**QUEEN'S
UNIVERSITY
BELFAST**

ALEA: Fine-Grain Energy Profiling with Basic Block Sampling

Mukhanov, L., Nikolopoulos, D. S., & de Supinski, B. R. (2015). ALEA: Fine-Grain Energy Profiling with Basic Block Sampling. In *Proceedings of the 24th International Conference on Parallel Architectures and Compilation Techniques (PACT)* (pp. 87) <https://doi.org/doi.org/10.1109/PACT.2015.16>

Published in:

Proceedings of the 24th International Conference on Parallel Architectures and Compilation Techniques (PACT)

Document Version:

Peer reviewed version

Queen's University Belfast - Research Portal:

[Link to publication record in Queen's University Belfast Research Portal](#)

Publisher rights

©2015 ACM

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee."

General rights

Copyright for the publications made accessible via the Queen's University Belfast Research Portal is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The Research Portal is Queen's institutional repository that provides access to Queen's research output. Every effort has been made to ensure that content in the Research Portal does not infringe any person's rights, or applicable UK laws. If you discover content in the Research Portal that you believe breaches copyright or violates any law, please contact openaccess@qub.ac.uk.

ALEA: Fine-grain Energy Profiling with Basic Block Sampling

Lev Mukhanov*, Dimitrios S. Nikolopoulos, and Bronis R. de Supinski

The School of Electronics, Electrical Engineering and Computer Science, Queen’s University Belfast,
Northern Ireland BT7 1NN, United Kingdom

KEY WORDS: energy profiling, sampling, energy efficiency, power measurement, ALEA

Abstract

Energy efficiency is an essential requirement for all contemporary computing systems. We thus need tools to measure the energy consumption of computing systems and to understand how workloads affect it. Significant recent research effort has targeted direct power measurements on production computing systems using on-board sensors or external instruments. These direct methods have in turn guided studies of software techniques to reduce energy consumption via workload allocation and scaling. Unfortunately, direct energy measurements are hampered by the low power sampling frequency of power sensors. The coarse granularity of power sensing limits our understanding of how power is allocated in systems and our ability to optimize energy efficiency via workload allocation.

We present ALEA, a tool to measure power and energy consumption at the granularity of basic blocks, using a probabilistic approach. ALEA provides fine-grained energy profiling via statistical sampling, which overcomes the limitations of power sensing instruments. Compared to state-of-the-art energy measurement tools, ALEA provides finer granularity without sacrificing accuracy. ALEA achieves low overhead energy measurements with mean error rates between 1.4% and 3.5% in 14 sequential and parallel benchmarks tested on both Intel and ARM platforms. The sampling method caps execution time overhead at approximately 1%. ALEA is thus suitable for online energy monitoring and optimization. Finally, ALEA is a user-space tool with a portable, machine-independent sampling method. We demonstrate three use cases of ALEA, where we reduce the energy consumption of a k-means computational kernel by 37%, an ocean modeling code by 33%, and a ray tracing code by 6% compared to high-performance execution baselines, by varying the power optimization strategy between basic blocks.

1 Introduction

Association of energy use with specific software abstractions and components enables the energy-efficient use of computing systems. Numerous energy profiling tools target platforms ranging from sensors, to smartphones, embedded systems, and high-end computing systems. These tools guide software-controlled energy optimization techniques such as dynamic voltage and frequency scaling, thread packing, and concurrency throttling. Emerging algorithmic energy models and metrics [1, 2] for high-level computation and communication abstractions make accurate energy accounting between software abstractions even more pressing.

Prior energy accounting tools can be broadly classified into two categories: Tools that measure energy by directly measuring power using on-board sensors or external instruments [3, 4, 5, 6, 7, 8]; and tools that model energy based on activity vectors of hardware performance counters, kernel event counters,

¹Correspondence to: The School of Electronics, Electrical Engineering and Computer Science, Queen’s University Belfast, Northern Ireland BT7 1NN, United Kingdom

finite state machines, or instruction counters in microbenchmarks [9, 10, 11, 12, 13, 14, 15, 16, 17, 18, 19]. All of these tools can associate energy measurements with software contexts via manual instrumentation, context tracing, or profiling.

Energy accounting tools based on direct power measurement can accurately measure both component-level and system-wide energy consumption, before and after the system’s power supply units. However, the time granularity of the sensors fundamentally limits these tools. State-of-the-art external instruments such as the Monsoon power meter have sampling rates of at most 5 kHz [20]. Some direct energy measurement and profiling tools use instruments with sampling rates as low as 1 Hz [5, 4]. Internal energy and power sensors such as Intel’s RAPL [21] or the sensors commonly found on ARM-based boards [22] have sampling frequencies between 1 and 3 kHz. The coarse granularity of direct power measurements limits their ability to account for the energy consumption of specific instructions or many software components such as basic blocks and most function instances, which typically execute for periods far shorter than the instrument sampling period.

Tools that model energy consumption from activity vectors can break the granularity barrier of direct energy measurements but suffer from several other shortcomings. Their accuracy may be limited and highly dependent on architectural variations between platforms and workload patterns [11, 12, 13, 18, 19]. The tools require extensive training and benchmarking processes that must be repeated per platform and workload, to calibrate platform parameters.

This paper presents a new method that directly measures power consumption in computing systems and accounts for energy consumption of fine-grain code blocks, including basic blocks with execution duration shorter than the minimum power consumption sampling period. We use the term coarse-grain for basic blocks of longer duration. Our energy accounting tool combines the accuracy of direct power measurements with the fine granularity of energy accounting between basic blocks. Our Abstraction-Level Energy Accounting (ALEA) tool uses the systematic sampling of physical power measurements and a probabilistic model to distribute energy between basic blocks of any granularity, while capturing the dynamic execution context of these blocks. ALEA achieves portability through a machine-independent sampling method that abstracts the details of the underlying architecture and power measurement instruments. We demonstrate its accuracy, efficiency and portability on two multicore platforms based on the Xeon Sandy Bridge and Samsung Exynos processors. We validate ALEA with 14 sequential and parallel applications. ALEA’s mean error for coarse-grain basic blocks, as well as for the whole program, is 1.4% on the Sandy Bridge server and 1.9% on the Exynos SoC. ALEA’s mean error for fine-grain basic blocks is 1.6% on the Sandy Bridge server and 3.5% on the Exynos SoC. We use ALEA to demonstrate the correlation between power consumption and cache accesses at the basic block level across our benchmark suite. Finally, we demonstrate three use cases of ALEA, where we reduce the energy consumption of a k-means computational kernel by 37%, an ocean modeling code by 33%, and a ray tracing code by 6% compared to high-performance execution baselines, by varying the power optimization strategy between basic blocks.

