# Accurate Energy Modelling on the Cortex-M0 Processor for Profiling and Static Analysis

Kris Nikov, Kyriakos Georgiou, Zbigniew Chamski, Kerstin Eder<sup>\*</sup> and Jose Nunez-Yanez<sup>†</sup>

January 31, 2023

#### Abstract

Energy modelling can enable energy-aware software development and assist the developer in meeting an application's energy budget. Although many energy models for embedded processors exist, most do not account for processor-specific configurations, neither are they suitable for static energy consumption estimation. This paper introduces a set of comprehensive energy models for Arm's Cortex-M0 processor, ready to support energy-aware development of edge computing applications using either profiling- or static-analysis-based energy consumption estimation. We use a commercially representative physical platform together with a custom modified Instruction Set Simulator to obtain the physical data and system state markers used to generate the models. The models account for different processor configurations which all have a significant impact on the execution time and energy consumption of edge computing applications. Unlike existing works, which target a very limited set of applications, all developed models are generated and validated using a very wide range of benchmarks from a variety of emerging IoT application areas, including machine learning and have a prediction error of less than 5%.

# 1 Introduction

One trillion new Internet of Things (IoT) devices are predicted to reach the market by 2035 [ARM] ushered by the increasingly expanding edge computing market. Typically, IoT devices are not part of a power grid but rather are scattered in the environment and powered by limited energy sources, such as batteries or energy harvesting. Thus, they are mostly based on small embedded processors with a tiny energy footprint, such as the Arm Cortex-M0. This kind of processor is inherently limited in processing power, making edge computing challenging. Developers must apply extreme optimisations to trim down the

<sup>\*</sup>K. Nikov, K. Georgiou, Z. Chamski and K. Eder are with the University of Bristol, UK. (e-mail: firstname.lastname@bristol.ac.uk)

<sup>&</sup>lt;sup>†</sup>J. Nunez-Yanez is with Linköping University, Sweden. (e-mail: jose.nunez-yanez@liu.se)

processing time, memory, and energy consumption of algorithms to enable their execution on such small embedded devices. A trending example is the streaming down of traditional machine learning algorithms to enable their execution on tiny IoT devices [NS].

The burden now lies with the software engineers to develop edge computing applications that can fit on the limited memory of the IoT embedded devices, execute within reasonable timeframes, and run within the available energy budget. Execution time and code size are easy to measure and well understood by the typical software developer. On the contrary, energy consumption information is not readily accessible, and something most software developers never had to account for. For edge computing, however, energy consumption feedback during the applications' development cycle is at least equally important as execution time and code size [GCAG<sup>+</sup>20, GKCE17].

Hardware measurements are the most accurate way of acquiring a program's energy consumption information, but they are not broadly supported by the hardware vendors and not within the know-how of typical software developers. Energy modelling and the integration of energy models into the development toolchains can solve both of these issues [GKCE17]. Once an accurate energy model has been developed for a particular platform, it can be integrated into a toolchain to allow for energy estimations with each compilation.

The literature offers a plethora of energy consumption models for embedded processors [PBH09, BCF11, NNY20, NMW<sup>+</sup>22, YJK<sup>+</sup>20]. For an energy model to be useful to the software developer, it must be able to convey energy consumption information at the source-code level. Thus, Instruction-Set-Architecture-based (ISA) energy models [TMWT96] became the most popular, because modelling at the ISA level allows for attributing energy costs to software components, such as ISA Control Flow Graph (CFG) basic blocks. Although ISA-based energy modelling approaches have benefits, extracting such models is time-consuming and challenging. It requires devising often complex energy measuring procedures to capture the energy consumption of each instruction in the ISA. On the other hand, energy modelling using Performance Monitoring Counters (PMCs), also named hardware event counters, is a more accessible approach compared to ISA-level modelling. It requires measuring the energy consumption of representative programs, collecting execution statistics from PMCs and then deducting energy consumption coefficients via mathematical analysis and machine learning techniques [NNY20, NYNEH20, NMW<sup>+</sup>22].

This paper demonstrates how to build PMC-based models for multiple embeddedprocessor configurations. The models can be used to attribute energy costs to software components and facilitate both profiling-based and static-analysisbased energy consumption estimation, similar to ISA-based models. Our main contributions are:

1. Due to limited support for PMCs on most IoT platforms, we customised an open-source Instruction Set Simulator (ISS) of the Arm thumb ISA, namely the Thumbulator [Thu] to produce accurate execution statistics useful for developing energy models.

