

Accurate and Stable Empirical CPU Power Modelling for Multi- and Many-Core Systems

Matthew J. Walker*, Stephan Diestelhorst†, Geoff V. Merrett* and Bashir M. Al-Hashimi*

*University of Southampton †Arm Ltd.

Abstract—Modern processors must provide an increasing level of performance, and are therefore including higher numbers of Heterogeneous Multi-Processing (HMP) elements. Intelligent run-time control of performance and power consumption is required to extend battery-life in mobile systems, reduce energy and cooling costs in data centres, and increase peak performance while respecting thermal and power constraints. Accurate online power estimation is essential in guiding run-time power management mechanisms and energy-aware scheduling decisions. We present a statistically-rigorous methodology for developing accurate and stable run-time power models and we experimentally demonstrate their ability to perform more accurately across a wider range of workloads. We highlight significant shortcomings in existing techniques and present an improved model formulation that also accounts for thermal effects. Moreover, we present the *Powmon*¹ software tools that automates our methodology, allowing power models to be developed for other platforms.

Accurate performance and power modelling is also essential in full-system simulation. We present the *GemStone*² open-source software tool, which automates the process of characterising hardware platforms; identifying sources of error in *gem5* performance models using machine learning techniques; applying the empirical power models to simulation data; and quantifying the effect of simulation errors on the performance, power and energy estimations, including their scaling across Dynamic Voltage-Frequency Scaling (DVFS) levels and HMP core types.

The presented work enables the development and implementation of smart run-time power management and energy-aware scheduling algorithms, as well as hardware-validated performance, power and energy simulation for design-space exploration and optimisation of future systems.

I. INTRODUCTION

Online power estimations are fundamental in effective control of power management policies and energy-aware scheduling (EAS) [1], which is required to improve energy consumption, extend lifetime reliability, and maximise peak performance, when required, while respecting thermal and power budgets.

Performance Monitoring Counters (PMCs), which are registers inside the CPU that count architectural and micro-architectural events, have been shown to be effective for estimating CPU power consumption with an Ordinary Least Squares (OLS) regression model [2]. The low overhead of accessing PMCs and low computational complexity of implementing the linear models makes this approach well-suited for complex systems utilising many HMP cores. Software packages have made it easy to calculate coefficients, however, such tools are often misused and assumptions of the algorithms

not respected. An online power model must make accurate estimations across a wide range of workloads and workload phases, and the most important attribute of a power model is therefore having stable model coefficients.

We identify key shortcomings in typical approaches and experimentally demonstrate the importance of model stability when validating across a large, diverse set of workloads. We focus on several key aspects of the methodology: identifying optimum input features, reducing multicollinearity between them and experimentally demonstrating their effect on model stability (Section III); correctly formulating the model using knowledge of CPU power consumption and demonstrating its effectiveness (Section IV); and adding thermal compensation to the power model (Section V).

Energy analysis is also required in design-space exploration where full-system modelling frameworks, such as *gem5*, are typically used. The stable empirical power models can be used as accurate reference models, with known and trusted accuracy. However, performance simulator models (estimating the execution time and modelled PMCs) have the problem of *specification error* [3], meaning that there are potentially significant sources of error preventing models responding in a representative manner to a proposed change. Section VI briefly presents a methodology for comparing full-system models to a hardware platform, using statistical and machine learning techniques to identify sources of error, and applying empirical PMC power models for energy analysis.

II. EXPERIMENTAL SETUP

The Hardkernel ODROID-XU3 development board, which contains a quad-core Arm Cortex-A15 CPU and a quad-core Cortex-A7 CPU, is used to demonstrate our approach. For the purpose of brevity, we will only consider the Cortex-A15 cluster in this paper. Critical to developing accurate power models is obtaining consistent experimental data across the full set of DVFS levels, cores, and across a large set of workloads. Obtaining PMC events on a mobile platform often has challenges and so we present the *Powmon* software tools that collects PMCs and automates the running of experiments.

