

Pac-Sim: Simulation of Multi-threaded Workloads using Intelligent, Live Sampling

CHANGXI LIU*, National University of Singapore, Singapore

ALEN SABU*, National University of Singapore, Singapore

AKANKSHA CHAUDHARI, University of Wisconsin-Madison, USA

QINGXUAN KANG, National University of Singapore, Singapore

TREVOR E. CARLSON, National University of Singapore, Singapore

High-performance, multi-core processors are the key to accelerating workloads in several application domains. To continue to scale performance at the limit of Moore's Law and Dennard scaling, software and hardware designers have turned to dynamic solutions that adapt to the needs of applications in a transparent, automatic way. For example, modern hardware improves its performance and power efficiency by changing the hardware configuration, like the frequency and voltage of cores, according to a number of parameters, such as the technology used or the workload running at the time. With this level of dynamism, it is essential to simulate next-generation multi-core processors in a way that can both respond to system changes and accurately determine system performance metrics. Currently, no sampled simulation platform can achieve these goals of dynamic, fast, and accurate simulation of multi-threaded workloads.

In this work, we propose a solution that allows for fast, accurate simulation in the presence of both hardware and software dynamism. To accomplish this goal, we present Pac-Sim, a novel sampled simulation methodology for fast, accurate sampled simulation that requires no upfront analysis of the workload. With our proposed methodology, it is now possible to simulate long-running dynamically scheduled multi-threaded programs with significant simulation speedups, even in the presence of dynamic hardware events. We evaluate Pac-Sim using the SPEC CPU2017, NPB, and PARSEC multi-threaded benchmarks with both static and dynamic thread scheduling. The experimental results show that Pac-Sim achieves a very low sampling error of 1.63% and 3.81% on average for statically and dynamically scheduled benchmarks, respectively. Pac-Sim also demonstrates significant simulation speedups as high as $523.5\times$ ($210.3\times$ on average) for the training input set of SPEC CPU2017 running eight threads.

CCS Concepts: • **Computing methodologies** → **Simulation tools; Simulation evaluation**; • **Computer systems organization** → **Multicore architectures**.

Additional Key Words and Phrases: Sampled Simulation, Multi-threaded Workloads, Dynamic Scheduling

1 INTRODUCTION

Computer architecture research heavily relies on simulations for design space exploration. However, microarchitectural simulation can be extremely time-consuming, particularly due to the ever-growing complexity of modern architectures. This is especially true in the *post-Dennard era*, where architectures are rapidly evolving to incorporate complex dynamic optimization techniques at both the hardware and software levels to improve system performance gains at runtime. Hardware-based dynamic techniques such as dynamic cache reconfiguration [3, 49, 50], DVFS [32, 41, 44], TurboBoost [19], and power management [10] techniques trigger optimizations based on dynamically identified hardware states (such as core frequency, cache reuse distance, etc.) to improve both energy-efficiency and overall performance of the system. Similarly, runtime information at the software level can be used to dynamically optimize code execution to further enhance the system performance. Some of the recent efforts on software-based optimization focus on dynamically scheduling tasks among threads [26, 28, 75] to ensure efficient resource utilization and employing

*Both authors contributed equally to this research.
New Paper, Not an Extension of a Conference Paper.

just-in-time (JIT) compilation techniques [21, 45, 69, 72] that generate high-performance instructions to optimize program execution during runtime. However, since these techniques utilize dynamic system state information in order to deploy optimizations at runtime, the execution behavior of an application (and, therefore, its performance) may vary greatly across multiple executions. This inherent variability may lead to an inaccurate performance evaluation when using existing simulation methodologies.

Conventionally, sampled simulation has served as a reliable and efficient technique to accelerate the performance estimation of multi-threaded workloads. In order to achieve low-latency, accurate results, most prior works relied on either (i) profile-driven sampling [17, 60, 65] or (ii) statistical sampling [42, 71] techniques. Profile-driven sampled simulation methodologies such as SimPoint [65], BarrierPoint [17], and LoopPoint [60] split the execution of an application into a series of repeatable regions and cluster them based on their execution features. A representative element from each cluster is then analyzed or simulated in order to extrapolate the performance of the entire application. However, these methodologies incur a significant cost in terms of the preprocessing effort that is needed to identify representative regions. These costs include the time required to profile and cluster the execution features of all application regions, along with the storage required. While it has been previously argued that these costs are a one-time investment and will be amortized over multiple runs, this argument does not necessarily hold for systems that optimize code execution dynamically. In such cases, the program execution paths followed by an application may vary considerably due to changes in hardware and software parameters that are being optimized. Therefore, the profiling information collected for one specific run would not necessarily extend to the program execution paths followed in the subsequent runs.

On the other hand, sampled simulation methodologies such as SMARTS [71] and PCantorSim[42] rely on statistical sampling techniques to speed up simulation-based performance measurements while meeting a given error bound. Unlike profile-driven sampling, these methodologies require minimal preprocessing and do not rely on the reproducibility of program execution paths. They are thus applicable to dynamically optimized systems. However, the simulation speedups achieved using these techniques are considerably lower than the profile-driven counterparts, and adjusting settings to achieve higher performance could lead to high errors.

For the above-mentioned reasons, it becomes challenging to sample and simulate generic multithreaded applications for dynamic hardware and software using existing methodologies. Architects need a simulation methodology that can dynamically adapt to changes in the system at runtime while accurately estimating the application’s performance without relying on the reproducibility of its execution. To this end, we propose Pac-Sim, a novel sampled simulation methodology that can, at runtime, efficiently analyze and sample the application to select the representative regions to be simulated in detail. The result is a methodology that enables both fast and accurate performance evaluation without the need for up-front analysis. We accomplish this by making use of an intelligent online¹ predictor and classifier that quickly and accurately decides whether the upcoming region needs to be simulated in detail.

In short, we make the following contributions:

- i. We propose Pac-Sim, a methodology that goes beyond prior sampled simulation techniques to be the first to allow for dynamic hardware and software support. The methodology² requires no upfront analysis and relies on an online predictor for sampling decisions, enabling the fast analysis of co-designed workloads.

1. We use the terms “online” and “offline” to distinguish between events that occur during and prior to the simulation of an application, respectively.

2. Pac-Sim has been open-sourced and can be found at <https://github.com/snipersim/snipersim>.

- ii. We experimentally demonstrate that Pac-Sim consistently improves performance in terms of speedup and accuracy over prior works that use offline profiling as Pac-Sim utilizes a lightweight but accurate online sampling technique.
- iii. We provide an extensive evaluation of Pac-Sim using standard benchmarks to compare against prior works and demonstrate best-in-class accuracy (average error of 1.63%). For the SPEC CPU2017 benchmarks (training inputs) running eight threads, we show a maximum serial speedup of 123.32 \times (26.09 \times on average) and a maximum parallel speedup of 523.5 \times (210.3 \times on average).
- iv. Finally, we showcase several case studies demonstrating that Pac-Sim is applicable to a number of research scenarios, including (but not limited to) the investigation of optimization techniques such as dynamically scheduled software and improving research into dynamic hardware and hardware-software co-design.

The rest of the paper is organized as follows. In Section 2, we discuss the relevant background and the challenges involved in the simulation of dynamic applications on modern architectures. Section 3 presents the Pac-Sim methodology in detail. We then describe the experimental infrastructure in Section 4, followed by an extensive evaluation of Pac-Sim in Section 5 along with case studies to demonstrate the applicability of the proposed methodology. Finally, we conclude the paper in Section 7.

2 SIMULATING MODERN ARCHITECTURES

In this section, we provide the necessary background and prior work on sampled simulation. Table 1 summarizes the widely used sampled simulation methodologies applicable to CPU workloads. We also discuss the challenges in simulating modern workloads and how the existing sampling methodologies are insufficient to address them.

Sampling Single-threaded Workloads. Sampling and workload reduction techniques are extensively utilized in computer architecture research for the purpose of program characterization and to reduce simulation time. Sampling

Table 1. Previously proposed sampled simulation methodologies for both single-threaded and multi-threaded applications. We categorize these methodologies into two main groups: profile-driven and statistical. The table also identifies the ‘Analysis Type’ used by each methodology. Notably, some methodologies may require only a lightweight analysis phase, whereas some require a lengthy profiling phase to extract application-specific characteristics. Additionally, the table indicates the methodologies that are amenable to parallel simulation, which determines the maximum speedup of the methodology. The field ‘Warmup’ shows the warmup technique used to reconstruct the microarchitectural state at the beginning of the detailed simulation.

	Methodology	Analysis Type	Parallel Simulation	Warmup	Applicability
Profile-driven	Simpoint [65]	○	✓	Prev Region	Single-threaded
	LiveSim [40]	○	✓	Checkpoint	Single-threaded
	BarrierPoint [17]	○	✓	Prev Region	Multi-threaded
	TaskPoint [37]	○	✓	Prev Region	Task-based
	LoopPoint [60]	○	✓	Prev Region	Multi-threaded
Statistical	SMARTS [71]	●	✗	Functional	Single-threaded
	SimFlex [70]	●	✗	Checkpoint	Multi-program
	Time-Based Sampling [4, 15]	●	✗	Functional	Multi-threaded
	Pac-Sim (this work)	●	✓	Statistical	Multi-threaded

● Online Profiling ● Offline Analysis ○ Offline Profiling

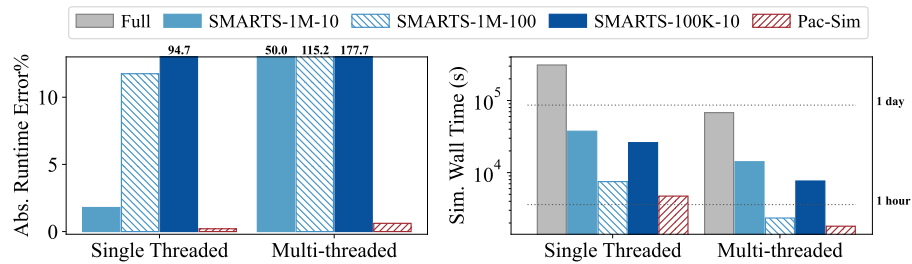


Fig. 1. Performance comparison of Pac-Sim with SMARTS [71] in different settings for the SPEC benchmark 644.nab_s.1 (multi-threaded version uses 8 threads). The left graph shows the runtime prediction errors compared to the full simulation using different sampled simulation techniques, whereas the right graph shows the overall simulation time (running on a parallel simulator). Both figures use lower-is-better metrics. *SMARTS-A-B* repeatedly switches between a single detailed simulation region of length A and B fast-forward regions of length A .

methodologies allow for the evaluation of a subset of the workload (a representative sample) in detail that can be used to reconstruct the performance of the whole workload accurately. These methodologies split the workload into different regions (or slices) based on predetermined conditions in order to identify a representative sample. Prior works that explored CPU workload sampling, like SimPoint [65] and SMARTS [71], tend to utilize fixed instruction counts to determine regions. However, instruction count-based techniques could lead to inconsistent and, therefore, invalid regions [1, 2, 15]. Some previous works [46, 73] proposed software phase markers that identify procedure and loop boundaries that correlate with phase changes to mark region boundaries instead of using fixed-sized regions.