- 2. We identified a set of PMCs that are both statically predictable at ISA basic block level and offer an energy consumption estimation error (a Mean Absolute Percentage Error (MAPE) of less than 5%) [Nik22].
- 3. We enhanced Thumbulator to include advanced configurations for the STM32F0xx family of processors [STM]. We tracked the use of the instruction PreFetch buffer (ON/OFF) which increases the efficiency of instruction fetching and the number of CPU WaitStates (0/1) required to correctly perform read operations from Flash memory (mandatory at higher CPU frequencies since flash memory latency is higher than the CPU clock speed).

# 2 Energy Modelling Methodology

### 2.1 Measurement Setup

Our proposed methodology involves a custom measurement set-up to extract energy consumption information from our target platform - the STM32F0-Discovery board a.k.a. the device under test (DUT) - and collating the data with PMC information from **Thumbulator** to obtain the full data used to generate the models. A diagram of the full set-up including the host PC and different components is presented in Figure 1.

We have used a custom measurement board, called MAGEEC [Mag], to intercept and sample the CPU power supply rails of the DUT. The samples are collected at a frequency of 10kHz and then converted to digital values. All of this is controlled via a python module called *pyenergy*, which is also used to flash and run the pre-compiled workloads on the target device. The workloads are compiled for bare-metal execution using GNU-GCC. The *pyenergy* control program runs on a host platform, connected to the measurement set-up via USB. All the physical DUT measurements are saved back on the host device as a series of .csv files.

The PMCs used for platform state characterisation and model generation are obtained using *Thumbulator*. The simulator has been modified to closely match the execution profile and memory set-up of the DUT. Further details about the modifications and the resulting accuracy are presented in Subsection 2.4. The two sets of binaries are required because the simulator does not fully handle access to off-core peripherals, e.g., PLL clock generators; these should be skipped in Thumbulator binaries. However, the same location and alignment of benchmark code for both types of binaries was maintained.

The aim of this work is to develop an accurate CPU model using ISS information, therefore the DUT peripherals and their interaction with the CPU are not included in the energy measurement collection and simulation. Wholesystem modelling is a very important topic, especially for embedded devices and IoT and remains an area for future research.



Fig. 1. Hardware and software harness for energy modelling of the Cortex-M0.

### 2.2 Benchmark Selection

Two sets of benchmarks were used for model characterisation and validation. First, the BEEBS benchmark suite [PHB13]; an open-source embedded-system benchmark suite designed for exploring the performance and energy consumption characteristics of embedded architectures. It features several categories of benchmarks, selected to represent real-world application areas such as Automotive, Consumer and Security. 76 out of the 88 BEEBS benchmarks have been used. The remaining twelve do not fit in the available memory of our STM32F051 target chip on the DUT. The selected benchmarks have a measured energy Coefficient of Variability (CoV) of 2.61, which shows very high heterogeneity. The second set of benchmarks is based on an industrial edge computing application, developed by Irida Labs [IRI]. The application uses a Convolutional Neural Network (CNN) and implements a smart monitoring system that can monitor, in real-time, a car parking lot with multiple parking slots to determine whether a slot is occupied or not. The different layers of the CNN, namely Convolutional, MaxPool, and Full-Connected, were isolated and configured with different hyper-parameters and optimisations, resulting in 154 distinct benchmarks, with a measured energy CoV of 1.31 indicating the diverse nature of the different CNN layers. Overall, a total of 230 benchmarks were used for the training and validation of our energy model with a measured energy CoV of 3.41, further highlighting the diverse profile of the workload set. This number goes significantly beyond the average number of used benchmarks reported for existing energy models of embedded processors [BSE13, RLE15, KCNL08]. In order to avoid over-fitting the model, we use 10-fold cross-validation to evaluate the model performance across a variety of workload configurations. Further details on model training are available in Subsection 2.5.

#### 2.3 PMC-based Code-level Energy Modelling

PMC-based energy consumption estimation models are typically obtained via multi-linear regression analysis, where coefficients,  $\beta_x$ , are determined for each counter,  $C_x$ , to predict the overall energy cost, i.e.,  $E = \sum_x (\beta_x \times C_x) + \alpha$ , with  $\alpha$  being the residual error term. The coefficients  $\beta_x$  are the constants in the energy model that are program independent while the counters  $C_x$  are the variables that depend on the program and its input. For a specific program with known counters, the energy model can be used to estimate the energy consumed during the program's execution.

