallel workload is dynamically distributed over the running threads. If that is not the case for the application, cores may run out of work and become idle while others might become overloaded and delay the end of the program's execution. Thus, the performance model may then make incorrect estimations as idle cores still consume power. This can lead to improper power predictions, resulting in inadequate energy predictions. Although, this is a limitation of the proposed approach, the advantages in terms of performance gain and energy reduction presented in this work is an argument for pursuing this type of workload balancing scheme.

3.2. Power Consumption Modelling

In this section, we devise a power model for an HMP platform with two clusters. Assuming that the transistors used on both types of cores have the same power consumption behaviour with respect to their operating frequency and voltages, we can model the power consumption of the big and LITTLE cores, as follows:

$$P_{\text{HMP}}(F_b, F_L, b, L) = P_b(F_b, b) + P_L(F_L, L), \quad (4)$$

where $P_b(F_b, b)$ is the power consumption of b big cores running at frequency F_b , $P_L(F_L, L)$ is the power consumption of the L LITTLE cores running at frequency F_L .

Consider the following equation to describe the power consumption of a processor running at frequency F [37–39].

$$P(F) = aCV^2F + VI_{\text{Leak}}, \quad (5)$$

where a is the average activity factor, C is the load capacitance, V is the supply voltage, and I_{Leak} is the average leakage current. Also, consider the equation that defines the maximum operating frequency for a transistor [40,41]:

$$F_{\text{Max}} = \kappa_1 \frac{(V - V_{\text{Th}})^h}{V}, \quad (6)$$

where V_{Th} denotes the threshold voltage, κ_1 is a constant, and h is a technology-dependent value often assumed to be within the range [2,3]. Since performance is proportional to the operating frequency, digital circuits often operate with voltage and frequency pairs that push frequency closer to the maximum. If maximum performance is not required, the voltage is scaled down with frequency to maintain this policy [42]. Usually, device vendors provide a table of discrete values of supply voltages that the processing chip can operate beneath by maximum frequencies [43]. Generally, the supply voltage V can be perceived as a linear function of the frequency depicted such as follows:

$$V = \kappa_2 F, \quad (7)$$

where κ_2 is any constant. Some authors make a similar approach [42,44] showing that it is reasonable and common in the literature to make such approximation. Furthermore, we reiterate that we intend to capture the trend of how a change in the operating frequency-voltage would affect the power consumption of the system. For this work, our target is to save energy consumption and not accuracy in power consumption prediction. Using Equation (7), we can rewrite Equation (5) as follows:

$$\begin{aligned} P(F) &= aC(\kappa_2 F)^2 F + \kappa_2 F I_{\text{Leak}}, \\ &= aC\kappa_2^2 F^3 + \kappa_2 F I_{\text{Leak}}. \end{aligned} \quad (8)$$

Considering I_{Leak} and a to be approximately constant, we can approximate Equation (8) by

$$P(F) = \alpha F^3 + \beta F, \quad (9)$$

where $\alpha = aC\kappa_2^2$ and $\beta = \kappa_2 I_{Leak}$ are considered constants that mainly abstract semiconductor technology attributes. These attributes should be application-independent and fixed since they rely on hardware design. Then, the power of each cluster type is modelled as follows:

$$P_b(F_b, b) = b\alpha_b F_b^3 + C_b \beta_b F_b, \quad (10)$$

$$P_L(F_L, L) = L\alpha_L F_L^3 + C_L \beta_L F_L, \quad (11)$$

where b and L are the numbers of active big and LITTLE cores, respectively, accounted for the dynamic power modelling component. The technical setup of our experimental platform did not permit disabling individual cores. Therefore, we consider a single term of leak current for each cluster. Then, C_b and C_L represent the total number of cores of each cluster type and account for the static power modelling component.

Note that this power model does not distinguish one application from another since the activity factor a is assumed to be constant. This is a very generic assumption, which may cause an inaccurate power consumption estimation. Nevertheless, for this approach, more critical than accurate estimations of power is to capture how a change of the operating frequency-voltage would affect power consumption. It is expected that the way in which this change affects the power is the same, regardless of the activity factor. In other words, as long as the activity factor of the target application does not change too much from one frequency configuration to another, this assumption will not adversely impact on the proposed power model.

By combining Equations (4), (10) and (11), we come to the power model for the whole HMP chip parameterized by the two frequencies of each cluster as follows:

$$P_{HMP}(F_b, F_L, b, L) = P_{HMP_p} = b\alpha_b F_b^3 + C_b \beta_b F_b + L\alpha_L F_L^3 + C_L \beta_L F_L. \quad (12)$$

Equation (12) represents the power consumption of a parallel fraction (P_{HMP_p}) of an parallel application. Also, we need to predict the power of the sequential part of a code that runs in one core.

$$P_{HMP_s}(F_b, F_L) = \begin{cases} \alpha_b F_b^3 + C_b \beta_b F_b + C_L \beta_L F_L, & \text{if } b > 0, \\ C_b \beta_b F_b + \alpha_L F_L^3 + C_L \beta_L F_L, & \text{otherwise.} \end{cases} \quad (13)$$

When there is, at least, one big core available, we consider the dynamic power of one big core and the static power of all cores. In the case that just LITTLE cores are available, then the sequential part runs on one LITTLE core. In this case, the power consumption of the sequential portion contains the static power of all cores and the dynamic component of one LITTLE core.

3.3. Energy Modelling

In this section, we devise an energy model for a parallel application running on a two-cluster HMP platform by aggregating the sequential and parallel energy consumption as follows:

$$E_{HMP}(F, F_b, F_L, b, L) = E_{HMP_s} + E_{HMP_p}. \quad (14)$$

By combining the parallel part of Equation (2) and the power model of the whole chip in Equation (12), we have the parallel portion energy model component:

$$E_{HMP_p} = \frac{T_L(F) f F (L\alpha_L F_L^3 + C_L \beta_L F_L + b\alpha_b F_b^3 + C_b \beta_b F_b)}{b \cdot \text{perf} \cdot F_b + L F_L}. \quad (15)$$

Using Equation (13) and the sequential part of the Equations (2) and (3), we have the sequential portion energy model component, which relies on the number of active big cores available:

$$E_{\text{HMP}_s} = \begin{cases} \frac{T_L(F)(1-f)F(C_L\beta_L F_L + \alpha_b F_b^3 + C_b\beta_b F_b)}{\text{perf} \cdot F_b}, & \text{if } b > 0, \\ \frac{T_L(F)(1-f)F(C_b\beta_b F_b + \alpha_L F_L^3 + C_L\beta_L F_L)}{F_L}, & \text{otherwise.} \end{cases} \quad (16)$$

Finally, combining Equations (15) and (16), we come to the consolidated energy model for the whole HMP chip parameterized by the frequencies and the number of active cores in each cluster:

$$E_{\text{HMP}}(F, F_b, F_L, b, L) = T_L(F)F \times \begin{cases} \frac{(1-f)(C_L\beta_L F_L + \alpha_b F_b^3 + C_b\beta_b F_b)}{\text{perf} \cdot F_b} + \frac{f(L\alpha_L F_L^3 + C_L\beta_L F_L + b\alpha_b F_b^3 + C_b\beta_b F_b)}{b \cdot \text{perf} \cdot F_b + L F_L}, & \text{if } b > 0, \\ \frac{(1-f)(C_b\beta_b F_b + \alpha_L F_L^3 + C_L\beta_L F_L)}{F_L} + \frac{f(L\alpha_L F_L^3 + C_L\beta_L F_L + C_b\beta_b F_b)}{L F_L}, & \text{otherwise.} \end{cases} \quad (17)$$

We can now estimate energy consumption and the performance of any configuration for a given parallel application using our energy model in Equation (17) and performance models in Equations (2) and (3), respectively. Considering fitted models, we can estimate all possible performance and energy consumptions of an application. Then, these estimations are used in the Pareto method to obtain the most optimal performance-energy trade-offs. In the next section, we are going to show the values obtained after the fitting process and the Pareto approach validation on a two-cluster heterogeneous board.

