

# UC Davis

## UC Davis Electronic Theses and Dissertations

### Title

Portable Targeted Sampling for Realistic Workload Evaluation in Computer Architecture

### Permalink

<https://escholarship.org/uc/item/5v4385sp>

### ISBN

9798189233500

### Author

Qiu, Zhantong

### Publication Date

2026-06-25

Peer reviewed|Thesis/dissertation

Portable Targeted Sampling for Realistic Workload Evaluation in  
Computer Architecture

By

ZHANTONG QIU

THESIS

Submitted in partial satisfaction of the requirements for the degree of

MASTER OF SCIENCE

in

Computer Science

in the

OFFICE OF GRADUATE STUDIES

of the

UNIVERSITY OF CALIFORNIA

DAVIS

Approved:

---

Jason Lowe-Power, Chair

---

John Owens

---

Tapti Palit

Committee in Charge

2026

Copyright © 2026 by

Zhantong Qiu

*All rights reserved.*

# CONTENTS

|                                                                          |           |
|--------------------------------------------------------------------------|-----------|
| List of Figures . . . . .                                                | iv        |
| List of Tables . . . . .                                                 | v         |
| Abstract . . . . .                                                       | vi        |
| Acknowledgments . . . . .                                                | vii       |
| <b>1 Introduction</b>                                                    | <b>1</b>  |
| <b>2 Motivation and Prior Work</b>                                       | <b>4</b>  |
| 2.1 Samples Are Expensive To Find . . . . .                              | 4         |
| 2.2 Samples Are Tied To A Single Executable Binary . . . . .             | 5         |
| 2.3 Validating The Selected Samples Is Infeasible . . . . .              | 6         |
| <b>3 Full-System LoopPoint</b>                                           | <b>8</b>  |
| 3.1 LoopPoint Background . . . . .                                       | 8         |
| 3.2 Full-System Implementation . . . . .                                 | 10        |
| 3.2.1 Phase Analysis . . . . .                                           | 10        |
| 3.2.2 Checkpoint Creation and Representative Region Simulation . . . . . | 12        |
| 3.3 Evaluation . . . . .                                                 | 12        |
| <b>4 The Nugget Framework</b>                                            | <b>16</b> |
| 4.1 Unit of Work . . . . .                                               | 18        |
| 4.2 Preparation . . . . .                                                | 19        |
| 4.3 Interval Analysis . . . . .                                          | 20        |
| 4.3.1 Interval Analysis LLVM Pass . . . . .                              | 21        |
| 4.3.2 Interval Analysis Executable . . . . .                             | 21        |
| 4.4 Nugget Creation . . . . .                                            | 23        |
| 4.4.1 Markers . . . . .                                                  | 23        |
| 4.4.2 Overhead of Hooks . . . . .                                        | 24        |
| 4.5 Sample Validation . . . . .                                          | 25        |

|          |                                                                                                           |           |
|----------|-----------------------------------------------------------------------------------------------------------|-----------|
| <b>5</b> | <b>Evaluation</b>                                                                                         | <b>26</b> |
| 5.1      | Interval Analysis Efficiency (Goal 1)                                                                     | 28        |
| 5.1.1    | Comparison with Functional Simulation                                                                     | 29        |
| 5.1.2    | Mix of Workload Types                                                                                     | 31        |
| 5.1.3    | Scaling with Number of Threads                                                                            | 32        |
| 5.1.4    | Summary                                                                                                   | 33        |
| 5.2      | Applying Sampling Methodologies to the Nugget Framework and Validating the Selected Samples (Goals 2 & 3) | 34        |
| 5.2.1    | Sampling Methodologies                                                                                    | 34        |
| 5.2.2    | Sample Validation                                                                                         | 35        |
| 5.2.3    | Hook Overhead                                                                                             | 37        |
| <b>6</b> | <b>Case Studies</b>                                                                                       | <b>40</b> |
| 6.1      | Nuggets in Simulation and on Hardware                                                                     | 41        |
| 6.1.1    | Experimental Setup                                                                                        | 41        |
| 6.1.2    | Results                                                                                                   | 42        |
| 6.2      | Model Accuracy Evaluation                                                                                 | 46        |
| <b>7</b> | <b>Conclusion</b>                                                                                         | <b>48</b> |
| 7.1      | LoopPoint Full-System Evaluation                                                                          | 48        |
| 7.2      | Nugget Framework Contributions                                                                            | 49        |
| 7.3      | Limitations                                                                                               | 49        |
| 7.4      | Future Work                                                                                               | 50        |

## LIST OF FIGURES

|     |                                                                          |    |
|-----|--------------------------------------------------------------------------|----|
| 3.1 | LoopPoint analysis flow in gem5 . . . . .                                | 11 |
| 3.2 | LoopPoint prediction error and speedup for NPB in full-system mode . .   | 14 |
| 4.1 | Nugget pipeline overview . . . . .                                       | 17 |
| 5.1 | Interval analysis overhead: Nugget vs. functional simulation . . . . .   | 30 |
| 5.2 | Interval analysis overhead across different workloads . . . . .          | 31 |
| 5.3 | Interval analysis overhead across thread counts . . . . .                | 33 |
| 5.4 | Prediction error by machine and sampling method . . . . .                | 35 |
| 5.5 | Marker hook executions for randomly selected nuggets . . . . .           | 38 |
| 6.1 | Speedup prediction error distribution: simulation vs. hardware . . . . . | 43 |
| 6.2 | Speedup prediction error: varying ISA, fixed microarchitecture . . . . . | 44 |
| 6.3 | Speedup prediction error: varying microarchitecture, fixed ISA . . . . . | 44 |
| 6.4 | Speedup prediction error: varying both ISA and microarchitecture . . . . | 45 |
| 6.5 | Memory instruction counts: gem5 microcodings vs. real hardware . . . .   | 47 |

## LIST OF TABLES

|     |                                                           |    |
|-----|-----------------------------------------------------------|----|
| 3.1 | gem5 full-system simulation configuration . . . . .       | 13 |
| 5.1 | Real Hardware Configuration . . . . .                     | 27 |
| 5.2 | Overview of Benchmark Suites Used in Evaluation . . . . . | 27 |
| 6.1 | gem5 simulation model configuration . . . . .             | 42 |

## ABSTRACT

### Portable Targeted Sampling for Realistic Workload Evaluation in Computer Architecture

We present Nugget, an open-source, portable framework for targeted sampling that defines representative samples at the LLVM IR level and analyzes and validates them on real hardware rather than in simulation. Because the unit of work is LLVM IR rather than machine instructions, a single interval analysis applies to every ISA, and the resulting samples compile into lightweight executables (*nuggets*) that validate natively and are then used in simulation. As a simulation baseline, we contribute and have upstreamed the first full-system implementation of LoopPoint in gem5, a methodology whose reliance on slow analysis, full-baseline validation, and binary-specific samples motivates Nugget’s design. Across SPEC CPU2017, NPB, and LSMS on AMD Ryzen and Ampere Altra hardware, Nugget’s interval analysis runs  $577\times$  faster than gem5 functional simulation, adding  $3\times$  overhead over native execution for single-threaded workloads and  $34\times$  for multi-threaded ones. We show that Nugget allows researchers to iterate quickly on sampling strategies using real hardware. We also find that consistency of prediction error across machines is a stronger indicator of sample quality than low absolute error on any single machine, and that microarchitectural differences affect prediction accuracy more than ISA differences. Beyond sampling, nuggets serve as realistic microbenchmarks that expose simulator inaccuracies and have revealed a microcoding error in gem5’s handling of Arm paired-memory instructions.

## ACKNOWLEDGMENTS

I would like to thank my advisor, Professor Jason Lowe-Power, for all of his support and guidance throughout my master’s thesis, my master’s program, my research, and my life. Professor Lowe-Power gave me my start in computer architecture research and showed me the fascinating world of simulation and performance evaluation, recognizing a potential in me that my transcript, at the time, did little to suggest. No one in my family had pursued education beyond middle school, so I had struggled to find my own way through school, learning through trial and error without anyone to guide me. As my advisor, Professor Lowe-Power gave me the guidance in education, science, and career that I had never had before. He did not just give me the opportunity to work on simulation and sampling methodology; he also opened the door for me to meet, learn from, and work alongside remarkable people whose intelligence and generous hearts continue to amaze me. As a researcher, he constantly showed me new ways to approach a problem, and he was always willing to discuss and debate ideas with me as an equal, never once talking down to me. He showed me what a great researcher looks like, and through his example, I learned how to become one. I will forever be grateful for the chance to work with him and to have him as my advisor.

Professor Christopher Batten at Cornell University welcomed me into his group and gave me the chance to work with him on hardware-software co-design in robotics during my year-long visit. Working in the Batten Research Group and the Computer Systems Laboratory (CSL) exposed me to a new field and a new research culture. There, I learned to build the infrastructure for co-simulating robotics and computer architecture, to help profile and evaluate new architectural ideas for the problems robotics faces. He was also candid about where I still need to grow, and I am still growing because of it. I am grateful to his students as well, who showed me new challenges and new ways of approaching problems. I feel fortunate to have spent that year being advised by Professor Batten and working alongside his group.

I would also like to thank my committee, Professors John Owens and Tapti Palit, for their time and their careful reading of this work. My thesis is not only about sampling

methodology and simulation; it also draws heavily on microarchitecture and compilers. Their feedback in these areas made my work more precise and complete, and their guidance on clarity, structure, and academic conventions strengthened it as a piece of writing. Beyond the thesis itself, they offered valuable suggestions for carrying this work forward.

Mahyar Samani has been my mentor and a close friend throughout my undergraduate internship and my master's program in the Davis Computer Architecture Research Group at UC Davis. He has never hesitated to help those around him. He taught me the basics of nearly everything in the early days of my research journey, and contributed enormously to my work later on. He was always willing to challenge my ideas and give me honest feedback, and some of our best research grew out of our lively discussions. He is as passionate about research as I am. I am fortunate to have had both his mentorship and his friendship, and I could not have come this far without him.

Dr. Alen Sabu, the author of the original LoopPoint paper, collaborated with me on my contribution to its gem5 implementation. An expert in sampling methodology working in the same narrow field, he understood the niche challenges I faced in a way few others could, and his support was invaluable as I found my footing as a researcher. He also offered generous guidance through my job search and interviews. I am glad that we ended up publishing together.

I would also like to thank everyone in the Davis Computer Architecture Research Group and the gem5 Development Group at UC Davis, whose discussions, feedback, and camaraderie made my time there both more productive and more enjoyable.

I would like to thank my parents and my brother for their support. While I have been growing as a student and researcher, I have spent most of these years away from them. I have missed all of their birthdays and most of our holidays together. Through all of it, my parents and my brother have been understanding, supportive of my choices, and a comfort whenever times were hard. My parents in particular have taught me the most important things by their own example: to work hard, to keep improving, and to learn from failure instead of being too scared to try. I am proud of my parents and my brother. I am the person I am because of them.

I met some of the kindest, smartest, and most remarkable people on this journey, and I will treasure it forever. I hope to pass on to others the same help and lessons they gave me.