For static-analysis-based energy consumption estimation, the overall energy consumption estimate of a piece of code is typically constructed from the estimates of the ISA basic blocks of the program [GKCE17]. Thus, a PMC-based energy model can enable energy consumption estimation via static analysis only if the counters used for the modelling and prediction can be statically predicted at the ISA basic block level.

In order to make the model scalable for block-level static analysis we have trained without using an intercept, so the residual is absorbed into the other event weights. This means that at time 0 the energy predicted is zero. We have also used a Non-Negative Least Squares (NNLS) solver to guarantee positive weights for all the events in the final model, thus always guaranteeing predictable energy consumption values from the model at discreet time slices.

# 2.4 Collection of Cortex-M0 Event Counters

Counter	Description		
$C_1$	Executed instructions (no Muls)		
$C_2$	Multiplication instructions - Muls		
$C_3$	Taken branches		
$C_4$	RAM data reads		
$C_5$	RAM writes		
$C_6$	Flash data reads		

Table 1: Statically predictable PMCs for energy-modelling.

The Cortex-M0 is a deeply embedded architecture with minimal resources available on-chip and it does not expose any PMCs. Thus, we modified an opensource ISS, namely Thumbulator [Thu], to extract the necessary event counters for our energy consumption modelling. The modifications wrt. the reference Thumbulator implementation [Thu] included four key aspects:

- Adaptation to reflect the memory organisation as well as the instruction fetch mechanism used in the STM32F0xx processor family.
- Implementation of a range of event counters and the associated reporting mechanism.
- Calibration and improvement of the timing behaviour of the simulation to match the hardware's behaviour.

The modified simulator can be used to simulate any of the processors in the STM32F0xx family [STM] and can collect a large number of event counters that represent various aspects of the architecture's runtime behaviour such as the effective RAM and Flash memory accesses, taken branches, per-opcode instruction execution statistics, and interactions between instruction- and datarelated memory accesses. The execution time model derived from event counts reported by Thumbulator is fully cycle-accurate wrt. hardware execution when the instruction PreFetch buffer is disabled or the WaitState count is 0. When the PreFetch buffer is enabled and the WaitState count is 1, the MAPE of the Thumbulator-based timing prediction is 1.55%. Theoretically this approach of using an ISS can be applied to other vendors or microprocessors, particularly where there is more available documentation about the micro-architectural implementation. This would allow even finer and quicker tuning of the ISS to match the DUT hardware.

Using the available architecture documentation and a series of modelling cycles, we constrained the number of event counters used for the modelling to the set of the counters that have the most significant impact on the energy consumption and are suitable for static analysis. Most notably all these PMCs can be statically predicted from code-block size using architecture models, which makes them suitable for use in energy analysis tools. These counters also yield the highest observed estimation accuracy compared to physical measurements when compared with the retrieved estimations of other event counter combinations. The selected counters are shown in Table 1.

### 2.5 Model Training and Validation

Hardware Config.	Energy Consumption Model [nJ]	Meas. Energy[J]	MAPE [%]
[20, OFF, 0]	$\mathbf{E} = 0.964258 \times C_1 + 1.652455 \times C_2 + 2.091986 \times C_3 + 1.109833 \times C_4 + 0.650563 \times C_5 + 0.633621 \times C_6$	221.4	2.80
[20, OFF, 1]	$\mathbf{E} = 1.282474 \times C_1 + 2.110668 \times C_2 + 2.191545 \times C_3 + 1.185609 \times C_4 + 0.416602 \times C_5 + 1.178991 \times C_6$	274.9	2.97
[20, ON, 0]	$\mathbf{E} = 1.003378 \times C_1 + 1.885309 \times C_2 + 1.802974 \times C_3 + 1.122833 \times C_4 + 0.849223 \times C_5 + 0.475831 \times C_6$	226.38	2.86
[20, ON, 1]	$\mathbf{E} = 0.895879 \times C_1 + 2.185851 \times C_2 + 2.001178 \times C_3 + 1.493364 \times C_4 + 1.076354 \times C_5 + 1.573758 \times C_6$	227.9	3.68
[24, OFF, 0]	$\mathbf{E} = 0.959172 \times C_1 + 1.888565 \times C_2 + 1.357556 \times C_3 + 1.089427 \times C_4 + 0.993145 \times C_5 + 0.562952 \times C_6 \times C_6$	214.62	3.22
[24, OFF, 1]	$\mathbf{E} = 1.178558 \times C_1 + 2.540429 \times C_2 + 2.042475 \times C_3 + 1.190892 \times C_4 + 0.979651 \times C_5 + 0.891088 \times C_6$	264.88	3.16
[24, ON, 0]	$\mathbf{E} = 0.985415 \times C_1 + 1.933276 \times C_2 + 1.448160 \times C_3 + 1.075671 \times C_4 + 1.011891 \times C_5 + 0.617510 \times C_6$	220.03	3.36
[24, ON, 1]	$\mathbf{E} = 0.883755 \times C_1 + 2.156046 \times C_2 + 1.633465 \times C_3 + 1.436556 \times C_4 + 1.152560 \times C_5 + 1.455166 \times C_6 \times C_6$	220.05	4.15
[48, OFF, 1]	$\mathbf{E} = 1.096677 \times C_1 + 2.364495 \times C_2 + 1.627854 \times C_3 + 1.173680 \times C_4 + 0.681475 \times C_5 + 0.652665 \times C_6 \times C_6$	243.44	3.65
[48, ON, 1]	$\mathbf{E} = 0.816331 \times C_1 + 2.014612 \times C_2 + 1.372157 \times C_3 + 1.402116 \times C_4 + 0.835035 \times C_5 + 1.250446 \times C_6$	202.5	4.33

