

Phase Characterization for Power: Evaluating Control-Flow-Based and Event-Counter-Based Techniques

Canturk Isci and Margaret Martonosi
Department of Electrical Engineering
Princeton University
{canturk,mrm}@princeton.edu

Abstract

Computer systems increasingly rely on dynamic, phase-based system management techniques, in which system hardware and software parameters may be altered or tuned at runtime for different program phases. Prior research has considered a range of possible phase analysis techniques, but has focused almost exclusively on performance-oriented phases; the notion of power-oriented phases has not been explored. Moreover, the bulk of phase-analysis studies have focused on simulation evaluation. There is need for real-system experiments that provide direct comparison of different practical techniques (such as control flow sampling, event counters, and power measurements) for gauging phase behavior.

In this paper, we propose and evaluate a live, real-system measurement framework for collecting and analyzing power phases in running applications. Our experimental framework simultaneously collects control flow, performance counter and live power measurement information. Using this framework, we directly compare between code-oriented techniques (such as “basic block vectors”) and performance counter techniques for characterizing power phases. Across a collection of both SPEC2000 benchmarks as well as mainstream desktop applications, our results indicate that both techniques are promising, but that performance counters consistently provide better representation of power behavior. For many of the experimented cases, basic block vectors demonstrate a strong relationship between the execution path and power consumption. However, there are instances where power behavior cannot be captured from control flow, for example due to differences in memory hierarchy performance. We demonstrate these with examples from real applications. Overall, counter-based techniques offer average classification errors of 1.9% for SPEC and 7.1% for other benchmarks, while basic block vectors achieve 2.9% average errors for SPEC and 11.7% for other benchmarks respectively.

1 Introduction

In recent years, phase behavior of applications has drawn a growing research interest for two main reasons. First, the increasing complexity and power demand of processor architectures mandate workload dependent dynamic management techniques. These techniques extensively benefit from tracking application phases to optimize power/performance trade-offs and to identify critical execution regions for management actions [1, 3, 9]. Second, in parallel with increasing processor complexities, architectural simulation studies have a growing need to research long execution timescales to capture the increasingly variable behavior of applications. These studies benefit from phase characterizations that summarize applica-

tion behavior with representative execution regions, alleviating the prohibitively high computational costs of large-scale simulations [32, 36].

Various prior studies demonstrated that phase behavior can be observed via different features of applications. Most of these approaches fall into two main categories: In the first category application phases are determined from the control flow of the applications or the program counter (PC) signatures of the executed instructions [9, 18, 24, 27, 32, 35, 36, 37]. In the second category, phases are determined based on the performance characteristics of the applications [3, 7, 12, 20, 39, 40].

Although there have been some previous efforts to compare or evaluate phase characterization techniques [2, 8, 26], they do not perform a direct comparison of the two main approaches. Moreover, there is generally a missing link between phase characterizations and their ability to represent power behavior, especially with real-system experiments. Such power characterization is very important for real systems, as a primary goal of phase characterization is dynamic power management of running systems.

Following from these motivations, in this work, we compare phase characterizations based on PC signatures and performance behavior of applications. Our study primarily evaluates these techniques for accurate power behavior characterization on a real-system. We compare these with respect to the actual, measured runtime power dissipation behavior of applications. Specifically, we look at phase analysis based on basic block vector (BBV) features of an application [36] to determine regions of similar power behavior. We compare this to phases determined by a particular set of performance monitoring counter (PMC) events that are chosen to reflect power dissipation [21]. We test the power characterization accuracy of these methods on 21 benchmarks from SPEC2000 suite and 9 other benchmarks derived from commonly used desktop and multimedia applications. We show that, in general, tracking performance metrics performs better than tracking control flow in identifying power phase behavior of applications. Additionally, we present specific examples from real applications demonstrating cases where power phase behavior cannot be deduced from code signatures.

There are three primary contributions of this work. First, we have designed an accurate, real-system method for synchronizing BBV signatures, performance events, and power measurements on running machines. This method allows us to study large-scale application behavior on running systems rather than being limited to simulation approaches. Second, utilizing this experimental framework, we evaluate how BBV and PMC based approaches perform from a real power characterization point of view. Compared to an uninformed phase characterization, both phase based techniques achieve significantly higher accuracies in identifying power phases, lead-

ing to 2-6X less errors for benchmarks with significant power variations. Last, we compare control flow (BBV) and performance (PMC) based approaches against each other for their power phase classification abilities. Overall tracking performance behavior leads to 30-40% fewer errors than tracking control flow in representing real power phase behavior.

The rest of the paper is organized as follows. Section 2 describes the goals of our phase analysis research and discusses the reasons why control flow and performance phases can differ. Section 3 describes our experimentation platform. Section 4 explains the collection of BBV and PMC information with our experimental setup. Section 5 describes our phase classification methods. Section 6 describes our quantitative evaluation and presents the power phase characterization results. Section 7 provides detailed observations from performed experiments. Section 8 provides a final discussion of BBV and PMC based approaches and presents recommendations for future research. Section 9 summarizes related work and Section 10 offers our conclusions.

2 Goals and Challenges of Phase Analysis

2.1 Phase Overview

Regardless of whether the phase characterization of an application is geared towards summarizing its execution behavior or towards identifying periods with different power/performance implications, the underlying goal is fundamentally the same. The principle purpose of phase characterization is to accurately classify execution behavior into self-similar operation regions based on the observed features. However, the choice of tracked features can be different for different endgoals. It might be desirable to have architecturally independent metrics to summarize execution for architectural exploration studies. In other cases, it is preferable to have metrics that reflect the different behavior under different architectures so that resulting phases closely track different power/performance behavior and correspond to different dynamic management opportunities.

In our research perspective, phase analysis lies as a layer between the architecture and the applications that make use of phase information. Based on the target application, there also exists a processing layer between phase analysis and the application, which helps interpret the phase behavior for application specific goals. An application can simply match observed phases with appropriate adaptive responses or can further process phase information to engage different dynamic management actions. For example, phase patterns can be used to detect recurrent behavior [22, 37] or to predict durations for certain modes of operation to amortize mode transition costs [23, 27]. Several system-level or architectural methods can benefit from phase information to guide dynamic management actions. Temperature aware scheduling [3] can benefit from detecting repetitive power phases to select among tasks with different power/temperature behavior to reduce performance degradation due to idling or throttling. Multicore power balancing and activity migration [16, 33] rely on application behavior to distribute or transfer activity among different components. Phases can provide both history and phase change information to decision policies of these techniques. Dynamic voltage and frequency scaling approaches [6, 30, 41] can evaluate costs and benefits of switching among operation modes at

runtime based on diversity and duration of different phases.

Specifically in this study, our focus is how well different phase characterizations—based on different features—represent workload power characteristics. We look at how previous control-flow based approaches perform for power characterization and compare this to our—power oriented—PMC based approach.

2.2 What Control Flow Information Does Not Show

Before delving into the details of our experimentation and phase characterization methodology, here we discuss the reasons why control flow and power/performance behavior of an application may disagree.

There are multiple aspects of application behavior that can cause the control flow and performance based approaches to reach different phase characterization conclusions. *Dynamic change in data locality* during an application’s execution can cause the power behavior to significantly change. While this change can be easily recovered from memory performance metrics, code signatures cannot reflect this as execution footprints are not altered. *Effectively same execution* represents the converse of the above effect. In various applications, multiple procedures or code segments perform similar processes, leading to practically identical power behavior. These are considered as fairly different phases in terms of control flow, which may result in many different phase clusters that do not reflect actual changes in program power. Typical examples for these are scientific or other iterative processing applications performing different tasks on an input with similar power/performance implications. *Operand dependent behavior* may result in similar effects as the first case, where power and latency of a unit depends on the input operands, despite the same control flow. Typical cases for these are overflow handling and scaling of execution based on the input operand values or widths [4].