# Chapter 1

## Introduction

Evaluating application performance on novel architectural ideas is key to computer architecture research. However, especially for hardware which does not yet exist, evaluating application performance can be very time-consuming. For instance, when using cycle-level software simulation, there can be over  $100,000\times$  slowdown. Even FPGA-accelerated simulation is significantly slower than wall-clock time [27]. At the same time, modern applications are growing in their computational demands from exascale applications [34] to machine-learning training (GPT-4 required 2,150,000 Exa- or  $2.15 \times 10^{25}$  operations to complete [39]).

To facilitate the evaluation of large applications, there has been a long history of reducing applications to smaller benchmarks, kernels, or short representative intervals. For instance, Livermore Loops provide a set of microbenchmarks that represent key computational kernels from scientific applications [18]. It is also common to pull out kernels from an application or to create “proxy applications” [4]. While these methods can reduce very large applications to much more manageable benchmarks, they require application experts to manually identify the important parts of applications and create new programs.

Another approach to reducing the size of applications is to *sample* the application to automatically find *representative intervals* of different application phases. Many techniques (e.g., SimPoint [43], BarrierPoint [14], and LoopPoint [40]) follow this approach.

In this thesis, we focus on targeted sampling techniques, which use profiling to auto-

---

Portions of this thesis have been published in [38] and [41].

matically select samples.

Prior targeted sampling techniques have three main drawbacks that limit their applicability to modern, realistically sized applications:

- **Samples are expensive to find:** Targeted sampling techniques require running the application from beginning to end once (typically via simulation or slow emulation) to define intervals, analyze program phases, and identify representative samples.
- **Samples are tied to a single executable binary:** These techniques generate samples that are specific to a particular binary, so any binary change necessitates repeating the entire sampling process.
- **Validating the selected samples is infeasible:** Targeted sampling techniques require running the application in detailed simulation to verify that the selected samples accurately represent the program.

This thesis takes a two-part approach. First, because existing sampling methodologies have mostly been validated in user-mode simulation, we implement the state-of-the-art LoopPoint methodology in gem5’s full-system mode and evaluate it with the NPB benchmark suite. This investigation confirms that LoopPoint can be extended to full-system simulation, but also highlights that the analysis remains time-consuming, that validating sampling accuracy still requires a full baseline simulation—the very bottleneck that sampling aims to avoid—and that prediction errors are inconsistent across ISAs because each ISA-specific binary produces different representative samples.

Motivated by these observations, we present *Nugget*, a framework for agile targeted sampling development. Because modern applications exhibit varied characteristics, we believe there is no one-size-fits-all sampling methodology or selection algorithm [16, 24, 28, 54]. Therefore, Nugget is designed to be flexible and extensible, allowing researchers to create new sample selection algorithms tailored to their targeted workload and input without the current limitations of targeted sampling techniques. Nugget’s cross-platform compatibility and ability to run on real hardware enable richer insights into both the selected samples and the systems that run them.

The main contributions of this thesis are:

- **A full-system LoopPoint implementation in gem5**, providing the first evaluation of LoopPoint with the NPB benchmark suite in full-system simulation, establishing a simulation-based baseline for sampling methodology validation.
- **Efficient and portable interval analysis** on real hardware using LLVM, unconstrained by the hardware’s microarchitecture or ISA.
- **An LLVM IR-level unit of work**, enabling sample selection algorithms to create and locate samples for a program in a cross-platform manner, agnostic of machine-level instructions, such as architecture-specific instructions and optimizations.
- **A sample creation methodology** that allows samples to run on any platform (e.g., real hardware or simulators).
- **An efficient validation methodology** that validates the selected samples for the target workload and input on native hardware, integrated into the development workflow.

# Chapter 2

## Motivation and Prior Work

Nugget is designed to address the following limitations of prior sampling techniques:

- Samples are expensive to find.
- Samples are tied to a single executable binary.
- Validating the selected samples is infeasible.

In this section, we discuss each of these limitations and how prior work in sampling has or has not addressed them.

### 2.1 Samples Are Expensive To Find

Targeted sampling techniques require analyzing the entire workload execution, which is a time-consuming process.

For example, the SPECspeed suite in the SPEC CPU2017 benchmark suite [11] has an average of 21.8 trillion dynamic instructions per benchmark for floating-point benchmarks with reference input size, which is the only input size designed to report time metrics [31]. Functional simulators, such as the gem5 atomic model, are often limited to a simulation rate of approximately 1 Million Instructions Per Second (MIPS) [5]. Therefore, it takes approximately 250 days of wall-clock time to analyze an average SPECspeed benchmark with reference input size. More recent workloads (e.g., the forthcoming SPEC CPUv8 [46],

HPC workloads, and machine-learning workloads) are significantly larger, making functional simulation infeasible.

For multi-threaded benchmarks, the analysis overhead is even worse. We must analyze the program once for every different thread configuration because the number of threads affects the program’s execution. Also, if the workload undergoes any changes (e.g., new compiler, different optimizations, etc.), we must redo the analysis.

Live sampling tools like Pac-Sim [32] parallelize the phase analysis and detailed simulation to reduce the overall phase analysis penalty. However, live sampling still relies on the simulator to analyze the program’s execution so this technique is constrained by the simulation speed. If a benchmark takes the simulator 250 days to analyze, as mentioned in the above example, then parallelizing the analysis and detailed simulation will still take at least 250 days.

Tools like DynamoRIO [10] and the Intel PinPoints toolkit [35] can analyze the program’s execution with dynamic binary instrumentation or binary translation. The instrumentation and profiling routines introduce overhead to the program’s execution, but this overhead is significantly smaller than that of simulation (e.g., 4-40 $\times$  for QEMU with some single-threaded workloads [49]). However, these tools often have restrictions on the architectures they can analyze or the host machines. Emerging architectures such as RISC-V, which must be analyzed across a wider range of benchmarks, often fall outside these tools’ support

**Goal 1:** Analyze the program’s execution quickly and without any restrictions on the architecture or host machines.

## 2.2 Samples Are Tied To A Single Executable Binary

Most prior work relies on binary-specific information, such as the program counter, to define intervals, identify program phases, and select samples. Any binary change, such as recompiling with different ISA extensions or compilers, alters the static instructions and the dynamic instruction stream, invalidating previously selected samples. As discussed in Section 2.1, the high overhead of this analysis makes repeating it for every binary revision

time-consuming.

Cross-binary simulation points [36] represent a notable attempt to decouple samples from a specific binary. However, they are limited to single-threaded programs and require significant manual effort, including mapping their control flow graphs and recalculating sample weights.

**Goal 2:** Automatically generate samples that are independent of the binary.

## 2.3 Validating The Selected Samples Is Infeasible

Sampling methodologies aim to predict overall program performance by measuring only a subset of representative samples and extrapolating to full-program behavior using statistical models. Most methodologies focus on predicting metrics that summarize overall hardware performance, such as instructions per cycle (IPC) or total runtime.

To evaluate the accuracy of a sample set selected by a sampling methodology, we compare its predicted performance metrics against ground-truth measurements. First, we measure the target metric from a full execution of the program on the experimental system, which serves as the ground truth. Next, we measure the same metric on the selected samples and apply statistical processing to estimate the full-program value. Finally, we compare the predicted value against the ground truth to quantify the accuracy of the selected samples.

The effectiveness of a sample-selection method depends on benchmark characteristics. Differences in control-flow phase behavior, working-set size and locality, and input sensitivity can cause a method that succeeds on one workload–input pair to misrepresent another [12, 15, 16, 19, 43, 51, 54]. For example, although SimPoint reports an average IPC error of only 3% on SPEC CPU2000 [44] programs [26], empirical studies such as MAV [12] report outliers where it fails to capture distinct memory-access behavior. Representative intervals should therefore be validated, yet the only way to validate them is to run the entire workload.

Prior sampling techniques are typically evaluated on small benchmarks or truncated runs because their validation depends on simulation [14, 22, 26, 30, 32, 36, 40, 42, 43, 52,

53]. What is missing is an agile path from methodology development to real-hardware validation. A sample can be measured natively only if its region boundaries are locatable on hardware, but these markers are methodology-specific. For example, LoopPoint [40] bounds intervals by a pair of program counters together with their global cross-thread execution counts, which `perf` cannot delimit directly. Locating such regions on hardware required a dedicated framework, ROIperf [41], which appeared four years after LoopPoint; this lag is evidence that no agile path from methodology development to hardware-level validation yet exists.

Validation therefore falls back on simulation, which introduces its own problems. Because the sample set is intended to approximate the behavior of the real system, validation by simulation is meaningful only if the simulator faithfully models the target; if it does not, the validation does not transfer. Simulation also makes methodologies hard to compare, since results are tied to different simulators and redundant engineering is required before any apples-to-apples comparison. For example, LoopPoint is validated on SPEC CPU2017 benchmarks with train-size inputs using Sniper [13], whereas FSA [42] is validated on the first 30 billion instructions of the SPEC CPU2006 [45] benchmarks using gem5 [9,33]. Such techniques can be validated only in a limited way, because each requires a full, detailed simulation of the workload from beginning to end. Detailed simulation can be 2–10 $\times$  slower than functional simulation, so validating full-sized workloads could take nearly a decade!

The lack of an agile validation workflow has meant that sampling techniques are rarely validated on large benchmarks. Evidence suggests that validation results from smaller benchmarks or partial executions do not generalize to larger benchmarks and realistic workloads [26, 31, 43, 51, 54]. Moreover, a workload’s behavior can change with its input [15]. Validation is therefore needed for each benchmark–input pair.

**Goal 3:** Quickly validate the selected samples for the targeted benchmarks and input size.

# Chapter 3

## Full-System LoopPoint

Because Nugget analyzes and evaluates sampling methodologies in a full-system environment (real hardware with an operating system), we first need to understand how prior sampling methodologies behave in such an environment. LoopPoint was originally validated in user-mode simulation with Sniper [13], which excludes OS-level effects such as system calls, interrupts, and context switches. Full-system simulation captures these effects and provides a more realistic evaluation environment, but it also introduces new challenges for sampling methodologies. By implementing and evaluating LoopPoint in gem5’s full-system mode, we both extend LoopPoint to a setting where it has not been tested and establish a simulation-based baseline against which Nugget’s real-hardware validation can be compared.

In this section, we describe the LoopPoint background (Section 3.1), present our full-system implementation in gem5 (Section 3.2), and evaluate the methodology with the NPB benchmark suite (Section 3.3).

### 3.1 LoopPoint Background

The LoopPoint methodology [40] is a targeted sampling technique for multi-threaded applications that contributes (1) repeatable, upfront application analysis, (2) a novel clustering approach that accounts for runtime parallelism, and (3) the use of loop-based simulation markers to divide execution into measurable regions, even in the presence of spin-loops.

The methodology proceeds in three stages: phase analysis, checkpoint creation, and representative region simulation. During phase analysis, LoopPoint profiles the application to identify loop-based region boundaries and constructs basic block vectors (BBVs) for each region. Regions are defined using (PC, count) pairs—a program counter value and the number of times that instruction has been executed—which locate execution points in multi-threaded applications. Synchronization library instructions are excluded from the analysis, as they do not reflect application-level phase behavior. The BBVs are then clustered offline to identify groups of similar regions, and a representative region is selected from each cluster.