4. Experimental Evaluation and Results

This section describes the experimental setup, including the specifications of the platform and the measurement methodology applied to collect data. Then, we will show the hardware design parameters and the values of the parallel fractions obtained through the non-linear regression. We validate our proposed methodology concerning the estimated Pareto frontier against direct measurements. We computed the error of the estimations by using the Mean Absolute Percentage Error (MAPE), defined as:

$$\text{MAPE} = \frac{1}{n} \sum_{t=1}^n \left| \frac{A_t - E_t}{A_t} \right|, \quad (18)$$

where n is the number of configurations, A_t is the actual measurement (execution time or energy consumption) and E_t is the estimated execution time or energy consumption.

4.1. Experimental Setup

The proposed methodology was validated using an ODROID-XU3 [45] board developed by Hardkernel co. with two core types of the same single-ISA, so both can execute the same compiled code. It uses a Samsung Exynos5422 System on a Chip (SoC), which utilizes ARM big.LITTLE technology and constitutes an ARM Cortex-A15 (big) quad-core cluster and a Cortex-A7 (LITTLE) quad-core cluster. The ODROID-XU3 has 19 available frequency levels ranging from 200 MHz up to 2 GHz on Cortex-A15 and 14 available frequency levels from 200 MHz to 1.5 GHz on Cortex-A7. Therefore, considering that each configuration is a combination of the number of cores in each cluster and a frequency pair that is available as a resource for an application to be executed, there are 6384 possible configurations to explore. To further clarify, consider that each cluster can select none to four cores, then $(\# \text{ of big cores}) \times (5 \times 19) \times (\# \text{ of LITTLE cores}) \times (5 \times 14) = 6650$. For each combination of active numbers of LITTLE and big cores, there are 19×14 arrangements of frequencies to choose. Since the combination of 0 big cores and 0 LITTLE cores is not possible, the number of available

configurations is $6650 - (19 \times 14) = 6384$. The experimental platform was set up with a Ubuntu Minimal 18.04 LTS running Linux Kernel LTS 4.14.43, and it did not permit disabling individual cores. Our software project can be found in the GitLab (<https://gitlab.com/lappsufrn/XU3EM>).

The ODROID-XU3 has four current and voltage sensors to measure the power dissipation of the cortex-A15 cores, cortex-A7 cores, GPUs and DRAMs individually. We sampled the clusters' power consumption readings every 0.5 s with timestamps. Then, by integrating the power samples, we computed the energy used by a given application. We consider the average power calculated by dividing total energy by execution time. Moreover, we used *cpuset* [46] subsystem to confine all the data collection instrumentation and operating system's processes/threads to the not active cores, when available, and assign the workload application appropriately to the cores. Thus, this mitigates the interference of all non-target application in the experiment.

4.2. Hardware Design Parameters

This section shows the values obtained for the parameters that represent the hardware design characteristics. As these parameters are related to the CPU design and they must be application-agnostic, we used a CPU stress tool, called *stress-ng* [47]. This tool can load and stress a computer system in various selectable ways and is available in any Linux distribution. This tool has an extensive range of specific stress tests that exercise floating-point, integer, bit manipulation, and control flow operations on the CPU.

We tested each method to find the one that maximizes the stress on the ODROID-XU3, i.e., the method that led to the highest power draw as cores executed their instructions (see Figure 2). Hence, after executing every method for 15 s on two big cores (1.9 GHz), and four LITTLE cores (1.5 GHz), we chose the *callfunc* stress method. Using all four big cores at the highest frequency would have overheated the board. High CPU frequencies, especially from the big cores, produce a large amount of heat, making the system protection mechanism against high temperatures halt the device to prevent damage. Moreover, running 15 s provides 30 power reading samples from the sensors. The average standard deviations regarding all methods are 0.0792 and 0.2395 for LITTLE and big power consumption, respectively.

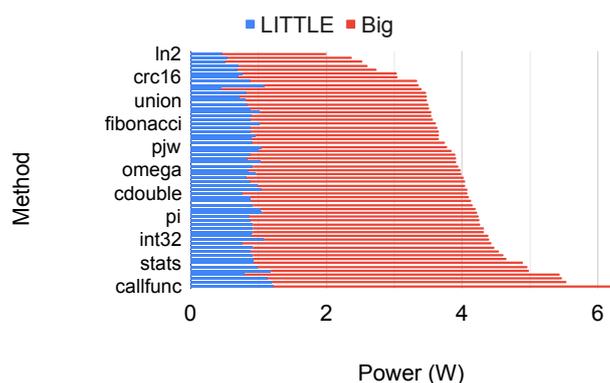


Figure 2. Power consumption from LITTLE and big clusters for every stress-ng method.

4.2.1. The Perf Parameter

The perf performance parameter represents the big core speedup in comparison with the LITTLE core. It is application-agnostic, relying only on the chip design. We executed the *stress-ng* with *callfunc* method on one core of each cluster type, fixing the number of operations at 87,555. This number of operations provides approximately one second to conclude in one big core with the highest frequency and up to 15 s for one LITTLE core with the lowest frequency.

Figure 3 shows the result of perf fitting. The left-hand side Y-axis shows the median execution time of five runs for one core of each cluster type and different frequencies. The frequency range used

is 200 MHz–1.5 GHz. Higher values are not possible due to the maximum LITTLE core frequency. The right-hand side Y-axis represents the speedup of one big core when comparing to one LITTLE core for each frequency. The average standard deviations of the runs regarding all frequencies are 0.0252 and 0.0404 for LITTLE and big core, respectively. The median of every speedup is 1.897 which is the obtained perf value.

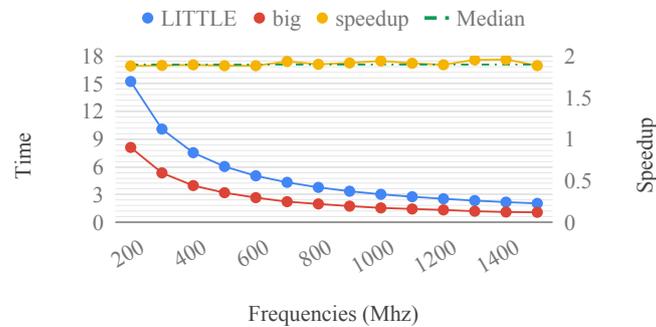


Figure 3. Performance parameter perf fitted for the architecture using a fixed number of operations of the stress-ng tool for one core of each cluster.

4.2.2. Power Consumption Parameters

We aim to obtain a power model that independent of the application, represents the power consumption of the two-clusters. As we want to model the total power consumption of the entire chip, we measured the whole chip to obtain P_{HMP} . We defined the number of processes and which cores would be active to execute the stress-ng by generating 95 evenly distributed Halton configurations. Each configuration is composed of the number of active big cores b , active LITTLE cores L , big's frequency cluster F_b and LITTLE's frequency cluster F_L .

We used the stress-ng with callfunc stress method and a fixed execution time of 75 s, with the command line `stress-ng -cpu N -cpu-method callfunc -t 75 -taskset t`. The `-cpu` parameter determines the number of processes on which the same method will be executed, and the parameter `-taskset` can be used to set a process's CPU affinity.