The rest of this paper is structured as follows. Section 2 presents related work. Section 3 describes our platforms and their direct energy measurement sensors. Section 4 details our energy sampling and profiling models and the key aspects of their implementation. Section 5 validates ALEA’s energy profiler. Section 6 presents a use case of ALEA in understanding the impact of memory accesses and thread synchronization on energy. Section 7 presents further use cases of ALEA for fine-grain energy optimization in parallel codes. Section 8 summarizes our findings.

2 Related Work

Statistical sampling of the execution context of a running program is an established method for performance profiling [23, 24, 25]. Sampling is also a state-of-the-art method for profiling large-scale data centers [26]. ALEA is the first tool to deploy basic block sampling and power sampling for fine-grain

energy profiling.

Several tools for energy profiling use manual instrumentation to collect samples of hardware event rates from hardware performance monitors (HPMs) [9, 11, 12, 13, 27]. These tools empirically model power consumption as a function of one or more activity rates that attempt to capture the utilization and dynamic power consumption of specific hardware components. HPM-based tools and their models have guided several power-aware optimizations. However, they often estimate power with low accuracy. Further, they rely on architecture-specific training and calibration.

PowerScope [4, 28], an early energy profiling mechanism, profiles mobile systems through direct hardware instrumentation. It samples power consumption, which it attributes to processes and procedures through post-processing. In contrast, ALEA profiles at a finer granularity.

Eprof [18, 19] models hardware components as finite state machines with discrete power states and emulates their transitions to attribute energy use to system calls. JouleUnit [17] correlates workload profiles with external power measurements to derive energy profiles across method calls. JouleMeter [6] uses post-execution event tracing to map measured energy consumption to threads or processes. These tools perform energy accounting at the granularity of functions or system calls, a limitation that ALEA overcomes. Fine-grained energy profiling enables more compile and run time opportunities for power-aware code optimization.

PowerPack [5] uses manual code instrumentation and platform-specific hardware instrumentation for component-level power measurement to associate power samples with functions. NITOS [7] measures energy consumption of mobile device components with a custom instrumentation device. Similarly, LEAP [8] measures energy consumption of code running on networked sensors with custom instrumentation hardware. These tools profile power at the hardware component level, thus capturing the power implications of non-CPU components, such as memories, interconnects, storage and networking devices. ALEA is complementary to these efforts. ALEA’s sampling method can account for energy consumed by any hardware component between basic blocks, while the statistical approach followed in ALEA overcomes the limitations of coarse and variable power sampling frequency in system components.

Other energy profiling tools build instruction-level power models bottom-up from gate-level models, or other hardware models extracted at design time to provide power profiles to simulators and prototyping environments [15, 16]. These inherently static models fail to capture the variability in instruction-level power consumption due to the context in which instructions execute in real programs. Similarly, using microbenchmarks [14] to estimate the energy per instruction (EPI) or per code block based on its instruction mix does not capture the impact of the execution context.

3 Platforms and Energy Measurement

The ALEA energy profiler builds on platform-specific substrates to measure or to model power at a fine granularity based on data constrained by the sampling rate of the underlying power sensors. In this paper we use two distinct platforms for power measurement, one based on Intel’s Running Average Power Limit (RAPL) apparatus on a Xeon Sandy Bridge server and a second based on integrated power sensors on an ARM Exynos board.

On the Sandy Bridge server, we directly measure energy consumption through on-chip energy counters, which we access through the RAPL interface [21]. RAPL allows us to account for the energy consumption of four components: `PKG`, which measures the energy consumed by the processor package, including the multicore processor; `PP0`, which measures the energy consumed by the power plane that powers the cores and the on-chip caches (L1/L2/L3); `PP1`, which measures the energy consumed by the on-chip graphics processor (for client platforms); and `DRAM`, which measures the energy consumed by memory DIMMs.

Client platforms can only access the `PKG`, `PP0` and `PP1` counters, while server platforms can access

Figure 1: Sampling process

Figure 2: Execution of a basic block in a program

the `PKG`, `PP0` and `DRAM` counters. Our Sandy Bridge server includes two Intel Xeon E5-2650 processors with eight cores per processor, 32KB/32KB I/D-Cache per core, 2MB shared L2 cache per 8 cores, and 20MB shared L3 cache per package. The system runs CentOS (release 6.5). The frequency of the system is up to 2 GHz. We disable the processor’s Turbo Boost and Hyperthreading options in our validation experiments.

Our second platform, an ODROID-XU+E board, has one Exynos 5 Octa processor. This ARM Big.LITTLE architecture has four Cortex-A15 cores and four Cortex-A7 cores, 32KB/32KB I/D-Cache per core, NEONv2 floating point support per core, VFPv4 support per core, one PowerVR SGX 544 MP3 GPU, and 2 GBytes of LPDDR3 DRAM. A 2 MByte L2 cache is shared between all Cortex-A15 cores and a 512 KByte L2 cache is shared between all Cortex-A7 cores. The ODROID board also includes power meters on each voltage plane to measure consumption for the following four sets of components: Cortex-A7 cores, including their shared L2 cache; Cortex-A15 cores, including their shared L2 cache; GPU; and DRAM. The system runs Ubuntu 14.04 LTS. In our experiments, we use the Cortex-A15 cores only at their maximum frequency of 1.6 GHz.

4 Profiling

Execution time profiling can use sampling or instrumentation [24, 26]. Compiler or binary instrumentation inserts profiling instructions that track dynamic execution counts and the execution time of code paths, as well as software or hardware events. Profilers based on sampling suspend binary execution to sample the execution state, typically the current program counter and possibly register contents or a stack traceback, and to correlate the sample with software events, hardware events, or metrics.

We use statistical sampling for fine-grained energy profiling and demonstrate that we can probabilistically estimate energy consumption at fine and coarse granularities. Our profiling approach simultaneously samples the currently executing basic block and takes power measurements, which it assigns to the basic block (Figure 1). We perform a one-pass sampling of power measurements during a single program execution. Our tool processes the profiling results off-line, using a probabilistic model to estimate the execution times and the mean power consumption for each basic block.

Figure 3: Random sampling

4.1 Execution time profiling model

To motivate the model, Figure 2 shows the iterative execution of a basic block that is executed k times. The model makes the simplifying assumption that the processor executes instructions from one basic block (bbm) in each clock cycle. The latency of each basic block ($latency_{bbm}^j$) may vary between iterations. For example, a basic block may execute the same load instruction with different latencies between iterations, depending on the level of the memory hierarchy that provides the requested data.