For simulation, LoopPoint creates checkpoints at the boundaries of selected representative regions. While fast-forwarding can also be used to reach region boundaries, checkpoints are preferable because they allow simulations to be distributed across simulation configurations and ensure that each configuration starts from the exact same execution state. Each checkpoint is restored and simulated in detailed mode with appropriate warmup, and the results are weighted by cluster membership to predict full-program performance.

In the original paper, LoopPoint achieves an average simulation error of 2.3% on multi-threaded SPEC CPU2017 benchmarks with the train input set, with a speedup of up to  $801\times$  over full detailed simulation. When representative regions are simulated in parallel, the speedup reaches  $11,587\times$  on average and up to  $31,253\times$ .

Notably, the original LoopPoint methodology uses Sniper [13] for both analysis and simulation, operating entirely in user-mode without modeling OS-level effects. A follow-up paper [41] extends the LoopPoint methodology and proposes ROIperf, a hardware-based framework for rapid correlation of representative regions, enabling performance measurement of full workloads directly on hardware. This thesis contributed to that follow-up work; the full-system gem5 implementation described in Section 3.2 served as the full-system simulation framework for that study.

## 3.2 Full-System Implementation

Full-system simulation models the entire software stack—from applications through the operating system down to microarchitecture and devices—providing deeper insights into system behavior within a realistic environment. However, the involvement of the operating system introduces new challenges for sampling methodologies. The address space must remain stable throughout the application’s execution to ensure accurate profiling, and kernel-level instructions should be excluded from application phase analysis.

We implement the LoopPoint methodology in gem5 [9, 33], a cycle-level simulator widely used in both academia and industry that supports full-system simulation. Following the three-stage pipeline described in Section 3.1, our implementation maps each stage to gem5 through two new probe modules: the LoopPoint Analysis module, which handles phase analysis, and the PcCount Tracker module, which drives checkpoint creation and representative region simulation. As of gem5 version 24.1, this LoopPoint functionality has been upstreamed and is readily available.

### 3.2.1 Phase Analysis

To accurately capture application phases, the LoopPoint methodology requires a stable address space and must exclude all kernel instructions [50]. Before starting the simulation, we disable Address Space Layout Randomization (ASLR) in the kernel and obtain the address mapping of the application. A stable address mapping is critical because LoopPoint uses virtual addresses to mark region boundaries and to identify synchronization library calls that should be excluded from phase analysis. We conduct the entire phase analysis in gem5’s functional mode using the atomic CPU.

The LoopPoint Analysis module filters out all kernel instructions to prevent them from affecting application phase analysis. Synchronization library instructions are also excluded, as they do not directly contribute to the application’s phase behavior (see Section 3.1). Only loops within the application code are considered as candidates for region markers. The analysis produces a list of regions with their basic block vectors (BBVs) and region boundaries expressed as (PC, count) pairs that serve as markers for each region.




Figure 3.1: LoopPoint analysis flow in gem5.

### 3.2.2 Checkpoint Creation and Representative Region Simulation

After clustering the BBVs and selecting representative regions, we create gem5 checkpoints using the respective (PC, count) pairs. These pairs serve as markers that identify the warmup start, simulation start, and end points of each representative region. The PcCount Tracker module maintains a list of these markers and monitors committed instructions during simulation, ensuring precise control over checkpoint creation and region boundaries.

Checkpoints are placed at the start of each warmup region. Upon restoring a checkpoint, gem5 executes the warmup instructions before collecting statistics in the representative region. By the time simulation reaches the region of interest, key microarchitectural structures (such as caches and branch predictors) have been warmed up, yielding more accurate performance measurements.

Since checkpoints are self-contained, simulations of representative regions can execute in parallel, significantly reducing overall simulation time. The PcCount Tracker module continuously monitors committed instructions to precisely detect the end of the warmup region and the end of each representative region, enabling accurate performance measurements.

## 3.3 Evaluation


| Component          | Config 1       | Config 2       |
|--------------------|----------------|----------------|
| ISA                | x86            | Arm            |
| Cores              | 8×Out-of-Order | 8×Out-of-Order |
| Frequency (GHz)    | 4              | 3              |
| L1 D/I (KiB)       | 64             | 64             |
| L2 (KiB)           | 1024           | 1024           |
| L3/SLC (MiB)       | –              | 32             |
| Memory clock (MHz) | 2400           | 2400           |
| Bandwidth (GiB/s)  | 38.4           | 76.8           |
| OS                 | Ubuntu 24.04   | Ubuntu 24.04   |

Table 3.1: The configuration used for the full-system simulation on gem5. L1 D/I is the per-core data and instruction cache size (equal here); L2 is per core and L3/SLC is shared. Both configurations use DDR4 memory.


To evaluate the effectiveness of the LoopPoint methodology in full-system simulation, we apply the complete sampling pipeline—using both the LoopPoint Analysis and PcCount Tracker modules—to the NPB benchmark suite. We use Class A input size with eight threads for our experiments, as larger input sizes (Class B and above) do not complete baseline full-system simulation within a reasonable time frame. In addition to evaluating in a full-system environment, we test the methodology across different ISAs and system configurations, as detailed in Table 3.1. We use Ubuntu 24.04 as the guest operating system in all experiments.

We conduct phase analysis and checkpoint creation in gem5’s functional mode. The checkpoints of selected representative regions are then simulated in detailed mode using the configurations shown in Table 3.1. We evaluate the accuracy of the LoopPoint methodology by comparing the predicted runtimes against baseline full simulations. Figure 3.2a shows the relative prediction error for the NPB Class A benchmarks.


The Class A input size does not produce enough regions for the methodology to



(a) Prediction error



(b) Serial Speedup



(c) Parallel Speedup

Figure 3.2: LoopPoint prediction error and speedup for NPB Class A with eight threads in full-system mode, compared across the x86 and Arm ISAs. For the same workload, the prediction error differs in both magnitude and direction between ISAs, and the speedups diverge as well. Because LoopPoint selects samples per binary, each ISA yields a different set of representative regions, introducing additional error when their predictions are used to compare performance across ISAs.

build an accurate sample set. For example, only eight total regions are discovered for *cg* (from which two representatives are selected) and three for *is* (from which one is selected). With so few regions, the clustering and selection process has limited statistical power. The prediction errors are not consistent across the x86 and Arm configurations. Because LoopPoint constructs BBVs from machine-level basic blocks and defines region boundaries with ISA-specific (PC, count) pairs, each ISA binary produces a different set of regions, clusters, and representative samples. As a result, using predicted runtimes to compare performance across ISAs can compound the error beyond the absolute error of each individual prediction. This underscores the need for ISA-independent samples—the same set of representative regions used across different ISAs—to avoid the additional error introduced by independently selecting samples for each configuration.

Figure 3.2 also shows the speedup achieved by sampled simulation. The speedup is calculated as the reduction in the number of instructions simulated in detailed mode compared to the full baseline. Due to the limited number of regions discovered for Class A benchmarks, the selected representatives constitute a large fraction of the total regions, resulting in modest speedup for this experiment. The speedup from the LoopPoint methodology primarily depends on the number of samples required to achieve a target accuracy, rather than on the simulator itself; larger workloads with more regions would yield greater reductions.

These results demonstrate that the LoopPoint methodology can be extended to full-system simulation in *gem5*, but they also highlight three key limitations. First, validating the methodology’s accuracy requires completing a full baseline simulation. Second, because LoopPoint’s analysis is tied to a specific binary, each ISA produces a different set of representative samples, making cross-ISA comparison unreliable. Third, interval analysis via functional simulation is time-consuming; in Section 5.1.1, we use the LoopPoint module to show that this overhead makes analysis impractical for realistic input sizes. These observations motivate the Nugget framework (Section 4), which performs interval analysis on real hardware, validates samples without simulation, and defines samples at the LLVM IR level so that the same set of samples can be used across different binaries.

# Chapter 4

## The Nugget Framework

Nugget<sup>1</sup> is designed to address the limitations of the prior sampling techniques, specifically focusing on 1) the high cost of finding samples, 2) the coupling of the samples to a specific executable binary, and 3) the infeasibility of validating the selected samples. Nugget leverages LLVM passes to statically insert annotations to analyze the program's execution on real hardware without the need for simulation, emulation, or binary translation. By using LLVM intermediate representation (IR) to define program progress and intervals, the samples become independent of the binary and can be reused across different executable versions of the same program, since LLVM IR abstracts away details of the final binary. Lastly, Nugget provides a method to create portable samples to be executed on any platform, such as real hardware or simulators, so that the samples can be validated in native time with real hardware before running them on a simulator.

---

<sup>1</sup>The Nugget framework is open source and publicly available at <https://github.com/studyztp/Nugget-LLVM-passes>.




Figure 4.1: The Nugget pipeline. A blue dragon marks a step that passes through the LLVM compiler, and a gear marks a component that is an executable. The pipeline is divided into two main phases: (1) preparation and analysis, and (2) nugget creation and validation. The preparation and analysis phase—base IR file creation and interval analysis—is performed only once per program. The nugget creation and validation phase can be repeated to refine and select the most representative set of samples for the workload.

In this section, we present an overview of Nugget as a pipeline (Figure 4.1), describing its workflow from preparation through interval analysis, nugget creation, and finally, sample selection verification. It is not the focus of this thesis to introduce a new sample selection algorithm, but rather to provide a framework that can be used to implement and validate diverse sample selection methodologies that can be tailored to specific applications and workloads.

We begin by defining program progress and intervals, which ensures that samples remain independent of the binary, as described in Section 4.1. Next, we describe the preparation required by Nugget to guarantee the LLVM IR of the program remains consistent across different binaries in Section 4.2. Then, we introduce our novel interval analysis methodology, which leverages LLVM passes, enabling comprehensive program analysis directly on real hardware without imposing hardware-specific constraints in Section 4.3. Following that, we present the nugget creation process, which allows for the generation of portable samples that can be executed on any platform, including real hardware and simulators, in Section 4.4. Finally, we describe the sample selection validation process, which enables efficient validation of the selected samples for specific workloads in native time, ensuring that the samples accurately represent the program’s execution in Section 4.5.

## 4.1 Unit of Work

A unit of work defines the metric by which progress toward completing a program’s execution is measured. It must remain consistent across identical runs of a program with the same input. While machine instructions could serve this purpose for a single binary, Nugget instead counts executed LLVM IR instructions as its unit of work. This ensures a uniform measurement of program progress across all binaries of the same program with the same input.

A unit of work directly determines the notion of intervals. For example, if the unit of work is the number of executed instructions, then each interval corresponds to a fixed number of executed instructions during runtime. A sample is a specific interval of a program’s execution; thus, the unit of work also defines what constitutes a sample.

The choice of unit of work influences not only how intervals are defined, but also the sample selection algorithm and the statistical models used in performance prediction. Therefore, the unit of work must be defined consistently across all binaries to ensure that samples remain independent of any specific binary. For example, if the unit of work is the number of executed machine-level instructions, then the sample will be tied to a specific binary because any binary changes can affect the number of machine-level instructions

executed, even if the control flow remains the same [36].