We revisit these effects after presenting our power phase characterization study. In Section 7, we show the differences that can arise between control flow and performance based phase tracking for power, with observations from real experimented applications.

3 Software and Hardware Measurement Platform

To collect synchronous PC, PMC and power information during an application’s execution, we use dynamic instrumentation via Pin [29]. Pin provides several flexible methods to dynamically instrument the binary at different granularities. This first step, *instrumentation*, simply decides where in the native code the additional procedures to analyze the application behavior should be inserted. Afterwards, whenever one of these instrumentation checkpoints are reached, Pin gains the control of the application and injects corresponding analysis routines. During execution, each time the instrumented locations are visited, their injected analysis routines also execute, providing the dynamic application information. This second phase of operation is called *analysis*. Pin utilizes a single executable, *Pintool*, to perform instrumentation and analysis on an application.

Figure 1 presents an overview of our experimental setup for power phase analysis with Pin. In our Pintool, we use trace level instrumentation to keep track of executed code traces.

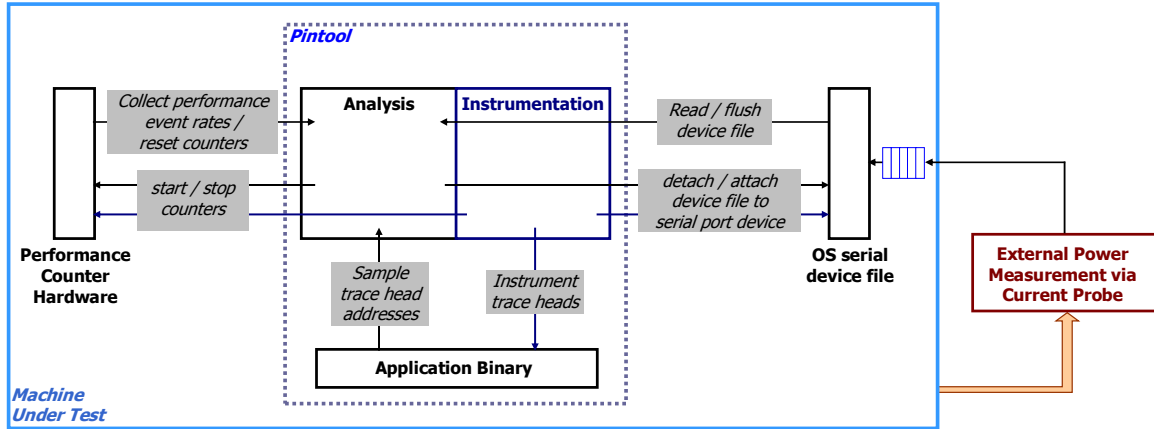


Figure 1. Experimental setup for power phase analysis with Pin.

Our analysis routine consists of three levels of hierarchy. The first level simply provides an account of executed instructions. This is implemented as an inlined conditional to improve performance and to avoid overwhelming power behavior. The second level samples one PC address approximately every 1 million instructions. The highest level analysis is invoked every 100 million instructions. This routine generates one BBV from the 100 PC samples, reads performance statistics from PMCs and logs the measured power history from the serial device file. These three sources of data collection are shown with the three incoming arrows to the analysis routine of our Pintool.

It is important to isolate application behavior from Pin operation. Pin provides application exclusive control flow information, however, performance monitoring and power measurements are out of Pin’s control. Therefore, we provide handles to our Pin routines to disable the logging of data for power and performance at routine entries, and to reenale data logging at routine exits. Under Pin execution, instrumentation and analysis are temporally intermixed. Therefore, we use these handles during both instrumentation and analysis.

We provide real power behavior feedback to our power phase characterizations via external, live power measurements. We perform power measurements by measuring the current flow into the processor with a current probe. This measurement information is then fed back to the measurement system over the serial port interface.

To isolate the application power behavior from Pin analysis and instrumentation, we use certain controls within the instrumentation and analysis routines of our Pintool. These handles detach/attach serial device driver from the device file at routine entries/exits via `termios` flags. This approach allows us to preserve previous application power history, while preventing further logging while inside an instrumentation or analysis routine. At the end of a 100 million instruction sampling period, the highest level analysis routine halts logging and reads the logged power history for the past sampling period. This history is then averaged and is assigned as the observed power for the past sampling quantum. Afterwards, the buffer is flushed and reenabled for logging at the start of the next sampling interval.

Similar to the power measurement method, we developed several handles to control PMC monitoring from within our

Pintool. At Pintool initialization, we first configure the events to be monitored. This is the most heavyweight operation, and is applied only once, before application execution commences. We selectively halt/start performance monitoring at instrumentation and analysis routine entries/exits. This is used to avoid polluting the PMC information with Pin execution. Although we provide the start/stop handles to all routines, after our initial experiments, we do not invoke them for instrumentation and the second level analysis routines, as their costs are comparable. Note that, this trade-off only affects PMC information without any effect on control flow information and power measurements. During the highest level analysis routine we read the past PMC statistics and reset the counters for the following sampling period.

With this experimental framework, we provide a valid matching between application execution flow, performance statistics and application specific power behavior. Inevitably, there can be error due to measurements, and due to transient operations that perform the control functions for selective logging and asynchronous operation of different data sources. However, in the experimented cases, our selective power collection and PMC sampling process produces power/performance information with good fidelity. Acquired power behavior and performance statistics are similar, both temporally and in terms of delta variations, to native executions of applications.

All of our experiments described in this paper are performed on a 1.4GHz Pentium 4 processor with Linux operating system, kernel 2.4.7-10. The experiments are carried out with the SPEC CPU 2000 benchmarks using reference datasets and other benchmarks derived from well-known suites and desktop applications. All benchmarks are compiled with `gcc` and `g77` compilers with base compiler flags.

4 Generating BBV and Performance Information from Pin/Hardware Structure

4.1 Program Counter Sampling and BBV Generation

To track control flow based application phases, we use the basic block vector (BBV) approach [36]. BBVs summarize application execution by tracking both which basic blocks of the application are touched and how many times each basic block is visited during a sampling interval. BBVs represent application execution behavior by providing both work-

ing set information and execution frequencies for different basic blocks [8]. BBVs are constructed from execution flow by mapping executed PC addresses to the basic blocks of an application binary. Originally, each component of a BBV is a specific basic block, and the magnitude of the component represents how often the corresponding basic block has been executed for a past sampling period. For practical purposes, BBVs are generally mapped into smaller dimensional vectors via random projection/hashing, component analysis or eliminating least significant dimensions [2, 13, 26, 36, 37].

In our implementation, we use Pin to sample the PC addresses at trace heads. As each trace head is also a basic block start address, each sampled PC actually corresponds to a specific basic block. Consequently, different sampled PCs represent different elements of the BBV and number of samples for a specific PC represents the execution frequency of the corresponding block. For sampling periods, we use previously published granularities [2]. We sample one PC every 1 million instructions and construct a BBV at every 100 million instructions. Thus, each BBV has an $L1$ -norm—sum of vector components—of 100. We perform static instrumentation of our experimented applications with gcc compiler to determine the dimensions of basic block profiles. Even after eliminating untouched basic blocks and libraries, applications exhibit large BBV dimensions ranging from 33000 (gcc) to 100 (swim). These lead to highly sparse and impractical to implement BBVs. Therefore, we also apply dimension reduction. For the reduced dimensions, we choose 32 buckets based on previous work [37]. We use a variation of Jenkins’ 32 bit integer hash function [25] to reduce the large and variable BBV dimensions into common 32 dimensional vectors.