Running 75 s provides 150 power samples. The average standard deviations regarding all 95 configurations are 0.0164 and 0.1478 for LITTLE and big core power consumption, respectively. As the configuration parameters and their total power consumption measurements are known, we fit the parameters $\alpha_L, \beta_L, \alpha_b, \beta_b$ from the Equation (12). The fitted values are: $\alpha_L = 5.953 \times 10^{-29}$, $\alpha_b = 2.914 \times 10^{-28}$, $\beta_L = 1.033 \times 10^{-10}$, $\beta_b = 9.342 \times 10^{-11}$.

4.3. Applications

To validate our approach, we used eight applications: the Black-Scholes, Bodytrack, Freqmine applications provided by PARSEC, Smallpt and x264 by Phoronix Test Suite and kmeans, Particle Filter and LavaMD from Rodinia benchmark. PARSEC [48] is a well-known benchmark suite of parallel applications, which focuses on emerging workloads and is designed to be representative of next-generation shared-memory programs for chip-multiprocessors. The Phoronix Test Suite [49] is a benchmarking platform that provides an extensible framework for carrying out tests in a fully automated manner from test installation to execution and reporting. Rodinia benchmark [50] is a set of applications designed for heterogeneous computing infrastructures with OpenMP [51], OpenCL [52], and CUDA [53] implementations.

We have chosen the applications because they were implemented using OpenMP. This allowed us to modify their workloads to be dynamically balanced using dynamic OpenMP schedule clause [54]. Moreover, the number of threads created is equal to the number of active cores available of a given application under a configuration. The thread master was bound to one big core in order to make sure

that the sequential part of the code would run on the higher performance core. The other threads were bound to one core each. This means that these applications follow the assumptions we made for the performance model.

The notable exception for the chosen criteria is the x264 encoder provided by the Phoronix test suite since it has only POSIX threads for parallelism. As higher video quality significantly depends on the encoding process, i.e., longer execution time and higher energy consumption, it is essential to improve the energy efficiency of video encoding on embedded devices. We therefore evaluate the behaviour of our methodology on this video encoder.

The characteristics of the applications which are used in this paper, according to [48,50,55–57], can be seen in Table 2. Most applications are CPU and memory intensive, so we expect a considerable use of the HMP chip and the memory system as well as the communication between them.

Nevertheless, some implicit hardware or software features, such as communication-to-computation ratio, cache sharing and off-chip traffic, are not explicitly included in our power and performance model. What we want to demonstrate is that a low-overhead, straightforward model is sufficient to estimate a Pareto frontier. Besides, as those features should limit the parallel speedup, it is expected that the sequential fraction of the application may be more significant for some applications.

We strive to capture the nuances of each configuration that will affect the way the application works on the hardware, producing different execution time. Therefore, we fit the performance model to find the specific parallel fraction f of each application by executing 50 Halton configurations (F_b, F_L, b, L). As Halton generates configurations that cover the domain more evenly than pseudo-random algorithms, 50 configurations provide sufficient data for us to build our performance model. One more configuration is needed to find $T_L(F)$, it is $b = 0, L = 1, F_b = 2$ GHz, $F = F_L = 800$ MHz. We calculated the median of the execution time of five runs for each configuration.

Considering the fitted perf, and the Equations (2) and (3), we use non-linear regression to find the parallel fraction f for every application, as explained in the Section 3. Table 3 shows the value of f for each application. Note that Black-Scholes and kmeans have a parallel fraction of 0.7743 and 0.6381, respectively, indicating a larger sequential portion of these applications when compared to other applications. The following section will show the validation of the Pareto frontier. It provides the evidence that our low-overhead models are sufficient for achieving our goals.

4.4. Pareto Frontier Validation

In this subsection, we validate our proposed methodology. Using the energy model in Equation (17) and the fitted performance model in Equations (2) and (3), we generated a pair of estimated performance and energy consumption values of all possible configurations for each application in Table 2. Then, the Pareto frontier method selects the performance and energy pairs that offer the optimal trade-off.

For all applications, the Pareto frontier selection step resulted in configurations that use all cores of the two clusters. The frequencies of the cluster big and LITTLE were the only parameters that have changed. This probably occurred because it was not possible to disable the unused cores during the fitting. As a result, the energy savings from configurations, which are not using all available cores, are not enough to justify the performance mitigation. Therefore, the optimal configurations become those that use all cores.

Figure 4 shows the estimated and measured Pareto Frontier for all applications w.r.t. all possible configurations. Each point describes a pair of performance and energy consumption of a configuration. The grey points are the estimated outcomes of our energy and performance models for all possible configurations. In total, there are 6384 available configurations w.r.t. the combinations of the number of cores of each cluster and pair frequency.

Table 2. Characteristics detail of each application based on [48,50,55–57].

| Benchmark | App. | Domain | Type | Problem Size | Description |
|-----------|-----------------|--------------------|--------------------------|--|---|
| PARSEC | Black-Scholes | Financial Analysis | CPU intensive | 10,000,000 options | Portfolio price calculation using Black-Scholes PDE |
| | Bodytrack | Computer Vision | CPU and memory intensive | 4 cameras, 261 frames, 4000 particles, 5 annealing layers | Computer vision, tracks 3D pose of human body |
| | Freqmine | Data Mining | CPU and memory intensive | Database composed of spidered collection of 250,000 web HTML documents | Array-based version of the FP-growth method |
| Rodinia | Kmeans | Data Mining | CPU and memory intensive | 1000,000 points, 34 dimensions, 5 clusters | Mean-based data partitioning method |
| | Particle Filter | Medical Imaging | CPU and memory intensive | Video resolution 128×128 , 10 frames, 10,000 particles | Tracks cells by statistically estimating the path in a Bayesian framework. |
| | LavaMD | Molecular Dynamics | CPU and Memory intensive | 100 boxes | Computes the interactions between particles in a three-dimensional cubic space |
| Phoronix | Smallppt | Image Renderer | CPU intensive | 1024×768 , 128 samples/ pixel | Generates a single image of a modified Cornell box rendered with full global illumination |
| | x264 | Media Processing | CPU and memory intensive | 1920×1080 pixels (HDTV resolution), 600 frames | H.264 video encoder |

Table 3. The fitted parallel fraction values of each application.

| Application | Black-scholes | Bodytrack | Freqmine | Smallpt | x264 | Kmeans | Particle Filter | LavaMD |
|-------------|---------------|-----------|----------|---------|-------|--------|-----------------|--------|
| <i>f</i> | 0.7743 | 0.9384 | 0.9343 | 0.9898 | 0.888 | 0.6381 | 0.9251 | 0.9961 |

The black circles depict the selected Pareto configurations from the estimated energy and performance pairs. Notice that the estimated Pareto have the least energy consumption concerning different performances. Each red point represents the actual values measured from the sensors using each selected Pareto configuration. Observe that the measured and estimated values are similar to each other but, as the execution time increases, the difference between the measured and estimated energy consumption also rises. The yellow triangle, the blue plus sign and the green hexagon are the performance, ondemand and powersave Linux governors, respectively. We compare our methodology with them in the next section.

Table 4 shows the total number of configurations selected by Pareto Frontier and the measured Pareto Frontier variations of energy and performance for each application. Our approach decreased the search space of 6384 available configurations for the given platform by approximately 99%. Considering the highest and least energy consumption and performance among the measured Pareto configurations, we observe energy savings of up to 68.23% and performance gains of up to 75.62% for kmeans application. The Freqmine application presented the least variation of energy, up to 54.35%, however, the performance gains were up to 85.41%. This demonstrates that the selected configurations by Pareto frontier can provide a good range of performance improvement and energy saving options for operating systems.