LLVM [29] is an open-source compiler infrastructure that transforms source code into a platform-independent intermediate representation (IR) before generating machine-specific binaries. LLVM IR instructions abstract away the underlying machine-level instructions, so as long as the control flow remains the same at the LLVM IR level, the definition of all intervals of program execution remains consistent across different binaries. Since LLVM IR instructions are an abstraction of machine-level instructions, the number of executed LLVM IR instructions reflects the progress of program execution just as the number of executed machine-level instructions does.

Therefore, Nugget’s unit of work is suitable for applications that have stable control flow across different platforms. Applications such as gradient descent, where floating-point precision might affect the control flow, are not suitable for Nugget or for any sampling methodologies that rely on the application having stable control flows across microarchitectures [14, 26, 30, 32, 36, 40, 43, 53].

## 4.2 Preparation

In the preparation stage, we generate an optimized LLVM IR representation of the program that is used consistently across all stages of the pipeline (Figure 4.1). This is necessary because Nugget relies on a stable LLVM IR representation as its unit of work, as introduced in Section 4.1.

Optimization matters for performance, but some IR level optimizations, such as loop unrolling, can significantly alter control flow. To keep this control flow consistent across the pipeline, Nugget applies these optimizations when it generates the base LLVM IR representation in the preparation stage.

Loop vectorization, which can also restructure control flow, is itself an IR level optimization, so it too is fixed in the base IR at this stage. After the base IR is generated, we run only `opt`, to insert Nugget’s block labels and hooks, and `llc`, to lower the result to each target. Target-specific machine instructions—including SIMD instructions such as AVX on x86 or NEON on Arm—and per-architecture branch lowering are selected by the

backend during code generation. The LLVM IR that defines the unit of work is therefore fixed before lowering and is identical for every binary.

Build-time modes such as LTO and PGO can alter LLVM IR, but they are disabled by default. In our experiments, the LLVM IR structure remained stable across optimization levels, CPU-feature flags, and LLVM versions. We can therefore safely apply backend optimizations during the final binary compilation.

Only code included in this base LLVM IR representation is subject to analysis. As a result, any linked libraries must also be processed through Nugget during the preparation stage; otherwise, their contents are excluded from the analysis. This selective inclusion also lets us ignore irrelevant or non-contributory operations, such as spin loops in multi-threaded libraries [40], that do not meaningfully affect program progress.

## 4.3 Interval Analysis

Nugget uses LLVM passes to automatically annotate programs at the LLVM IR level, enabling interval analysis to be performed directly on real hardware. Interval analysis is necessary to discover and profile all intervals in a program’s execution as defined by the chosen unit of work. However, profiling an entire program can be time-consuming if done via simulation or emulation (Section 2.1).

While tools that rely on dynamic instrumentation and binary translation exist to perform such analyses, they often impose restrictions on the ISA and the host platform. By embedding interval analysis directly into the program’s execution, Nugget achieves near-native speed without introducing hardware-specific constraints: any machine capable of running the program can execute the interval analysis.

Because our unit of work is defined at the LLVM IR level and is independent of the final binary, all discovered intervals remain consistent regardless of the host machine’s ISA or microarchitecture. For example, we can analyze a program on an x86 system and then use the resulting intervals to drive sampling on an Arm system.

In this section, we describe (1) how Nugget uses LLVM passes to generate an interval analysis executable, and (2) how this executable is used to discover and profile intervals.

### 4.3.1 Interval Analysis LLVM Pass

The interval analysis LLVM pass, a modular compiler transformation that operates directly on LLVM IR, identifies all basic blocks (IRBBs) in the program’s IR. For each IRBB, it records the number of LLVM IR instructions and inserts a hook at the end of the block. Each hook is uniquely associated with its IRBB and encodes both the block’s ID and instruction count, allowing us to retrieve LLVM IR-level information during binary execution.

The hook is a lightweight function that counts occurrences and invokes other functions when specified conditions are met. This design allows us to flexibly adjust what data to collect by simply changing the functions the hook calls. After compilation with this LLVM interval analysis pass, the resulting binary executes the hook every time an IRBB is completed, enabling runtime data collection at the LLVM IRBB level. This hook-based approach also lets us automatically map LLVM IR basic blocks to their corresponding machine-level basic blocks. We discuss the overheads introduced by the hooks in Section 4.4.2.

### 4.3.2 Interval Analysis Executable

The interval analysis process has two main goals: (1) to partition the program’s execution into a sequence of intervals, and (2) to generate a characteristic signature (e.g., a basic block vector) for each interval. In Nugget, an interval is defined by a fixed number of executed LLVM IR instructions (Section 4.1). Thus, discovering intervals involves recording the current execution point whenever the LLVM IR instruction count is reached and marking the current execution point as the end of a newly identified interval.

Once intervals are identified, each is characterized by a signature. Whereas the interval’s definition specifies its location in the program’s execution, the signature captures its behavioral characteristics. For example, SimPoint defines intervals by instruction count but represents each interval’s signature as a basic block vector that tracks how often each basic block is executed.

A signature is closely tied to the chosen sample selection algorithm. Since this thesis does not propose a new selection algorithm, we do not introduce a new signature. How-

ever, we demonstrate how our hook-based approach can flexibly collect different types of signatures by constructing both LLVM IRBB vectors and count-stamp vectors. The LLVM IRBB vector tracks the frequency of entering each IRBB, while the count-stamp vector records the global IR instruction counter value when each IRBB is last entered. These will later be used to create markers (Section 4.4).

**Discovering Intervals** Discovering intervals is managed by the hooks. Each hook knows the ID and instruction count of its IRBB. Whenever a hook executes, it increments a global IR instruction counter by the IRBB’s instruction count and checks whether this counter has reached the defined interval size. If so, a new interval is marked by assigning it a sequential ID; otherwise, execution continues as normal.

All discovered intervals are binary independent because their definitions rely on the LLVM IR-level. Thus, regardless of the ISA or microarchitecture of the machine running the interval analysis executable, the discovered intervals remain consistent for the final compiled binary.

**Generate A Characteristic Signature** Our hook-based design makes it easy to specify what data to collect during execution by simply modifying the functions called by each hook. Because each hook has access to the LLVM IR basic block (IRBB) information, we can build an *IRBB vector* by incrementing a global vector based on the IRBB ID whenever a block is entered. This vector effectively records how many times each IRBB is executed within an interval, capturing the frequency profile of the program’s control flow. It is recorded and reset whenever a new interval is detected. Since the LLVM IR control flow remains consistent across all final binaries (Section 4.2), the IRBB vector is inherently binary independent.

In contrast, the *count-stamp vector* records the global IR instruction count each time an IRBB is entered, marking when during the program’s execution each IRBB was last encountered. This provides a timestamped trace of the control flow relative to instruction execution. The count-stamp vector is also recorded and reset at the end of each interval.

In Section 4.4.1, we show how the IRBB vector is used to create execution markers and how the count-stamp vector helps reduce the hook overhead. This flexible design allows

researchers to explore additional signatures for sampling methodologies with minimal development effort.

## 4.4 Nugget Creation

After selecting a set of samples with a sampling algorithm, Nugget generates *nuggets* based on these samples. We define a *nugget* as an executable that marks the selected sample with a start point and an end point, enabling execution on any platform capable of running the original program, including both real hardware and simulators. Thus, every *nugget* requires a cross-platform start marker and end marker to indicate the boundaries of the sample execution. In this section, we describe how Nugget creates these markers and mitigates their overhead.

### 4.4.1 Markers

In Nugget, a *marker* is a mechanism for identifying a specific point in program execution, defined according to our unit of work. Because progress is measured by the number of executed LLVM IR instructions (Section 4.1), our markers are directly tied to the IR instruction count.

We designate the last executed LLVM IR basic block (IRBB) that completes the target instruction count as a *marker*. This marker consists of two components: the IRBB itself and the number of times it must execute. The end marker for an interval is determined using the IRBB vector and count-stamp vector gathered during interval analysis (Section 4.3.2). Specifically, we use the count-stamp vector (Section 4.3.2) to identify the last executed IRBB of the interval, which becomes the basic block for the end marker, and the IRBB vector to determine how many times that block must execute to reach the interval’s endpoint. Because the start of one interval corresponds to the end of the previous interval, we reuse the end marker of the preceding interval as the start marker for the next.

Nugget inserts hooks at these marker IRBBs during the nugget creation LLVM pass. These hooks track IRBB executions and trigger specific actions, such as resetting statistics in a simulator or starting a timer on real hardware.

When using *nuggets* in simulation, we can also mark the start of a warmup interval. The warmup interval does not belong to the selected sample but is a period of execution used to warm up the microarchitectural state in simulation models, assuming functional simulation is used to fast-forward execution before the sample starts [20]. Any contiguous and antecedent intervals of the sample interval can be used as the warmup interval, so one can choose how many intervals to use for warmup.

#### 4.4.2 Overhead of Hooks

Hooks inserted at marker IRBBs introduce overhead due to frequent counting and condition checks, which can affect measurements if executed often. This overhead appears when running *nuggets* on real hardware, as the hooks add extra instructions.

For single-threaded programs, the overhead is typically small because hooks perform only simple operations, such as counting and checking thresholds. Moreover, hooks stop all counting and checks once the threshold is reached, and modern branch predictors generally handle these conditional checks efficiently.

For multi-threaded programs, the overhead can be more significant because counting and threshold checks require synchronization (e.g., atomic operations). Although hooks still cease checking after meeting the threshold, the overhead from the end marker’s hook is unavoidable.

To mitigate this, we propose two strategies:

- Identifying a lower-overhead marker.
- Using the program counter to track the marker instead of hook instrumentation when running in simulation.

**Lower Overhead Marker** To reduce overhead on real hardware, we can trade some precision by selecting an IRBB near the end of the interval that executes less frequently than the final IRBB. Because hook overhead is tied to execution frequency, this choice reduces cost. We do this by defining a search distance in IR instructions from the interval’s end, identifying candidate IRBBs within this distance using the count-stamp vector, and selecting the least frequently executed one using the IRBB vector. This effectively marks

a point slightly before the interval’s true end, reducing overhead at the cost of some precision.

**Program Counter Tracking in Simulation** While hook instrumentation is necessary on real hardware, when running *nuggets* in a simulator we can leverage simulator capabilities to eliminate hooks entirely. The LLVM nugget creation pass labels the marker IRBBs in the assembly, and we use the program counter addresses of these labels as inputs to the simulator to track markers. This completely removes marker overhead in simulation, ensuring measurements remain unaffected.

## 4.5 Sample Validation

Nugget addresses the limitations of using simulation to validate selected samples (Section 2.3) by generating *nuggets* that can run directly on both real hardware and simulation. Therefore, we can validate the selected samples on real hardware, and use them for simulation. This approach allows for fast, scalable validation of sampling accuracy without the high overhead of simulation, even for large applications.

Because nuggets are **binary-independent**, they can be validated across different systems with minimal constraints. Validating the same samples on multiple platforms enhances confidence in their *representativeness and robustness*. As shown later in Section 6.1, we found that having *consistent prediction error across platforms* is a stronger indicator of sample quality than achieving low error on a single platform.