If we sample the program counter once during program execution at a random point in time, we define the random variable X_{bbm} as:

$$X_{bbm} = \begin{cases} 1, & \text{if } bbm \text{ is the sampled basic block} \\ 0, & \text{otherwise} \end{cases} \quad (1)$$

In our probabilistic model, CPU clock cycles (ticks) correspond to the units of the finite population (U) and a sample during a specific clock cycle instantiates X_{bbm} [29]. The probability that bbm is sampled is:

$$p_{bbm} = P(X_{bbm} = 1) = \frac{C_{t_{bbm}}^1}{C_{t_{exec}}^1} = \frac{\sum_{j=1}^k latency_{bbm}^j}{t_{exec}} \quad (2)$$

$$\frac{\sum_{j=1}^k latency_{bbm}^j}{t_{exec}} = \frac{t_{bbm}}{t_{exec}} \quad (3)$$

where t_{bbm} is the total execution time of instances of bbm , t_{exec} is the total execution time of the program, and C_S^1 is a 1-combination of a set S . We measure time in ticks and represent it in seconds by dividing it by the CPU frequency. Equation 2 captures the observation that the probability of sampling a basic block at a random clock cycle is equal to the ratio of its execution time to the program's total execution time. If the probability p_{bbm} and the total execution time are known then t_{bbm} is:

$$t_{bbm} = p_{bbm} \cdot t_{exec} \quad (4)$$

We assume that X_{bbm} follows a Bernoulli distribution because it is binary, random, and p_{bbm} is a constant in our model. By applying random sampling (see Figure 3), we can estimate the probability as the maximum likelihood estimator of parameter p_{bbm} in the Bernoulli distribution for $X_{bbm} = 1$ [30, 31]:

$$\hat{p}_{bbm} = \frac{n_{bbm}}{n} \quad (5)$$

In Equation 5, n_{bbm} is the number of samples of some instruction from bbm , and n is the total number of samples. Thus, we estimate the execution time of any basic block as:

$$\hat{t}_{bbm} = \hat{p}_{bbm} \cdot t_{exec} = \frac{n_{bbm} \cdot t_{exec}}{n} \quad (6)$$

We measure the total execution time t_{exec} of an application during the profiling run.

4.2 Energy profiling model

We apply the same probabilistic approach to profile power and energy. Similarly to the execution time profiling model, we consider power consumption as a random variable (pow , Figure 3) and an implementation of this variable at a clock cycle as a characteristic associated with the clock cycle. We simultaneously take samples of the program counter and power consumption, which we assign to the sampled basic block even though power consumption likely includes power that instructions outside that basic block consume.

Assuming n_{bbm} samples of block bbm , we estimate its mean power consumption as [31]:

$$\hat{pow}_{bbm} = \frac{1}{n_{bbm}} \cdot \sum_{i=1}^{n_{bbm}} pow_{bbm}^i \quad (7)$$

In Equation 7, pow_{bbm}^i is the power consumption associated with the i -th sample of block bbm .

We estimate the energy consumption of bbm as:

$$\hat{e}_{bbm} = \hat{pow}_{bbm} \cdot \hat{t}_{bbm} \quad (8)$$

4.3 Bounds and Confidence

If p_{bbm} is not too close to 0 or 1 and n is relatively large ($n \cdot p_{bbm} > 5$, $n \cdot (1 - p_{bbm}) > 5$) [31], then we can construct the confidence interval with upper and lower bounds on p_{bbm} :

$$\hat{p}_{bbm}^u = \hat{p}_{bbm} + z_{\alpha/2} \sqrt{\frac{1}{n} \cdot \hat{p}_{bbm} \cdot (1 - \hat{p}_{bbm})} \quad (9)$$

$$\hat{p}_{bbm}^l = \hat{p}_{bbm} - z_{\alpha/2} \sqrt{\frac{1}{n} \cdot \hat{p}_{bbm} \cdot (1 - \hat{p}_{bbm})} \quad (10)$$

$$\hat{p}_{bbm}^l \leq p_{bbm} \leq \hat{p}_{bbm}^u \quad (11)$$

In Equations 9, 10 and 11, z_{α} is the $1 - \alpha/2$ percentile of the standard normal distribution, and $1 - \alpha$ is a confidence level. The interval in Equation 11 includes the true value of p_{bbm} with probability $1 - \alpha$. According to Equation 6, by multiplying the lower and upper bounds of p_{bbm} with the total execution time t_{exec} , we obtain an interval in which the true execution time t_{bbm} of bbm lies:

$$\hat{p}_{bbm}^l \cdot t_{exec} \leq t_{bbm} \leq \hat{p}_{bbm}^u \cdot t_{exec} \quad (12)$$

We can similarly build a confidence interval for power [31]:

$$\hat{pow}_{bbm}^u = \hat{pow}_{bbm} + z_{\alpha/2} \frac{s}{\sqrt{n_{bbm}}} \quad (13)$$

$$\hat{pow}_{bbm}^l = \hat{pow}_{bbm} - z_{\alpha/2} \frac{s}{\sqrt{n_{bbm}}} \quad (14)$$

$$s = \sqrt{\frac{1}{n_{bbm} - 1} \cdot \sum_{i=1}^{n_{bbm}} (pow_{bbm}^i - \hat{pow}_{bbm})^2} \quad (15)$$

$$\hat{pow}_{bbm}^l \leq pow_{bbm} \leq \hat{pow}_{bbm}^u \quad (16)$$

where s is the corrected sample standard deviation. Using confidence intervals for execution time and power, we can derive a confidence interval for energy consumption:

$$\hat{p}_{bbm}^l \cdot t_{exec} \cdot \hat{pow}_{bbm}^l \leq e_{bbm} \leq \hat{p}_{bbm}^u \cdot t_{exec} \cdot \hat{pow}_{bbm}^u \quad (17)$$

If we increase the total number of samples, we reduce the width of the confidence intervals as they are inversely proportional to the square root of the number of samples (time: $\sim \frac{const}{\sqrt{n}}$, power: $\sim \frac{const}{\sqrt{n_{bbm}}}$). Thus, the accuracy of the energy estimates should increase with increasing total number of samples (n) and the given basic block samples (n_{bbm}). Because n_{bbm} is strongly correlated with n , the accuracy of the energy estimates is primarily affected by the total number of samples (n).

4.4 Profiling of parallel applications

We employ the same execution time and energy profiling models for multithreaded applications. The essential difference is that each sample is a vector of program counters simultaneously sampled across all threads. Thus, we distribute the execution time and energy across combinations of basic blocks, which are executed on different threads:

$$\hat{t}_{comb} = \hat{p}_{comb} \cdot t_{exec} = \frac{n_{comb} \cdot t_{exec}}{n} \quad (18)$$

$$\hat{pow}_{comb} = \frac{1}{n_{comb}} \cdot \sum_{i=1}^{n_{comb}} pow_{comb}^i \quad (19)$$

$$comb = bb_{thread_1}, bb_{thread_2}, \dots, bb_{thread_l} \quad (20)$$

where $comb$ corresponds to a combination of basic blocks that were sampled on different threads (l threads).