Sampling Multi-threaded Workloads. SimFlex [70], which selects sampling units for the simulation of server throughput workloads, does not appear to be generally extensible to synchronizing multi-threaded workloads [15]. In the presence of synchronizing threads, the application performance tends to vary more frequently [2, 15]. Sampling methodologies such as SMARTS, SimFlex, and PCantorSim [42] rely on statistical confidence, and adjusting settings to achieve higher performance could lead to high errors, as shown in Figure 1. Time-Based Sampling methodologies [4, 15] are the first to address the problem of sampling synchronizing multi-threaded applications. These methodologies are generally applicable and are suitable for the sampled simulation of dynamic systems. However, Time-Based Sampling techniques are the slowest of all the sampling techniques, and as a result, they are not practical for handling long-running workloads. On the other hand, profile-driven sampling methodologies for multi-threaded applications like BarrierPoint [17], TaskPoint [37], and LoopPoint [60] select specific program constructs, such as barrier synchronization primitives, task instances, and loops, respectively, to identify periodic behavior. This enables the utilization of representative-sized regions for simulation, regardless of the program’s length.

Feature Vectors. Profiling captures feature vectors to characterize the execution behavior of an application across regions. Previous works have introduced several microarchitecture-independent feature vectors, of which basic block vectors (BBVs) [58, 65] are the most widely used for performance characterization. Lau et al. [47] showed a strong correlation between BBVs and region performance. Apart from BBVs, Shen et al. [64] introduced LRU stack distance vectors (LDVs) [48] to summarize program behavior for different regions. BarrierPoint [17] combines BBVs and LDVs into a signature vector (SV) in an attempt to represent more accurate features of multi-threaded applications. Furthermore, Cotson [5] and Dynamic sampling [35] record statistics such as the number of instructions executed, memory accesses,

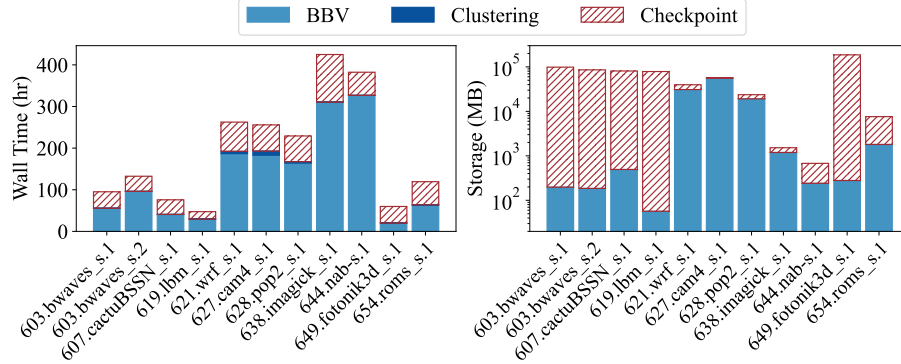


Fig. 2. The figure shows the resource utilization of a previously proposed multi-threaded sampled simulation technique, LoopPoint, for the SPEC CPU2017 benchmarks with the ref inputs running eight OpenMP threads. The graph on the left shows the time required to generate the profiling data (with checkpoints stored as pinballs [55]), whereas the graph on the right shows the amount of storage required.

exceptions, bytes read or written, etc., in order to plot the feature of a given region. Unfortunately, none of these offline techniques can handle runtime optimizations that impact applications.

Overheads. Figure 2 illustrates the overhead of profiling data for LoopPoint (the evaluation was performed using the LoopPoint tools [61]) methodology, indicating that profile-driven methodologies incur significant overheads. When it is required to emulate an architecture (for example, simulating ARM or RISC-V binaries on x86) during profiling, it is necessary to resort to functional simulation to gather feature vectors, which can be a time-consuming process. For instance, Sandberg et al. [63] demonstrated that it took up to a month to generate profile data for SPEC CPU2006 benchmarks using simulators like gem5. Prior works [17, 60] argue that this overhead is amortized over multiple runs as the profiling results will be reused. However, this assertion does not hold in the case of dynamic software and hardware where certain performance optimization decisions are made using runtime information. For example, profiling for asymmetric cores, such as the *big.LITTLE* cores is challenging as the operating frequency (and other dynamic hardware settings) of each core may not be known during profiling. Handling and storing simulation architecture/microarchitecture checkpoints can be another daunting task. For instance, x86 architecture checkpoints like ELFies [56] require a significant amount of storage space. Checkpoints are often specific to a simulator or are tied to particular software/hardware configurations. Microarchitectural checkpoints, requiring detailed hardware information like cache states, are specific to the underlying hardware configuration.

Hardware and Software Dynamism. Researchers have introduced several dynamic optimization techniques in hardware and software to achieve higher performance and reduce power consumption. Techniques such as dynamic voltage and frequency scaling (DVFS) and cache reconfiguration have been developed to adjust the hardware state in response to executed instructions and active processes. Software optimization techniques [21, 45, 69, 72] generate optimized code sequences at runtime. Additionally, dynamic scheduling techniques [28] have been developed for multi-threaded applications. In such cases, profile-driven sampling methodologies may end up with different performance results for each execution. Methodologies such as trace-based simulations [13] or deterministic replay platforms [57] can guarantee consistent performance across multiple executions but demand extensive profiling and large storage resources. Dynamic hardware events, such as changes in core frequency, cache size, etc., can be unknown during

profiling. These events, when they are performance and power-dependent, become difficult to predict. Sherwood et al. [66] utilize a Markov predictor to predict the phase behavior at runtime. Kihm et al. [43] propose switching to the detailed simulation mode whenever the BBV variance exceeds a specified threshold. However, these methods have only been demonstrated using single-threaded applications as the phase behavior of synchronizing multi-threaded applications varies frequently due to the interaction of threads [2, 68].

Requirements for Fast and Accurate Simulation. Sampled simulation techniques that do not require upfront application analysis demonstrate significant potential under dynamic software and hardware constraints. The inherent variability of dynamic software behavior renders a single analysis insufficient, while the unpredictable nature of modern hardware compromises the reliability of upfront analysis. Additionally, the overhead due to the detailed application analysis becomes a bottleneck for researchers engaged in fields like hardware-software co-design. Therefore, it is imperative to leverage the best aspects of SimPoint-like and SMARTS-like methodologies to achieve optimal simulation efficiency and accuracy. In our approach, we integrate application analysis to guide sampled simulations, similar to SimPoint-like methodologies, but without the need for upfront preprocessing, as seen in SMARTS-like methodologies. In this work, we capture the dynamic information of the software and hardware to make intelligent simulation decisions through online learning. Therefore, our methodology is capable of handling hardware state changes, software dynamism, and other factors influencing application performance. To achieve optimal performance with online analysis, efficient and lightweight profiling, clustering, and warmup techniques are essential.

In short, to quickly estimate the performance of multithreaded applications running on next-generation hardware, a sampled simulation methodology is needed that can dynamically adapt to changes in the system at runtime while accurately determining relevant performance metrics. In Section 3, we will provide a thorough discussion of these aspects and present our solution for fast and accurate simulation.

3 THE PAC-SIM METHODOLOGY

In this section, we describe our proposal for an end-to-end sampled simulation methodology, Pac-Sim (see Figure 3), that supports both dynamic hardware and software without requiring up-front workload analysis. Pac-Sim consists of five main stages: Marker Detection, Region Profiling, Clustering, Prediction, and Simulation, which are all carried out online. We have carefully designed each of these stages to minimize the runtime overhead of the methodology while maintaining the sampling accuracy. An important advantage of an online sampled simulation methodology like Pac-Sim is its ability to accurately determine the execution profile of an application without relying on the reproducibility of a program’s execution paths. This characteristic allows Pac-Sim to accurately analyze and evaluate dynamic multi-threaded applications, accounting for any performance variability that may occur at runtime.

Pac-Sim operates by making use of the program structure and runtime hardware state to identify the regions and their boundaries online. Each of these region boundaries or *markers* defines the ending of the current region and the beginning of the next region (Section 3.1). Once a marker is identified, Pac-Sim collects the profiling data and simulation results of the current region (Section 3.2) and clusters it with the previously identified regions to determine its cluster ID (Section 3.3). This cluster ID is added to the program execution history, which is then used by the *Predictor* (Section 3.4) along with the current marker and hardware state to predict whether the next region needs to be simulated in detailed mode or fast-forward mode.

While we only demonstrate the effectiveness of Pac-Sim in estimating the performance of synchronizing multi-threaded workloads in this work, our methodology has the potential to support a variety of modern workload classes, such as cloud and mobile applications, and could also be implemented for full system simulations. However, in such

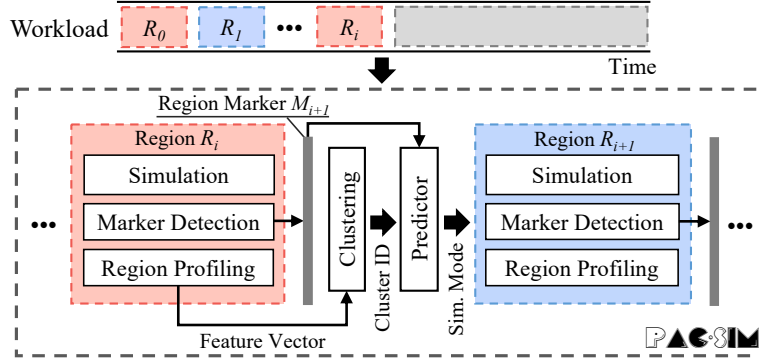


Fig. 3. The overall workflow of Pac-Sim methodology. At any given time, the regions of a multi-threaded workload till R_i are identified (as shown above). First, Pac-Sim monitors the application code structure to determine an appropriate region marker M_{i+1} , which marks both the end of the region R_i and the start of the region R_{i+1} . Next, the feature vector and simulation results for R_i are collected, and a prediction mechanism determines the simulation mode for region R_{i+1} . Finally, region R_{i+1} will be simulated, either in detail or in fast-forward mode.

cases, various factors must be taken into consideration, such as kernel and driver performance, which can significantly impact the overall efficiency of the workloads. In this work, we focus on user-space workloads, and enabling support for the above-mentioned use cases is out-of-scope in this context which we leave for future work.