Table 2: Energy models for selected Cortex-M0 hardware configurations – Hardware Configuration Format: [Frequency (MHz), PreFetch (ON/OFF), Wait-State (0/1)] and MAPE: Mean Absolute Percentage Error

When using regression modelling, it is critical to include as broad and representative a training sample as possible in the training phase. This ensures that the model is as generic as possible and can capture a large part of the space being modelled. Thus, instead of splitting our data into predefined training and testing sets, we included all data into the training, and we used k-fold crossvalidation to ensure the retrieved models avoid over-fitting and selection bias. In our case, we used 10-fold cross-validation and we used the  $R^2$  to evaluate the performance of each of the ten models for each of the modelling configuration, shown in Table 2. The 10-fold cross-validation yielded an  $R^2$  mean value of close to 0.99 for all configurations, with a standard deviation of around 0.2%, where an  $R^2$  value close to 1 indicates an excellent prediction. This demonstrates that the counters selected for the model are accurately capturing the energy consumption of a variety of programs. For the final model coefficients and results, all the data points were used in the training.

Energy models for the different hardware configurations and their accuracy are listed in Table 2. For all models the MAPE is less than 5%, compared to hardware energy measurements. Compared to other relevant works our models achieve lower error, while being trained and validated on a much larger variety of benchmarks using only statically predictable events suitable for code-block-level analysis [BSE13, RLE15, KCNL08].

Analysing the calculated model weights for the PMCs across the different hardware configurations shows a high variation, however some interesting general deductions can be made. For example, when the WaitState is 1 there is a higher cost associated with *Flash reads*, due to the fact that the processor stays idle while waiting for data from the memory. Also, the cost for *RAM data reads*  is close to or higher than *RAM writes* and *Flash reads*, because there are more than twice as many *RAM data reads* operations than the other two and the NNLS solver associates a large part of the energy consumption to them, even if the operation itself uses much less energy. Introducing a WaitState clearly increases energy consumption and thus the PMC coefficients for the entire DUT (however it is needed for correct functionality at higher frequencies). When the WaitState is 0, turning on the PreFetch results in slightly higher energy consumption for the frequencies that support WaitState 0. Consequently, when the WaitState is 1 and PreFetch is ON, there is a significantly reduced overall energy consumption with lower model weights for arithmetic PMCs and branches, but higher model weights for data movement PMCs.

# 3 Conclusion and future work

This paper offers an open-source, ready-to-use energy model for the Arm Cortex-M0 processor [Nik22]. The model can be used for profiling-based analysis to accurately estimate the total energy consumption of a program and in static analysis to predict the energy budget of a particular block of code with a MAPE of less that 5%. The models also account for various frequency and flash instructionbuffer configurations of the processor that can significantly affect the execution time and energy consumption of an application. Our customised open-source ISS [Thu] is also readily available to profile the execution time and energy consumption of edge computing applications for any of the STM32F0xx family of processors. This allows developers to choose the hardware configuration that can meet the resource requirements.

# Acknowledgement

This research has been supported by the European Union's Horizon 2020 Research and Innovation Programme under grant agreement No. 779882, Team-Play (Time, Energy and security Analysis for Multi/Many-core heterogeneous PLAtforms).