We consider all threads of an application running on the same processor package collectively during sampling, because they share resources and because resource sharing contributes additional energy consumption due to contention between threads. Shared resources include caches, buses and network links, all of which can significantly increase power consumption under contention. We could apportion power between threads based on dynamic activity vectors that measure the occupancy of shared hardware resources per thread [10]. However, these vectors are difficult to collect on real hardware, as current monitoring infrastructures cannot distinguish between the activity of different threads on shared resources. As such, per-thread energy apportioning cannot be accurately validated on real hardware.

We can still correlate power consumption with basic blocks with this approach. For example, we can investigate how the energy profile of a basic block changes between stand-alone execution and execution with different co-runners, to capture contention for shared resources. Further, our methodology helps us understand how synchronization can decrease power consumption, which in turn reveals opportunities for reducing energy consumption in the runtime system by applying dynamic concurrency throttling [27].

4.5 Power measurements

We measure processor power consumption on our Sandy Bridge server for a given sample (pow_{bbm}^i) by dividing the energy consumed since the last sample by the length of the sampling period. Our analysis of sampling overhead and accuracy, which we present in the following sections, led to a 10 ms sampling period. This approach conforms to RAPL, which provides running energy but not power measurements.

Figure 4: Overhead and energy estimate error

Our Exynos platform has TI INA231 power meters, which directly sample power consumption for the system-on-chip averaged over a user-defined period. We used the minimum feasible period on the Exynos, which is 280 microseconds.

In general the sampling period used in our model is different than the platform power sampling period. Our method estimates the energy consumption of basic blocks of any duration, including ones that run for less than the sampling period, under a probabilistic model of the fraction of program execution time that each given basic block consumes and the average power consumption due to execution of that basic block.

4.6 Implications of systematic sampling

Systematic sampling, which approximates random sampling, selects units from an ordered population with the same sampling period. It selects the first unit of a sample randomly from the bounded interval $[1, \text{length of sampling period}]$. We use systematic sampling for time and energy profiling, in which units correspond to CPU clock cycles and the user sets the sampling period [29].

Systematic sampling can be inefficient with populations that exhibit a periodic variation that is an integral multiple of the sampling period. For example, if the same basic block is executed with a period equal to the sampling period then theoretically, we will only sample that basic block. In practice, the precise size of a sampling period in CPU clock cycles varies randomly between samples due to the inaccuracy of the timer and variance in the execution length of the sampling code itself. We find that on the Sandy Bridge and Exynos platforms, the variation in the delay between samples may be up to hundreds of microseconds. This random variation obviates the need to add deliberate randomization during the sampling process.

4.7 Sampling period

The accuracy of our sampling estimates improves with an increasing number of samples. However, sampling incurs overhead, which biases execution time and energy estimates. This overhead increases linearly or superlinearly with the number of samples, since the program must be interrupted for each sample. Thus, the estimation error is composed of random error, which is introduced by sampling, and systematic error, which is introduced by profiling overhead. If we increase the number of samples, then the random error decreases but the systematic error increases.

We use our benchmark suite to capture basic blocks with diverse execution times and power consumption to find the best sampling period in terms of energy estimation accuracy and execution time overhead. As an example, the `streamcluster` benchmark from the Rodinia suite includes basic blocks with latency varying between 1 and 30 ms on the Sandy Bridge platform. Figure 4 shows the trade off between the length of the sampling period, overhead and accuracy of energy estimates for the

Figure 5: Average error in execution time and energy estimates, compared with direct measurements (Sandy Bridge)

Sandy Bridge and Exynos platforms, using both sequential and parallel executions of the benchmark. We observe similar results in all benchmarks, pointing to a sampling period of 10 ms as a good compromise between energy estimation error and runtime overhead. A fixed sampling period helps deployment of ALEA as a continuous, online application energy profiler with capped overhead. However, we can select an application-specific sampling frequency since the tool exposes the sampling interval as a user-defined parameter.

4.8 Implementation

ALEA uses a separate control process to obtain the current instruction pointer of the profiled application and to take power measurements. We use the `ptrace` interface, which allows one process to retrieve the contents of registers in another process or thread. Thus, the profiled program does not execute any additional code, unlike sampling schemes based on signals [23]. Instead, the control process captures context information and energy/power measurements. This approach reduces system overhead because system call interfaces are offloaded from the profiled program’s critical path to the control process. However, this approach still incurs performance and energy overhead because processes or threads of the profiled program are suspended while the control process reads the registers via the `ptrace` interface. ALEA currently executes on a dedicated core that the profiled application does not use.

5 Validation

We use 14 benchmarks (sequential and parallel) from four suites (SPEC 2000, Parsec, Rodinia, SPEC OMP) to validate the accuracy of ALEA’s execution time and energy consumption estimates. We use a range of benchmarks to achieve good coverage of basic block features such as execution time, including fine-grain and coarse-grain blocks, and energy consumption, including blocks with distinct power profiles and/or power variations between their samples. We use the `native` input data set for benchmarks from Parsec and standard input for benchmarks from other suites.

We measure whole program execution time and energy. We also measure the execution time and energy of those basic blocks with latency that exceeds the sampling period (10 ms) in isolation. Further, in isolation, we measure the execution time and energy of fine-grain basic blocks that have shorter latency than the sampling period, but are enclosed in innermost loops such that the overall loop latency exceeds this period. Overall, direct per-basic block measurements covers 81% of the execution time of each benchmark on average. We compare ALEA’s execution time and energy consumption estimates to

per-basic block direct measurements. For basic blocks that are not captured by direct measurements, we compare whole program measurements to the sum of execution time and energy consumption estimates for all basic blocks sampled by ALEA at least once during program execution.

We execute each benchmark at least six times. The first run directly measures energy and time. The other runs use ALEA to estimate the execution time and energy consumption of each basic block. We use at least five ALEA runs and as many more as needed (up to 20 total) to bring the 95% confidence interval of the time, power and energy measurements within 5% of the mean. We compile all benchmarks using gcc with `-O1` and `-ffast-math`, which inlines mathematical and other functions when possible. For validation, we use the `-O1` optimization level instead of `-O3` to increase latencies of some basic blocks to the minimum needed to take direct measurements.