III. FEATURE SELECTION

The selection of features (model inputs) is critical to the stability of the model coefficients. Only seven PMC events can be simultaneously monitored on our platform and care must be taken to ensure the maximum amount of information useful to predicting power consumption is obtained, without duplicating any information, which causes multicollinearity

¹See <http://powmon.ecs.soton.ac.uk>

²See <http://gemstone.ecs.soton.ac.uk>

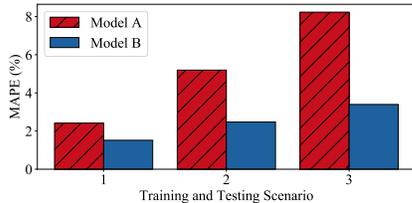


Fig. 1. Demonstration of stability on MAPE with three different training and testing methods

and inflated coefficients. Multicollinearity does not necessarily degrade the fit to the training data, but can reduce the accuracy of the model when tested on observations outside of the training set. We use the *Variance Inflation Factor* (VIF) to measure multicollinearity. The VIF indicates how much the variance (square of standard error) of an input feature has been inflated due to the presence of multicollinearity, compared to the situation where no multicollinearity is present.

We use Hierarchical Cluster Analysis (HCA) to group similar behaving PMC events together and then analyse how each event in a cluster correlates to the power consumption. We propose an automated method of choosing PMC events using a forward stepwise selection process that considers the coefficient of determination (R^2 , indicating goodness-of-fit), the p -values (indicating statistical significance), and the VIF to select events. Once the events have been chosen, the VIFs are further analysed to guide transformations between the events to further reduce multicollinearity (full methodology presented in [2]). Without this last transformation step, only four PMC events can be chosen in our Cortex-A15 before the VIF rises to an unacceptable level. After the transformations have been made, all seven events can be used and the model Mean Absolute Percentage Error (MAPE) is reduced from $> 5\%$ to $< 3\%$.

We experimentally demonstrate how the selection of PMC events affects the model stability, which can be observed by plotting the MAPE when different training and testing workload sets are used (Fig. 1). When training and tested with 20 workloads from a benchmark suite (Scenario 1), a (optimistic) low MAPE is achieved. When trained and tested on a more diverse set of workloads (Scenario 2), the PMC event selection of Model A is not able to capture this diversity as well but a reasonable (and optimistic) MAPE is achieved. However, when testing on a larger set of workloads, including ones that are not present in the training set, Model A performs poorly, while the stable coefficients of Model B, enable it to achieve low error of $< 3\%$. Furthermore, the maximum error of Model A is $> 45\%$ while the maximum error of Model B is $< 15\%$.

IV. FORMULATION

A fundamental assumption of a linear regression model is correct specification of the model. Previous works insert features into OLS solvers without considering the relationship between them and the power consumption (e.g. (1), adapted from [4]). This subsequently leads to the incorrect conclusion

that an accurate model using a single set of coefficients (with voltage, V , and frequency, f , inputs) for every DVFS level is not achievable and that a *per-frequency* model is required. However, the equation used does not correctly specify how the inputs relate to the power consumption, and the model will therefore perform poorly compared to a *per-frequency* model.

$$P = const. + \beta_1 V + \beta_2 f + \beta_3 T + \beta_4 IPC + \beta_5 \frac{INT}{No.Inst.} + \beta_6 \frac{VPF}{No.Inst.} + \dots + \beta_{15} SoftIRQ \quad (1)$$

Our power model is formulated using knowledge of CMOS (complementary metal-oxide-semiconductor) power consumption and breaks down the static power and dynamic power:

$$P_{cluster} = \underbrace{\left(\sum_{n=0}^{N-1} \beta_n E_n V^2 f \right)}_{\text{dynamic activity}} + \underbrace{\beta_b V^2 f}_{\text{BG dynamic}} + \underbrace{g(V, f)}_{\text{static}} \quad (2)$$

where E_0 to E_{N-1} are the chosen PMC event *rates* (and are also divided by f), β_0 to β_{N-1} , β_b are estimated coefficients and *BG dynamic* is a constant dynamic power component.