3.1 Online Region Detection

Previous research [17, 46, 60] has shown that certain program constructs, such as barriers or loops, can be utilized to characterize the phase behavior of multi-threaded applications by splitting them into a series of individually analyzable regions. Since barriers represent the global synchronization points within a program execution, all threads align at these points, making them natural boundaries for application regions. However, relying solely on barriers to split an application may not be ideal, especially in the presence of large inter-barrier regions, as this can lead to low simulation speedups as representatives can still be too large to complete detailed simulation in a reasonable amount of time. In contrast, loops offer a finer level of granularity, allowing for greater control over the size of regions. Typically, multi-threaded applications consist of both loops and barriers in varying proportions. The online *Marker Detector* combines both of these program constructs to effectively split multi-threaded applications into regions with sizes that are well-suited for clustering while also avoiding aliasing [18]. The Marker Detector uses the following approach in order to identify the barrier- and loop-based markers online:

Barriers. Typically, a multi-threaded region begins with a fork call, which spawns additional worker threads and ends with a join call, which terminates the current thread and synchronizes with other threads. A new region is triggered at events of thread creation and termination, as regions with different active threads have different performances. For multi-threaded programs that use the OpenMP library, special function names are generated depending on the compiler used. We utilize this information in the online Marker Detector to quickly and efficiently detect barriers with low overhead.

Loops. Both loop and conditional statements use conditional branch instructions, with the target address usually given as an offset from the instruction pointer. The key difference between the two statements is that the offset of the branch instructions in a loop statement is usually negative, whereas that in a conditional statement is positive. In

most cases, selecting conditional branches with negative offsets is sufficient to identify loop markers [73]. In the rare case of exceptions, Pac-Sim predicts the simulation mode of the next region to be detailed mode. We also make sure to disregard spinloops from our analysis.

As an application executes, the Marker Detector identifies markers online, splitting the application into multiple regions. While doing so, it also monitors the region sizes to ensure they fall approximately within the bounds of δ_{min} and δ_{max} instructions. A minimum number of instructions, δ_{min} , is necessary to capture the frequent variations in the multi-threaded program behavior and accurately cluster the obtained regions. Whenever the Marker Detector chooses barrier-based markers as region boundaries, the size of the region can be as small as δ_{min} instructions but no larger than δ_{max} instructions. Otherwise, the Marker Detector chooses the first loop-based marker it encounters beyond δ_{max} instructions as the next region boundary. For loop-bounded regions, it is necessary to keep region sizes large enough to avoid aliasing [15]. In our experiments with fixed region sizes of 10 million, 20 million, 50 million, and 100 million instructions, the SPEC CPU2017 benchmarks showed average error rates of 6.9%, 3.3%, 1.8%, and 1.8%, respectively. We set the lower bound δ_{min} to be 20 million to ensure sampling accuracy and the upper bound δ_{max} to be 50 million for better performance.

Hardware State. The Marker Detector also monitors the hardware state of the simulated system. If it detects changes, the current region is ended at the next marker so that each region has a consistent hardware state. Once a marker is detected, the program counter (PC) and the hardware state of the simulated system are collected and stored corresponding to the marker. The collected hardware state includes the system parameters, like processor frequency, cache size/configuration, power management techniques, etc., that can be configured during runtime.

3.2 Online Region Profiling

Conventionally, BBVs have been used to characterize the execution behavior of code regions, as they have been shown to exhibit a strong correlation with the region’s performance [47]. BBVs record the execution counts of each basic block (i.e., code blocks with single entry and exit points) within a given code region. The number of dimensions for a BBV depends on the number of basic blocks executed, which could range anywhere from thousands to even millions for very large applications. This presents a major challenge for the online analysis of BBVs as the time and effort required for this stage would significantly increase as the vector dimensionality increases. SimPoint [65] uses random linear projection [22] to overcome this problem. However, this method is not suitable for our online algorithm as the matrix-vector multiplication operations involved could introduce significant runtime overheads.

To overcome these issues, we propose a fast online BBV generation technique (illustrated in Figure 4). Rather than creating a fixed-size BBV for each region, we use an online projection technique to generate fixed-size vectors BBV'_i for each basic block BB_i , where the elements of BBV'_i are computed by multiplying the instruction count of a basic block with the hash results of its program counter (PC) value. We use the hash function `drand48()`, which generates pseudo-random numbers for an integer value input. The initial four dimensions of the online BBV are determined using the hash values utilizing inputs PC, PC+1, PC+2, and PC+3, respectively. The values of the subsequent four dimensions are generated using the output of the preceding four dimensions as inputs to the hash function. We experimentally determined that using 16 dimensions adequately captures the representation of a region using the online BBV. The resultant BBV'_i vectors are then accumulated to obtain the per-thread BBV (BBV'_{online}) for the given region, which can be represented as:

$$BBV'_{online} = \sum_i BBV'_i = \sum_i (BBV_i \cdot M_{proj}),$$




Fig. 4. The figure shows the workflow of online BBV generation. Whenever a basic block BB_i is encountered, a corresponding execution fingerprint BBV_i is generated using hash functions applied to the program counter of BB_i and the number of instructions it contains. $hash_1$ to $hash_d$ are d distinct hash functions, where d is the dimension of the BBV. The BBV for each region is obtained by accumulating all BBV_i s that belong to the region.

where the values of the elements in M_{proj} are generated using hash functions as mentioned above. This BBV'_{online} for a region is analogous to the BBV utilized in SimPoint, which is obtained through random linear projection. The projected down BBV used in SimPoint, $BBV'_{offline}$, is obtained from the dot product of the actual BBV of the region and projection matrix M_{proj} : $BBV'_{offline} = BBV \cdot M_{proj} = \sum_i (BBV_i \cdot M_{proj})$. We then normalize these per-thread BBVs and concatenate them into a single global-BBV to represent the software feature of a given multi-threaded code region. In Pac-Sim, it is necessary to maintain the online BBV to capture the dynamic program behavior. Using online BBVs to represent regions eliminates the need to perform computationally intensive dimensionality reduction techniques during simulation.

3.3 Determining Region Similarity

Pac-Sim employs an online clustering mechanism to group regions with similar execution behavior based on the feature vectors collected for each region in the online profiling stage. The clustering, which is done at the end of simulating each region, is required for the learning process of the Predictor. Prior works, like SimPoint [65], cluster feature vectors using the k-means algorithm [36]. However, k-means uses an iterative refinement technique that is computationally intensive, and therefore, a more efficient algorithm might be better suited for online analysis.

In order to reduce this computational load and enable real-time clustering, we devise an alternative technique for clustering feature vectors (i.e., global-BBVs) in Pac-Sim. In our technique, we maintain two separate queues: (i) detailed queue and (ii) fast-forward queue. The detailed queue includes the BBVs corresponding to the regions that have been simulated in detail, while the fast-forward queue includes those corresponding to the regions that have been fast-forwarded. When a new BBV is recorded, it is first compared with the BBVs in the detailed queue. If its distance from any of these BBVs is less than the specified threshold θ , then we return the cluster ID of the closest region. If there is no region whose distance is less than θ , we repeat the same procedure with the regions in the fast-forward queue. If we still do not find similar regions, we assign a new cluster ID for the current region and insert it into the BBV queue corresponding to its simulation mode. In our experiments, we set $\theta = 0.05$ to ensure a reasonable simulation accuracy while maintaining high speedups. To further improve the efficiency of our clustering technique, we incorporate the triangle inequality optimization [31] into our algorithm, which can skip redundant BBV distance calculations. We use Euclidean distance for all BBV distance calculations.

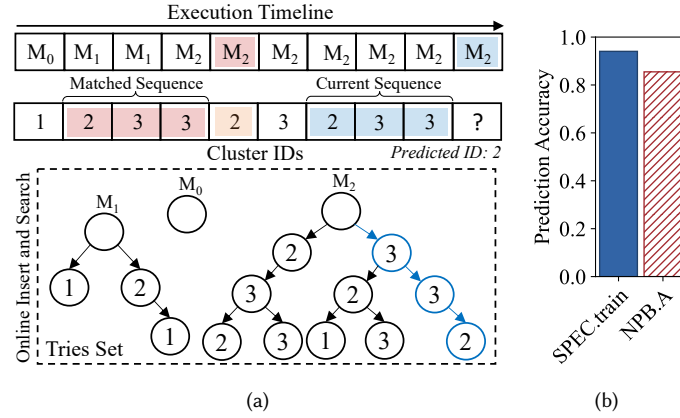


Fig. 5. The predictor utilizes the trie [11] data structure to quickly predict the cluster ID of the next region by searching for a similar history with the same region start marker M_i . In this example, the cluster ID of the next region is predicted to be 2 since the prior region with the cluster ID of 2 has the same start marker M_2 and the longest matching sequence ($3 \rightarrow 3 \rightarrow 2$). Plot (b) shows the accuracy of the predictor for different benchmark suites.

3.4 Prediction Mechanism

Pac-Sim employs a *Predictor* – an online prediction mechanism that leverages region markers, execution history, and hardware state to predict the phase behavior of the next region in an application and decide its simulation mode at runtime.

Region Markers. The Marker Detector identifies PC-based region markers that act as the boundaries of the regions. In certain cases, using region markers to classify regions is effective for applications where the same part of the code displays similar phase behavior, as in the case of 619.lbm_s.1 and 644.nab_s.1 using train inputs.

Execution History. When executing the same part of the source code, differences in memory access patterns, branching, etc., can result in varying phase behavior at runtime. We, therefore, make use of execution history, which is a sequence of the cluster IDs of prior regions, to predict these differences in the phase behavior among applications.

Hardware State. Pac-Sim takes into account the state of the simulated system, such as core frequency, while executing each region. *Predictor* predicts the next region to be detailed mode if there are no prior similar regions with the same hardware state.

The Predictor decides the cluster ID of the next region by choosing the cluster ID of the previous region with the same region marker and has the longest matching sequence. For the regions that do not have a previous region with the same start marker or the same history, Pac-Sim enables detailed execution for that region. Then, it decides the simulation mode of the next region by checking whether prior regions with cluster ID and the current hardware state are simulated in detailed mode. This history is learned online and is updated every time Pac-Sim finishes simulating a region.

To accelerate this stage for large application lengths, we further optimize our clustering algorithm to reduce its average search time complexity from $\mathcal{O}(n^2)$ to $\mathcal{O}(n)$. This is achieved by maintaining the execution history in a *trie* [11] data structure, with a maximum depth of 16, which allows for more efficient *search* and *insert* operations. In Pac-Sim, we




Fig. 6. The graph shows the regions identified using Pac-Sim for the NPB benchmark ft, grouped together with the respective cluster they belong to. The shaded portion represents the regions that are simulated in detail.

utilize the trie data structure to maintain the execution history of the application being simulated and quickly predict the cluster ID of the next region based on this information.

Figure 5a illustrates the usage of tries to predict the cluster ID of the next region by considering the example of a hypothetical execution sequence. *Insert*: The cluster ID of the current region is inserted into the trie. In Figure 5a, when the online clustering of the fifth region is finished, we insert the current cluster ID 2 for both branches corresponding to the three histories: 3, 3 \rightarrow 3, and 2 \rightarrow 3 \rightarrow 3. *Search*: Once Insert of the current region is completed, the cluster ID of the next region is predicted by searching the trie for a matching cluster ID sequence. The search operation ends when the sequence matches one of the leaf node paths. Note that two regions having the same marker do not necessarily mean that the regions belong to the same cluster.