The ALEA profiler executes on a core that is not in use by the profiled application, to minimize interference. Specifically, ALEA runs on a separate Sandy Bridge socket but on the same Exynos four-core Cortex A15 cluster since our Odroid board does not allow co-execution on both of the A15 and A7 clusters. We present results from experiments using up to eight threads on one socket of the Sandy Bridge platform and up to two threads of the A15 cluster on the Exynos platform for the execution of parallel benchmarks. Running the profiler on a separate core keeps the overhead under 1% on both platforms. We also experimented with running the profiler on the same core as one of the threads of each profiled program and observed the overhead to increase to up to 10% (not shown). This overhead can be mitigated by reducing the sampling frequency (Figure 4). Halving the sampling frequency halves the overhead and keeps the ALEA average energy estimation error at a manageable 5% (Exynos) to 6% (Sandy Bridge).

5.1 Sandy Bridge results

Figure 5 presents the average error of ALEA’s execution time and energy consumption estimates for basic blocks on the Sandy Bridge platform. The average error is 1.3% for the execution time estimates and 1.4% for the energy consumption estimates. 99% of the execution time and energy measurements lie within 95% confidence intervals. For those fine-grain basic block sets enclosed in loops that allow us to measure time and energy directly, the average error in ALEA’s energy estimate is 1.6% (1.3% for execution time). For coarse-grain basic blocks, the ALEA profiling error is 1.4% for both execution time and energy consumption. The average errors of the ALEA execution time and energy estimates for parallel benchmarks (Figure 5) are 3.1% and 2.6%. Our average whole program absolute error across all benchmarks is 1.1% for execution time and 1.4% for energy.

5.2 Exynos results

While RAPL supports direct energy measurements on the Sandy Bridge platform, we can only directly measure power on the Exynos platform. We thus follow a different approach to validate energy profiling between basic blocks on it. We again instrument the benchmarks to perform execution time profiling. However, in each instrumented basic block, we sample the power consumption using the system timer and corresponding signal handler. We set the Exynos TI power meters to compute average power over the minimum feasible period of 280 microseconds. This instrumentation has higher overhead than direct energy measurements on the Sandy Bridge platform because it enforces one interrupt per sample. This higher overhead introduces a bias in energy measurements, which leads to higher error.

The average error in ALEA’s energy estimates (not shown due to space limitations) is 2.6% (also 2.6% in execution time estimates) for sequential benchmarks and 3.6% (2.8% in execution time estimates) for parallel benchmarks. 99% of all time and energy measurements lie within 95% confidence intervals. The average error in ALEA’s energy estimate for fine-grain basic blocks is 3.5% (3.7% for execution time) and 1.9% (1.8% for execution time) for coarse-grain basic blocks. The average error of total execution time estimates is 1.4% and that of total energy estimates is 1.9%.

Basic Block A Art(Spec2000) / 10.10 Watts	Basic Block B Heartwall(Rodinia) / 8.80 Watts
<pre> mov rdx,[r8+rax*1] movsd xmm0,[r9+rax*8+0x28] add rax,0x8 cmp rax,rdi mulsd xmm0,[rdx+rsi*1] addsd xmm2,xmm0 movsd QWORD PTR [rcx],xmm2 jne 402a80 <match+0x690> </pre>	<pre> mov ecx,esi sub ecx,eax movsxd rcx,ecx lea edx,[rax+rdi*1] movsxd rdx,edx movss xmm1,[rbx+rcx*4-0x4] mulss xmm1,[rbp+rdx*4-0x4] addss xmm0,xmm1 add eax,0x1 cmp eax,r8d jle 4018cb <kernel+0x796> </pre>

Figure 6: art and heartwall basic blocks (Sandy Bridge)

Figure 7: Power,energy and execution time measurements taken for microbenchmarks

6 Impact of Memory Instructions and Synchronization on Energy

We can optimize a program’s energy consumption by reducing its execution time or power consumption. However, reducing execution time often increases power consumption. We use ALEA to investigate the causes of increased power consumption in optimized programs. Our experiments indicate that the power consumption may vary considerably between basic blocks. Figure 6 shows a basic block from *art* (BBA) and a basic block from *heartwall* (BBB). On the Sandy Bridge platform BBA consumes 10.10W (98.39J in total), while BBB consumes 8.80W (278.63J in total). Our experimental study shows that the power consumption of a basic block is primarily affected by the cache access intensity and does not vary considerably with the type of executed instructions. In our example, BBA accesses approximately 7 MB of data during its execution (which fits in the L3 cache), while BBB accesses only 36KB of data (which fits in the L1 cache). The Exynos platform exhibits similar behavior.

To confirm the effect of cache accesses, we develop microbenchmarks based on BBA. We create a basic block with the same set of instructions and context for both processors. We divide its instructions into two groups: memory access instructions and arithmetic/logic instructions. We use these groups to implement different versions of BBA (Table 1). We then add a basic block with a single *nop* instruction, which does not use the floating point units (FPUs). We limit the size of the accessed data so that the data fits in the L2 cache.

Figure 7 shows the power, execution time and energy measurements for our experimental set of basic blocks on the Sandy Bridge platform (the basic blocks are sorted by power consumption). The *NoMem* and *NoMem(L1,load)* blocks consume almost the same power even though the second block occupies the FPU. In contrast, the difference in power consumption between the *Mem* and *NoMem* blocks is more than 1.5W. Similarly to the Sandy Bridge platform, the *NoMem* and *NoMem(L1,load)* basic blocks show the same power

Block	Description
Basic block A	Copy of BBA
Mem	Only memory access instructions of BBA
NoMem	Only arithmetic/logic instructions of BBA
Mem (L2)	Mem block with the size of accessed data limited to 2MB (L2 cache size on Exynos)
Mem (L1)	Mem block with the size of accessed data limited to 2KB (L1 cache size on Exynos)
Mem (load)	Mem block with load instructions only
Mem (store)	Mem block with store instructions only
Mem (L2, load)	Mem (L2) block with loads only
Mem (L2, store)	Mem (L2) block with stores only
Mem (L1, load)	Mem (L1) block with loads only
Mem (L1, store)	Mem (L1) block with stores only

Table 1: Versions of BBA

consumption on the Exynos platform, while the Mem (L2) block consumes more power than does the NoMem block (Figure 7). Thus, the increase in power consumption on both platforms is primarily due to data cache accesses and not the type of instructions executed.