Table 4. Pareto Frontier energy and performance variations concerning the measured values and the total number of Pareto configurations.

| | Application | | | | | | | |
|--|---------------|-----------|----------|---------|--------|--------|-----------------|--------|
| | Black-scholes | Bodytrack | Freqmine | Smallpt | x264 | Kmeans | Particle Filter | LavaMD |
| Number of Pareto configurations | 93 | 71 | 72 | 66 | 78 | 108 | 74 | 57 |
| Variation of Energy | 60.20% | 64.71% | 54.35% | 59.43% | 57.74% | 68.23% | 58.28% | 54.48% |
| Variation of Performance | 80.47% | 84.84% | 85.41% | 89.53% | 85.55% | 75.62% | 83.97% | 88.58% |

Figure 5 gives a zoom-in in Figure 4. The grey points that represent all possible estimated configurations' outcomes are not displayed. Moreover, the measured and the related estimated Pareto configurations are linked using a dotted line, thus, showing the configurations' results from the models against the measured data. Depending on the application, we could not execute some configurations to avoid overheating the board. As anticipated in Section 4.2, high CPU frequencies produce high temperatures that make the system halt. For instance, the Freqmine results presented in Figure 5c show four estimated Pareto points with the frequencies (big,LITTLE) being (1500,2000), (1500,1900), (1400,1900) and (1500,1800) without their respective measured points. Note that two of them are overlapping. The same occurred with the Smallpt, LavaMD, x264 and particle Filter applications (see Figure 5d,e,g,h).

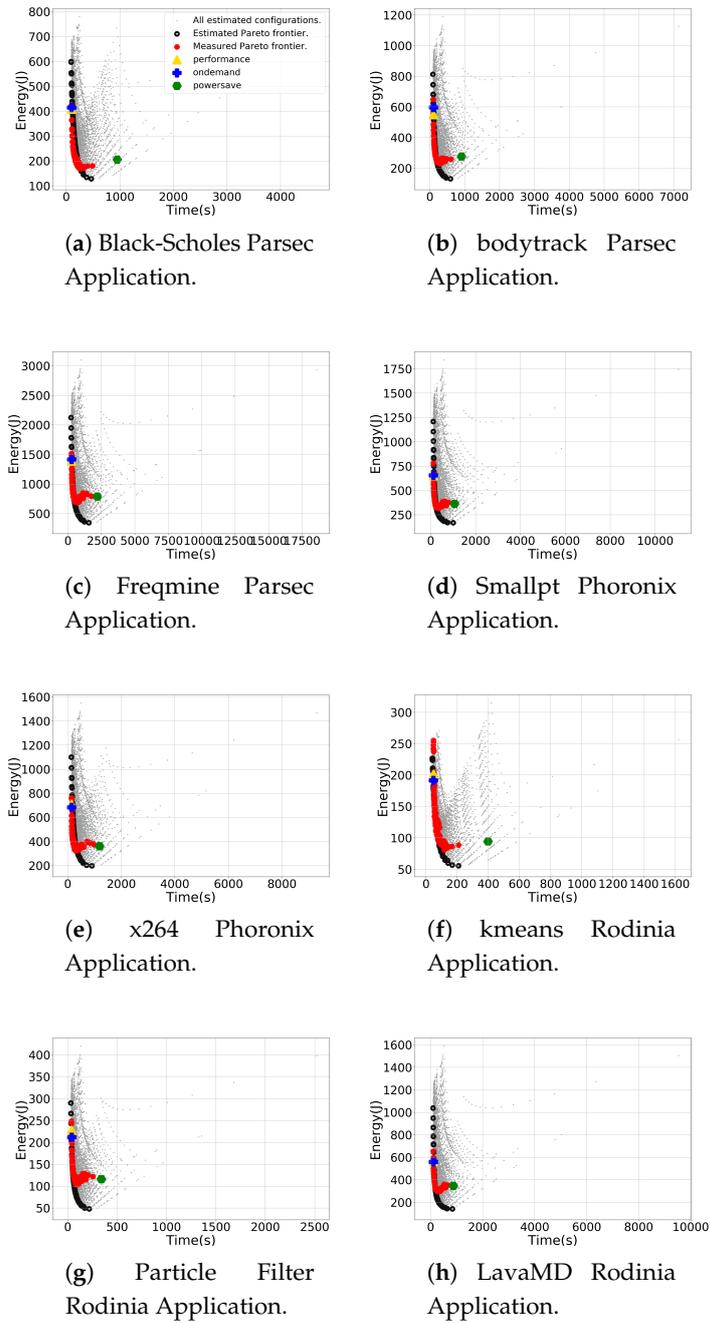
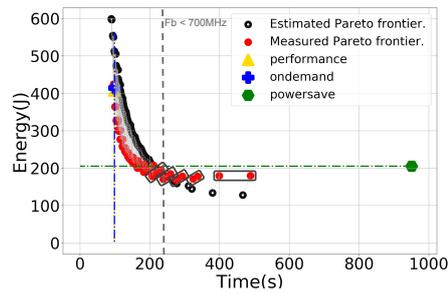
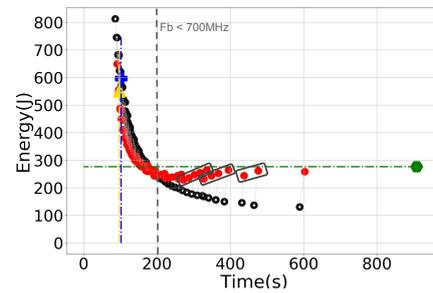


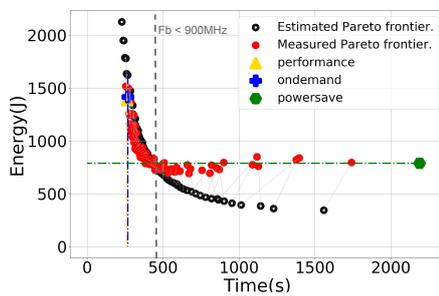
Figure 4. Estimated and measured Pareto frontier compared with the default governors, and all possible modeled configurations.



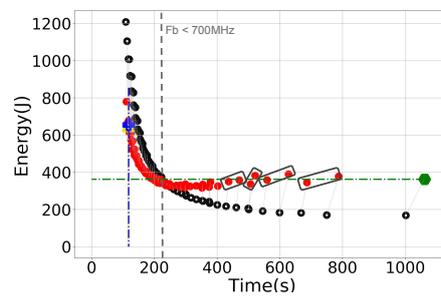
(a) Black-Scholes Parsec Application.



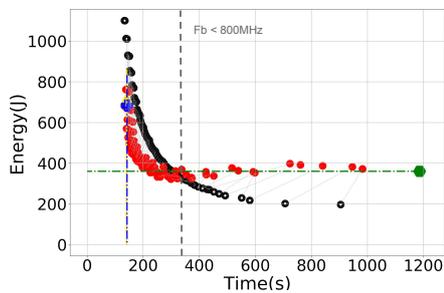
(b) bodytrack Parsec Application.



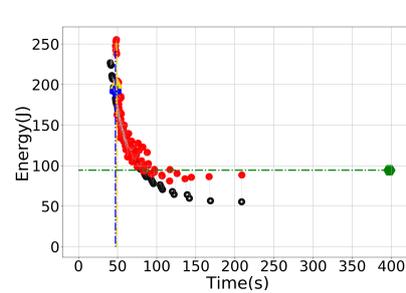
(c) Freqmine Parsec Application.