# References

- [ARM] ARM. A trillion devices A trillion dollars. Accessed: 2020-11-02.
- [BCF11] C. Brandolese, S. Corbetta, and W. Fornaciari. Software energy estimation based on statistical characterization of intermediate compilation code. In Low Power Electronics and Design (ISLPED) 2011 International Symposium on, pages 333–338, Aug 2011.
- [BSE13] Mostafa Bazzaz, Mohammad Salehi, and Alireza Ejlali. An accurate instruction-level energy estimation model and tool for embed-

ded systems. *IEEE transactions on instrumentation and measure*ment, 62(7):1927–1934, 2013.

- [GCAG<sup>+</sup>20] Kyriakos Georgiou, Zbigniew Chamski, Andres Amaya Garcia, David May, and Kerstin Eder. Lost In Translation: Exposing Hidden Compiler Optimization Opportunities. *The Computer Journal*, Aug 2020.
- [GKCE17] Kyriakos Georgiou, Steve Kerrison, Zbigniew Chamski, and Kerstin Eder. Energy transparency for deeply embedded programs. ACM Trans. Archit. Code Optim., 14(1), Mar 2017.
- [IRI] Irida Labs. Accessed: 2020-12-28.
- [KCNL08] Vasilios Konstantakos, Alexander Chatzigeorgiou, Spiridon Nikolaidis, and Theodore Laopoulos. Energy consumption estimation in embedded systems. *IEEE Transactions on instrumentation and measurement*, 57(4):797–804, 2008.
- [Mag] MAGEEC measuring setup. Accessed: 2020-08-14.
- [Nik22] Kris Nikov. Stm32f0discovery cortexm0 energy model data. https://doi.org/10.5523/bris.v5vp60ddmvu02rukwb5vtguy9, 2022.
- [NMW<sup>+</sup>22] Kris Nikov, Marcos Martinez, Simon Wegener, Jose Nunez-Yanez, Zbigniew Chamski, Kyriakos Georgiou, and Kerstin Eder. Robust and accurate fine-grain power models for embedded systems with no on-chip pmu. *IEEE Embedded Systems Letters*, 2022.
- [NNY20] Krastin Nikov and Jose Nunez-Yanez. Intra and inter-core power modelling for single-isa heterogeneous processors. International Journal of Embedded Systems, 12(3):324–340, 2020.
- [NS] Danny Loh Naveen Suda. Machine Learning on Arm Cortex-M Microcontrollers — White paper. Accessed: 2020-11-02.
- [NYNEH20] Jose Nunez-Yanez, Kris Nikov, Kerstin Eder, and Mohammad Hosseinabady. Run-time power modelling in embedded gpus with dynamic voltage and frequency scaling. In Proceedings of the 11th Workshop on Parallel Programming and Run-Time Management Techniques for Many-core Architectures/9th Workshop on Design Tools and Architectures for Multicore Embedded Computing Platforms, pages 1–6, 2020.
- [PBH09] S. Penolazzi, L. Bolognino, and A. Hemani. Energy and Performance Model of a SPARC Leon3 Processor. In Digital System Design, Architectures, Methods and Tools, 2009. DSD '09. 12th Euromicro Conference on, pages 651–656, Aug 2009.

- [PHB13] James Pallister, Simon Hollis, and Jeremy Bennett. Beebs: Open benchmarks for energy measurements on embedded platforms. arXiv preprint arXiv:1308.5174, 2013.
- [RLE15] Priit Ruberg, Keijo Lass, and Peeter Ellervee. Microcontroller energy consumption estimation based on software analysis for embedded systems. In 2015 Nordic Circuits and Systems Conference (NORCAS): NORCHIP & International Symposium on Systemon-Chip (SoC), pages 1–4. IEEE, 2015.
- [STM] STM32F030x4/x6/x8/xC and STM32F070x6/xB advanced ARMbased 32-bit MCUs - Reference Manual. Accessed: 2020-12-28.
- [Thu] The Thumbulator Git repository modified for collecting performance monitoring counters for the STM32F0-Discovery board. branch: prefetch-model, Git-tag: teamplay-D4.5. Accessed: 2020-11-09.
- [TMWT96] V. Tiwari, S. Malik, A. Wolfe, and M. Tien-Chien Lee. Instruction level power analysis and optimization of software. Journal of VLSI Signal Processing Systems for Signal, Image, and Video Technology, 13(2-3):223-238, 1996.
- [YJK<sup>+</sup>20] Yahya H. Yassin, Magnus Jahre, Per Gunnar Kjeldsberg, Snorre Aunet, and Francky Catthoor. Fast and Accurate Edge Computing Energy Modeling and DVFS Implementation in GEM5 Using System Call Emulation Mode. 36(8), may 2020.