As has been discussed in previous studies [26], sampling always incurs some amount of information loss. However, for any practical implementation of runtime control flow tracking, sampling is inevitable. Our observations show that our sampled PC information offers good similarity information for large scale control flow behavior. We compare full-blown BBVs, constructed from complete PC information, to our sampled BBVs with similarity matrices [36]. Both methods reflect the major phase content in terms of execution flow similarity and lead to similar phases for small numbers of target phase clusters.

4.2 Using Performance Counters to Generate PMC Vectors

In order to track power phases, we use a set of 15 performance counters that are good proxies for power estimation. The chosen counters track metrics such as CPU instruction counts, L1 and L2 access rates and bus utilizations for memory behavior. The complete list of chosen performance counters are shown in Table 1 together with the applied mask configurations that define the particular event subsets we choose to track. The performance monitoring method is similar to prior research [21], but streamlined to avoid counter rotations. The final set of 15 PMC events can be monitored simultaneously without conflicts. Therefore, no PMC configuration is required except at the initial Pintool startup.

Every 100 million instructions, we collect the performance event counts and cycle count for the past sampling period. We then convert these event counts into per-cycle rates. These 15 event rates are then used to construct a 15 dimensional PMC

PMC Event	Mask	Description
IOQ_allocation	0x0EFE1	I/O Queue and Bus Sequence Queue allocations from all agents
BSQ_cache_ref	0x0507	L2 cache read and write accesses
FSB_data_activity	0x03F	Front Side Bus utilization for reading, driving or reserving the bus.
ITLB_reference	0x07	ITLB translations performed
uop_queue_writes	0x07	All μ ops written to the μ op queue
TC_deliver_mode	0x038	Number of cycles the processor is building traces from instruction decode
uop_queue_writes	0x04	μ ops written to the μ op queue by microcode ROM
x87_FP_uop	0x08000	All x87 floating point μ ops executed
LD_port_replay	0x02	Number of replays at the load port
x87_SIMD_moves	0x018	Executed x87, MMX, SSE and SSE2 load, store and register move μ ops
ST_port_replay	0x02	Number of replays at the store port
branch_retired	0x0F	All branches retired
uops_retired	0x03	Number of μ ops retired
front_end_event	0x03	Number of loads and stores retired
uop_type	0x06	Tags load and stores (Does not count)

Table 1. The set of chosen performance counter events and mask configurations.

vector, which gauges the similarity of execution samples in a similar manner as BBVs.

5 Phase Classification

We cluster BBV and PMC vector samples into phases with multiple clustering algorithms. First, we develop a fast, but less accurate method based on the descriptions of previous work [20]. This method is more suitable for runtime analysis as it assigns samples to phases as they are observed. We call this method *First Pivot Clustering*. To corroborate the observed characterization results are not due to the choice of clustering, we also experiment with a very computationally expensive method, *Agglomerative Clustering* [11]. We use two variations of this method: *complete linkage* and *average linkage*. Patil et al. [32] show that SPECint and SPECfp lead to on average 4 and 5 phases respectively. Therefore, in this study, to provide consistent results and error metrics across all applications, we target towards 5 final phases for all benchmarks. Afterwards, we show that observed results are consistent as the target number of phases changes.

5.1 First Pivot Clustering

First Pivot Clustering uses *pivot* samples to represent different phases. In the original description of this method, a new gathered sample is compared to all previous pivots, i.e. starters of different phases. If the current sample is within a specified threshold distance of a pivot, it is assigned to that phase. If it is not within the similarity distance of any of the pivots, it starts a new phase and is added to the list of pivots as the representative sample for the new phase. By this way, The original description can assign samples to phases at runtime. This approach provides an upper bound to the distance within each phase, but it does not guarantee a fixed number of phases.

We change this to an iterative process, where the threshold is changed dynamically based on both the acquired and target number of phases. With this modification, we classify both BBVs and PMC vectors into 5 final phases after a few iterations.

5.2 Agglomerative Clustering

Agglomerative clustering is a tedious bottom-up approach to clustering samples into phases. In this approach, clustering algorithm starts with an initial clustering solution of N clusters, where N is the number of samples. At each iteration, the algorithm compares all pairwise combinations of the current set of clusters and finds the best candidate pair of clusters to combine into a single cluster. The pairs are compared based on a *linkage* criterion, which determines the best candidates. This iterative process continues until a final target number of clusters are reached or a distance threshold among clusters is exceeded. For agglomerative clustering, we experiment with two types of linkages, complete and average linkage. We describe these below.

5.2.1 Average Linkage

Average linkage compares the average distance between all sample pairs belonging to two different clusters. For two clusters with i and j samples respectively, it computes the distance between all the $i \cdot j$ pairs and finds the average distance between the clusters. Performing this operation for all cluster combinations, it chooses to combine two clusters with the minimum average distance. This leads to clusters with similar ranges in all dimensions, but can result in significantly different ranges for different clusters.

5.2.2 Complete Linkage

Complete linkage does similar comparisons as average linkage. However, it compares the maximum pairwise sample-distance among clusters. It combines the clusters with the least maximum distance among all their pairs. Consequently, the final set of clusters have similar ranges among most of their samples, although the range across each dimension can be different.

In all our analyses we use $L1$ —Manhattan—distance, as our measure of distance between two samples. For BBV based clustering, we compute the $L1$ -distance between the two corresponding 32 dimensional BBVs. For PMC based clustering, we use the two 15 dimensional PMC vectors to gauge the similarity between points. We apply the above three clustering methods and evaluate clustering criteria based on these distances.

6 Power Phase Characterization: Evaluation of Techniques and Results

We apply our described power phase classification methods to several benchmarks. Using both control flow and performance features, we cluster each benchmark into 5 phases with multiple clustering methods. Here, we discuss first how we evaluate the fidelity of these phases in terms of power behavior characterization. Afterwards, we provide the complete set of results based on these evaluations. With the demonstrated results, we show how code signatures and PMC phases perform in identifying power behavior characteristics with respect to a “gold standard” phase classification as our lower bound and an “uninformed” classification as the upper bound. We also present a direct comparison between BBV phases and PMC phases for power characterization.

6.1 Evaluating the Error of Power Phase Characterization

We evaluate the quality of generated phase clusters by comparing the measured power at each sample to the aggregate power for the whole cluster the sample belongs to. For a benchmark with N samples, each sample i ($i = 1, \dots, N$) is an element of one of the final phase sets P_j ($j = 1, \dots, 5$). Each sample has a corresponding set of data $[bbv_i, pmc_i, pwr_i]$, where bbv_i and pmc_i are the corresponding BBV and PMC vectors used during phase clusterings, and pwr_i is the measured power value during sample i 's execution. For each phase P_j , we compute a “representative power”, R_j , as the arithmetic average of the power values for the total N_j samples belonging to that phase. Then, for each sample i , we compute the squared difference between the sample's actual power value pwr_i and the representative power R_j for its owner phase P_j . We denote R_j values corresponding to each sample i with R_{ji} . Afterwards, we compute the rooted average of these squared differences over all samples for our final RMS error figure E_{RMS} . We summarize this error computation in Equation 1.

$$R_j = \frac{\sum_{i \in P_j} pwr_i}{N_j} \quad (j = 1, \dots, 5)$$

$$E_{RMS} = \sqrt{\frac{\sum_{i=1}^N (pwr_i - R_{ji})^2}{N}} \quad (1)$$

This error value represents the quality of power phase characterization for a given phase classification method on the evaluated benchmark. The methods are the combinations of tracked feature (BBVs or PMCs) and clustering algorithm (first pivot, agglomerative with average or complete linkage). We use this error measure to gauge the effectiveness of BBV and PMC based features in representing power phase behavior of applications in our experiments with various benchmarks.