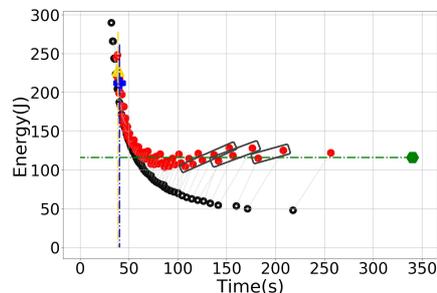
(d) Smallpt Phoronix Application.



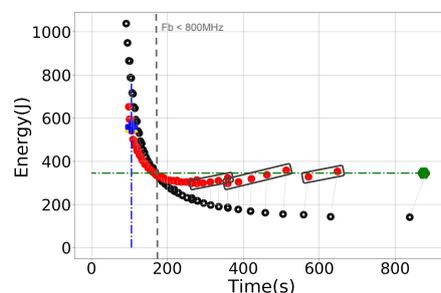
(e) x264 Phoronix Application.



(f) kmeans Rodinia Application.



(g) Particle Filter Rodinia Application.



(h) LavaMD Rodinia Application.

Figure 5. Estimated and measured Pareto frontier compared with the default governors removing all possible modeled configurations.

Notice that all applications except the Particle Filter and kmeans (see Figure 5g,f), have a point when the measured energy consumed is higher than the estimated. A vertical line represents at which time and big cluster frequency this occurs (see Figure 5). As the energy is also influenced by the performance modeling, the measured execution time in some cases takes longer when compared with the expected performance consuming more energy.

Notice the highlighted configurations (see Figure 5) that have similar measured energy. Those points have the same big frequency, and, as the LITTLE frequency scales down, the execution time increases, and so does energy consumption. Further, when the big cluster frequency decreases, there is an abrupt drop in energy consumption. For instance, this can be seen in Figure 5b of the Bodytrack application in points above 200 s. This shows that the big cores play an important role in energy consumption.

Some applications do not have a workload as well-balanced as we expected. This results in the predicted performance not being well correlated with the measured data. Even though Freqmine (see Figure 5c) and kmeans (see Figure 5f) have the dynamic clause included in their parallel loops, they do not present consistent performance. Additional work will be required to understand why these applications are not workload balanced. In particular, the x264 application (see Figure 5e) requires more refined modelling since video processing is much more complex, and the application phases tends to vary according to the frames from the input video. Our approach shows reasonable energy savings, even though fixing a Pareto-optimal configuration to the whole application execution.

Table 5 shows the mean absolute percentage error (MAPE) between the measured and estimated performance and energy consumption of the Pareto Frontier, including the average standard deviation of the measured data. The highest energy and performance errors are from Black-Scholes (38.92%) and x264 (11.10%) applications, respectively. The lowest errors are observed in kmeans and LavaMD with 10.48% and 1.75% of energy and performance errors, respectively. It is important to notice that the x264 had a reasonable accuracy, considering its complexity. On average, we achieved performance and energy errors of 5.53% and 22.30%, respectively.

Table 5. The average standard deviation and mean absolute percentage error of performance and energy consumption of the estimated Pareto frontier against measured data.

| | | Application | | | | | | | |
|-------------|-------------------|---------------|-----------|----------|---------|--------|--------|-----------------|--------|
| | | Black-scholes | Bodytrack | Freqmine | Smallpt | x264 | kmeans | Particle Filter | LavaMD |
| Performance | MAPE | 4.88% | 2.55% | 6.8% | 2.16% | 11.10% | 6.96% | 8% | 1.75% |
| | Ave. Std. Deviat. | 2.4536 | 0.8468 | 17.534 | 0.6111 | 6.8414 | 2.4148 | 0.2099 | 0.34 |
| Energy | MAPE | 38.92% | 23.97% | 17.49% | 27.42% | 19.64% | 10.48% | 17.96% | 22.52% |
| | Ave. Std. Deviat. | 1.2244 | 1.7626 | 11.7941 | 4.9319 | 1.9157 | 7.2793 | 0.6401 | 1.1927 |

It is important to recall that our concern is the trend of the measured Pareto frontier. Therefore, even without an exact match between measured and estimated values, the Pareto configurations are reasonable alternatives that give each application suitable trade-off choices. In the next section, we will compare our approach against the Linux governors.

4.5. Comparison against DVFS Governors

All applications were executed under the performance, ondemand, and powersave Linux governors (see Figure 5). The performance and powersave governors set the CPU statically to the highest and lowest frequency, respectively, within the borders of available minimum and maximum frequencies. Moreover, the ondemand governor scales the frequency dynamically according to the current load. It boosts to the highest frequency and then likely decreases as the idle time raises.

We compared our approach to these three governors as they provide a consistent variety of optimization options and are implemented on numerous mobile devices, making them competitive baselines.

The green line in Figure 5 shows the measured points that have higher performance and less energy consumption when compared with the powersave governor to all applications. Notice that there are many configurations pointed from the Pareto frontier better than the powersave governor. The blue and yellow lines show the measured points that have higher or similar performance when compared with the ondemand and performance governor, respectively. Notice that there are few configurations at the Pareto frontier better or similar to those governors.

Normalized Comparison

Figures 6 and 7 show, respectively, the normalized energy consumption and the normalized performance of all the benchmarks compared with the performance, ondemand and powersave governors. We chose the measured Pareto configurations that give the least energy consumption (MPLE) and the highest performance (MPHP) to normalize the energy and the performance, respectively, for each application.

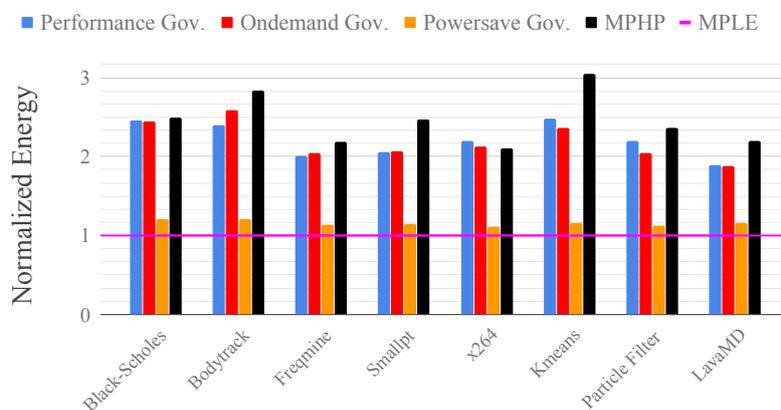


Figure 6. Normalized energy consumption by the measured Pareto configuration that provides the least energy consumption (MPLE) for each application.

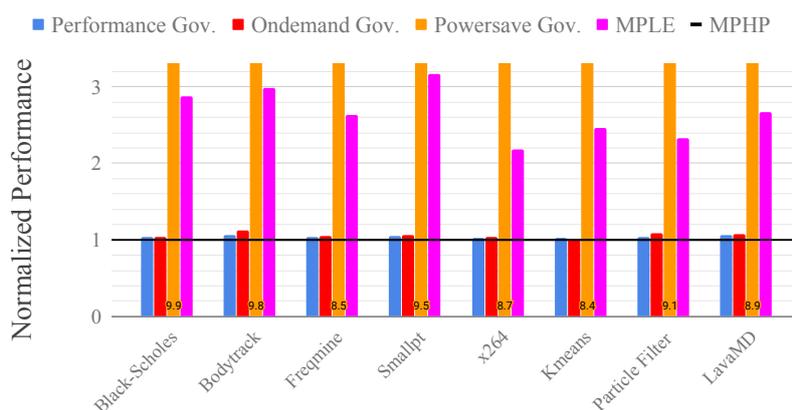


Figure 7. Normalized performance by the measured Pareto configuration that provides the highest performance (MPHP) for each application.