Figure 5b shows the average accuracies of the online predictor for the benchmarks of SPEC CPU2017 and NPB are 94% and 85%, respectively, ensuring the sampling accuracy and performance of Pac-Sim. The accuracy of the predictor is determined by comparing the predicted cluster ID prior to simulating the region with the actual cluster ID obtained through clustering after simulation. Figure 6 shows the results of the Predictor in clustering different regions as identified by Pac-Sim simulating the ft benchmark from the NPB benchmark suite using eight threads. We observe that the majority of regions from each cluster are simulated in detail (shaded portions). This is in accordance with the learning phase of our algorithm where Pac-Sim works to establish a comprehensive understanding of the phase behavior of the application.

3.5 Simulation by Application Reconstruction

Previously proposed multi-threaded sampling methodologies [17, 60] rely fully on offline analysis to determine the regions that need to be simulated in detail. Pac-Sim assumes no prior knowledge about the nature of the workload that it is about to simulate. Instead, it (a) samples regions online during the simulation and (b) uses the detailed simulation results of previous regions to estimate the performance of the current fast-forwarded region by applying the four different methods described below successively until convergence is reached.

- i. Use the detailed performance metrics of a region that belongs to the same cluster and has the same start marker as the current region.
- ii. Use the detailed performance metrics of a region that belongs to the same cluster.
- iii. Use the performance of a region with the closest BBV ($\theta = 1$) and the same number of active threads.
- iv. Use the average performance of the regions that have the same number of active threads.




Fig. 7. The workflow of Pac-Sim when the representative regions are simulated in parallel. Pac-Sim starts in the emulation mode, collecting feature vectors and MTR [7] warmup data online, and then predicts the simulation mode of the next region. For regions predicted for detailed mode, Pac-Sim forks new processes to perform warmup and detailed simulation.

We extrapolate the runtime T_r of a fast-forwarded region r using the region r' by $T_r = T_{r'} \frac{insn_r}{insn_{r'}}$, where $T_{r'}$ is the runtime of the previous region r' identified above, and $insn_r$ and $insn_{r'}$ are the maximum instruction counts among all threads for the regions r and r' , respectively.

Runtime Hardware Events. Pac-Sim takes into account the state of the simulated system while estimating the performance of the fast-forwarded region. As runtime hardware events can happen at any time, we do not guarantee the regions that are divided by those events to be large enough. In such cases, we estimate the performance of these regions using the closest previous region with the same hardware state, as these regions are too small to be clustered. Moreover, the impact of these regions on the overall application performance is typically negligible as the regions are too short.

3.6 Sampled Simulation in Parallel

Pac-Sim is primarily targeted for runtime varying scenarios using live sampling. However, for statically scheduled multi-threaded applications, Pac-Sim can support sampled simulation in parallel, similar to checkpoint-based mechanisms, to further speed up the sampled simulation. The workflow of Pac-Sim for parallel simulation is shown in Figure 7. Previous methods, like LiveSim [40] and LoopPoint [60], require offline analysis and store checkpoints for sampled simulation. A huge amount of storage is required for these methods, as mentioned in Section 2. Pac-Sim starts in emulation mode, collecting feature vectors and warmup data online, and then predicts the simulation mode of the next region. For regions predicted for detailed mode, Pac-Sim forks new processes, which run in parallel, to perform warmup and detailed simulation. Pac-Sim reconstructs the performance of the entire application once the whole application is emulated and the simulation of all regions is completed.

3.7 Microarchitectural Warmup

One of the major challenges of sampled simulation is to choose the right warmup technique that can directly build up an accurate microarchitectural state prior to the detailed simulation of a region. Methodologies like SMARTS [71] and Time-Based Sampling techniques [4, 15] keep functional warming enabled during the entire sampled simulation, leading to large slowdowns. We find that the statistical warmup techniques [7, 39, 52, 67] can reconstruct the accurate microarchitectural state of a simulated system online. We choose the memory time-stamp record (MTR) [7] technique to be used with Pac-Sim. MTR can rapidly collect memory reference patterns during the fast-forward mode and reconstruct the cache state before switching to detailed simulation. In this work, we limit the simulation infrastructure to explicit cache warming, as the smaller structures tend to be warmed relatively quickly.

We employ a cache warmup strategy by populating the cache structures with the N most recent unique memory accesses, where N represents the total number of cache lines being simulated. We maintain a minimum region length of 20 million instructions to mitigate the impact of smaller structures on sampling accuracy. While prior research explores warming up various microarchitectural structures [23, 30], our experiments demonstrate that focusing solely on last-level cache warmup is sufficient for high sampling accuracy. Pac-Sim can be configured with other warmup techniques, but evaluating their effectiveness is beyond the scope of this work.

4 EXPERIMENTAL SETUP

In this section, we describe the experimental setup used to evaluate Pac-Sim. We begin by providing the specifics of the simulation framework, comprising the simulator used and details of the simulated architecture employed in our experiments. We then describe the different workloads that are used to evaluate the performance of our methodology, including SPEC CPU2017 [12], PARSEC [8], and NAS Parallel Benchmarks (NPB) [6] with different multi-threaded programming models, namely OpenMP [53] and OmpSs [28]. The default parameters of Pac-Sim used in our experiments are listed in Table 2.

Table 2. The default parameters of Pac-Sim used in our experiments.

Parameters	Values
Min. Region Length (instructions)	20,000,000
Max. Region Length (instructions)	50,000,000
Dimensions of Online BBV	16
Max. Depth of Trie used in Predictor	16
Clustering Threshold	0.05

4.1 Simulation Tools

In this work, we use a modified version of the Sniper multi-core simulator [14, 16] (version 7.4), which is updated to support loop-based and barrier-based region specifications in order to evaluate Pac-Sim. Sniper is a many-core simulator using high-level abstract models and is widely used for architectural evaluation and design space exploration. Note that our methodology does not utilize any features specific to the Sniper simulator. Therefore, porting the methodology to other simulators, such as gem5 [9] or ZSim [62], should be relatively straightforward. To demonstrate that Pac-Sim is indeed a microarchitecture-independent methodology, we experimentally evaluate it by running simulations upon two different processor configurations that mimic the performance/behavior of Intel’s Gainestown, Skylake [27], and Sunnycove [54] microarchitectures using Sniper. The configuration details for each of these models are listed in Table 3.

To enable high-performance simulation, Pac-Sim intelligently switches among the three simulation modes supported by Sniper, namely, fast-forward mode, cache-only mode, and detailed simulation mode. The fast-forward mode is used to reach a particular point in an application during simulation without enabling the performance models. The cache-only mode performs the functional warming of the caches, whereas the detailed simulation mode is the default simulation mode that enables the timing model for performance estimation. For Pac-Sim, we capitalize on the split execution and

Note that Gainestown is the latest microarchitecture available on Sniper simulator that has been validated against hardware. We made modifications to the back-end of Sniper to support the contention model and instruction latencies for Skylake and Sunnycove architectures.

Table 3. The configuration parameters we used for Gainestown, Skylake, and Sunnycove microarchitectures on Sniper.

Component	Gainestown Parameters	Skylake Parameters	Sunnycove Parameters
Processor	1, 8 cores	1, 8 cores	1, 8 cores
Core	2.66 GHz, 128-entry ROB	2.66 / 3.7 GHz, 224-entry ROB	3.60 GHz, 352-entry ROB
L1-I / L1-D	32 KB, 4 / 8 way, LRU	32 KB, 8 / 8 way, LRU	32 / 48 KB, 8 / 12 way, LRU
L2 cache	256 KB, 8 way, LRU	1 MB, 16 way, LRU	1.25 MB, 20 way, LRU
L3 cache	8 MB (shared), 16 way, LRU	22 MB (shared), 12 way, LRU	16 MB (shared), 16 way, LRU

timing model of Sniper to fast-forward in the front-end of the simulator so that the simulation wall time is further minimized.

Every time Pac-Sim switches from fast-forward to detailed simulation mode, the cache state is reconstructed at the beginning of the region using the memory time-stamp record (MTR) [7] technique. We implement MTR in Sniper to collect the cache line information accessed by each *Load* and *Store* instruction during simulation, ordered in LRU fashion per set and then inject the requests into the cache in the correct order to rebuild the appropriate cache state.

4.2 Benchmarks Used

To demonstrate the applicability of Pac-Sim, we experimentally evaluate the methodology using multiple benchmark suites such as (i) the SPEC CPU2017 benchmark suite [12], (ii) the NAS Parallel Benchmarks (NPB) [6] version 3.4.2, and (iii) the PARSEC [8] version 3.0 benchmark suite. Note that these are multi-threaded benchmarks that synchronize frequently and share memory.

We configure these benchmarks to use two different multi-threaded programming models, namely OpenMP [53] and OmpSs [28]. OpenMP [53] provides a set of compiler directives, library routines, and environment variables that help developers to parallelize their code. On the other hand, OmpSs [28] extends OpenMP, and it is able to dynamically manage and schedule tasks to maximize multi-threaded application performance. We set up the multi-threaded benchmarks to use *passive* thread wait policy, meaning that the threads will sleep while waiting for other threads at a synchronization point.

SPEC CPU2017 is a collection of benchmarks used for performance evaluation in computer architecture research. In our experiments, we use the *speed* version of multi-threaded SPEC CPU2017 benchmarks that are parallelized with OpenMP. The benchmarks are compiled using GCC 6.4.0 and GFortran with the `-O3` compiler flag for x86-64 architecture. We configure these benchmarks to run with eight threads and evaluate them using the *training* (train) input set. Note that the full detailed simulation of these benchmarks using *reference* (ref) inputs for sample verification is not practical. NAS Parallel Benchmarks (NPB) [6] is another set of benchmarks widely used to evaluate the performance of highly parallel systems in computer architecture. The reference implementations of these benchmarks are available in the two most commonly used programming models, i.e., MPI and OpenMP. In our experiments, we use the OpenMP-based implementation with input *class A* and generate the binaries using `icc` compiler (with `-O2` flag) as part of the Intel oneAPI (version 2022.0.2) toolkit. We also present experimental evaluations of Pac-Sim using PARSEC, which is another standard benchmark suite consisting of computationally intensive applications designed to facilitate the study of multi-core systems with shared memory. PARSEC implementations are available in both OpenMP and OmpSs [20] versions. In our experiments, we use both these versions with the *simlarge* input set.