6.2 Error Boundaries

To gauge the ability of the phase classification techniques in discerning application power behavior, we also provide the error boundaries that can be achieved with perfect knowledge of power information—lower bound—as well as without any knowledge of application behavior—upper bound.

To compute lower error bounds, we look directly at the measured power, which is the independent target experiment parameter in all other analyses. We apply all three clustering algorithms to each benchmark's power information and for each case choose the smallest error value achieved. We refer to this “gold standard” measure as *baseline error* in our results.

For the upper error bounds, we design a separate clustering method, which assigns each sample to any of the final target phases randomly, without using any application behavior information. We refer to the results of this “uninformed” phase characterization as *random error*. We show the results achieved with these approaches for each benchmark. These demonstrate opportunities for improvement that remain and how much improvement each tested phase analysis feature brings to power characterization.

6.3 Experimented Benchmarks

For our power phase analysis experiments, we obtain control flow, performance and power characteristics for several benchmarks on our test machine. We look at 11 SPECint benchmarks—all except `perlbmk` due to compilation problems—and 10 SPECfp benchmarks—excluded are F90 benchmarks. We experiment with all reference datasets for the 21 SPEC benchmarks leading to a total of 37 different experiments.

In addition to SPEC, we also use 9 other benchmarks from previous studies and derived from well-known applications. These benchmarks are `ghostscript`, `dvips`, `gimp`, `lame`, `cjpeg`, `djpeg`, `mesh`, `stream` and `mdbnch`. For some cases, we alter the dataset or iterations for the benchmarks to achieve longer execution times. We describe these benchmarks and any modifications here.

In the first category, `ghostscript` and `dvips` are conversion utilities commonly used in document creation. Their behavior depends on the nature and layout of the input document. Next, `gimp`, `lame`, `cjpeg` and `djpeg` are media processing tools used to convert among formats or manipulate media files. Last, `mesh`, `stream` and `mdbnch` are iterative applications with multiple sequential functions similar to many scientific computation tools.

For `ghostscript` and `dvips` we use a large document of 190 pages, with different size images in the middle of the document. `ghostscript` converts a postscript input to pdf, and `dvips` converts dvi input to postscript.

`Gimp` is an image manipulation tool [15]. We use `gimp` in batch mode to perform several image processing operations such as blurring, filtering and applying digital effects. Depending on the computation and memory intensity of the applied functions, they can lead to different power behavior. We use the `lame` MP3 encoder [38] to encode a wave file under varying quality settings. Both power levels and the total execution increase with the quality settings. `Cjpeg` and `djpeg` are image compression and decompression programs from MediaBench [28]. We use `cjpeg` to encode a very large (110 MB) ppm image file into jpeg and `djpeg` to decode the jpeg file into ppm. Their power behavior also changes during execution and with input data.

`Mesh` is a well-known program used in dynamic program optimization studies [10, 34]. It performs various computations over the input mesh edges and faces, with sequentially executed repetitive functions. Our mesh input consists of 10K nodes and 60K edges, leading to very quick iterations. To emphasize the execution of separate functions, we alter the original mesh code to repeat each function 100-200 times. `Mdbnch` is a relatively older, scalar molecular dynamics benchmark [14]. It performs seven different molecular dynamics tasks with different sizes or complexities. To extend its execution, we increase the number of time steps for each task by 4-50x. Both `mesh` and `mdbnch` have similar iterative properties of scientific computation. Although they iterate within different control paths, each task usually has similar computation properties—except for changes in memory intensity. These lead to fairly flat behavior with small data footprints. `Stream` is actually a synthetic benchmark, commonly used to measure sustainable memory bandwidth [31]. It iterates over four small tasks doing different computations. Similar to the above

two applications, `stream` also exhibits a stable power behavior during normal operation. However, it has a loop carried positive feedback that eventually overflows the inputs for its tasks, resulting in a drastic change in power behavior. For our `stream` experiments, we use an iteration count of 275 and data size of 2 million entry arrays.

6.4 Power Phase Characterization Results

We show the overall results for our experiments in Figures 2-4. Three figures show phase characterization errors for the three clustering algorithms. In each figure, we show the upper—*random*—and lower—*baseline*—error bounds for each application and the achieved error with BBV and PMC based approaches. We also show the average accuracies for SPECint, SPECfp and other experimented benchmarks.

First, obtained characterization results are consistent, independent of the applied clustering algorithm. In general Figure 2 shows relatively higher errors due to the cheaper clustering method. However, the general accuracy relation between BBVs and PMCs are preserved.

Comparing among the three sets of applications, SPECfp applications lead to relatively low errors even with random phase clustering for some cases. This is due to the generic flat power behavior of these benchmarks (`applu`, `art`, `sixtrack`, `wupwise`). In some other cases, benchmarks go through specific initialization (i.e. `equake`) or periodic (i.e. `ammp`) phases with significant changes in all control flow, performance and power features. In these cases, both BBVs and PMCs achieve very good power characterizations approaching baseline errors.

SPECint shows significantly higher errors for all approaches due to higher variations in behavior. In many of the shown cases, BBVs and PMCs are seen to have significant improvement over random clustering. This shows the benefits of phase tracking for power behavior characterization.

Most of the other experimented benchmarks show significantly higher error ranges due to their high power variability based on input data characteristics and functional behavior. In these cases, applying phase analysis, especially with PMCs, proves to be very useful in identifying similar power behavior.

Overall, for the three benchmark sets, BBVs achieve errors that are on average 48% less than random clustering errors, for benchmarks with non-flat power behavior. PMC phases lead to 66% less errors than random clustering. For PMC based approach, power characterization accuracies vary between 2-6X improvements over random clusterings for these benchmarks. Performing same comparisons with respect to baseline errors show, BBVs on average achieve 2.9X higher errors compared to baseline, while PMC errors are 1.8X of baseline figures. These comparisons show, BBV and PMC phase analyses have significant benefit in characterizing power behavior. However, there still exist opportunities to improve power phase behavior characterization of applications.

As above measures also indicate, in almost all experimented cases, PMC based phase analysis performs better than BBV based approach for representing power behavior. PMCs lead to 2.2% and 1.4% errors for SPECint and SPECfp, while BBVs achieve 3.4% and 1.5% errors. For the other experimented benchmarks, PMCs and BBVs have 7.1% and 14.7% average errors respectively. For most of the benchmarks PMCs achieve 30-40% less errors than BBVs with an average




Figure 2. Power characterization errors (absolute) for BBV and PMC phases with first pivot clustering.




Figure 3. Power characterization errors (absolute) for BBV and PMC phases with agglomerative clustering-average linkage.




Figure 4. Power characterization errors (absolute) for BBV and PMC phases with agglomerative clustering-complete linkage.

of 33%. This direct comparison between BBVs and PMCs shows, although both techniques provide useful features to identify power phase behavior, in general PMCs features are better candidates for identifying power phases.

6.5 Sensitivity to Different Target Number of Phases

We have presented our complete analysis for a fixed target number of 5 phases for consistency. However, we have also experimented with different numbers of target phases to verify the reliability of our results. We show these in Figure 5.