Even though the NoMem block merely omits the memory access instructions of BBA, these blocks have nearly the same execution time on both platforms because pipelining hides the data access latencies of BBA. Thus, its execution time does not increase despite the energy used for the data accesses. Pipelining can lead to significant errors in energy consumption estimates based on EPI [14], which ALEA mitigates. For example, BBA is a union of instructions from Mem and NoMem blocks. On the Sandy Bridge, according to an EPI model, BBA, which consumes 1,474J, should consume the sum of the energy consumed by Mem (955J) and NoMem (1,245J) blocks, which is 2,200J or over $1.5\times$ more than the actual energy consumption. On the Exynos platform, the energy consumption of BBA is $1.29\times$ less than the sum of energy consumption of the NoMem and Mem blocks.

Our experiments show that the power consumption of basic blocks executed in parallel applications depends on the form of each thread’s activity. For example, the `ammp` (SPEC OMP) benchmark contains a basic block with 564 instructions that correspond to a loop body in the `mm_fv_update_nonbon` procedure (`rectmm.c`, line 1210). This basic block includes regular accesses to caches. When four threads execute this block in parallel, the Sandy Bridge processor consumes 19.07W (1153J). However, if only one thread executes this basic block while the other threads wait in synchronization, power consumption drops to 13.19W (513J). Results on the Exynos platform are similar.

7 Use cases

We present three use cases of how basic block level energy profiling can be used in energy-aware program optimization. Our first use case analyzes hot spots to uncover opportunities for energy optimizations in a single dominant basic block, based on techniques that adapt the degree of parallelism in the program [32, 33, 34, 35]. Our second and third use cases explore fine-grain optimization and power capping opportunities across multiple basic blocks.

7.1 Hotspot energy optimization

Our first use case applies ALEA to optimize hot spot energy use in the `k-means` benchmark of the Rodinia suite using one socket on our Sandy Bridge platform. ALEA runs on one core of the other socket. We scaled up the standard input set $6\times$ to model realistic runs of the benchmark. Profiling of the sequential version shows that 56% of the total execution time is spent on the basic block

Figure 8: Profiling results of k-means (Sandy Bridge)

that corresponds to the loop that calculates the multidimensional spatial Euclidean distance square (`euclid_dist_2` function). We use the `-O3` compilation flag as a default option. However `unroll` and `auto-vectorization` optimizations are, surprisingly, not applied to the basic block. We use compiler hints (C-extensions: parameter and function attributes) to force the compiler to apply unrolling. We also use parameter attributes to align and to restrict pointers so the compiler recognizes the proper context for auto-vectorization. Finally, the `-ffast-math -mavx` flag enables floating-point arithmetic transformations and the use of AVX-256 instructions. We refer to this set of optimizations as *hints*.

Figure 8 shows execution time, power, energy, energy-delay and energy-delay² estimates for the key k-means basic block optimized with `-O3` and with `-O3+hints`. The energy-delay and energy-delay² measurements of the latter version are shown in separate charts to assist the reader, because the optimization hints reduce these metrics by two to three orders of magnitude. We also measure the corresponding metrics for the entire k-means program. Our optimizations reduce execution time of the dominant basic block by up to 8× when running with one or two threads but the impact of these optimizations on performance is less pronounced with more threads, due to memory contention that limits scalability. The speedup of the full benchmark when running with more cores is limited by the significant percentage of sequential execution time spent on I/O operations (up to 55% after optimizations). The optimization hints that significantly accelerate the dominant basic block actually reduce the speedup from using more cores.

The impact of optimizations on energy consumption is considerably different from that on execution time. Power consumption increases disproportionately when optimizations and additional concurrency are applied to the benchmark. Energy consumption is not minimized with the set of optimizations or the degree of concurrency that minimizes execution time. A combination of unrolling, vectorization and maximal concurrency (eight threads) achieves peak performance for the benchmark (18.51 seconds), while energy consumption is minimized with optimizations turned on but using only two cores, at a 20% performance loss. Overall, optimizing the dominant basic block for energy consumption yields 37% energy savings for the entire program, compared to the high-performance baseline (eight cores, `-O3 + hints`).

	Baseline		Energy-optimal				
	Time(s)	Energy (J)	Time (s)	Energy (J)	Threads	Frequency	Manual optimization
bb1,jacobcalc2.C:301	2.03	8.48	1.87	6.03	4	1500 MHz	No
bb2,slave2.C:641	1.54	6.70	1.31	4.16	2	1600 MHz	Yes
bb3,laplacalc.C:83	2.02	9.53	2.55	7.98	2	1500 MHz	No
bb4,multi.C:253	2.17	7.22	2.62	6.52	2	1500 MHz	No
bb5,multi.C:235	2.36	7.88	3.29	5.56	1	1500 MHz	No
bb6,multi.C:290	2.67	9.23	3.23	5.46	1	1500 MHz	No
program	29.93	108.64	26.88	72.84	2.0 (avg.)	1516 MHz (avg.)	Yes

Table 2: Time and energy impact of basic-block level optimization for `ocean_cp` on Exynos

The `k-means` example exhibits clear trade-offs between performance and energy consumption. Optimization criteria that place heavier emphasis on performance (execution time, energy-delay²), when applied to the dominant basic block, indicate preference for the highest concurrency and manual code optimization via hints. Optimization criteria that place heavier emphasis on power and energy opt for lower concurrency. Further, we should apply a different optimization strategy for the whole of the program, compared to the strategy followed for the dominant basic block (see *EDP* and *ED2P* in Figure 8, configurations are annotated). This result motivates fine-grain energy accounting.

7.2 Fine-grain power optimization across basic blocks

We use the `ocean_cp` benchmark from the PARSEC suite to explore whether ALEA exposes different energy optimizations for basic blocks in the same code, in order to achieve better whole-program energy-efficiency. Such an optimization strategy would motivate ALEA’s fine-grain profiling. We use the `native` input data set and modify the time between relaxations to increase the overall execution time of the benchmark in order to achieve stable and repeatable results. Time profiling of `ocean_cp` indicates that more than 50% of the total execution time is spent executing six basic blocks (Table 2), to which we refer as `bb1` through `bb6`. We initially compile this benchmark for highest performance using the flags: `-O3, -mfpu=neon-vfpv4, -mtune=cortex-a15, -ffast-math, -funroll-loops, -ftree-vectorize, -fprefetch-loop-arrays`.

Motivated by our experimental analysis of the power implications of memory instructions (Section 6), we disable optimizations that could increase cache access rates to reduce power. The disabled optimizations are prefetching, for `bb3`, and the combination of unroll and vectorization, for `bb1` and `bb2`. By disabling these optimizations for those basic blocks, we reduce power consumption by up to 14% for `bb2`, 10% for `bb1`, and 4% for `bb3`. Further code inspection of `bb4`, `bb5` and `bb6` reveals that the compiler inserts additional stack access instructions before each of these basic blocks, due to the predictive commoning optimization, which has no effect on performance, but increases power consumption. By disabling this optimization we reduce power consumption for these three basic blocks by between 3% to 10%.