To run a nugget on real hardware, the system must still execute the portion of the program leading up to the sample’s start marker (Section 4.4.1). For samples that are in later stages of a large program’s execution, this can introduce some validation overhead.

In the future, we plan to reduce this overhead further by exploring techniques such as creating software checkpoints using CRIU [7] or BLCR [23] to resume execution closer to the sample start point.

# Chapter 5

## Evaluation

In this section, we revisit the goals of Nugget and evaluate how well we achieve them.

1. Analyzing program execution quickly and without restrictions on architecture or host machine.
2. Automatically generating samples that are independent of the binary.
3. Quickly validating the selected samples for the targeted benchmarks and input size.

We begin by evaluating the first goal through measurements of the analysis overhead introduced by Nugget, comparing it to both uninstrumented execution and to functional simulation (Section 5.1). We then address the second and third goals by employing two simple sample selection methodologies, random sampling and LLVM IR basic block (IRBB) vector-based k-means sampling, and evaluating their effectiveness (Section 5.2).

| Attribute          | AMD Ryzen Threadripper<br>3960X | Ampere Altra     |
|--------------------|---------------------------------|------------------|
| ISA                | x86                             | Arm              |
| Cores              | 24×Out-of-Order                 | 160×Out-of-Order |
| Frequency (GHz)    | 3.8                             | 2.8              |
| L1 D/I (KiB)       | 32                              | 64               |
| L2 (KiB)           | 512                             | 1024             |
| L3/SLC (MiB)       | 128                             | 64               |
| Memory clock (MHz) | 2133                            | 3200             |
| Bandwidth (GiB/s)  | 136.5                           | 409.6            |

Table 5.1: Real Hardware Configuration. L1 D/I is the per-core data and instruction cache size (equal on both systems); L2 is per core and L3/SLC is shared. Both systems use DDR4 memory.

| Suite        | Workload                              | Parallelism             | Lang.           |
|--------------|---------------------------------------|-------------------------|-----------------|
| SPEC CPU2017 | Desktop and server workloads          | Single-threaded         | C, C++, Fortran |
| NPB          | HPC kernels                           | Multi-threaded (OpenMP) | Fortran, C      |
| LSMS         | First-principles materials simulation | Single-threaded         | C++             |

Table 5.2: Overview of Benchmark Suites Used in Evaluation

Table 5.1 presents the detailed configurations of the real hardware used in our evaluation. For the simulation models, we employ a basic generic functional model using the ATOMIC CPU in gem5, running a guest Ubuntu 24.04 environment, to perform interval analysis on both x86 and Arm architectures.

We evaluate Nugget using the SPEC CPU 2017 benchmark suite, the NAS Parallel

Benchmarks (NPB) suite [8], and the LSMS benchmark [48], covering a wide range of realistic workloads as summarized in Table 5.2.

We use the “-O2” optimization level for the compilation. The input chosen for each benchmark is too large to be simulated in feasible time (except for NPB class “A”).

For SPEC CPU2017, we use the reference input sizes, which require between 198 billion and 9,897 billion instructions. For NPB, we use class “C” and “D” inputs, with class “C” requiring between 24 billion and 2,714 billion instructions and class “D” ranging from 1,217 billion to 37,533 billion instructions. For LSMS, we use the Fe input, which requires between 28,484 billion and 198,885 billion instructions (one binary per ISA).

We used the SPEC CPU2017 Integer suite, except for “619.lbm\_s” and “644.nab\_s”, which are from the Floating-Point suite. We do not include CPU2017 benchmarks that failed to compile with the original SPEC CPU2017 GCC build system (6 benchmarks) or encountered Fortran linking issues (5 benchmarks). We did not encounter the same Fortran linking issues with any NPB Fortran benchmarks.

To accurately measure hardware performance, we employ *cpuset* to strictly bind workloads to specific cores and memory nodes. We also disable CPU frequency scaling and fix the CPU at a constant frequency. The LLVM passes are implemented on top of LLVM 19.0.0git.

## 5.1 Interval Analysis Efficiency (Goal 1)

In this section, we evaluate the efficiency of Nugget’s interval analysis. We define efficiency using two criteria: (1) *low slowdown*, quantified as the ratio between the time spent on interval analysis and the time required for executing the full workload; and (2) *flexibility*, meaning the ability to perform the analysis across different host machines.

During interval analysis, we identify all valid intervals and collect two key artifacts: the LLVM IR Basic Block (IRBB) vector and the count-stamp vector (Section 4.3.2). These outputs are later used to generate nuggets.

For all workloads, we execute the complete preparation and analysis pipeline shown in Figure 4.1. Additionally, we compile an unmodified baseline binary from the same LLVM

IR to isolate and measure the overhead introduced solely by the interval analysis.

We begin by comparing the overhead of interval analysis when performed via functional simulation versus Nugget. Next, we evaluate how overhead varies across a range of workload types and use different host machines to demonstrate the flexibility of Nugget. Finally, we study how overhead scales with the number of threads in the workload, noting that synchronization requirements during interval analysis can amplify slowdown at higher thread counts.


### 5.1.1 Comparison with Functional Simulation

For this experiment, we use the NPB suite with Class “A” inputs and 4 threads. We select Class “A” because larger input sizes, such as those used in Figure 5.2, would require months to analyze using functional simulation. We use multi-threaded workloads to demonstrate a worst-case overhead scenario, since interval analysis must track synchronization events between threads. Additionally, multi-threaded workloads reflect current trends in modern computing, making them a relevant evaluation target.


To evaluate the efficiency of Nugget, we compare it against LoopPoint [40], a recent state-of-the-art sampling methodology. LoopPoint is selected for comparison because it supports multi-threaded workload sampling and has integrated support within the gem5 simulator for full-system analysis [37]. This is particularly relevant, as all nuggets in our framework are generated on real hardware, a full-system environment.

As shown in Figure 5.1, we observe an average  $577.65\times$  reduction in interval analysis overhead when using Nugget instead of gem5’s functional simulation. The overhead on the order of tens of thousands in functional simulation highlights its impracticality for interval analysis in large applications. On average, Nugget incurs a slowdown of  $54\times$ , with a maximum of  $118\times$  for the “cg” benchmark. By comparison, LoopPoint’s average overhead is  $31,343\times$ , peaking at  $55,497\times$  for “cg-x86.”

Overall, multi-threaded workloads amplify interval analysis overhead because their baseline execution times are significantly shorter (due to parallelism) while the analysis itself incurs greater cost from synchronization requirements. The single-threaded workloads in Figure 5.2 exhibit much lower overheads, which we discuss further in Section 5.1.2.



(a) The LoopPoint interval analysis slowdown using gem5 functional simulation.



(b) The Nugget interval analysis slowdown.

Figure 5.1: Interval analysis overhead for Nugget and gem5 functional simulation. Both analyses use the NPB Class A input running with four threads. The slowdown is the wall-clock time of running the analysis divided by the wall-clock time of running the unmodified program. On average, Nugget performs interval analysis  $577.65\times$  faster than the gem5 functional simulation used by LoopPoint.

The “cg” benchmark has the shortest execution time in this experiment (0.28 seconds), which further amplifies the measured overhead for both methods. The “is” benchmark exhibits the lowest overhead under functional simulation but still incurs a  $95\times$  slowdown with Nugget, primarily due to the frequent executions of interval analysis hooks, as discussed in Section 5.2.3. The “mg” and “ft” benchmarks show relatively high overhead for both methods because their parallel regions increase synchronization costs during interval




Figure 5.2: Nugget interval analysis overhead across a mix of workload types: NPB with Class C input and four threads, SPEC CPU2017 with the **ref** input on a single thread, and LSMS with the **Fe** input. The slowdown is the wall-clock time of the analysis divided by that of the unmodified workload. Overhead stays low for the single-threaded SPEC CPU2017 and LSMS workloads but rises for the multi-threaded NPB workloads, which require synchronization across threads.

analysis, which in turn adds further simulation overhead. We examine the relationship between synchronization and overhead in more detail in Section 5.1.3.

In addition to its lower overhead, Nugget offers a key practical advantage: it requires only a single interval analysis run to generate nuggets that can be executed across multiple architectures. In contrast, LoopPoint must repeat interval analysis for every binary variant or target architecture, significantly increasing its total cost.

### 5.1.2 Mix of Workload Types

We also evaluate overhead across a mix of workload types. For single-threaded workloads, we use the SPEC CPU2017 suite with the **ref** input size and the LSMS benchmark with the **Fe** input size. For multi-threaded workloads, we use the NPB suite with Class “C” in this experiment.

Figure 5.2 shows the interval analysis slowdown across the evaluated workloads and demonstrates that **the overhead varies significantly depending on the level of synchronization**. The average observed overhead is  $3\times$  for SPEC CPU2017,  $2\times$  for LSMS, and  $34\times$  for NPB. The overhead remains low for SPEC CPU2017 and LSMS because they are single-threaded workloads, which simplifies interval analysis by eliminating the need for synchronization. In contrast, the high overhead for NPB arises from

the synchronization required during multi-threaded interval analysis, particularly due to aggregation across threads. As synchronization demands increase, so does the overhead, a relationship we explore further in Section 5.1.3.

**To demonstrate the flexibility of our approach across host machines, we repeat the interval analysis on two distinct systems listed in Table 5.1 and find consistent results across both platforms for single-threaded workloads, with one corner case.** For SPEC CPU2017, interval analysis results are identical across the two systems. For NPB, the multi-threaded nature of the benchmarks introduces variability between runs (even on the same host), a well-known challenge in multi-threaded sampling that is beyond the scope of this work [14, 40]. For LSMS, we observe minor differences in the execution frequency of some IRBBs, which we attribute to differences in floating-point precision across CPUs. In particular, the number of iterations in certain data-driven loops in LSMS is influenced by the underlying architecture’s precision. Such differences could affect nugget representativeness if a marker IRBB is sensitive to these variations. Nevertheless, as shown in Section 5.2.2, nuggets generated from interval analysis on one system still achieve low prediction error when validated on the other system for LSMS.

### 5.1.3 Scaling with Number of Threads

Finally, we study how interval analysis overhead scales with the number of threads. For this experiment, we use the NPB suite with Class “D” inputs and vary the thread count across  $\{1, 4, 8, 16\}$ . To ensure fair comparisons, all overheads are normalized to the execution time of the original workload running with a single thread.

As shown in Figure 5.3, most workloads exhibit increasing overhead with higher thread counts, consistent with our expectation that more threads lead to greater synchronization overhead. However, this growth is not uniform across all workloads because not all of them require the same level of synchronization during interval analysis, as not all hooks are executed within parallel regions. For “bt”, “cg”, “lu”, and “sp”, the consistently low slowdown as the thread count increases suggests that hooks are not frequently executed within parallel regions for these benchmarks. In contrast, “mg” and “ft” show a significant increase in overhead with more threads, indicating that these benchmarks require more




Figure 5.3: Nugget interval analysis overhead as the thread count varies across  $\{1, 4, 8, 16\}$ , using NPB with Class D inputs. Unlike Figure 5.1b and Figure 5.2, the slowdown here is the wall-clock time of the analysis divided by that of the unmodified *single-threaded* workload. Normalizing to a single-threaded baseline isolates the overhead added by synchronization, so the bars are not skewed by a well-parallelized baseline whose execution time shrinks as threads are added.

synchronization during interval analysis—likely due to the complexity of their parallel regions.