Each component of the power model is found to be statistically significant ($p < 0.0001$ for all components) and the total 10-fold cross-validated MAPE across all 2160 observations is 2.8%. To further demonstrate the benefit of coefficient stability and carefully formulating the model, we train the model with half the number of workloads (30) and only run the workloads at a single DVFS level (collecting single observations at the others). Training time is reduced from 40 hours to 25 minutes and a MAPE of 3.8% across all 2160 is achieved.

Other overlooked model development stages include inspection of the residuals, which identify the inherent problem of heteroscedasticity in CPU power modelling, which we address using a heteroscedasticity-consistent estimator. The model formulation and high coefficient stability mean that each individual coefficient of the model accurately understands how it uniquely contributes to the total power consumption, enabling high accuracy across a wide range of diverse observations, even if they are not captured in the training set (Fig. 2). The *Powmon* software tools implementing the full methodology, raw experimental data, and an online results visualiser are available at <http://powmon.ecs.soton.ac.uk>.

V. THERMAL COMPENSATION

The $g(V, f)$ term in the model equation (2) includes polynomial values of V and f to absorb the affects of temperature on the static power consumption due to increased voltage and switching speeds at different DVFS levels. However, it is possible to add data from on-board temperature sensors to account for changes in ambient temperature. We achieve this by first removing existing thermal compensation components from the model and analysing the residuals to derive an equation relating the error to the temperature sensor data. We apply this to the power model equation and achieve a MAPE of 3.7% across 45 workloads, 8 DVFS levels and three different thermal environments (CPU temperature variation between $31^\circ C$ and $91^\circ C$) [6].

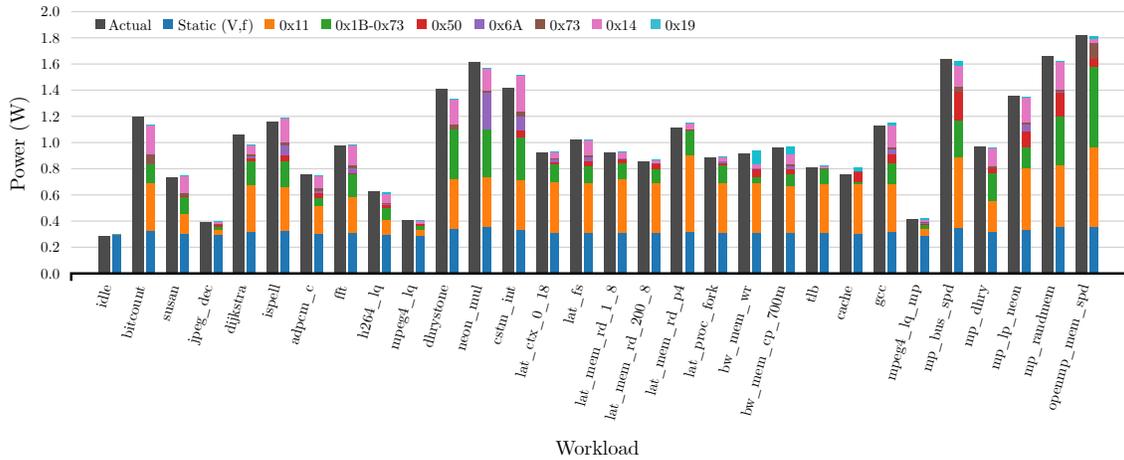


Fig. 2. Actual (measured) power vs. the predicted power for half of the testing workloads, with the predicted dynamic activity power broken down into its constituting parts (Cortex-A15 CPU model). Hex values in the key indicate the event IDs of chosen events (see [5]).