5 EVALUATION

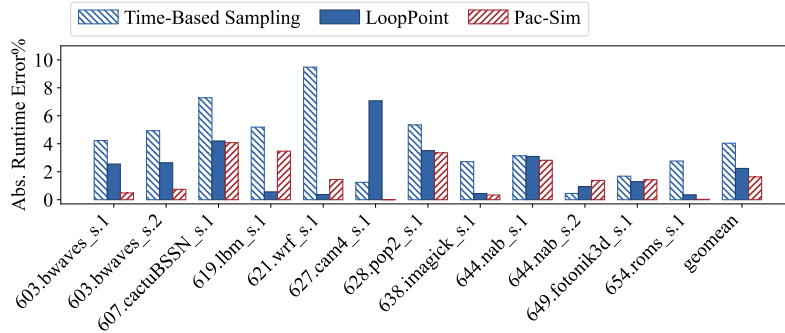


Fig. 8. A comparison of the absolute runtime prediction error using different methodologies, namely, Time-Based Sampling, LoopPoint, and Pac-Sim for 8-threaded SPEC CPU2017 benchmarks using train inputs. On average, Pac-Sim achieves better accuracy compared to Time-Based Sampling and LoopPoint.

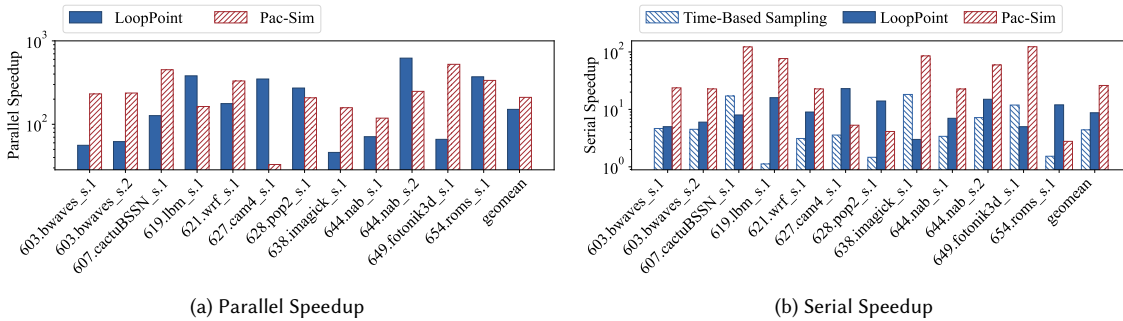


Fig. 9. The parallel and serial speedups achieved using Pac-Sim for 8-threaded SPEC CPU2017 benchmarks using train inputs. For speedup calculations, the simulation walltime corresponding to Pac-Sim includes both online analysis and simulation time, whereas, for LoopPoint, we consider only the checkpoint simulation time, excluding the time required for offline profiling and checkpoint generation. Pac-Sim outperforms both Time-Based Sampling and LoopPoint in terms of speedup achieved. Note that Time-Based Sampling techniques are not suitable for sampled simulation in parallel.

In this section, we first present a comprehensive evaluation of Pac-Sim, comparing its efficacy with the current state-of-the-art. Additionally, we provide experimental evidence showing that Pac-Sim is indeed a hardware-independent methodology. Finally, we present case studies that demonstrate the applicability and effectiveness of Pac-Sim in estimating workload performance in dynamic, multi-threaded hardware and software environments. Note that, throughout this paper, the term runtime refers to the simulated runtime of the application, whereas the term wall-time refers to the actual time taken by the simulator to finish the run.

Evaluation metrics. In order to evaluate the effectiveness of any simulation methodology, it is crucial to quantitatively measure its performance in terms of two critical metrics: *accuracy* and *speedup*. In our experiments, we define these metrics in the following manner:




Fig. 10. The absolute differences in predicting L2 cache misses per kilo instructions (MPKI) using Pac-Sim as compared to the full detailed simulation. In this experiment, we use the NPB benchmarks with class A inputs running eight threads. The geometric mean of the absolute differences in predicting L2 MPKI is 0.23.

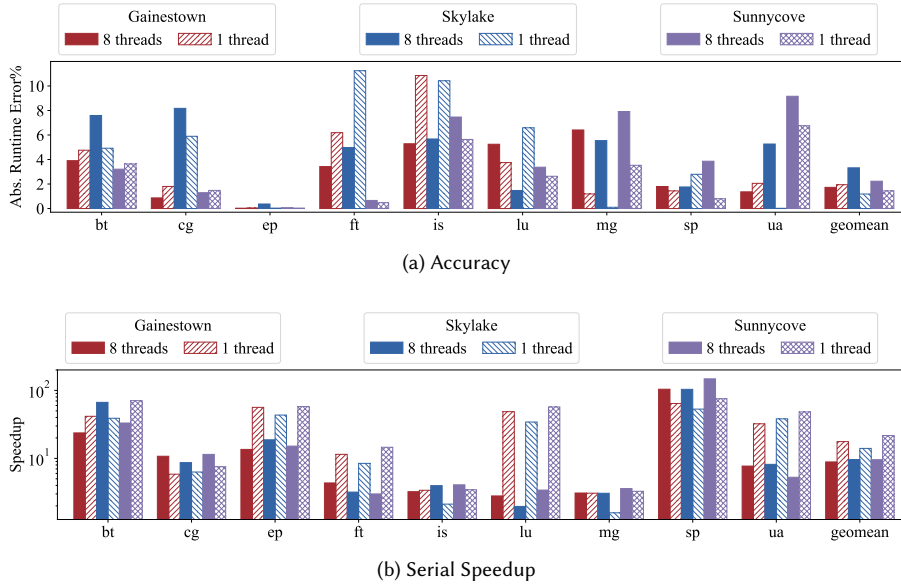


Fig. 11. The accuracy and serial speedup achieved for Pac-Sim methodology when simulated using three different microarchitectures, namely, Gainestown, Skylake, and Sunnycove for NPB benchmarks with class A inputs running eight threads and one thread.

Accuracy: We assess the accuracy of our proposed methodology by comparing the simulation runtime obtained from the full simulation and the sampled simulation in terms of absolute runtime prediction error Δ_{time} , which is defined as

$$\Delta_{time} = \frac{|T_{full} - T_{sample}|}{T_{full}},$$

where T_{full} represents the simulation runtime obtained from the full run, and T_{sample} represents the simulation runtime extrapolated from the sampled simulation. It is important to note that in our evaluation, we use the runtime (execution time as inferred from simulation) of the application as the performance metric to measure the accuracy of sampling. This is because time-per-program is the gold-standard performance measure, and IPC is not a valid performance metric for multi-threaded applications [2].

Speedup: In our experiments, we calculate the speedup by taking the ratio of the wall-clock time for the full simulation to that of the sampled simulation and the average speedup by computing the geometric mean of the speedups across

all benchmarks. Serial speedup is defined as the speedup achieved when all representative regions are simulated sequentially, while parallel speedup is obtained when the representative regions are simulated in parallel, assuming infinite resources.

5.1 Comparison with the State-of-the-Art

In this section, we evaluate the performance of Pac-Sim in comparison to the state-of-the-art profile-driven sampled simulation methodology, LoopPoint [60]. While several other profile-driven methodologies exist, LoopPoint provides the benefit of being applicable across a variety of applications and synchronization types. It has also been shown to outperform other multithreaded sampled simulation methodologies (such as BarrierPoint) in terms of speedup and accuracy, thus serving as a strong baseline for our evaluations. We also compare Pac-Sim with Time-Based Sampling [4, 15] techniques, which repeatedly alternate between detailed simulation and fast-forwarding of regions. For a fair comparison, we adopt Pac-Sim’s approach of injecting the warmup state at the beginning of each detailed simulation region. We now report the results of our simulation experiments evaluating and comparing the performance of these methodologies using the SPEC CPU2017 benchmarks.

Accuracy. Figure 8 shows a comparison of absolute runtime prediction errors for Pac-Sim, Time-Based Sampling, and LoopPoint obtained for the 8-threaded SPEC CPU2017 benchmarks using train inputs. Our analysis reveals that, in most cases, Pac-Sim performs comparably with LoopPoint in predicting the runtime of the applications, with the individual errors differing by no more than 2 to 3%. The relatively higher errors for some applications, such as `619.lbm_s.1` is because Pac-Sim relies on online extrapolation to estimate application performance using the limited profile data that is available from regions that have already been simulated. Meanwhile, methodologies like LoopPoint rely on offline profiling and, therefore, utilize information about the whole application. We also evaluate the accuracy of Pac-Sim to determine the L2 cache misses per kilo instructions (MPKI), as shown in Figure 10. The final results show that the average absolute difference of MPKI for all benchmarks evaluated is 0.23, demonstrating Pac-Sim accurately extrapolates microarchitectural metrics of the whole application from the selected samples.

Speedup. Figure 9 shows the speedup comparison of Pac-Sim with previously proposed techniques, Time-Based Sampling and LoopPoint, for the SPEC CPU2017 benchmarks using train inputs running eight threads. Figure 9a shows the parallel speedup for which Pac-Sim outperforms LoopPoint in most cases (7 out of 12 benchmarks). The primary reason for this is that Pac-Sim uses smaller regions as compared to LoopPoint. Although Pac-Sim requires emulation of the entire application, the online analysis overhead is minimized, and therefore, the average parallel speedup for SPEC CPU2017 benchmarks (train inputs) using Pac-Sim is 210.3 \times , which is larger than that obtained for LoopPoint (150.97 \times). Figure 9b shows the serial speedup, and we observe Pac-Sim outperforms LoopPoint in most cases, attaining a maximum serial speedup of 123.32 \times . Specifically, Pac-Sim outperforms LoopPoint by 1.8 \times and 1.4 \times for the serial and parallel speedup, respectively. While the online analysis can introduce runtime overheads, the performance advantages of Pac-Sim seem to outweigh these overheads in most cases. We observe that the performance of Pac-Sim surpasses that of Time-Based Sampling across all evaluated benchmarks. This is because Pac-Sim selects regions for detailed simulation through online learning, in contrast to Time-Based Sampling. Furthermore, the use of program constructs like barriers and loops in Pac-Sim for region selection enables identifying repetitive program behavior in multi-threaded workloads. However, there are some cases where LoopPoint performs better than Pac-Sim, such as for `627.cam4_s.1` and `628.pop2_s.1` benchmarks in Figure 9b. This is mainly because Pac-Sim uses a small clustering threshold (0.05) for the online clustering algorithm in order to maintain higher accuracy, compared with Time-Based sampling and LoopPoint.

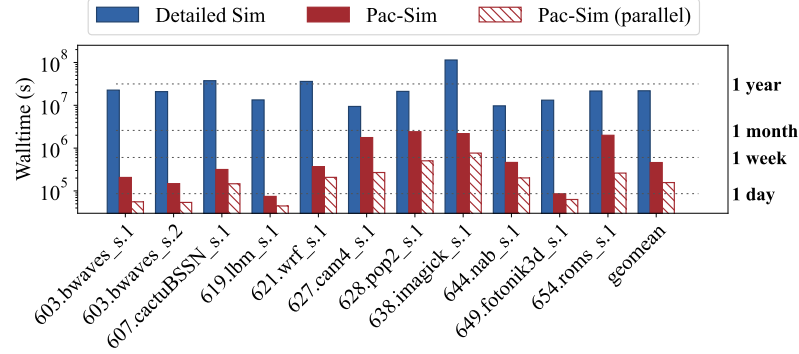


Fig. 12. A comparison of the estimated walltime for fully detailed simulation and sampled simulation using the serial and parallel versions of Pac-Sim for 8-threaded SPEC CPU2017 benchmarks using ref inputs. The estimated walltime includes the time required for online analysis, warmup, and simulation.