Table 2 shows selected results from an experimental campaign to understand how to minimize the energy consumption of the six dominant basic blocks in `ocean_cp`. The baseline for this campaign is execution of the code using the maximum number of cores on an Exynos cluster (four) and the maximum frequency (1600 MHz). Besides execution time and energy of the baseline case, we show execution time and energy of the energy-optimal configuration, as well as details of the program and system configurations that achieve energy minimization, including clock frequency, number of threads and use or no use of the three manual power optimizations considered: unrolling, vectorization and predictive commoning.

The table reveals several findings that motivate the ALEA approach to fine-grain profiling. First, fine-grain energy optimization at the basic block level yields substantial energy savings, ranging from 10% for `bb4` to 41% for `bb6`; and 33% for the program as a whole compared to the baseline. Second, the factor that

catalyzes energy minimization varies between basic blocks: most basic blocks are more energy-efficient when running at slightly lower than the maximum frequency (1500 vs. 1600 MHz); most basic blocks run most efficiently with one or two, not all four, cores on the chip, suggesting that system bottlenecks such as memory contention dominate energy consumption; and at least one basic block (bb2) requires manual optimization to achieve maximum energy-efficiency. Third, fine-grain power optimization implies the ability to perform fine-grain power capping and more efficient power-constrained execution beyond that afforded by voltage and frequency scaling. For example a 10% reduction of the power cap in Exynos can be met by reducing frequency by one step but also by concurrency throttling and manual or compiler-driven code optimization. The latter two options show better energy savings potential.

7.3 Optimization of fine-grain basic blocks in acyclic regions

Loops enclose all basic blocks considered in our other use cases. However, applications, such as the Raytrace benchmark from the PARSEC suite, often contain hot basic blocks in acyclic regions. With the `simlarge` input, the `SphPeIntersect` function, which contains two hot blocks in an acyclic region (lines 323–328, lines 333–335, `sph.C`) consumes about 50% of the total execution time on the Exynos platform. The compiler optimizes these blocks poorly, leading to redundant memory accesses and indirect addressing instructions. We manually modified the generated code to remove redundant instructions, which reduced total energy consumption of the sequential version by 6.1% (2.8% for the parallel version). We cannot directly profile the targeted basic blocks due to the latency of hardware energy measurements. The execution time of the `SphPeIntersect` function is no more than 200 cycles on average. ALEA’s probabilistic model was the only viable option to profile and to optimize these basic blocks.

8 Conclusion

We presented a probabilistic approach for fine-grained energy profiling, implemented in ALEA, an energy profiling tool based on statistical sampling. We demonstrated that fine-grain energy accounting provides better insight into the power implications of microarchitectural and memory structures to support energy-aware code optimization. ALEA importantly overcomes the fundamental limitation of the low sampling frequency of power sensors, which is common across computing platforms. The tool operates entirely in user space and is portable across architectures.

We demonstrated ALEA’s high accuracy and low overhead on an Intel and an ARM platform with radically different architectural characteristics. ALEA achieved both functional and performance portability. We used ALEA to demonstrate the strong correlation between power consumption and memory access rates, as well as a clear impact of shared cache contention on power consumption. We presented use cases of ALEA where we applied new energy optimizations of individual basic blocks, using different strategies and achieved whole-program energy savings of up to 37%. These use cases motivated fine-grain energy accounting and uncovered the complex interplay between code optimization, multi-core execution and energy consumption.

We will pursue three directions for future work in ALEA. The first direction is to evolve ALEA into a production-strength energy accounting tool that maps energy consumption to source code and data structures, along the lines of tools such as Intel’s Vtune and HPCToolkit. The second direction is to extend ALEA’s capabilities to provide binary-level energy accounting of legacy programs running on virtualized software stacks. The third direction is to use ALEA for constructing a new library of code optimizations for power-constrained environments.

Acknowledgments

This research has been supported by the UK Engineering and Physical Sciences Research Council (EPSRC) through grant agreements EP/L000055/1 (ALEA), EP/L004232/1 (ENPOWER), and EP/K017594/1

(GEMSCCLAIM) and by the European Commission under the Seventh Framework Programme, grant agreement FP7-610509 (NanoStreams) and FP7-323872 (SCORPIO).

References

- [1] Hsu CH, Kuehn JA, Poole SW. Towards Efficient Supercomputing: Searching for the Right Efficiency Metric. *Proceedings of the 3rd ACM/SPEC International Conference on Performance Engineering, ICPE '12*, ACM: New York, NY, USA, 2012; 157–162.
- [2] Alonso P, Dolz MF, Mayo R, Quintana-Ortí ES. Energy-Efficient Execution of Dense Linear Algebra Algorithms on Multi-core Processors. *Cluster Computing* Sep 2013; **16**(3):497–509.
- [3] Chang F, Farkas KI, Ranganathan P. Energy-Driven Statistical Sampling: Detecting Software Hotspots. *Proceedings of the 2nd International Conference on Power-Aware Computer Systems, PACS'02*, Springer-Verlag: Berlin, Heidelberg, 2003; 110–129.
- [4] Flinn J, Satyanarayanan M. PowerScope: A Tool for Profiling the Energy Usage of Mobile Applications. *2nd IEEE Workshop on Mobile Computing Systems and Applications*, 1999.
- [5] Ge R, Feng X, Song S, Chang HC, Li D, Cameron KW. PowerPack: Energy Profiling and Analysis of High-Performance Systems and Applications. *IEEE Transactions on Parallel and Distributed Systems* 2010; **21**(5):658–671.
- [6] Kansal A, Zhao F. Fine-Grained Energy Profiling for Power-Aware Application Design. *SIGMETRICS Perform. Eval. Rev.* Aug 2008; **36**(2):26–31.
- [7] Keranidis S, Kazdaridis G, Passas V, Igoumenos G, Korakis T, Koutsopoulos I, Tassioulas L. NITOS Mobile Monitoring Solution: Realistic Energy Consumption Profiling of Mobile Devices. *Proceedings of the 5th International Conference on Future Energy Systems, e-Energy '14*, ACM: New York, NY, USA, 2014; 219–220.
- [8] McIntire D, Stathopoulos T, Kaiser W. Etop: Sensor Network Application Energy Profiling on the LEAP2 Platform. *Proceedings of the 6th International Conference on Information Processing in Sensor Networks, IPSN '07*, ACM: New York, NY, USA, 2007; 576–577.
- [9] Bertran R, Gonzalez Tallada M, Martorell X, Navarro N, Ayguade E. A Systematic Methodology to Generate Decomposable and Responsive Power Models for CMPs. *IEEE Trans. Comput.* Jul 2013; **62**(7):1289–1302.
- [10] Manousakis I, Zakkak FS, Pratikakis P, Nikolopoulos DS. TProf: An Energy Profiler for Task-Parallel Programs. *Sustainable Computing: Informatics and Systems* 2014; doi:<http://dx.doi.org/10.1016/j.suscom.2014.07.004>.
- [11] Contreras G, Martonosi M. Power Prediction for Intel XScale Processors Using Performance Monitoring Unit Events. *Proceedings of the 2005 International Symposium on Low Power Electronics and Design, ISLPED '05*, ACM: New York, NY, USA, 2005; 221–226.
- [12] Isci C, Martonosi M. Runtime Power Monitoring in High-End Processors: Methodology and Empirical Data. *Proceedings of the 36th Annual IEEE/ACM International Symposium on Microarchitecture, MICRO 36*, IEEE Computer Society: Washington, DC, USA, 2003; 93–.
- [13] Curtis-Maury M, Shah A, Blagojevic F, Nikolopoulos DS, de Supinski BR, Schulz M. Prediction Models for Multi-dimensional Power-Performance Optimization on Many Cores. *Proceedings of the 17th International Conference on Parallel Architectures and Compilation Techniques, PACT '08*, ACM: New York, NY, USA, 2008; 250–259.