Here, we show the effect of target phases with agglomerative clustering/complete linkage. For all the benchmarks, we perform clusterings for final phase numbers varying from 1 to 5000. We show the achieved errors as both RMS and maximum observed values. For each benchmark, we compute the RMS and maximum error figure for each target phase count. Afterwards, we average these values over all benchmarks to reach a single error figure for each target phase count.

Intuitively, for a single final phase, both BBVs and PMCs will reach the same error, equivalent to the standard deviation of all the power samples of the benchmark. Afterwards, as the number of phases increase, errors for both methods will decrease with different slopes. As number of target final phases grows towards infinity, both error curves will converge to 0, i.e. where each phase is a singleton sample.

In Figure 5, we show the behavior up to 100 phases for demonstration purposes. As phase counts grow beyond 100, both curves approach 0. For all practical purposes, PMC based phases perform consistently better, independent of the number of final phase clusters.

7 Observations from Experimented Applications

Initially we discussed some of the possible reasons that can cause control flow information and performance statistics to arrive at different conclusions about application power behav-




Figure 5. Variation of errors with respect to number of final phases.

ior. Here, we show our observations from actual applications that we experimented on. We demonstrate the effects for two of the sources of disagreement, *operand dependent behavior* and *effectively same execution*.

7.1 Operand Dependent Behavior

Here, we show an example of operand dependent behavior with the *stream* benchmark. *Stream* performs four repetitive operations with simple vector kernels. It operates on three vectors, *a*, *b* and *c*. The four operations are *copy* ($c[j] = a[j]$), *scale* ($b[j] = scalar * c[j]$), *add* ($c[j] = a[j] + b[j]$) and *triad* ($a[j] = b[j] + scalar * c[j]$). It targets at measuring sustainable memory bandwidth with vectors larger than cache sizes and by avoiding data reuse. We use this application to show operand dependent behavior and its implications on power. There exists a positive feedback between each iteration of the four described operations. This causes the the FP operations to overflow at iteration 261, where first vector *a* overflows at *triad*. This is then propagated to vectors *b* and *c* in the next iteration. This overflow causes the three FP kernels to experience a slowdown larger than 10x, while the *copy* operation is not significantly effected. Consequently, power dissipation experiences a drastic phase change, while execution path is still conserved.

In Figure 6, we show the resulting behavior in terms of power, BBV signatures and PMC signatures. Figure 6(a) shows, the power (top) and BBV signatures (bottom) with respect to executed instructions. We show BBV signatures as stacked vector sample bars, where magnitude of each vector component adds on top of the stack. Here, we see the repetitive BBV vector patterns throughout the execution, corresponding to the 4 different operations repeated 275 times. As the control flow is repetitive, the sudden power drop goes undetected with BBVs. In Figure 6(b), we show the same execution with power (top) and few of the PMC vector samples (bottom). Shown PMC metrics represent instructions per cycle (IPC), L2 cache access rates (L2) and memory access rates (MEM). Here, we show the execution with respect to cycles, to emphasize the actual effect of overflow on elapsed time in different power phases. While, the lower power phase occupies less than 6% of executed instructions, the time spent in this phase is more than half of the total execution. Tracking PMCs easily identifies this power phase change resulting from operand dependent behavior of *stream*.

7.2 Effectively Same Execution

Phase characterizations of applications have two related outcomes. First, phase characterizations provide feedback for

identifying phase changes in program behavior. Second, they classify applications into similar regions of execution. These two aspects have an inverse relation, which can be considered in terms of *similarity* and *granularity* [17]. Dictating more restrictive similarity features within each phase results in higher number of phases with smaller granularity. These may, then, lead to numerous false alarms for spurious phase transitions, as many of the small variations in tracked features do not reflect in application (power) behavior. Thus, a desired property for phase characterization is to lead to high granularity phases that capture major application behavior; balancing similarity and granularity.

Effectively same execution represents a characteristic behavior when PMC and BBV approaches perform differently in achieving this balance. In many occurrences, applications walk through different code paths, while performing similar computational tasks. These lead to different code signatures, indicating different phases, while actual power phase behavior is similar.

We demonstrate the impact of this effect with the *mesh* benchmark. During its execution, *mesh* first reads an input mesh configuration and performs various tasks on the input mesh. Most of these tasks have computationally similar properties, leading to effectively same execution behavior—while in different execution address spaces. In Figure 7, we show part of the execution characteristics for *mesh*. In the figure, we first show the measured power behavior. We can easily separate *mesh* execution into three power phases by observing the power trace. We label these “actual” power phases as *H*, *L* and *M* on the power trace, representing phases with high, low and medium power consumption. Underneath the power trace, we show the corresponding BBV vector patterns for each sample. Again, we present the 32 dimensional BBVs as stacked bars, where each vector component adds up to the stack based on its magnitude. Several distinct control flow phases are observable from the BBV patterns. We separate each of these regions with vertical dotted lines. These correlate well with *mesh* tasks. The first high power phase corresponds to the sorting task after reading nodes and initialization. This task sorts nodes based on their types. It operates mainly in L1 cache and performs several arithmetics. The following low power phase, results from *SetBoundaryData* task which sets the values for boundary nodes. This task mostly accesses L2, and has low overlapping computation, which leads to less power. After this task, *mesh* repetitively operates on three computation tasks, namely, *ComputeForces()*, *ComputeVelocityChange()* and *SmoothenVelocity()*. These constitute the medium power phase of *mesh*. All these tasks also make significant L2 accesses. However,

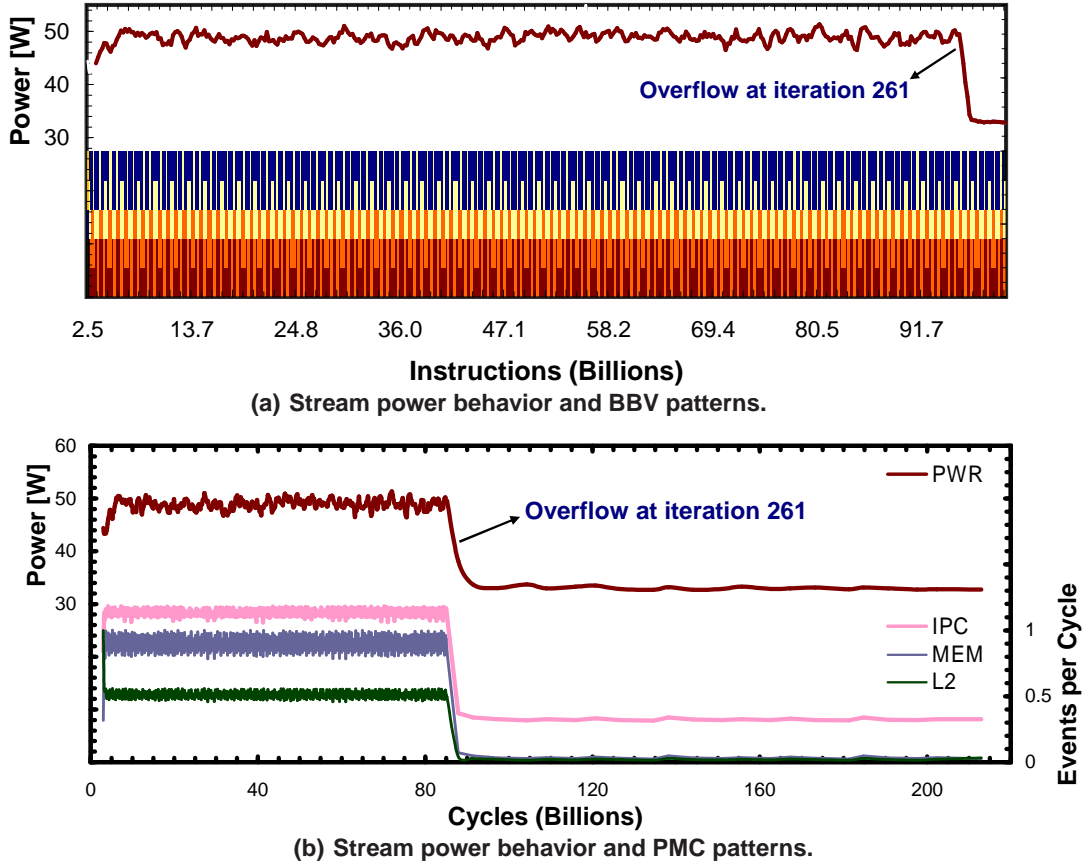


Figure 6. Power phase change at overflow condition for `stream` benchmark. (a) shows BBV signatures, unable to detect the phase change, (b) shows PMCs detecting the change. (b) is drawn with respect to elapsed cycles to show the actual time behavior.

their overlapping FP computations lead to relatively higher power.