Efficacy in Evaluating Realistic Workloads. The full detailed simulation of SPEC CPU2017 benchmarks with reference inputs takes an extremely long time – about a year on average using multi-core simulators like Sniper. Instead, we estimate their simulation walltime by considering the instruction count of the benchmark using reference inputs along with the average simulation rate of the benchmark using train inputs. The walltime of Pac-Sim includes the time required for online analysis and emulation of the entire workload along with the time for detailed simulation of the representative regions. Figure 12 shows that Pac-Sim takes less than a week, on average, to run the entire application sequentially, while the parallel version of Pac-Sim takes about 1.8 days on average. In experiments where the microarchitecture structures like cache size are adjusted or when the application itself undergoes instruction-level modifications, it is necessary to regenerate the checkpoints. In such cases, Pac-Sim is more appropriate as LoopPoint takes 6.2 days on average (shown in Figure 2) to complete its preprocessing before simulation.

Microarchitecture-agnostic sampling. In addition to achieving high accuracy and speedups, Pac-Sim also provides the advantage of being a microarchitecture-independent methodology. We experimentally demonstrate this by evaluating our methodology with three different processor configurations, namely the Sunnycove, Skylake, and Gainestown microarchitectures, for the NPB benchmarks that run using one thread and eight threads. The accuracy and speedup numbers obtained in our experiments are plotted in Figure 11a and Figure 11b, respectively. From Figure 11a, we can observe that the absolute runtime errors estimated by Pac-Sim for all NPB applications are quite low (all under 8%) and are similar for both these processor configurations (differing by 5% at most). Moreover, the speedups obtained for both configurations are similar for most benchmarks, as observed in Figure 11b. Hence, the choice of a target microarchitecture for evaluation does not affect the efficacy of Pac-Sim.

Wall-time Distribution. We show the time spent by Pac-Sim in different stages of sampled simulation. Figure 13 shows the average time spent in the online analysis stage for NPB (class A inputs) and SPEC CPU2017 (train inputs) benchmarks is 7.88% and 11.20%, respectively. This is the result of the optimizations described in Section 3, which are applied to the analysis part. Moreover, Pac-Sim spends 8.25% and 16.00% of the execution time on warmup for NPB and SPEC CPU2017 benchmarks, respectively. This is because Pac-Sim needs to reconstruct the memory access patterns at the beginning of the detailed simulation of a region. Note that Pac-Sim reduces the time spent in both profiling and analysis of the benchmarks significantly as compared to prior profile-driven methodologies for sampled simulation.




Fig. 13. The graph shows the percentage of time that Pac-Sim spends at each phase during the sampled simulation of each benchmark suite (average across all benchmarks). The Analysis component includes online marker detection, region profiling, clustering, and prediction.

Memory Overhead. The memory overhead of Pac-Sim is minimal (no disk access) and allocated once per simulation. For the online analysis, Pac-Sim needs to record the BBV information of the previously simulated regions. The space complexity of the online analysis is $\mathcal{O}\left(\frac{insn_{full}}{region_length}\right)$, where $insn_{full}$ is the total dynamic instructions executed by the application and $region_length$ is the average length (size) of the regions identified. The MTR warmup injection technique keeps counters for each cache line accessed (num_cache_lines), and the space complexity is $\mathcal{O}(num_cache_lines)$.

5.2 Case Studies

We showcase the versatility of Pac-Sim through several compelling case studies. Firstly, we demonstrate that our methodology remains agnostic to dynamic thread scheduling decisions made during runtime, highlighting its robustness and adaptability. Next, we provide examples of how Pac-Sim operates seamlessly in the presence of various runtime hardware events, further cementing its reliability. Finally, we exhibit the applicability of the proposed methodology in hardware-software co-design studies, showcasing its potential to facilitate more efficient and effective design processes.

5.2.1 Dynamically Scheduled Software. The advent of multi-core and many-core architectures has necessitated the efficient parallel execution of dynamically scheduled multi-threaded applications to maximize system performance. However, the non-determinism resulting from the execution of such applications on multi-core platforms often leads to notable performance variability across multiple runs. This variability can be attributed to software-level factors such as dynamic job scheduling by the operating system, thread migration between cores, load balancing optimizations, and contention for shared resources at runtime.

Table 4. Table shows the IPC of freqmine benchmark from the PARSEC benchmark suite using the *simlarge* input for threads 0 through 7. Pac-Sim shows the details of dynamically scheduled software whose IPC and thread mapping differ across two runs.

Thread ID	0	1	2	3	4	5	6	7	Aggregate
IPC _{run1}	0.15	0.09	1.75	0.43	0.07	0.07	0.10	0.09	2.75
IPC _{run2}	0.15	0.09	1.76	0.07	0.44	0.07	0.09	0.10	2.76

Table 4 illustrates the thread-level differences in terms of instructions per cycle (IPC) for two different runs of the OpenMP-parallelized freqmine application from the PARSEC benchmark suite. There are variations in the per-thread IPCs between the two runs, particularly for thread IDs 3 and 4. To investigate the impact of these variations on conventionally used sampling techniques, we conducted two independent profiling runs of freqmine using LoopPoint. The experimental results showed significant variability, with 14% of the regions clustered differently between runs. This




Fig. 14. Figure shows the average error rates (from five different runs) and error bars in predicting the runtime of dynamically scheduled benchmarks. We use PARSEC benchmarks with the *simlarge* input using OmpSs and OpenMP, and NPB benchmarks with *class A* inputs using OpenMP runtime.

presents a challenge for sampled simulation, which relies on profiling data from a prior execution to guide simulation in subsequent runs. Dynamically scheduled applications, by nature, have profiling data that fluctuates across executions. To address this challenge caused by runtime variability, Pac-Sim performs online profiling and simulation within the same run. This allows Pac-Sim to capture any performance variations that might occur during execution.

To demonstrate the effectiveness of Pac-Sim in this regard, we now present an experimental study of dynamically scheduled multi-threaded versions of PARSEC with *simlarge* inputs and NPB with class A inputs. While the per-thread behavior varies for dynamically scheduled applications, the global execution time and global IPC remain consistent across multiple runs. In Table 4, we observe that while there are some variations in per-thread behavior, the aggregate IPCs across the two runs remain nearly unchanged. Figure 14 demonstrates the average runtime prediction errors of Pac-Sim simulating dynamically scheduled multi-threaded applications. We run the benchmarks multiple times in full detailed mode and using Pac-Sim. The errors are calculated by comparing the runtime obtained using Pac-Sim with the average runtime obtained from the full detailed simulations. The results show that Pac-Sim achieves a very low error in predicting the runtime of dynamically scheduled software (3.81% on average). The benchmark, *freqmine*, which shows the largest IPC variation (without spinloops) maintains an average error of 11.43%. Moreover, Pac-Sim demonstrates speedups of up to $43.96\times$ ($6.29\times$ on average) for all dynamically scheduled benchmarks.

5.2.2 Dynamic Hardware Events. Dynamic event-based hardware optimizations help improve performance gains and energy efficiency in modern architectures. DVFS [32, 41, 44] is one of the most widely employed dynamic hardware event-based optimization techniques. It monitors core frequencies and load variations in order to match the system power consumption with the required level of performance by triggering voltage and frequency optimizations at runtime. These optimizations may lead to a diverse range of dynamic hardware states (i.e., core frequency, power configurations) over a given run, consequently resulting in a significant degree of performance variability for a given workload across different executions.

Pac-Sim deals with this performance variability by monitoring the simulated hardware events during execution. While prior sampled simulation methodologies can be modified to support dynamic hardware events triggered only at region boundaries (to maintain consistent hardware state within a region), Pac-Sim allows hardware events at any point during the application execution. Each time an event occurs, Pac-Sim triggers a new region to ensure hardware state consistency within that region. The predictor then speculates the cluster ID of the next region and checks the

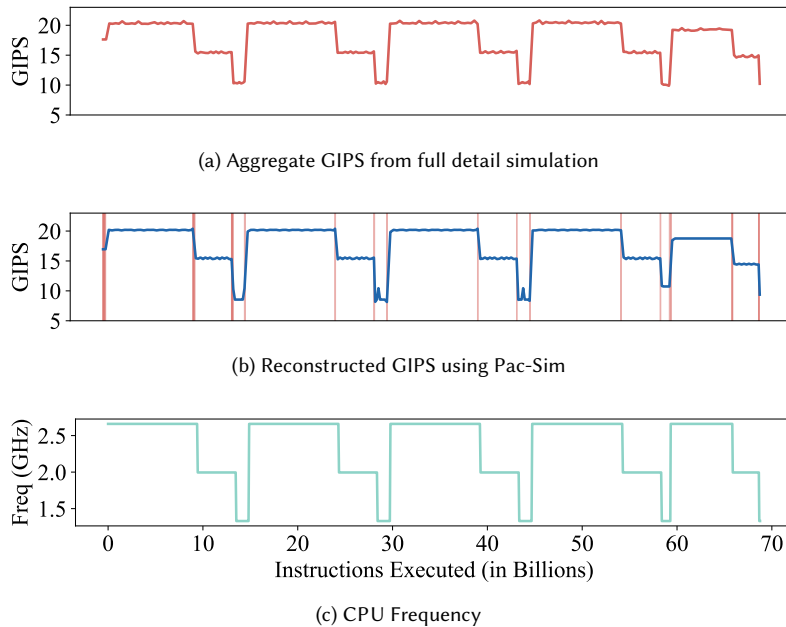


Fig. 15. The aggregate giga (billion) instructions per second (GIPS) of the full run (a), reconstructed GIPS using Pac-Sim (b), and the varying CPU frequency for all CPUs (c) 644.nab_s.1 benchmark with train inputs running eight threads. The shaded regions in (b) represent the regions simulated in detail. The figures share the same x-axis.

execution history to determine whether similar regions (i.e., regions with the same cluster ID and hardware state) were previously encountered. If a match is found, the region is fast-forwarded; otherwise, a detailed simulation is triggered.

We now present an experimental study demonstrating the effectiveness of Pac-Sim in handling the variability caused by dynamic hardware events by specifically considering the case of DVFS-optimized workloads. In our experiments, we evaluate the performance of the benchmarks by comparing the results of Pac-Sim with the baseline while changing the frequency at predetermined intervals; however, just like in actual DVFS-optimized executions, the information on the frequency changes is not available to the simulator a priori. In order to evaluate the performance of Pac-Sim, we consider a DVFS scenario in which the processor frequency f switches among a fixed range of values, i.e., $f \in \{1.33 \text{ GHz}, 2.00 \text{ GHz}, 2.66 \text{ GHz}\}$ as shown in Figure 15c.