Even at higher thread counts (up to  $16\times$ ), Nugget maintains overheads that are orders-of-magnitude lower than those of functional simulation-based interval analysis (Figure 5.1).

#### 5.1.4 Summary

Our experiments show that Nugget achieves orders-of-magnitude lower interval analysis overhead compared to functional simulation, with an average reduction of over  $577.65\times$  relative to LoopPoint for multi-threaded workloads. It also produces consistent analysis results across different host machines, demonstrating its portability.

By enabling fast, binary-independent interval analysis, even for large workloads, Nugget successfully meets Goal 1.

In the following sections, we build upon this foundation to demonstrate how Nugget supports a variety of sampling methodologies (Goal 2) and enables rapid validation of selected samples on real hardware (Goal 3), further establishing its effectiveness for sampling methodology research.

## 5.2 Applying Sampling Methodologies to the Nugget Framework and Validating the Selected Samples (Goals 2 & 3)

The primary objective of Nugget is to support diverse sample selection methodologies while providing efficient interval analysis, binary-independent samples, and scalable validation capabilities. In this section, we leverage these features to execute the selected samples, *nuggets*, directly on real hardware, enabling us to assess whether different sampling strategies can accurately predict full workload performance *without relying on simulation*.

Because this thesis does not propose new sampling algorithms, we focus on two straightforward approaches: random sampling and LLVM IRBB vector-based k-means clustering. We show that Nugget seamlessly implements these methodologies and facilitates validating the selected samples on *real hardware*. We also show that different workloads may benefit from different sampling strategies; Nugget enables discovering and applying methods tailored to specific workload characteristics. Moreover, we highlight the importance of validating sampling methodologies across multiple machines, especially when the target system is not yet available. Finally, we discuss the overhead introduced by the hooks used to mark sample intervals within nuggets.

### 5.2.1 Sampling Methodologies

For the random sampling approach, we randomly select a set of intervals from the program execution to estimate its average or cumulative properties. This selection method is based on statistical sampling methods [52].

For the clustering approach, we use the k-means algorithm to cluster the IRBB vectors, select “k” samples to represent the program phases of the workload, and use the cluster weights to estimate the average or cumulative properties of the program execution. This approach is based on basic block vector (BBV) phase-identification techniques [16] (e.g., SimPoint) and on the fact that the IRBB vector is an abstraction of the actual BBV.

We refer to the random approach as “Random” and the clustering approach as “K-means”. We use 50 samples as the target number for both methodologies, which is




Figure 5.4: Runtime prediction error of the Random and K-means sampling methods across the two real-hardware machines in Table 5.1, for NPB and SPEC CPU2017. The prediction error is the difference between the runtime predicted from the nuggets and the runtime of the full workload. **All results are obtained by executing the nuggets directly on real hardware**, without simulation.

statistically large enough to approximate the population distribution [25]. We apply the Silhouette score with the K-means algorithm to determine the optimal number of clusters for each workload, enforcing the constraint  $\#clusters \leq 50$  (with  $k = 50$  as the upper bound). We assign equal weights to all intervals in the “Random” approach, whereas the “K-means” approach weighs each interval according to the size of its cluster.

### 5.2.2 Sample Validation

We use the analysis information collected in Section 5.1.2 to select samples with the “Random” and “K-means” approaches. We then generate nuggets for the selected samples and validate them as described in the Sample Creation and Validation step in Figure 4.1.

The nuggets are used to predict the runtime of each workload, with the runtime of the full execution serving as the baseline for calculating prediction error. We choose runtime as the metric because it is a common target of sampling methodologies [40], and IPC has been shown to be unsuitable for multi-threaded workloads [2].

Figure 5.4 presents the prediction errors of nuggets generated by the Random and K-means approaches for each workload on the real-hardware machines listed in Table 5.1. As shown, many workloads exhibit significant errors, often larger than the errors reported by prior methodologies such as SimPoint (3% for single-threaded workloads) and LoopPoint (2.87% for 8-threaded workloads).

Several factors contribute to these elevated errors. (1) The sampling methodologies are naively designed. For example, the Random approach relies heavily on luck and the specific composition of the interval pool. (2) Prior work rarely validates samples on real hardware [14, 22, 26, 30, 32, 36, 40, 42, 43, 52, 53]; instead, it mostly relies on user-mode-only simulation, which is typically less noisy due to more controllable thread interleaving compared to actual hardware [1, 3, 9, 17]. In Section 6.1, we further analyze the sources of prediction error using smaller workloads in simulation. (3) Prior work rarely validates large workloads as we used in our experiments (e.g., reference input sizes for CPU2017). Larger and more realistic applications exhibit increasingly complex behaviors [26, 31, 40, 43, 51, 54], which demand more representative sample sets, as discussed in Section 2.3. In such cases, large program phases can amplify prediction error. (4) Hardware-based nuggets incur additional overhead for the begin and end interval markers, particularly in multi-threaded contexts, which further inflates prediction error. We provide a detailed discussion of this overhead in Section 5.2.3.

As the figure shows, the sampling approach that yields the lowest prediction error varies by workload. For example, for “sp,” the “Random” approach with 50 samples achieves the lowest prediction error across all methods, and this holds on both machines, suggesting that this sample set is architecturally and microarchitecturally independent. In contrast, for “lbm,” the “K-means” approach consistently yields the lowest prediction error.

Although the absolute runtime prediction errors can be large for some workloads, a more practically relevant question is whether these errors are consistent across machines. To investigate this, we use the predicted runtimes from our nuggets to compute predicted speedup between the two hardware platforms and compare the resulting speedup prediction error against the individual runtime prediction errors. The average speedup prediction error is significantly smaller than the average runtime prediction error, even when the outlier workload “is” is included in the data set. For example, the K-means approach yields an average runtime prediction error of 70.13% but a speedup prediction error of only 7.58%, while the Random approach with 50 samples yields 126.90% for runtime but only 8.92% for speedup. This occurs because when the prediction error is consistent across machines—that is, the predicted runtime on both machines deviates from the true runtime in the same direction and by a similar proportion—the errors cancel out in the speedup ratio, yielding an accurate relative performance prediction. This observation indicates that *consistency* of prediction error across platforms matters more than minimizing the absolute prediction error on any single machine.

We focus on demonstrating Nugget’s ability to validate samples on distinct machines in this section. In Section 6.1, we will further analyze how to utilize the hardware validation data with a case study. We also find that maintaining consistency across machines is a stronger indicator of sample quality than minimizing prediction error on any single platform. This insight further underscores the importance of efficient, cross-platform validation in the design of robust sampling methodologies.

### 5.2.3 Hook Overhead

Nugget uses hook instructions to mark sample intervals, which can introduce runtime overhead during workload execution. In Figure 5.5, we present the number of times the marker hooks are executed for the randomly selected nuggets shown in Figure 5.4. The number of hook executions is normalized to the total number of IRBB executions during all nugget executions for each workload and input pair. As shown, “is” exhibits a significantly higher number of hook executions compared to other workloads, indicating that the prediction error observed for “is” in Figure 5.4 is influenced by the excessive




Figure 5.5: Marker hook executions for the randomly selected nuggets from Figure 5.4, normalized to the total IRBB executions across all nuggets for each workload and input pair. A higher fraction indicates that more of the nugget’s execution is spent on interval marking rather than useful work, inflating the measured prediction error.

overhead from hook executions. Because “is” frequently executes only a small number of IRBBs within a parallel region, it triggers an excessive number of hook invocations. “is” was also one of the workloads with the largest overheads for analysis as shown in Figure 5.1. This leads to substantial overhead and results in inaccurate prediction errors. Hook overhead can only slow execution, so it can cause false negatives (rejecting good samples) but is unlikely to cause false-positive conclusions.

For other workloads, the number of hook executions does not show a clear correlation with the prediction error observed in Figure 5.4. Cleanly separating the impact of hook overhead from other sources of prediction error is still challenging. The number of hook executions and the prediction errors observed suggest that a good marker should execute fewer than 10% of total IRBBs for single-threaded workloads and 1% for multi-threaded ones. Multi-threaded workloads are more sensitive to hook overhead due to the synchronization required for counting and threshold checks, as discussed in Section 4.4.2.

This behavior highlights a limitation of the current hook-based interval marking approach. Future work could explore alternative mechanisms to reduce this overhead or to decouple it from the validation results. However, addressing this issue is nontrivial, especially for multi-threaded workloads: overhead in parallel regions can differ significantly from that in serial regions, and parallelism may sometimes mask or amplify hook-related costs, complicating the isolation of their impact.

It is important to note that this overhead is unique to real hardware execution. In simulation, as discussed in Section 4.4.2, sample intervals can be marked directly by the simulator without introducing runtime overhead.

# Chapter 6

## Case Studies

The evaluation in the previous chapter showed that hook-based markers introduce runtime overhead on real hardware (Section 5.2.3), which makes it difficult to determine whether prediction error originates from poor sample selection or from instrumentation artifacts. Without separating these two sources of error, users cannot confidently assess the quality of their selected samples. By running nuggets in simulation, where the simulator tracks interval markers directly without runtime overhead, we can isolate sampling accuracy from marker overhead.

Additionally, hand-crafted microbenchmarks are labor-intensive to develop, exercise only a narrow slice of the design space, and require ongoing maintenance as architectures evolve. Compiler optimizations can also remove essential parts of microbenchmark code, producing misleading results. Nuggets, by contrast, are automatically derived from real workloads and capture the diverse instruction mix and control flow present in actual applications, making them a scalable alternative for evaluating simulation model fidelity.

Together, these two case studies demonstrate that Nugget is not only a sampling framework but also a practical tool for understanding both the quality of selected samples and the accuracy of simulation models. In this section, we present two case studies. First, we separate the overhead of nugget markers from the accuracy of the sampling selection process by running the nuggets both on hardware and in simulation. Then, we illustrate how nuggets can serve as lightweight microbenchmarks to assess simulation model accuracy, eliminating the need to develop dedicated microbenchmark suites.

## 6.1 Nuggets in Simulation and on Hardware

The goal of computer architecture analysis is often to compare the performance of two systems. Therefore, to investigate the usefulness of nuggets, we run both the entire workload and the nuggets and compare the speedup between different systems. We compare four different systems *in simulation* to eliminate the potential inaccuracies from the interval selection marker in hardware. We also compare four different hardware systems. With these eight systems, we can draw two conclusions. (1) We can use hardware to provide increased confidence in the selected intervals, but we cannot prove or bound the error due to the overheads discussed in Section 6.1.2. (2) We find that Random and basic-block vector approaches to interval selection are not microarchitecturally independent. In fact, we find that microarchitectural differences affect the accuracy of the selection methodology more than the ISA.

In this thesis we report the *error in predicted speedup*, not the error in absolute runtime. We found that in cases with modest absolute error the predicted speedup (or relative performance) is often much more accurate.