In the lower two plots of Figure 7, we show the phase classifications performed by BBVs and PMCs. We apply agglomerative clustering with complete linkage and use target phase numbers, N , of 5—our base choice—and 3 for a more restrictive case. In these plots, y axis shows different phases ranging from 1 to 5 for the first case and 1 to 3 for the second. For each sample, we add a tick mark above the horizontal line corresponding to its phase assigned by BBV classification. We also add a tick mark below the horizontal line that corresponds to each sample’s PMC phase. These marks then form the bands of phases seen in these plots. For example, for the case with 5 phases, low power phase of mesh is classified into phase “1” by BBVs and phase “3” by PMCs.

These plots show the significant impact of effectively same execution in phase classification. For $N = 5$, PMCs correctly identify the three actual power phases. BBVs on the other hand, collapse the high and low power phases into a single phase, leading to a false characterization. This is because, BBVs identify several different large-scale control flow phases. Clustering starts to overlap these based on their $L1$ -distances, and these result in combining the high and low phases of power. The three repetitive control flow phases with effectively same power behavior are seen as the more different phases by BBVs, and are assigned to different clusters. These

indicate several false alarms of spurious phase changes. For $N = 3$, BBV phases still show more sensitivity to the three repetitive tasks of medium power phase and assign them to three different phases. In this case, all high, low and parts of medium power phases are assigned to same phase (“1”) by BBVs. In contrast, PMCs show very good fidelity. They successfully identify three power regions and assign them to different phases.

This example demonstrates the clear impact of effectively same execution on control-flow-based phase characterization. It is important to note that, this effect has implications for not only phase characterizations, but also runtime phase detection. Various control-flow phases with similar power behavior can cause a detection framework to produce several false alarms for phase transitions.

There are also other cases where differences between PMC and BBV approaches arise including some SPEC benchmarks such as `mcf`. We do not present these here for brevity. Nonetheless, overall both BBV and PMC phases provide a good account of application power phase behavior; in many cases showing good correlation between power and both control flow and performance measures. PMCs usually show a better mapping to power behavior due to both their proximity to the actual flow of power in the processor, as well as due to these discussed sources of disagreement between power and code signatures.

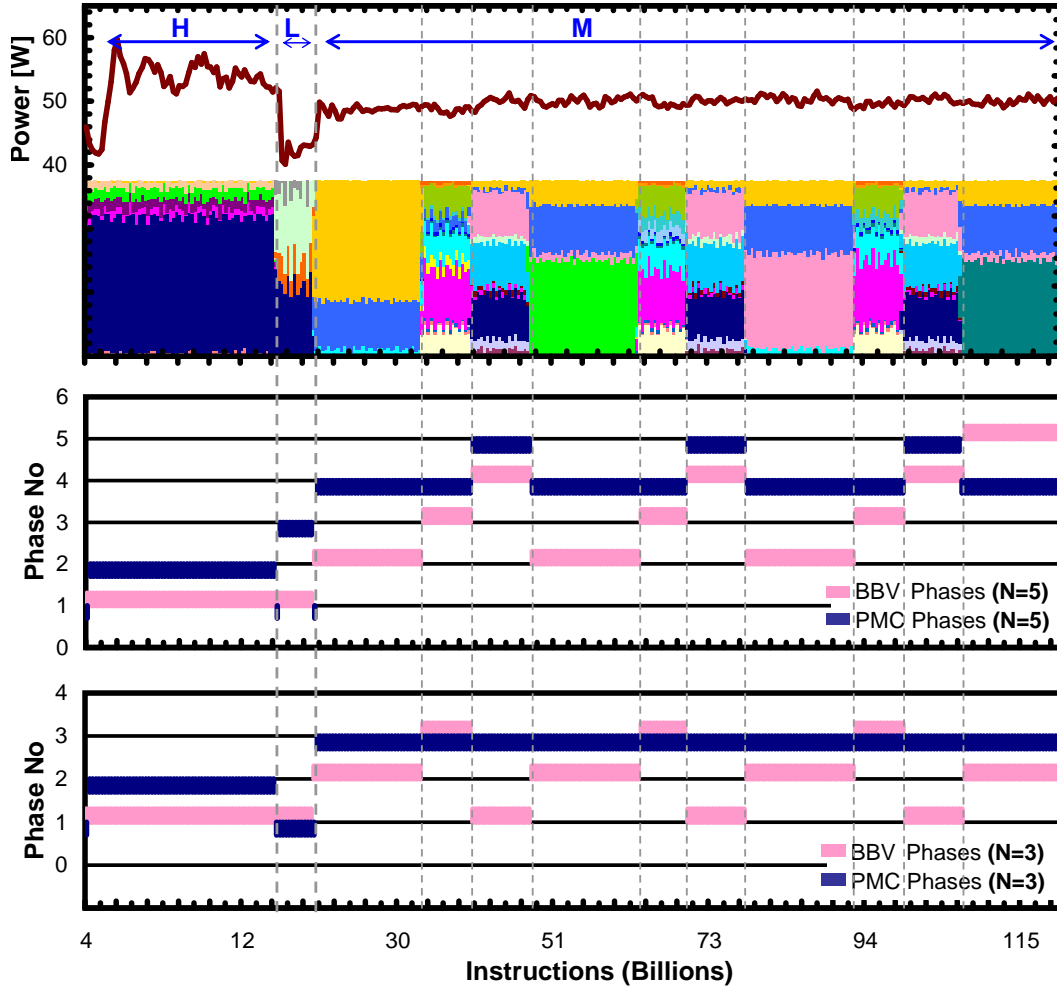


Figure 7. Mesh power and BBV signatures (top) and generated PMC and BBV phases with target cluster numbers of 5 (middle) and 3 (bottom). Multiple control flow phases with effectively same power characteristics disguise actual power phases in BBV based classification. Actual power phases are labeled as H , L and M , for high, low and medium power dissipation regions.

8 Summary and Recommendations for Future Research

Here, we first make a final comparison of BBVs and PMCs for power phase characterization. We summarize different pros and cons of the two approaches. Afterwards, we discuss how the strengths of these two features can be conceptually combined for application to dynamic management techniques. Finally, we present the directions for future research following from these discussions.

BBVs are known to have several benefits for summarizing application performance or tracking application phases. The most important advantage of BBVs is the repeatability of the observed phase behavior. Tracked code signatures do not change due to system effects or with the application of dynamic management actions that affect system power and performance.