VI. FULL-SYSTEM SIMULATION

Full-system simulation tools, such as gem5, are used extensively to evaluate new research ideas and proposals. With energy-efficiency becoming a primary design constraint, they are often used in conjunction with a power simulation framework, such as McPAT (a power, area and timing framework for multi- and many-core architectures). While these tools provide design flexibility, they are known to contain significant sources of error which can impact the results and conclusions drawn from works of research. In many applications, a researcher or system designer requires a baseline, or reference, model on which to implement the proposed (hardware or software) change. If the baseline model is not accurate, it may not respond in a representative way to the changes under test. Existing works have identified that *specification error* to be the most significant contributor [3], which is caused by a lack of (publicly available) detailed knowledge on the CPUs being modelled, preventing correct model parameters from being set. We first validated existing models against a hardware platform and identified a MAPE of 59%, which motivated the development of a methodology that uses statistical and machine learning techniques to identify the key sources of error in gem5 models, without the need for detailed CPU specifications [7]. Our methodology identified a branch predictor problem to be the key source of error, which reduced to 18% once fixed. Other smaller sources of error were also identified (TLB hierarchy, classification of floating-point and SIMD operations, and how the L1I cache is accessed). The open-source GemStone tool implements the methodology, allowing gem5 models to be improved, extended to other CPUs, validated after changes, and applicability tested for specific use-cases.

The empirical power models are adapted for use in gem5 and the GemStone tool also allows Powmon models to be applied to gem5 simulations results. GemStone also evaluates the effect of errors in the gem5 model on performance, power and energy, including how they scale across DVFS level and different core types. Hardware-validated, empirical power

models provide significantly higher accuracy on baseline CPU models over more flexible power modelling frameworks.

VII. CONCLUSION

We have presented a methodology for developing accurate and stable empirical power models. We demonstrate the methodology on a quad-core Arm Cortex-A15 CPU and achieve a MAPE of 2.8%. We extend this model to include thermal compensation and achieve a MAPE of 3.7% when testing across different ambient temperature conditions. We presented a methodology for identifying sources of error in full-system simulators and applying the empirical power models for hardware-validated performance, power and energy analysis of multi- and many-core systems. Furthermore, we have presented the Powmon and GemStone software tools that automate the power modelling and full-system error identification methodologies, respectively.

ACKNOWLEDGEMENTS

This work was supported by Arm Ltd. and EPSRC Grant EP/K034448/1 (the PRIME Programme)

REFERENCES

- [1] Arm Developer, “Energy aware scheduling,” 2018, [Online; accessed 08-March-2018]. [Online]. Available: <https://developer.arm.com/open-source/energy-aware-scheduling>
- [2] M. J. Walker, S. Diestelhorst, A. Hansson, A. K. Das, S. Yang, B. M. Al-Hashimi, and G. V. Merrett, “Accurate and stable run-time power modeling for mobile and embedded CPUs,” *IEEE TCAD*, vol. 36, no. 1, pp. 106–119, Jan 2017.
- [3] A. Butko, F. Bruguier, A. Gamati, G. Sassatelli, D. Novo, L. Torres, and M. Robert, “Full-system simulation of big.LITTLE multicore architecture for performance and energy exploration,” in *IEEE Int. Symp. Embedded Multicore/Many-core Systems-on-Chip*, Sept 2016, pp. 201–208.
- [4] K. Nikov, J. L. Nunez-Yanez, and M. Horsnell, “Evaluation of hybrid run-time power models for the arm big.little architecture,” in *2015 IEEE 13th International Conference on Embedded and Ubiquitous Computing*, Oct 2015, pp. 205–210.
- [5] Arm Ltd., “Cortex-A15 MPCore, r3p3,” 2013.
- [6] M. J. Walker, S. Diestelhorst, A. Hansson, D. Balsamo, G. V. Merrett, and B. M. Al-Hashimi, “Thermally-aware composite run-time CPU power models,” July 2016.
- [7] M. Walker, S. Bischoff, S. Diestelhorst, G. Merrett, and B. Al-Hashimi, “Hardware-validated cpu performance and energy modelling,” in *IEEE ISPASS*, April 2018.