### 6.1.1 Experimental Setup

We use the selected nuggets to conduct experiments on both gem5, using full-system simulation, and real hardware. In addition to the systems listed in Table 5.1, we include two additional real hardware platforms: an Intel i9-9900X system and an Intel i7-9700 system. The simulation models used are listed in Table 6.1 and run a guest Ubuntu 24.04 kernel, matching the software environment of the real hardware.

Sample sets are selected using the “Random” and “K-means” methods described in Section 5.2.1. We set  $k = 20$  in this experiment, as the use of NPB Class A with a single-threaded configuration results in a relatively small pool of intervals available for sampling. Single-threaded workloads also help eliminate confounding factors such as thread interleaving, which could otherwise affect our measurement of prediction error.

In simulation, we use gem5 to track sample interval markers directly, thereby avoiding the runtime overhead introduced by hook-based marking mechanisms. To simulate nuggets efficiently, we first use KVM to fast-forward execution to the warmup marker

| Attribute          | Config 1    | Config 2    | Config 3   | Config 4   |
|--------------------|-------------|-------------|------------|------------|
| ISA                | x86         | Arm         | x86        | Arm        |
| Core               | 4×Large OoO | 8×Small OoO | 4×Base OoO | 8×Base OoO |
| Frequency (GHz)    | 4           | 3           | 4          | 4          |
| L1 D/I (KiB)       | 64          | 64          | 64         | 64         |
| L2 (KiB)           | 1024        | 1024        | 1024       | 1024       |
| L3/SLC (MiB)       | –           | 32          | –          | –          |
| Memory clock (MHz) | 2400        | 2400        | 2400       | 2400       |
| Bandwidth (GiB/s)  | 38.4        | 76.8        | 38.4       | 76.8       |

Table 6.1: gem5 Simulation Model Configuration. The large, small, and base out-of-order cores are internal models that target different performance characteristics. L1 D/I is the per-core data and instruction cache size (always equal here); L2 is per core and L3/SLC is shared. All configurations use DDR4 memory.

and then create a gem5 checkpoint [21]. This allows us to resume nugget execution from that point using different microarchitectural configurations, subject to certain constraints (e.g., the number of cores must remain fixed). For each nugget, we restore from the checkpoint and begin measurement only after reaching the start marker, allowing the microarchitectural state to warm up naturally.

## 6.1.2 Results

**Hardware-Based Validation Offers Scalability and Early Signal** Validating samples on real hardware is significantly faster and more scalable than using simulation. Several simulation runs (e.g., “bt”, “lu”, and “sp”), even with small inputs such as Class A, remain incomplete after two weeks, making simulation impractical for extensive validation. By contrast, hardware validation enables us to quickly assess the variability in speedup prediction across platforms.

Figure 6.1 shows the distribution of absolute speedup prediction errors for selected benchmarks. Each point is the absolute value of the error of one pair of configurations’




Figure 6.1: Distribution of absolute speedup prediction errors using K-means samples across simulation (blue) and hardware (orange) comparisons for selected benchmarks. “0.1” means 10% difference in predicted speedup.

(in simulation) or pair of hardware systems’ speedup compared to the true speedup of those two systems determined by running the entire workload. The hardware results allow us to quickly identify benchmarks with high variability in prediction error, such as “cg”, “ep”, and “is”, which exhibit wide error ranges across different platforms. For these workloads with large variations or large errors on hardware, we find the intervals are *not representative of the true speedup*. However, we cannot use the error on hardware to bound the error in simulation (e.g., the error for “ft” in simulation is larger than on hardware). The main reason we cannot use the hardware error to bound the simulation error is that the error is strongly affected by the underlying microarchitecture. In the next section, we analyze the sources of prediction error in more detail using simulation data and controlled case studies.

**ISA vs. Microarchitecture: Case Study** To isolate the sources of prediction error, we compare three types of platform pairs: (1) systems with different ISAs but small microarchitecture differences, (2) systems with significantly different microarchitectural configurations, but the same ISA, and (3) systems that differ in both ISA and microarchitecture.

Figure 6.2 shows the speedup prediction error when the main difference between platforms is the ISA (Config 3 vs. Config 4). In this experiment, the ISA has a noticeable impact on the binaries: for instance, the Arm binaries have vectorization enabled, whereas




Figure 6.2: Speedup: Random vs. K-means samples when ISA varies and microarchitecture is fixed, Config 3 vs. Config 4.




Figure 6.3: Speedup between Random and K-means samples when microarchitecture varies and ISA is fixed: Config 4 over Config 2 for Arm, and Config 3 over Config 1 for x86.

the x86 binaries do not. The average absolute speedup prediction error is 5% for Random-selected nuggets and 2.0% for K-means-selected nuggets. Among all benchmarks, “ft” exhibits the highest K-means error (6.5%), likely due to an unrepresentative sample set. A linear regression of nugget runtimes across the two ISAs explains 94.5% of the variance, indicating that the selected samples capture similar execution phases across architectures. The K-means speedup error for “ft” is consistent with its runtime prediction errors on Config 3 (4.9%) and Config 4 (−1.5%), suggesting the overestimation stems from sample imbalance. This implies that certain performance-critical phases were not captured by the selected nuggets. The Random-selected samples show a similar pattern for both “ft” and “is”.

Figure 6.3 shows the speedup prediction error when the microarchitecture varies: Config 4 over Config 2 for Arm and Config 3 over Config 1 for x86. The average absolute prediction error increases to 8.1% for Random-selected nuggets and 5.9% for K-means-




Figure 6.4: Speedup comparison between Random and K-means samples when both ISA and microarchitecture vary across simulation models. “K” indicates K-means samples, and “R” indicates Random samples.

selected nuggets. The comparison between Config 4 and Config 2 for Arm results in higher error in predicted speedup using the selected nuggets than between Config 3 and Config 1 for x86, reflecting greater microarchitectural differences in the Arm systems despite running the same benchmarks and sample sets.

Figure 6.4 presents the speedup prediction error when both ISA and microarchitecture vary across simulation models. The average absolute error rises to 10.8% for Random-selected nuggets and 7.9% for K-means-selected nuggets. The lowest errors occur in comparisons between Config 1 and Config 4 (7.7% for Random, 5% for K-means), where microarchitectural differences are relatively minor. In contrast, we observe larger errors in comparisons involving Config 2, which diverges significantly from other configurations in both ISA and microarchitecture. These findings highlight the sensitivity of prediction accuracy to microarchitectural differences more than to ISA variation.

**In summary**, while this thesis does not propose a new method for selecting representative intervals, we show the importance of comparing different interval selection techniques across a wide variety of architectures and microarchitectures. Nugget enables future work to have apples-to-apples comparisons across architectures and fast comparisons on hardware, opening this new line of research.

## 6.2 Model Accuracy Evaluation

In this section, we present how we used nuggets to identify sources of errors in gem5’s models. Identifying sources of errors in the simulation model is key to establishing baseline performance for evaluating architectural ideas.

Microbenchmarks are a common tool for fine-tuning architectural simulators. To build a useful model of the Ampere Altra processor, we developed a collection of models including CPU, cache, and memory models using publicly available information on the hardware [6]. We used a set of microbenchmarks and fine-tuned the configuration of the models to lower the error in simulated cycles [47]. While we were able to reduce cycle count error to below 5% for many of the microbenchmarks, the models struggled to accurately simulate real applications such as NAS Parallel Benchmarks—with errors exceeding 600% [8].

Using microbenchmarks has two main limitations. (1) In many cases the compiler optimizations can remove essential parts of the code, resulting in false positives. (2) Even when preserved after compiler optimizations, these microbenchmarks often exercise a narrow slice of the design space, missing the complex interactions between microarchitectural components present in real programs.

While it is possible to develop more sophisticated microbenchmarks, this approach is labor-intensive and does not scale well as architectures develop. Nuggets are an organic alternative to microbenchmarks; they capture the diverse instruction mix and control flow found in real workloads and they are small enough to provide a focused window for manual inspection by users.

To improve our model’s accuracy, we collected hardware performance counters for 234 nuggets and compared the hardware counters to gem5 statistics for each nugget to identify the nuggets with the largest error. We found that the data cache statistics explained most of the runtime error, and for more than 97% of the nuggets, the error in the number of accesses to the data cache could be explained by differences in the number of memory instructions. We isolated the problem by inspecting execution traces from a small set of the nuggets. We observed that the nugget with the largest error was dominated by




Figure 6.5: Comparison of the number of memory instructions for two different microcodings in gem5 vs. real hardware.

Arm’s paired memory instructions (e.g., ldp, stp), while fewer were present in samples with small error.

Figure 6.5 shows a comparison of the number of memory instructions in our selected nuggets for three different microcodings. The orange bars show the number of memory instructions with the default microcoding of gem5; the blue bars show the number of memory instructions if the paired memory instructions were microcoded as one microop; while the green bars show the number of memory instructions measured on native hardware. The blue bars closely track the green bars, suggesting that the microcoding of paired memory instructions in gem5 differs from that of the native hardware.

# Chapter 7

## Conclusion

This thesis addressed three limitations of existing targeted sampling techniques: the high cost of finding samples via simulation, the coupling of samples to a specific binary, and the lack of a practical path to validation without full-program simulation. We took a two-part approach: first implementing and evaluating the LoopPoint methodology in gem5’s full-system mode to establish a baseline, then presenting Nugget, a portable framework for agile targeted sampling development. Nugget uses LLVM IR to perform binary-independent interval analysis on real hardware, generates cross-platform executable samples—nuggets—that can be validated natively, and enables researchers to iterate on sampling methodologies at scale without relying on simulation.

### 7.1 LoopPoint Full-System Evaluation

We implemented the LoopPoint methodology in gem5’s full-system mode and evaluated it with the NPB benchmark suite using class A inputs and eight threads on both x86 and Arm configurations (Section 3). This evaluation confirmed that LoopPoint can be extended to full-system simulation, but it also highlighted three limitations that motivated the Nugget framework. First, interval analysis via functional simulation is time-consuming, requiring tens of thousands of times slowdown relative to native execution, which makes analyzing large workloads infeasible. Second, because LoopPoint constructs basic block vectors from machine-level instructions and defines region boundaries with ISA-specific (PC, count) pairs, each binary produces a different set of representative sam-

ples, leading to inconsistent prediction errors across ISAs. Third, validating the accuracy of the selected samples requires completing a full baseline simulation—the very bottleneck that sampling aims to avoid.

## 7.2 Nugget Framework Contributions

Nugget’s LLVM-based interval analysis achieves an average  $577\times$  speedup over functional simulation for multi-threaded workloads, with only  $3\times$  overhead for single-threaded SPEC CPU2017 benchmarks (Section 5.1). This efficiency enables analysis of reference-sized workloads that would otherwise take months or years to profile via simulation.

Because nuggets are defined at the LLVM IR level, the same set of representative samples can be compiled and validated across different ISAs and machines without repeating the analysis (Section 5.2). Our evaluation showed that cross-platform consistency in prediction error is a stronger indicator of sample quality than achieving low error on any single platform, underscoring the value of binary-independent, cross-platform validation.