The biggest disadvantage of BBVs lies in runtime applicability. It is impractical to collect full blown BBV information during application runtime. Sampling methods, as applied in this study, provide acceptable resolution, but BBV generation

still requires mapping PC samples into control flow blocks. Another related issue is the high dimensionality of BBVs that requires processing for dimension reduction. In addition to these, false alarms due to changes only in control flow are an important consideration for a runtime detection system. Finally, the indifference of BBVs to varying data locality and operand dependent behavior can be a significant impediment in power/performance phase characterizations for some real applications [2].

The important advantages of PMCs are their straightforward runtime applicability and their proximity to processor power consumption. PMCs are easily accessible at runtime with lightweight interfaces, which makes them good candidates for dynamic applications on real systems. Several PMCs show good correlations with processor power behavior, so they don't suffer significantly from false alarms. Also as the number of simultaneously monitored PMCs is less than 20, they require no dimension reduction during phase characterizations.

The most important issue with PMCs is repeatability. As PMC data comes from several event counts over the proces-

sor, the values are not identical among repetitions of phases. A phase detection method that utilizes PMCs is required to consider event count ranges or has to track deltas together with events to detect phases or phase changes. Our previous studies show quantization can be unreliable, while tracking changes produces higher fidelity [22]. The PMC based approach also requires range considerations. Different dimensions of PMC vectors are not of similar strength. For example, memory access counts and instructions issued have different orders of magnitude. Therefore, scaling of vector components or normalizations may be necessary to match the impact of certain events.

Our quantitative results show that PMCs have relatively higher fidelity in characterizing power phase behavior. However, we believe a better solution can be achieved by combining the strengths of BBVs with the PMC approach. For a general phase tracking method, we envision a hierarchical approach between BBVs and PMCs. We consider using PMC based phase tracking as the global mechanism to identify phase changes and using BBVs to track the repetitive execution progress. In terms of decision hierarchy, PMCs provide confidence to phase changes detected from control flow and provide the final decision whether this is an actual or spurious phase change. On the other hand, BBVs enhance the repeatability of observed PMC phases, by informing the PMC method when a repetitive control flow is detected.

As future research, we consider application of control-flow feedback to PMC-based phase tracking for runtime dynamic management. For example, we plan to use a DVFS enabled system and process phase information to detect repetitive behavior under different DVFS settings. Using both PMC and control flow features, we expect to achieve better power/performance trade-offs than a simple reactive response. In addition to this, other future research avenues include applying these power phase analysis methods to global CMP power management and to dynamic thermal management as mentioned in Section 2.

9 Related Work

Several previous studies investigate phase behavior of applications for adaptation and characterization purposes. Most of these research studies focus on either control flow or performance characteristics of applications. Iyer and Marculescu [24], Dhodapkar and Smith [9], Sherwood et al. [36, 37], Huang et al. [19] and Lau et al. [27] analyze control flow behavior of applications via different features such as subroutines, working sets and basic block profiles. These studies use simulation based methods to identify application phases for summarizing performance and architectural studies. Patil et al. [32] also look at control flow phases with real-system experiments. They use similar dynamic instrumentation to identify BBV phases of applications. Their work uses basic block profiles of applications to find representative execution points, while we look at power characterizations with BBV and PMC phases.

Cook et al. [7] show the repetitive performance phase characteristics of different applications using simulations. Todi [39], Weissel and Bellosa [40] and Duesterwald et al. [12] utilize performance counters to identify performance based phases. They use performance statistics to guide dynamic optimizations and metric predictions. These works do not con-

sider power behavior of applications. Our earlier work [20] employed runtime power measurements and power estimation with performance counters to identify phases of applications. Chang et al. [5] apply process power profiling to determine software power breakdowns. While these studies also look at power behavior, they do not investigate control flow approaches. Hu et al. [18] describe a compile time methodology to find basic block phases at runtime for power studies. This study looks at control flow information from a compiler perspective, while we investigate runtime power phase behaviors of both control flow and performance statistics.

There are also previous studies that compare or evaluate phase characterization techniques. Dhodapkar and Smith [8], perform a comparison between different control flow techniques, working set signatures and BBVs. Annavaram et al. [2] sample executed program counters as a proxy to control flow and show the correlations between code signatures and application performance. They show that, control flow does not always correlate well with application performance. Lau et al. [26] also look at control flow and performance of applications to show a strong correlation can be established by linking program counter to procedures and loops of applications via profiling. In comparison, our work looks at the direct comparison of two phase characterization features, BBVs and PMCs with runtime measurement feedback for real power evaluation on a real-system.

10 Conclusion

Phase analysis is increasingly important for computer systems first because simulation-based techniques rely on phase-directed sampling to reduce simulation time, and second, because real-life adaptive hardware and software mechanisms rely on dynamic phase-directed readjustments.

With power being such a pressing constraint in current processors, it becomes important to understand not just the phases of performance metrics, but also of their related power counterparts. Observing power phase behavior on real systems is particularly important because the real-system phases show the impact of a comprehensive range of systems effects typically excluded from simulations.

This work has explored methods for real-system power phase generation. Drawing on prior work, we have developed an experimental framework for comparing both control-flow-based and performance-monitoring-based phase techniques, and for comparing against live power measurements. Our results show that both control-flow and performance statistics provide useful hints to power phase behavior. In general, performance-based phase tracking leads to approximately 33% less power characterization errors than code signatures.

In some cases where power behavior depends on aspects other than control flow (e.g. data locality, operand values, or other characteristics), phases based on control flow can “miss” some transitions. In other cases, control flow phase classification can result in “extra” phases, where applications perform different tasks with effectively the same execution characteristics. These effects lead to both false alarms for phase changes and incorrect power phase classifications.

Overall, the results presented here show a roadmap to effective power phase analysis in real systems. Control-flow techniques offer a good base, but may well be best applied as hybrid techniques together with performance counters that

can more closely track the details of program behavior, needed for detection of power phases with high fidelity.

Acknowledgments

We would like to thank Gilberto Contreras, Belma Dogdas, James Donald, Qiang Wu and the members of the Pin mailing group for their help during the development of this work and for several useful discussions. We also thank Chen Ding for his help with the benchmarks, and the anonymous reviewers for their useful suggestions. This research was supported by NSF grants CCR-0086031 (ITR) and CNS-0410937. Martonosi's work is also supported in part by Intel, IBM, and SRC.