Figure 6 shows that the MPLE saved more energy when compared to all governors for every application. Also, we can observe that MPHP saved energy when compared to the performance and ondemand governors for the x264 application. Figure 7 shows that MPHP has a small speedup when compared to performance and ondemand Linux governors. On average, notice that MPLE has

7× higher speedup when compared with the powersave governor and it is only 2× slower when compared to the performance governor.

Table 6 shows the percentage of performance gains and energy savings when compared to the MPHP and MPLE, respectively, for all applications. Our methodology saved, on average, 54.38%, 53.99% and 13.67% of energy w.r.t. the performance, ondemand, and powersave Linux governors, respectively. Also, we observed 4.23%, 5.52% and 89.03% of speedup w.r.t. the performance, ondemand, and powersave Linux governor, respectively.

Table 6. Percentage of performance gains and energy savings when compared to the measured Pareto configurations that give the highest performance (MPHP) and the least energy consumption (MPLE), respectively, for all applications.

| | | Linux Governors | | |
|-----------------|------|-----------------|----------|-----------|
| | | Performance | Ondemand | Powersave |
| Black-Scholes | MPHP | 3.69% | 3.37% | 89.96% |
| | MPLE | 59.33% | 59.19% | 17.70% |
| Bodytrack | MPHP | 6.42% | 10.63% | 89.94% |
| | MPLE | 58.29% | 61.38% | 17.15% |
| Freqmine | MPHP | 3.17% | 5.24% | 88.39% |
| | MPLE | 50.11% | 51.15% | 12.40% |
| Smallpt | MPHP | 4.53% | 5.93% | 89.54% |
| | MPLE | 51.30% | 51.68% | 12.75% |
| x264 | MPHP | 2.67% | 4.07% | 88.48% |
| | MPLE | 54.62% | 52.88% | 10.63% |
| kmeans | MPHP | 2.86% | −0.01% | 88.10% |
| | MPLE | 59.77% | 57.67% | 14.07% |
| Particle Filter | MPHP | 4.18% | 8.05% | 89.09% |
| | MPLE | 54.54% | 51.09% | 10.67% |
| LavaMD | MPHP | 6.29% | 6.86% | 88.78% |
| | MPLE | 47.05% | 46.86% | 14.02% |
| Average | MPHP | 4.23% | 5.52% | 89.03% |
| | MPLE | 54.38% | 53.99% | 13.67% |

5. Conclusions and Future Work

We presented a novel methodology to estimate optimal performance and energy trade-off configurations for parallel applications on Heterogeneous Multi-Processing (HMP) systems. We devised an analytical low-overhead straightforward performance model for a given multi-thread application and an analytical application-agnostic power model for a specific two-cluster HMP system. These models, when combined, generate the energy model which can assess all available configurations to predict an application's energy consumption. The Pareto frontier uses the provided offline performance and energy models to select, from all available options, the optimal performance-energy trade-off.

We validated our methodology on an ODROID-XU3 board which we used to fit our models and to validate the Pareto frontier configurations. Moreover, we compared the measured Pareto frontier with the performance, powersave and ondemand Linux governors. Our approach achieved a reduction of the configuration search space of approximately 99%, significantly decreasing the number of options available to identify an optimal performance-energy trade-off configuration. On average, the performance and energy absolute percentage errors between the measured and the estimated Pareto frontier for all applications are 5.53% and 22.30%, respectively. The average variation of performance and energy concerning all applications are 84.25% and 59.68%, respectively. Also, we obtained 13.67% of energy savings regarding the powersave governor with higher or similar performance. These results encourage future research in performance- and energy-aware schedulers using our methodology. Furthermore, we can apply our models to predict optimal energy and performance trade-offs.

The proposed approach can be used to run parallel applications that have previously been characterized, minimizing the overhead and energy waste associated with runtime characterization. Also, offline characterization can be made as precise as necessary since resource limitation is often not

an issue. As future work, we intend to reduce the performance restrictions resulting from data-content dependent workloads, the size of the problem's input and application phase changing. Furthermore, the accuracy of the power prediction may be remodelled, considering the chip voltage as a non-linear relationship with the maximum operating frequency.

Author Contributions: Funding acquisition, K.E. and S.X.-d.-S.; Project administration, K.E. and S.X.-d.-S.; Software, D.A.M.C.; Supervision, S.X.-d.-S. and K.E.; Validation, D.A.M.C.; Writing—original draft, D.A.M.C.; Writing—review and editing, D.D.S., A.F.L., K.G., J.N.-Y., K.E. and S.X.-d.-S. All authors have read and agreed to the published version of the manuscript.

Funding: This work was financed in part by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior-Brasil (CAPES)-Finance Code 001, and in part by the Royal Society-Newton Advanced Fellowship award no. NA160108. It is also supported in part by the European Union's Horizon 2020 Research and Innovation Programme under grant agreement no. 779,882, TeamPlay (Time, Energy and security Analysis for Multi/Many-core heterogeneous PLATforms).

Acknowledgments: The authors would like to thank the anonymous reviewers.

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

References

1. Pham, D.; Aipperspach, T.; Boerstler, D.; Bolliger, M.; Chaudhry, R.; Cox, D.; Harvey, P.; Harvey, P.; Hofstee, H.; Johns, C.; et al. Overview of the architecture, circuit design, and physical implementation of a first-generation Cell processor. *IEEE J. Solid State Circuits* **2006**, *41*, 179–196. [CrossRef]
2. Corp. A. big.LITTLE Technology—ARM. Available online: <https://www.arm.com/products/processors/technologies/biglittletprocessing.php> (accessed on 6 April 2019).
3. Chitlur, N.; Srinivasa, G.; Hahn, S.; Gupta, P.K.; Reddy, D.; Koufaty, D.; Brett, P.; Prabhakaran, A.; Zhao, L.; Ijth, N.; et al. QuickIA: Exploring heterogeneous architectures on real prototypes. In Proceedings of the IEEE International Symposium on High-Performance Comp Architecture, New Orleans, LA, USA, 25–29 February 2012; pp. 1–8. doi:10.1109/HPCA.2012.6169046. [CrossRef]
4. Wang, G.; Li, W.; Hei, X. Energy-aware real-time scheduling on Heterogeneous Multi-Processor. In Proceedings of the 2015 49th Annual Conference on Information Sciences and Systems, CISS 2015, Baltimore, MD, USA, 18–20 March 2015; pp. 1–7, doi:10.1109/CISS.2015.7086884. [CrossRef]
5. Koufaty, D.; Reddy, D.; Hahn, S. Bias scheduling in heterogeneous multi-core architectures. In Proceedings of the 5th European Conference on Computer Systems—EuroSys '10, Paris, France, 13–16 April 2010; ACM Press: New York, NY, USA, 2010; p. 125, doi:10.1145/1755913.1755928. [CrossRef]
6. Saez, J.C.; Pousa, A.; Castro, F.; Chaver, D.; Prieto-Matias, M. Towards completely fair scheduling on asymmetric single-ISA multicore processors. *J. Parallel Distrib. Comput.* **2017**, *102*, 115–131, doi:10.1016/j.jpdc.2016.12.011. [CrossRef]
7. Sawalha, L.; Wolff, S.; Tull, M.P.; Barnes, R.D. Phase-Guided Scheduling on Single-ISA Heterogeneous Multicore Processors. In Proceedings of the IEEE 2011 14th Euromicro Conference on Digital System Design, Oulu, Finland, 31 August–2 September 2011; pp. 736–745, doi:10.1109/DSD.2011.98. [CrossRef]
8. Jin, C.; de Supinski, B.R.; Abramson, D.; Poxon, H.; DeRose, L.; Dinh, M.N.; Endrei, M.; Jessup, E.R. A survey on software methods to improve the energy efficiency of parallel computing. *Int. J. High Perform. Comput. Appl.* **2017**, *31*, 517–549, doi:10.1177/1094342016665471. [CrossRef]
9. Endrei, M.; Jin, C.; Dinh, M.N.; Abramson, D.; Poxon, H.; DeRose, L.; de Supinski, B.R. Energy Efficiency Modeling of Parallel Applications. In Proceedings of the International Conference for High Performance Computing, Networking, Storage, and Analysis, Dallas, TX, USA, 11–16 November 2018; IEEE Press: Piscataway, NJ, USA, 2018; pp. 212–224.
10. De Sensi, D.; Torquati, M.; Danelutto, M. A Reconfiguration Algorithm for Power-Aware Parallel Applications. *ACM Trans. Archit. Code Optim.* **2016**, *13*, 1–25, doi:10.1145/3004054. [CrossRef]
11. Gupta, U.; Patil, C.A.; Bhat, G.; Mishra, P.; Ogras, U.Y. DyPO: Dynamic Pareto-Optimal Configuration Selection for Heterogeneous. *ACM Trans. Embed. Comput. Syst.* **2017**, *16*, 1–20, doi:10.1145/3126530. [CrossRef]