- [14] Shao YS, Brooks D. Energy Characterization and Instruction-Level Energy Model of Intel's Xeon Phi Processor. *Proceedings of the 2013 International Symposium on Low Power Electronics and Design, ISLPED '13*, IEEE Press: Piscataway, NJ, USA, 2013; 389–394.
- [15] Tsoi KH, Luk W. Power Profiling and Optimization for Heterogeneous Multi-core Systems. *SIGARCH Comput. Archit. News* Dec 2011; **39**(4):8–13.
- [16] Tu CH, Hsu HH, Chen JH, Chen CH, Hung SH. Performance and Power Profiling for Emulated Android Systems. *ACM Trans. Des. Autom. Electron. Syst.* Mar 2014; **19**(2):10:1–10:25.
- [17] Wilke C, Götz S, Richly S. JouleUnit: A Generic Framework for Software Energy Profiling and Testing. *Proceedings of the 2013 Workshop on Green in/by Software Engineering, GIBSE '13*, ACM: New York, NY, USA, 2013; 9–14.
- [18] Pathak A, Hu YC, Zhang M. Where is the Energy Spent Inside My App?: Fine Grained Energy Accounting on Smartphones with Eprof. *Proceedings of the 7th ACM European Conference on Computer Systems, EuroSys '12*, ACM: New York, NY, USA, 2012; 29–42.
- [19] Schubert S, Kostic D, Zwaenepoel W, Shin KG. Profiling Software for Energy Consumption. *2012 IEEE International Conference on Green Computing and Communications 2012*; **0**:515–522.
- [20] Brouwers N, Zuniga M, Langendoen K. NEAT: A Novel Energy Analysis Toolkit for Free-Roaming Smartphones. *Proceedings of the 12th ACM Conference on Embedded Network Sensor Systems, SenSys '14*, ACM: New York, NY, USA, 2014; 16–30.
- [21] Intel Corporation. *Intel® 64 and IA-32 Architectures Software Developer's Manual*. 253669-033US, 2009.
- [22] Cao T, Blackburn SM, Gao T, McKinley KS. The Yin and Yang of Power and Performance for Asymmetric Hardware and Managed Software. *Proceedings of the 39th Annual International Symposium on Computer Architecture, ISCA '12*, IEEE Computer Society: Washington, DC, USA, 2012; 225–236.
- [23] Tallent NR, Mellor-Crummey JM, Adhianto L, Fagan MW, Krentel M. Diagnosing Performance Bottlenecks in Emerging Petascale Applications. *Proceedings of the Conference on High Performance Computing Networking, Storage and Analysis, SC '09*, ACM: New York, NY, USA, 2009; 51:1–51:11.
- [24] Tallent NR, Mellor-Crummey JM, Fagan MW. Binary Analysis for Measurement and Attribution of Program Performance. *Proceedings of the 2009 ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI '09*, ACM: New York, NY, USA, 2009; 441–452.
- [25] Adhianto L, Banerjee S, Fagan M, Krentel M, Marin G, Mellor-Crummey J, Tallent NR. HPCToolkit: Tools for Performance Analysis of Optimized Parallel Programs. *Concurr. Comput.: Pract. and Exper.* Apr 2010; **22**(6):685–701.
- [26] Ren G, Tune E, Moseley T, Shi Y, Rus S, Hundt R. Google-Wide Profiling: A Continuous Profiling Infrastructure for Data Centers. *IEEE Micro* Jul 2010; **30**(4):65–79.
- [27] Li D, de Supinski BR, Schulz M, Nikolopoulos DS, Cameron KW. Strategies for Energy-Efficient Resource Management of Hybrid Programming Models. *IEEE Trans. Parallel Distrib. Syst.* Jan 2013; **24**(1):144–157.
- [28] Flinn J, Satyanarayanan M. Energy-Aware Adaptation for Mobile Applications. *Proceedings of the Seventeenth ACM Symposium on Operating Systems Principles, SOSP '99*, ACM: New York, NY, USA, 1999; 48–63.
- [29] Lohr S. *Sampling: Design and Analysis*. Brooks/Cole, 1999.

- [30] Casella G, Berger R. *Statistical Inference*. Duxbury Press Belmont, Calif, 1990.
- [31] Montgomery D, Runger G. *Applied Statistics and Probability for Engineers*. Wiley, 2002.
- [32] Curtis-Maury M, Dzierwa J, Antonopoulos CD, Nikolopoulos DS. Online power-performance adaptation of multithreaded programs using hardware event-based prediction. *Proceedings of the 20th Annual International Conference on Supercomputing, ICS '06*, ACM: New York, NY, USA, 2006; 157–166.
- [33] Curtis-Maury M, Blagojevic F, Antonopoulos CD, Nikolopoulos DS. Prediction-based power-performance adaptation of multithreaded scientific codes. *IEEE Trans. Parallel Distrib. Syst.* Oct 2008; **19**(10):1396–1410.
- [34] Jeon M, He Y, Elnikety S, Cox AL, Rixner S. Adaptive parallelism for web search. *Proceedings of the 8th ACM European Conference on Computer Systems, EuroSys '13*, ACM: New York, NY, USA, 2013; 155–168.
- [35] Sridharan S, Gupta G, Sohi GS. Adaptive, efficient, parallel execution of parallel programs. *Proceedings of the 35th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI '14*, ACM: New York, NY, USA, 2014; 169–180.