For this scenario, we measure the aggregate giga/billion instructions per second (GIPS) values obtained from both the full detailed simulation and Pac-Sim over the entire execution. The findings of our experiment are presented in Figure 15. We observe that the GIPS values obtained from both the full simulation (Figure 15a) and Pac-Sim (Figure 15b) exhibit a great deal of similarity, indicating Pac-Sim’s effectiveness in estimating the performance of a dynamically optimized workload with a high level of accuracy. Furthermore, our findings reveal that Pac-Sim simulates only a small fraction of the entire application in detail (depicted by shaded regions in Figure 15b). Notably, most of the detailed simulation occurs either at points of change in the phase behavior of the application or hardware states. This demonstrates that Pac-Sim can use this information to identify a minimal representative subset for applications using online analysis.




Fig. 16. The figure shows the absolute difference in performance (in terms of runtime) for NPB benchmarks using class A inputs and eight threads with (w/) and without (w/o) SSE2 simulated in detailed mode and with Pac-Sim.

5.2.3 Hardware-Software Co-design. Hardware-software co-design is an emerging field of study that optimizes the system performance by concurrently designing the compiler and hardware components of a system to exploit the synergy between the two. Prior works [34, 38, 74] have investigated several directions in this context. To identify the most effective strategies, hardware-software co-design research relies on fast and accurate architectural simulation methodologies to explore the design space efficiently. However, among existing methodologies, the profile-driven methodologies [17, 65] incur significant profiling and preprocessing costs, as shown in Figure 2, whereas the statistical sampling methodologies [42, 71] (which do not rely on preprocessing) have low speedups.

Pac-Sim addresses these issues by sampling and analyzing the regions online. Thus, it incurs no additional profiling cost if new compilers are used or new applications are generated, enabling fast and efficient exploration of hardware-software co-design space. To demonstrate the effectiveness of Pac-Sim in this regard, we now present a performance evaluation study of the NPB benchmarks under different compiler optimization techniques. We study the impact of SIMD (Single Instruction, Multiple Data) optimizations on the generated binaries using both Pac-Sim and full detailed simulations. SIMD-enabled processors are equipped with special-purpose registers that can simultaneously load multiple machine words and perform operations on them in parallel in order to improve processor performance. For instance, the Streaming SIMD Extensions 2 (SSE2) instruction set uses 128-bit XMM registers to process packed data elements at once.

In our experiments, we measure the performance improvement (in terms of runtime) obtained by enabling SSE2 and compare it against the baseline. The results of our simulations are presented in Figure 16. We observe that the average difference in the performance improvements obtained from full detailed mode and Pac-Sim is 3.65%. Specifically, Pac-Sim reveals the performance effects of SIMD instructions. For example, some benchmarks achieve a significant speedup over the baseline as these applications meet the `icc` vectorization criteria [25]. `ft` calculates a 3D fast Fourier transform, and its innermost loop consists of multiply-add statements with contiguous memory accesses and no data dependency. On the other hand, `is`, which uses the quick sort algorithm, is hard to be vectorized. The SIMD overheads resulting from register transfer costs exacerbate the overall application performance.

6 RELATED WORK

In Section 2, we comprehensively discussed the most relevant prior works. Sampled simulation has been an active research area for several decades, and several techniques were proposed [4, 5, 15, 17, 37, 40, 59, 60, 63, 65, 70, 71] in this direction for different workload classes primarily for the reduction of simulation time and resources. In order to simulate the identified sample correctly, it is important to reconstruct the microarchitectural state of the system

using techniques like statistical warming [7, 29, 39, 51], checkpoint-based warming [52, 67], or by enabling functional simulation. Analytical modeling is yet another solution to evaluate a complex workload quickly. Prior works proposed analytical models to derive the performance of processors [24, 33], cache miss rates [30], branch miss rates [23], etc. However, analytical performance modeling can be limited in supporting new hardware designs, requiring new models for each design.

7 CONCLUSION

In this work, we propose a novel methodology, Pac-Sim, that allows for the sampled simulation of dynamic software that responds to workload and run-time execution conditions. Pac-Sim is the first, to the best of our knowledge, to propose a sampling solution that simulates these dynamic conditions in both a fast (up to 523.5× speedup, 210.3× on average) and accurate way (average errors of 1.63% and 3.81% for statically and dynamically scheduled benchmarks, respectively).

ACKNOWLEDGMENTS

We would like to thank each of the people who took the time to review and provide detailed feedback for this work – it has been greatly improved because of each of these suggestions. In particular, we would like to acknowledge the help of Lieven Eeckhout, Harish Patil, Wim Heirman, and Andreas Diavastos. We also thank the anonymous reviewers for their valuable feedback. This work was funded, in part, by Intel, VMware, and Huawei.

REFERENCES

- [1] A.R. Alameldeen and D.A. Wood. 2003. Variability in architectural simulations of multi-threaded workloads. In *International Symposium on High-Performance Computer Architecture (HPCA)*.
- [2] A.R. Alameldeen and D.A. Wood. 2006. IPC Considered Harmful for Multiprocessor Workloads. *IEEE Micro* (2006).
- [3] D.H. Albonesi. 1999. Selective cache ways: on-demand cache resource allocation. In *International Symposium on Microarchitecture (MICRO)*.
- [4] E. K. Ardestani and J. Renau. 2013. ESESC: A fast multicore simulator using Time-Based Sampling. In *International Symposium on High Performance Computer Architecture (HPCA)*.
- [5] Eduardo Argollo, Ayose Falcón, Paolo Faraboschi, Matteo Monchiero, and Daniel Ortega. 2009. COTSon: infrastructure for full system simulation. *SIGOPS Operating Systems Review* (2009).
- [6] David Bailey, Tim Harris, William Saphir, Rob Van Der Wijngaart, Alex Woo, and Maurice Yarrow. 1995. *The NAS Parallel Benchmarks 2.0*. Technical Report.
- [7] Kenneth C Barr, Heidi Pan, Michael Zhang, and Krste Asanovic. 2005. Accelerating multiprocessor simulation with a memory timestamp record. In *International Symposium on Performance Analysis of Systems and Software (ISPASS)*.
- [8] Christian Bienia. 2011. *Benchmarking Modern Multiprocessors*. Ph.D. Dissertation.
- [9] Nathan Binkert, Bradford Beckmann, Gabriel Black, Steven K. Reinhardt, Ali Saidi, Arkaprava Basu, Joel Hestness, Derek R. Hower, Tushar Krishna, Somayeh Sardashti, Rathijit Sen, Korey Sewell, Muhammad Shoab, Nilay Vaish, Mark D. Hill, and David A. Wood. 2011. The Gem5 Simulator. *SIGARCH Computer Architecture News* (2011).
- [10] William Lloyd Bircher and Lizy John. 2010. Predictive power management for multi-core processors. In *International Symposium on Computer Architecture (ISCA)*.
- [11] Ferenc Bodon and Lajos Rónyai. 2003. Trie: an alternative data structure for data mining algorithms. *Mathematical and Computer Modelling* (2003).
- [12] James Bucek, Klaus-Dieter Lange, and JÓakim v. Kistowski. 2018. SPEC CPU2017: Next-Generation Compute Benchmark. In *International Conference on Performance Engineering (ICPE)*.
- [13] Anastasia Butko, Rafael Garibotti, Luciano Ost, Vianney Lapotre, Abdoulaye Gamatie, Gilles Sassatelli, and Chris Adeniyi-Jones. 2015. A trace-driven approach for fast and accurate simulation of manycore architectures. In *Asia and South Pacific Design Automation Conference (ASP-DAC)*.
- [14] T. E. Carlson, W. Heirman, and L. Eeckhout. 2011. Sniper: Exploring the level of abstraction for scalable and accurate parallel multi-core simulation. In *International Conference for High Performance Computing, Networking, Storage and Analysis (SC)*.
- [15] T. E. Carlson, W. Heirman, and L. Eeckhout. 2013. Sampled simulation of multi-threaded applications. In *International Symposium on Performance Analysis of Systems and Software (ISPASS)*.