12. Manumachu, R.R.; Lastovetsky, A. Bi-Objective Optimization of Data-Parallel Applications on Homogeneous Multicore Clusters for Performance and Energy. *IEEE Trans. Comput.* **2018**, *67*, 160–177, doi:10.1109/TC.2017.2742513. [[CrossRef](#)]
13. Tzilis, S.; Trancoso, P.; Sourdis, I. Energy-Efficient Runtime Management of Heterogeneous Multicores using Online Projection. *ACM Trans. Archit. Code Optim.* **2019**, *15*, 1–26, doi:10.1145/3293446. [[CrossRef](#)]
14. Loghin, D.; Teo, Y.M. The time and energy efficiency of modern multicore systems. *Parallel Comput.* **2018**, *86*, 1–10, doi:10.1145/3229710.3229714. [[CrossRef](#)]
15. Vasilakis, E.; Sourdis, I.; Papaefstathiou, V.; Psathakis, A.; Katevenis, M.G. Modeling energy-performance tradeoffs in ARM big.LITTLE architectures. In Proceedings of the 2017 27th International Symposium on Power and Timing Modeling, Optimization and Simulation, PATMOS 2017, Thessaloniki, Greece, 25–27 September 2017; pp. 1–8, doi:10.1109/PATMOS.2017.8106950. [[CrossRef](#)]
16. De Sensi, D. Predicting Performance and Power Consumption of Parallel Applications. In Proceedings of the 24th Euromicro International Conference on Parallel, Distributed, and Network-Based Processing, PDP 2016, Heraklion, Greece, 17–19 February 2016; pp. 200–207, doi:10.1109/PDP.2016.41. [[CrossRef](#)]
17. Aalsaud, A.; Shafik, R.; Rafiev, A.; Xia, F.; Yang, S.; Yakovlev, A. Power-Aware Performance Adaptation of Concurrent Applications in Heterogeneous Many-Core Systems. In Proceedings of the 2016 International Symposium on Low Power Electronics and Design, San Francisco, CA, USA, 8–10 August 2016; ACM: New York, NY, USA, 2016; pp. 368–373, doi:10.1145/2934583.2934612. [[CrossRef](#)]
18. Gustafson, J.L. Reevaluating Amdahl’s Law. *Commun. ACM* **1988**, *31*, 532–533. [[CrossRef](#)]
19. Amdahl, G.M. Validity of the single processor approach to achieving large scale computing capabilities. In Proceedings of the AFIPS 1967 Spring Joint Computer Conference, Atlantic City, NJ, USA, 18–20 April 1967; pp. 483–485.
20. Woo, D.H.; Lee, H.H.S. Extending Amdahl’s law for energy-efficient computing in the many-core era. *Computer* **2008**, *41*, 24–31, doi:10.1109/MC.2008.494. [[CrossRef](#)]
21. Herbert, S.; Marculescu, D. Analysis of dynamic voltage/frequency scaling in chip-multiprocessors. In Proceedings of the 2007 International Symposium on Low Power Electronics and Design (ISLPED ’07), Portland, OR, USA, 27–29 August 2007; pp. 38–43, doi:10.1145/1283780.1283790. [[CrossRef](#)]
22. 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, doi:10.1109/92.845896. [[CrossRef](#)]
23. Singh, A.K.; Shafique, M.; Kumar, A.; Henkel, J. Mapping on multi/many-core systems: Survey of current and emerging trends. In Proceedings of the 2013 50th ACM/EDAC/IEEE Design Automation Conference (DAC), Austin, TX, USA, 29 May–7 June 2013; pp. 1–10, doi:10.1145/2463209.2488734. [[CrossRef](#)]
24. Donyanavard, B.; Mück, T.; Sarma, S.; Dutt, N. SPARTA: Runtime task allocation for energy efficient heterogeneous manycores. In Proceedings of the 2016 International Conference on Hardware/Software Codesign and System Synthesis (CODES+ISSS), Pittsburgh, PA, USA, 1–7 October 2016; pp. 1–10.
25. Sarma, S.; Muck, T.; Bathen, L.A.D.; Dutt, N.; Nicolau, A. SmartBalance: A sensing-driven Linux load balancer for energy efficiency of heterogeneous MPSoCs. In Proceedings of the 2015 52nd ACM/EDAC/IEEE Design Automation Conference (DAC), San Francisco, CA, USA, 8–12 June 2015; pp. 1–6, doi:10.1145/2744769.2744911. [[CrossRef](#)]
26. Kim, M.; Kim, K.; Geraci, J.R.; Hong, S. Utilization-aware load balancing for the energy efficient operation of the big.LITTLE processor. In Proceedings of the Design, Automation & Test in Europe Conference & Exhibition (DATE), Dresden, Germany, 24–28 March 2014; pp. 1–4, doi:10.7873/DATE2014.236. [[CrossRef](#)]
27. Endrei, M.; Jin, C.; Dinh, M.N.; Abramson, D.; Poxon, H.; DeRose, L.; de Supinski, B.R. A Bottleneck-Centric Tuning Policy for Optimizing Energy in Parallel Programs. In Proceedings of the Parallel Computing is Everywhere. In Proceedings of the International Conference on Parallel Computing, ParCo 2017, Ologna, Italy, 12–15 September 2017; pp. 265–276, doi:10.3233/978-1-61499-843-3-265. [[CrossRef](#)]
28. Balaprakash, P.; Tiwari, A.; Wild, S.M. Multi objective optimization of HPC kernels for performance, power, and energy. *Lect. Notes Comput. Sci. (Incl. Subser. Lect. Notes Artif. Intell. Lect. Notes Bioinform.)* **2014**, *8551*, 239–260, doi:10.1007/978-3-319-10214-6_12. [[CrossRef](#)]
29. Sen, R.; Wood, D.A. Pareto Governors for Energy-Optimal Computing. *ACM Trans. Archit. Code Optim.* **2017**, *14*, 1–25, doi:10.1145/3046682. [[CrossRef](#)]