References

- [1] D. Albonesi, R. Balasubramonian, S. Dropsho, S. Dwarkadas, E. Friedman, M. Huang, V. Kursun, G. Magklis, M. Scott, G. Semeraro, P. Bose, A. Buyuktosunoglu, P. Cook, and S. Schuster. Dynamically Tuning Processor Resources with Adaptive Processing. *IEEE Computer*, 36(12):43–51, 2003.
- [2] M. Annavaram, R. Rakvic, M. Polito, J.-Y. Bouguet, R. Hankins, and B. Davies. The Fuzzy Correlation between Code and Performance Predictability. In *Proceedings of the 37th annual International Symp. on Microarchitecture*, 2004.
- [3] F. Bellosa, A. Weissel, M. Waitz, and S. Kellner. Event-Driven Energy Accounting for Dynamic Thermal Management. In *Proceedings of the Workshop on Compilers and Operating Systems for Low Power (COLP'03)*, New Orleans, Sept. 2003.
- [4] D. Brooks and M. Martonosi. Dynamically exploiting narrow width operands to improve processor power and performance. In *Proceedings of the 5th International Symposium on High Performance Computer Architecture*, Jan. 1999.
- [5] F. Chang, K. Farkas, and P. Ranganathan. Energy driven statistical profiling: Detecting software hotspots. In *Proceedings of the Proceedings of the Workshop on Computer Systems*, 2002.
- [6] K. Choi, R. Soma, and M. Pedram. Dynamic Voltage and Frequency Scaling based on Workload Decomposition. In *Proceedings of International Symposium on Low Power Electronics and Design (ISLPED)*, Aug. 2004.
- [7] J. Cook, R. L. Oliver, and E. E. Johnson. Examining performance differences in workload execution phases. In *Proceedings of the IEEE International Workshop on Workload Characterization (WWC-4)*, 2001.
- [8] A. Dhodapkar and J. Smith. Comparing Program Phase Detection Techniques. In 36th International Symp. on Microarchitecture, 2003.
- [9] A. Dhodapkar and J. Smith. Managing multi-configurable hardware via dynamic working set analysis. In 29th Annual International Symposium on Computer Architecture, 2002.
- [10] C. Ding and K. Kennedy. Improving Cache Performance in Dynamic Applications through Data and Computation Reorganization at Run Time. In *Proceedings of ACM SIGPLAN Conference on Programming Language Design and Implementation*, 1999.
- [11] R. O. Duda, P. E. Hart, and D. G. Stork. *Pattern Classification. Second Edition*. Wiley Interscience, New York, 2001.
- [12] E. Duesterwald, C. Cascaval, and S. Dwarkadas. Characterizing and Predicting Program Behavior and its Variability. In *IEEE PACT*, pages 220–231, 2003.
- [13] L. Eeckhout, R. Sundareswara, J. Yi, D. Lilja, and P. Schrater. Accurate Statistical Approaches for Generating Representative Workload Compositions. In *Proceedings of the IEEE International Symposium on Workload Characterization*, Oct. 2005.
- [14] F. Ercolessi. MDBNCH - A molecular dynamics benchmark. International School for Advanced Studies in Trieste. www.fisica.uniud.it/ercolessi/mdbnch.html.
- [15] GIMP. GNU Image Manipulation Program. <http://www.gimp.org/>.
- [16] S. Heo, K. Barr, and K. Asanovic. Reducing Power Density through Activity Migration. In *Proceedings of International Symposium on Low Power Electronics and Design (ISLPED)*, Seoul, Korea, Aug. 2003.
- [17] M. J. Hind, V. T. Rajan, and P. F. Sweeney. Phase Shift Detection: A Problem Classification. IBM Research Report RC-22887, IBM T. J. Watson, Aug. 2003.
- [18] C. Hu, D. Jimenez, and U. Kremer. Toward an Evaluation Infrastructure for Power and Energy Optimizations. In *Workshop on High-Performance, Power-Aware Computing*, 2005.
- [19] M. Huang, J. Renau, and J. Torrellas. Positional Adaptation of Processors: Application to Energy Reduction. In *Proceedings of the International Symp. on Computer Architecture*, 2003.
- [20] C. Isci and M. Martonosi. Identifying Program Power Phase Behavior using Power Vectors. In *Proceedings of the IEEE International Workshop on Workload Characterization (WWC-6)*, 2003.
- [21] C. Isci and M. Martonosi. Runtime Power Monitoring in High-End Processors: Methodology and Empirical Data. In *Proceedings of the 36th International Symp. on Microarchitecture*, Dec. 2003.
- [22] C. Isci and M. Martonosi. Detecting Recurrent Phase Behavior under Real-System Variability. In *Proceedings of the IEEE International Symposium on Workload Characterization*, Oct. 2005.
- [23] C. Isci, M. Martonosi, and A. Buyuktosunoglu. Long-term Workload Phases: Duration Predictions and Applications to DVFS. *IEEE Micro: Special Issue on Energy Efficient Design*, 25(5):39–51, Sep/Oct 2005.
- [24] A. Iyer and D. Marculescu. Power aware microarchitecture resource scaling. In *Proceedings of Design Automation and Test in Europe, DATE*, Mar. 2001.
- [25] R. Jenkins. Hash functions. *Dr. Dobbs's Journal*, 9709, Sept. 1997.
- [26] J. Lau, J. Sampson, E. Perelman, G. Hamerly, and B. Calder. The Strong Correlation between Code Signatures and Performance. In *IEEE International Symposium on Performance Analysis of Systems and Software*, Mar. 2005.
- [27] J. Lau, S. Schoenmackers, and B. Calder. Transition Phase Classification and Prediction. In *11th International Symposium on High Performance Computer Architecture*, 2005.
- [28] C. Lee, M. Potkonjak, and W. H. Mangione-Smith. MediaBench: A Tool for Evaluating and Synthesizing Multimedia and Communications Systems. In *Proceedings of the 30th annual International Symposium on Microarchitecture*, 1997.
- [29] C. Luk, R. Cohn, R. Muth, H. Patil, A. Klauser, G. Lowney, S. Wallace, V. Reddi, and K. Hazelwood. Pin: Building Customized Program Analysis Tools with Dynamic Instrumentation. In *Programming Language Design and Implementation (PLDI)*, June 2005.
- [30] G. Magklis, M. L. Scott, G. Semeraro, D. H. Albonesi, and S. Dropsho. Profile-based Dynamic Voltage and Frequency Scaling for a Multiple Clock Domain Processor. In *Proceedings of the 30th International Symposium on Computer Architecture*, June 2003.
- [31] J. McCalphin. STREAM: Sustainable Memory Bandwidth in Current High Performance Computers. Technical report, University of Virginia, 1995.
- [32] H. Patil, R. Cohn, M. Charney, R. Kapoor, A. Sun, and A. Karunanidhi. Pinpointing Representative Portions of Large Intel Itanium Programs with Dynamic Instrumentation. In *Proceedings of the 37th annual International Symp. on Microarchitecture*, 2004.
- [33] M. Powell, M. Gomaa, and T. N. Vijaykumar. Heat-and-run: Leveraging SMT and CMP to manage power density through the operating system. In *Eleventh International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS XI)*, 2004.
- [34] X. Shen, Y. Zhong, and C. Ding. Locality Phase Prediction. In *Eleventh International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS XI)*, Oct. 2004.
- [35] T. Sherwood, E. Perelman, and B. Calder. Basic block distribution analysis to find periodic behavior and simulation points in applications. In *International Conference on Parallel Architectures and Compilation Techniques*, Sept. 2001.
- [36] T. Sherwood, E. Perelman, G. Hamerly, and B. Calder. Automatically characterizing large scale program behavior, 2002. In Tenth International Conference on Architectural Support for Programming Languages and Operating Systems, October 2002. <http://www.cs.ucsd.edu/users/calder/simpoint/>.
- [37] T. Sherwood, S. Sair, and B. Calder. Phase tracking and prediction. In *Proceedings of the 28th International Symposium on Computer Architecture (ISCA-30)*, June 2003.
- [38] Sourceforge.net. The LAME Project. <http://www.mp3dev.org/>.
- [39] R. Todi. Speclite: using representative samples to reduce spec cpu2000 workload. In *Proceedings of the IEEE International Workshop on Workload Characterization (WWC-4)*, 2001.
- [40] A. Weissel and F. Bellosa. Process cruise control: Event-driven clock scaling for dynamic power management. In *Proceedings of the International Conference on Compilers, Architecture and Synthesis for Embedded Systems (CASES 2002)*, Grenoble, France., Aug. 2002.
- [41] Q. Wu, V. Reddi, Y. Wu, J. Lee, D. Connors, D. Brooks, M. Martonosi, and D. W. Clark. A Dynamic Compilation Framework for Controlling Microprocessor Energy and Performance. In *Proceedings of the 38th annual International Symp. on Microarchitecture*, 2005.