- [16] Trevor E Carlson, Wim Heirman, Stijn Eyerman, Ibrahim Hur, and Lieven Eeckhout. 2014. An evaluation of high-level mechanistic core models. *ACM Transactions on Architecture and Code Optimization (TACO)* 11, 3 (2014), 1–25.
- [17] T. E. Carlson, W. Heirman, K. Van Craeynest, and L. Eeckhout. 2014. BarrierPoint: Sampled simulation of multi-threaded applications. In *International Symposium on Performance Analysis of Systems and Software (ISPASS)*.
- [18] Marc Casas, Harald Servat, Rosa M. Badia, and Jesús Labarta. 2012. Extracting the optimal sampling frequency of applications spectral analysis. *Concurrency and Computation: Practice and Experience* (2012).
- [19] James Charles, Preet Jassi, Narayan S Ananth, Abbas Sadat, and Alexandra Fedorova. 2009. Evaluation of the Intel® core™ i7 turbo boost feature. In *International Symposium on Workload Characterization (IISWC)*.
- [20] Dimitrios Chasapis, Marc Casas, Miquel Moretó, Raul Vidal, Eduard Ayguadé, Jesús Labarta, and Mateo Valero. 2015. PARSECS: Evaluating the impact of task parallelism in the PARSEC benchmark suite. *Transactions on Architecture and Code Optimization (TACO)* (2015).
- [21] Wen-Ke Chen, Sorin Lerner, Ronnie Chaiken, and David M Gillies. 2000. Mojo: A dynamic optimization system. In *Workshop on Feedback-Directed and Dynamic Optimization*.
- [22] Sanjoy Dasgupta. 2000. Experiments with random projection. In *Conference on Uncertainty in artificial intelligence (UAI)*.
- [23] Sander De Pestel, Stijn Eyerman, and Lieven Eeckhout. 2015. Micro-architecture independent branch behavior characterization. In *International Symposium on Performance Analysis of Systems and Software (ISPASS)*.
- [24] S. De Pestel, S. Van den Steen, S. Akram, and L. Eeckhout. 2019. RPPM: Rapid Performance Prediction of Multithreaded Workloads on Multicore Processors. In *International Symposium on Performance Analysis of Systems and Software (ISPASS)*.
- [25] M Deilmann et al. 2012. A guide to vectorization with Intel C++ compilers. *Intel Corporation* (2012).
- [26] Andreas Diavastos and Pedro Trancoso. 2017. SWITCHES: A lightweight runtime for dataflow execution of tasks on many-cores. *Transactions on Architecture and Code Optimization (TACO)* (2017).
- [27] Jack Doweck, Wen-Fu Kao, Allen Kuan-yu Lu, Julius Mandelblat, Anirudha Rahatekar, Lihu Rappoport, Efraim Rotem, Ahmad Yasin, and Adi Yoaz. 2017. Inside 6th-generation intel core: New microarchitecture code-named skylake. *IEEE Micro* (2017).
- [28] Alejandro Duran, Eduard Ayguadé, Rosa M Badia, Jesús Labarta, Luis Martinell, Xavier Martorell, and Judit Planas. 2011. Ompss: a proposal for programming heterogeneous multi-core architectures. *Parallel processing letters* (2011).
- [29] Lieven Eeckhout, Yue Luo, Koen De Bosschere, and Lizy K John. 2005. BLRL: Accurate and efficient warmup for sampled processor simulation. *Comput. J.* (2005).
- [30] David Eklov and Erik Hagersten. 2010. StatStack: Efficient modeling of LRU caches. In *International Symposium on Performance Analysis of Systems and Software (ISPASS)*.
- [31] Charles Elkan. 2003. Using the triangle inequality to accelerate k-means. In *International Conference on Machine Learning (ICML)*.
- [32] Stijn Eyerman and Lieven Eeckhout. 2011. Fine-grained DVFS using on-chip regulators. *Transactions on Architecture and Code Optimization (TACO)* (2011).
- [33] Stijn Eyerman, Lieven Eeckhout, Tejas Karkhanis, and James E Smith. 2009. A mechanistic performance model for superscalar out-of-order processors. *Transactions on Computer Systems (TOCS)* (2009).
- [34] Stijn Eyerman, Wim Heirman, Sam Van den Steen, and Ibrahim Hur. 2021. Enabling Branch-Mispredict Level Parallelism by Selectively Flushing Instructions. In *International Symposium on Microarchitecture (MICRO)*.
- [35] Ayose Falcón, Paolo Faraboschi, and Daniel Ortega. 2007. Combining simulation and virtualization through dynamic sampling. In *International Symposium on Performance Analysis of Systems and Software (ISPASS)*.
- [36] Edward W Forgy. 1965. Cluster analysis of multivariate data: efficiency versus interpretability of classifications. *Biometrics* (1965).
- [37] T. Grass, A. Rico, M. Casas, M. Moreto, and E. Ayguadé. 2016. TaskPoint: Sampled simulation of task-based programs. In *International Symposium on Performance Analysis of Systems and Software (ISPASS)*.
- [38] Ali Hajiabadi, Andreas Diavastos, and Trevor E. Carlson. 2021. NOREBA: A Compiler-Informed Non-Speculative out-of-Order Commit Processor. In *International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*.
- [39] John W Haskins and Kevin Skadron. 2003. Memory reference reuse latency: Accelerated warmup for sampled microarchitecture simulation. In *International Symposium on Performance Analysis of Systems and Software (ISPASS)*.
- [40] Sina Hassani, Gabriel Southern, and Jose Renau. 2016. LiveSim: Going live with microarchitecture simulation. In *International Symposium on High Performance Computer Architecture (HPCA)*.
- [41] Canturk Isci, Alper Buyuktosunoglu, Chen-Yong Cher, Pradip Bose, and Margaret Martonosi. 2006. An analysis of efficient multi-core global power management policies: Maximizing performance for a given power budget. In *International Symposium on Microarchitecture (MICRO)*.
- [42] Chuntao Jiang, Zhibin Yu, Hai Jin, Chengzhong Xu, Lieven Eeckhout, Wim Heirman, Trevor E. Carlson, and Xiaofei Liao. 2013. PCantorSim: Accelerating Parallel Architecture Simulation through Fractal-Based Sampling. *Transactions on Architecture and Code Optimization (TACO)* (2013).
- [43] Joshua L Kihm, Samuel D Strom, and Daniel A Connors. 2007. Phase-guided small-sample simulation. In *International Symposium on Performance Analysis of Systems and Software (ISPASS)*.
- [44] Wonyoung Kim, Meeta S Gupta, Gu-Yeon Wei, and David Brooks. 2008. System level analysis of fast, per-core DVFS using on-chip switching regulators. In *International Symposium on High Performance Computer Architecture (ISCA)*.
- [45] Neeraj Kulkarni, Feng Qi, and Christina Delimitrou. 2019. Pliant: Leveraging approximation to improve datacenter resource efficiency. In *International Symposium on High Performance Computer Architecture (HPCA)*.

- [46] Jeremy Lau, Erez Perelman, and Brad Calder. 2006. Selecting Software Phase Markers with Code Structure Analysis. In *International Symposium on Code Generation and Optimization (CGO)*.
- [47] J. Lau, J. Sampson, E. Perelman, G. Hamerly, and B. Calder. 2005. The Strong correlation Between Code Signatures and Performance. In *International Symposium on Performance Analysis of Systems and Software (ISPASS)*.
- [48] R.L. Mattson, J. Gecsei, D. R. Slutz, and I. L. Traiger. 1970. Evaluation techniques for storage hierarchies. *IBM Systems Journal* (1970).
- [49] Sparsh Mittal, Yanan Cao, and Zhao Zhang. 2013. Master: A multicore cache energy-saving technique using dynamic cache reconfiguration. *Transactions on very large scale integration systems (VLSI)* (2013).
- [50] Sparsh Mittal, Zhao Zhang, and Jeffrey S Vetter. 2013. FlexiWay: A cache energy saving technique using fine-grained cache reconfiguration. In *International Conference on Computer Design (ICCD)*.
- [51] Nikos Nikolieris, Lieven Eeckhout, Erik Hagersten, and Trevor E. Carlson. 2019. Directed Statistical Warming Through Time Traveling. In *International Symposium on Microarchitecture (MICRO)*.
- [52] Nikos Nikolieris, Andreas Sandberg, Erik Hagersten, and Trevor E Carlson. 2016. CoolSim: Statistical techniques to replace cache warming with efficient, virtualized profiling. In *2016 International Conference on Embedded Computer Systems: Architectures, Modeling and Simulation (SAMOS)*.
- [53] OpenMP [n. d.]. OpenMP 3.1 API C/C++ Syntax Quick Reference Card. <https://www.openmp.org/wp-content/uploads/OpenMP3.1-CCard.pdf>.
- [54] Irma Esmir Papazian. 2020. New 3rd Gen Intel® Xeon® Scalable Processor (Codename: Ice Lake-SP). In *Hot Chips Symposium (HCS)*.
- [55] Harish Patil and Trevor E. Carlson. 2014. Pinballs: portable and shareable user-level checkpoints for reproducible analysis and simulation. In *Workshop on Reproducible Research Methodologies (REPRODUCE)*.
- [56] Harish Patil, Alexander Isaev, Wim Heirman, Alen Sabu, Ali Hajiabadi, and Trevor E Carlson. 2021. ELFies: Executable Region Checkpoints for Performance Analysis and Simulation. In *International Symposium on Code Generation and Optimization (CGO)*.
- [57] Harish Patil, Cristiano Pereira, Mack Stallcup, Gregory Lueck, and James Cownie. 2010. PinPlay: A Framework for Deterministic Replay and Reproducible Analysis of Parallel Programs. In *International Symposium on Code Generation and Optimization (CGO)*.
- [58] E. Perelman, M. Polito, J.-Y. Bouguet, J. Sampson, B. Calder, and C. Dulong. 2006. Detecting phases in parallel applications on shared memory architectures. In *International Parallel Distributed Processing Symposium (IPDPS)*.
- [59] Alen Sabu, Changxi Liu, and Trevor E. Carlson. 2024. Viper: Utilizing Hierarchical Program Structure to Accelerate Multi-Core Simulation. *IEEE Access* (2024).
- [60] Alen Sabu, Harish Patil, Wim Heirman, and Trevor E. Carlson. 2022. LoopPoint: Checkpoint-driven Sampled Simulation for Multi-threaded Applications. In *International Symposium on High Performance Computer Architecture (HPCA)*.
- [61] Alen Sabu, Harish Patil, Wim Heirman, and Trevor E. Carlson. 2022. LoopPoint Tools. <https://github.com/nus-comparch/looppoint>.
- [62] Daniel Sanchez and Christos Kozyrakis. 2013. ZSim: Fast and Accurate Microarchitectural Simulation of Thousand-Core Systems. In *International Symposium on Computer Architecture (ISCA)*.
- [63] Andreas Sandberg, Nikos Nikolieris, Trevor E. Carlson, Erik Hagersten, Stefanos Kaxiras, and David Black-Schaffer. 2015. Full Speed Ahead: Detailed Architectural Simulation at Near-Native Speed. In *International Symposium on Workload Characterization (IISWC)*.
- [64] Xipeng Shen, Yutao Zhong, and Chen Ding. 2004. Locality phase prediction. *International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*.
- [65] Timothy Sherwood, Erez Perelman, Greg Hamerly, and Brad Calder. 2002. Automatically Characterizing Large Scale Program Behavior. In *International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*.
- [66] Timothy Sherwood, Suleyman Sair, and Brad Calder. 2003. Phase tracking and prediction. *International Symposium on Computer Architecture (ISCA)*.
- [67] Michael Van Biesbrouck, Brad Calder, and Lieven Eeckhout. 2006. Efficient sampling startup for SimPoint. *IEEE Micro* (2006).
- [68] M. Van Biesbrouck, T. Sherwood, and B. Calder. 2004. A co-phase matrix to guide simultaneous multithreading simulation. In *International Symposium on Performance Analysis of Systems and Software (ISPASS)*.
- [69] Michael J Voss and Rudolf Eigemann. 2001. High-level adaptive program optimization with ADAPT. In *Symposium on Principles and Practices of Parallel Programming (PPoPP)*.
- [70] T. F. Wenisch, R. E. Wunderlich, M. Ferdman, A. Ailamaki, B. Falsafi, and J. C. Hoe. 2006. SimFlex: Statistical Sampling of Computer System Simulation. *IEEE Micro* (2006).
- [71] Roland E. Wunderlich, Thomas F. Wenisch, Babak Falsafi, and James C. Hoe. 2003. SMARTS: Accelerating Microarchitecture Simulation via Rigorous Statistical Sampling. In *International Symposium on Computer Architecture (ISCA)*.
- [72] Xin You, Changxi Liu, Hailong Yang, Pengbo Wang, Zhongzhi Luan, and Depei Qian. 2022. Vectorizing SpMV by Exploiting Dynamic Regular Patterns. In *International Conference on Parallel Processing (ICPP)*.
- [73] C. Yount, H. Patil, and M. S. Islam. 2015. Graph-matching-based simulation-region selection for multiple binaries. In *International Symposium on Performance Analysis of Systems and Software (ISPASS)*.
- [74] Jianping Zeng, Hongjune Kim, Jaejin Lee, and Changhee Jung. 2021. Turnpike: Lightweight Soft Error Resilience for In-Order Cores. In *International Symposium on Microarchitecture (MICRO)*.
- [75] Weifeng Zhang, B. Calder, and D.M. Tullsen. 2005. An event-driven multithreaded dynamic optimization framework. In *International Conference on Parallel Architectures and Compilation Techniques (PACT)*.