30. Bailey, P.E.; Marathe, A.; Lowenthal, D.K.; Rountree, B.; Schulz, M. Finding the Limits of Power-constrained Application Performance. In Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis, Austin, TX, USA, 15 November 2015; ACM: New York, NY, USA, 2015; pp. 1–12, doi:10.1145/2807591.2807637. [CrossRef]
31. Coutinho, D.A.M.; Georgiou, K.; Eder, K.I.; Nunez-Yanez, J.; Xavier-de-Souza, S. Performance and Energy Efficiency Trade-Offs in Single-ISA Heterogeneous Multi-Processing for Parallel Applications. In Proceedings of the 2019 IFIP/IEEE 27th International Conference on Very Large Scale Integration (VLSI-SoC), Cuzco, Peru, 6–9 October 2019; pp. 232–233.
32. Morokoff, W.J.; Cafilisch, R.E. Quasi-Random Sequences and Their Discrepancies. *SIAM J. Sci. Comput.* **1994**, *15*, 1251–1279, doi:10.1137/0915077. [CrossRef]
33. Hill, M.D.; Marty, M.R. Amdahl’s Law in the Multicore Era. *Computer* **2008**, *41*, 33–38, doi:10.1109/MC.2008.209. [CrossRef]
34. Barros, C.; Silveira, L.; Valderrama, C.; Xavier-de Souza, S. Optimal processor dynamic-energy reduction for parallel workloads on heterogeneous multi-core architectures. *Microprocess. Microsystems* **2015**, *39*, 418–425, doi:10.1016/j.micpro.2015.05.009. [CrossRef]
35. Morad, T.; Weiser, U.; Kolodny, A.; Valero, M.; Ayguade, E. Performance, Power Efficiency and Scalability of Asymmetric Cluster Chip Multiprocessors. *IEEE Comput. Archit. Lett.* **2006**, *5*, 4, doi:10.1109/L-CA.2006.6. [CrossRef]
36. Furtunato, A.F.A.; Georgiou, K.; Eder, K.; de Souza, S.X. When parallel speedups hit the memory wall. *arXiv* **2019**, arXiv:1905.01234.
37. Chandrakasan, A.P.; Brodersen, R.W. Minimizing power consumption in digital CMOS circuits. *Proc. IEEE* **1995**, *83*, 498–523, doi:10.1109/5.371964. [CrossRef]
38. Alonso, P.; Dolz, M.F.; Mayo, R.; Quintana-Ortí, E.S. Modeling power and energy of the task-parallel Cholesky factorization on multicore processors. *Comput. Sci. R D* **2014**, *29*, 105–112, doi:10.1007/s00450-012-0227-z. [CrossRef]
39. Kim, N.S.; Austin, T.; Blaauw, D.; Mudge, T.; Flautner, K.; Hu, J.S.; Irwin, M.J.; Kandemir, M.; Narayanan, V. Leakage Current: Moore’s Law Meets Static Power. *Computer* **2003**, *36*, 68–75, doi:10.1109/MC.2003.1250885. [CrossRef]
40. Suleiman, D.; Ibrahim, M.; Hamarash, I. Dynamic voltage frequency scaling (DVFS) for microprocessors power and energy reduction. In Proceedings of the 4th International Conference on Sustainable Energy and Environmental Engineering (ICSEEE 2015), Shenzhen, China, 20–21 December 2005; pp. 5–9.
41. Venkatchalam, V.; Franz, M. Power reduction techniques for microprocessor systems. *ACM Comput. Surv.* **2005**, *37*, 195–237, doi:10.1145/1108956.1108957. [CrossRef]
42. Cardoso, J.M.; Coutinho, J.G.F.; Diniz, P.C. Chapter 2—High-performance embedded computing. In *Embedded Computing for High Performance*; Cardoso, J.M., Coutinho, J.G.F., Diniz, P.C., Eds.; Morgan Kaufmann: Boston, MA, USA, 2017; pp. 17–56, doi:10.1016/B978-0-12-804189-5.00002-8. [CrossRef]
43. Tang, A.; Sethumadhavan, S.; Stolfo, S.; Tang, A.; Stolfo, S. CLKSCREW: Exposing the Perils of Security-Oblivious Energy Management. In Proceedings of the 26th USENIX Security Symposium, Vancouver, BC, Canada, 16–18 August 2017; pp. 1057–1074.
44. 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 EIT 2013, Rapid City, SD, USA, 9–11 May 2013; pp. 1–5, doi:10.1109/EIT.2013.6632716. [CrossRef]
45. ODROID-XU3 - Hardkernel. Available online: <https://www.hardkernel.com/shop/odroid-xu3/> (accessed on 12 April 2019).
46. Cpuset. Available online: <https://www.kernel.org/doc/Documentation/cgroup-v1/cpusets.txt> (accessed on 15 April 2019)
47. Stress-ng. Available online: <https://github.com/ColinIanKing/stress-ng> (accessed on 15 April 2019)
48. Bienia, C. Benchmarking Modern Multiprocessors. Ph.D. Thesis, Princeton University, Princeton, NJ, USA, 2011.
49. Suite, P.T. Open-Source, Automated Benchmarking. Available online: <https://www.phoronix-test-suite.com/> (accessed on 17 April 2019).

50. Che, S.; Boyer, M.; Meng, J.; Tarjan, D.; Sheaffer, J. W.; Lee S.; Skadron, K. Rodinia: A benchmark suite for heterogeneous computing. In Proceedings of the 2009 IEEE International Symposium on Workload Characterization (IISWC), Austin, TX, USA, 4–6 October 2009; pp. 44–54, doi:10.1109/CASES.2013.6662519. [CrossRef]
51. OpenMP. Available online: www.openmp.org (accessed on 15 April 2019)
52. Team, O. Available online: <https://opencv.org/opencv/> (accessed on 28 October 2019).
53. Corporation N. Available online: <https://developer.nvidia.com/cuda-zone>, (accessed on 28 October 2019).
54. Vivek Kale. Loop Scheduling in OpenMP. Available online: https://www.openmp.org/wp-content/uploads/SC17-Kale-LoopSchedforOMP_BoothTalk.pdf (accessed on 28 October 2019).
55. Goodrum, M.A.; Trotter, M.J.; Aksel, A.; Acton, S.T.; Skadron, K. Parallelization of particle filter algorithms. *Lect. Notes Comput. Sci. (Incl. Subser. Lect. Notes Artif. Intell. Lect. Notes Bioinform.)* **2012**, *6161*, 139–149, doi:10.1007/978-3-642-24322-6_12. [CrossRef]
56. Özisikyilmaz, B.; Narayanan, R.; Zambreno, J.; Memik, G.; Choudhary, A. An architectural characterization study of data mining and bioinformatics workloads. In Proceedings of the 2006 IEEE International Symposium on Workload Characterization, IISWC-2006, San Jose, CA, USA, 25–27 October 2006; pp. 61–70, doi:10.1109/IISWC.2006.302730. [CrossRef]
57. Che, S.; Sheaffer, J.W.; Boyer, M.; Szafaryn, L.G.; Wang, L.; Skadron, K. A characterization of the Rodinia benchmark suite with comparison to contemporary CMP workloads. In Proceedings of the IEEE International Symposium on Workload Characterization, IISWC'10, Atlanta, GA, USA, 2–4 December 2010, doi:10.1109/IISWC.2010.5650274. [CrossRef]



© 2020 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 (<http://creativecommons.org/licenses/by/4.0/>).