In a case study, we found that microarchitectural differences affect prediction accuracy more significantly than ISA differences, underscoring the need to validate sampling methodologies across diverse microarchitectures. Beyond sampling validation, nuggets serve as lightweight microbenchmarks for assessing simulation model accuracy (Section 6.2). In a separate case study, we used nuggets from the NPB suite to identify a microcoding discrepancy in gem5’s handling of Arm paired memory instructions—a modeling error that traditional hand-crafted microbenchmarks had failed to expose.

## 7.3 Limitations

The hook-based interval markers used for real-hardware validation introduce runtime overhead, particularly for multi-threaded workloads with frequent, short parallel regions (Section 5.2.3). For example, the “is” benchmark triggers an excessive number of hook invocations because it executes only a small number of IRBBs within each parallel region, inflating prediction error. While hook overhead can only slow execution and thus cannot produce false-positive conclusions about sample quality, it remains a confounding factor that is difficult to isolate cleanly from other sources of error.

The Random and K-means sampling methodologies evaluated in this thesis are deliberately simple and show elevated prediction errors compared to prior approaches such as SimPoint and LoopPoint (Section 5.2). These elevated errors reflect the naivety of the selection algorithms rather than a limitation of the framework itself; Nugget is designed to support the development of more sophisticated methodologies.

Validating samples that are located in later stages of a large program’s execution still requires executing the program prefix up to the sample’s start marker (Section 4.5), which can be time-consuming for very large workloads.

## 7.4 Future Work

To reduce validation overhead for samples located late in a program’s execution, we plan to explore software checkpointing techniques such as CRIU [7] or BLCR [23]. These tools can create a checkpoint of the process state at or near a sample’s start marker, allowing subsequent validation runs to resume directly from that point rather than replaying the entire prefix.

We can also explore using nuggets to generate realistic modeling microbenchmarks from state-of-the-art benchmark suites, giving the community more realistic options and again leveraging Nugget’s portability.

Further, we can use Nugget to locate and analyze specific regions of a program, because Nugget can, by design, pinpoint a specific execution point within a program across platforms—something that, as discussed in Section 2.3, is otherwise difficult to do.

Finally, Nugget’s fast analysis and validation cycle makes it practical to develop and iterate on more sophisticated sample-selection algorithms tailored to specific workload characteristics. For example, we can define a set of IR-level features—such as LLVM IR data-instruction counts per IR basic block (IRBB), loop hints from IR instructions, and user-input metadata—and test each as a signature for selecting representative samples, then choose the set that performs most consistently across all experimental platforms as our final samples. Alternatively, rather than using the IR instruction as our unit of work, we could use the number of IR data instructions executed as the unit of work—a

direction that was previously infeasible to explore because of the difficulty of the sampling-methodology development flow, but that Nugget now makes practical. Future work can leverage the framework to explore phase-aware, workload-adaptive, or hybrid selection strategies and validate them at scale on real hardware.

## REFERENCES

- [1] A. Akram and L. Sawalha, “A survey of computer architecture simulation techniques and tools,” *IEEE Access*, vol. 7, pp. 78 120–78 145, Jun. 2019. doi:10.1109/ACCESS.2019.2917698
- [2] A. R. Alameldeen and D. A. Wood, “IPC considered harmful for multiprocessor workloads,” *IEEE Micro*, vol. 26, no. 4, pp. 8–17, Jul. 2006. doi:10.1109/MM.2006.73
- [3] A. R. Alameldeen and D. A. Wood, “Variability in architectural simulations of multi-threaded workloads,” in *Proceedings of the 9th International Symposium on High-Performance Computer Architecture*, ser. HPCA ’03. USA: IEEE Computer Society, Feb. 2003, pp. 7–18. doi:10.1109/HPCA.2003.1183520
- [4] F. Alexander, A. Almgren, J. Bell, A. Bhattacharjee, J. Chen, P. Colella, D. Daniel, J. DeSlippe, L. Diachin, E. Draeger, A. Dubey, T. Dunning, T. Evans, I. Foster, M. Francois, T. Germann, M. Gordon, S. Habib, M. Halappanavar, S. Hamilton, W. Hart, Z. (Henry) Huang, A. Hungerford, D. Kasen, P. R. C. Kent, T. Koley, D. B. Kothe, A. Kronfeld, Y. Luo, P. Mackenzie, D. McCallen, B. Messer, S. Mniszewski, C. Oehmen, A. Perazzo, D. Perez, D. Richards, W. J. Rider, R. Rieben, K. Roche, A. Siegel, M. Sprague, C. Steefel, R. Stevens, M. Syamlal, M. Taylor, J. Turner, J.-L. Vay, A. F. Voter, T. L. Windus, and K. Yelick, “Exascale applications: skin in the game,” *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 378, no. 2166, p. 20190056, Mar. 2020. doi:10.1098/rsta.2019.0056
- [5] M. Alian, “Sustainable gem5 simulations,” <https://www.gem5.org/assets/files/workshop-isca-2023/slides/sustainable-gem5-simulations.pdf>, presented at the gem5 Workshop, ISCA 2023, Accessed: 2023-10-01.
- [6] Ampere Computing. Ampere Altra family product brief. [Online]. Available: <https://amperecomputing.com/briefs/ampere-altra-family-product-brief>
- [7] F. Andrijauskas, I. Sfiligoi, D. Davila, A. Arora, J. Guiang, B. Bockelman, G. Thain, and F. Wurthwein, “CRIU – checkpoint restore in userspace for computational simulations and scientific applications,” Feb. 2024. [Online]. Available: <https://arxiv.org/abs/2402.05244>
- [8] D. H. Bailey, E. Barszcz, J. T. Barton, D. S. Browning, R. L. Carter, L. Dagum, R. A. Fatoohi, P. O. Frederickson, T. A. Lasinski, R. S. Schreiber, H. D. Simon, V. Venkatakrisnan, and S. K. Weeratunga, “The NAS parallel benchmarks,” *Int. J. High Perform. Comput. Appl.*, vol. 5, no. 3, pp. 63–73, Sep. 1991. doi:10.1177/109434209100500306
- [9] N. Binkert, B. Beckmann, G. Black, S. K. Reinhardt, A. Saidi, A. Basu, J. Hestness, D. R. Hower, T. Krishna, S. Sardashti, R. Sen, K. Sewell, M. Shoaib, N. Vaish,

- M. D. Hill, and D. A. Wood, “The gem5 simulator,” *SIGARCH Comput. Archit. News*, vol. 39, no. 2, pp. 1–7, Aug. 2011. doi:10.1145/2024716.2024718
- [10] D. L. Bruening and S. Amarasinghe, “Efficient, transparent, and comprehensive run-time code manipulation,” Ph.D. dissertation, Massachusetts Institute of Technology, USA, 2004, aAI0807735.
- [11] J. Bucek, K.-D. Lange, and J. v. Kistowski, “SPEC CPU2017: next-generation compute benchmark,” in *Companion of the 2018 ACM/SPEC International Conference on Performance Engineering*, ser. ICPE ’18. New York, NY, USA: Association for Computing Machinery, Apr. 2018, pp. 41–42. doi:10.1145/3185768.3185771
- [12] S. Caculo, M. Madhav, and J. Baxter, “Memory access vectors: improving sampling fidelity for CPU performance simulations,” Jun. 2025. [Online]. Available: <https://arxiv.org/abs/2506.02344>
- [13] T. E. Carlson, W. Heirman, and L. Eeckhout, “Sniper: exploring the level of abstraction for scalable and accurate parallel multi-core simulation,” in *SC ’11: Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis*, Nov. 2011, pp. 1–12. doi:10.1145/2063384.2063454
- [14] T. E. Carlson, W. Heirman, K. Van Craeynest, and L. Eeckhout, “BarrierPoint: sampled simulation of multi-threaded applications,” in *2014 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*, Mar. 2014, pp. 2–12. doi:10.1109/ISPASS.2014.6844456
- [15] E. Colp, J. N. Amaral, C. Benedicto, R. E. Rodrigues, and E. Borin, “Report: the 523.xalancbmk.r benchmark (June 2016 development kit),” Department of Computing Science, University of Alberta, Technical Report, Jan. 2018, published January 18, 2018. Part of “The Alberta Workloads for the SPEC CPU2017 Benchmark Suite.” [Online]. Available: [https://webdocs.cs.ualberta.ca/~amaral/AlbertaWorkloadsForSPEC2017/reports/xalancbmk\\_report.html](https://webdocs.cs.ualberta.ca/~amaral/AlbertaWorkloadsForSPEC2017/reports/xalancbmk_report.html)
- [16] K. Criswell and T. Adegbiya, “A survey of phase classification techniques for characterizing variable application behavior,” *IEEE Transactions on Parallel and Distributed Systems*, vol. 31, no. 1, pp. 224–236, Jan. 2020. doi:10.1109/TPDS.2019.2929781
- [17] J. Devietti, B. Lucia, L. Ceze, and M. Oskin, “DMP: deterministic shared memory multiprocessing,” in *Proceedings of the 14th International Conference on Architectural Support for Programming Languages and Operating Systems*, ser. ASPLOS XIV. New York, NY, USA: Association for Computing Machinery, Mar. 2009, pp. 85–96. doi:10.1145/1508244.1508255
- [18] J. Dongarra and P. Luszczek, *Livermore Loops*. Boston, MA: Springer US, 2011, pp. 1041–1043. doi:10.1007/978-0-387-09766-4\_161

- [19] L. Eeckhout, J. Sampson, and B. Calder, “Exploiting program microarchitecture independent characteristics and phase behavior for reduced benchmark suite simulation,” in *IEEE International. 2005 Proceedings of the IEEE Workload Characterization Symposium, 2005.*, Oct. 2005, pp. 2–12. doi:10.1109/IISWC.2005.1525996
- [20] L. Eeckhout, *Computer architecture performance evaluation methods*, 1st ed. Morgan & Claypool Publishers, 2010. doi:10.2200/S00273ED1V01Y201006CAC010
- [21] gem5 Project. (2025) Checkpoints — gem5 documentation. [Online]. Available: [https://www.gem5.org/documentation/general\\_docs/checkpoints/](https://www.gem5.org/documentation/general_docs/checkpoints/)
- [22] T. Grass, A. Rico, M. Casas, M. Moreto, and E. Ayguadé, “TaskPoint: sampled simulation of task-based programs,” in *2016 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*, Apr. 2016, pp. 296–306. doi:10.1109/ISPASS.2016.7482104
- [23] P. Hargrove and J. Duell, “Berkeley Lab checkpoint/restart (BLCR) for Linux clusters,” *Journal of Physics: Conference Series*, vol. 46, p. 494, Sep. 2006. doi:10.1088/1742-6596/46/1/067
- [24] J. W. Haskins and K. Skadron, “Memory reference reuse latency: accelerated warmup for sampled microarchitecture simulation,” in *2003 IEEE International Symposium on Performance Analysis of Systems and Software. ISPASS 2003.*, Mar. 2003, pp. 195–203. doi:10.1109/ISPASS.2003.1190246
- [25] R. V. Hogg, E. A. Tanis, and D. L. Zimmerman, *Probability and statistical inference*, 10th ed. Boston: Pearson, 2020.
- [26] T. Huffmire and T. Sherwood, “Wavelet-based phase classification,” in *2006 International Conference on Parallel Architectures and Compilation Techniques (PACT)*, Sep. 2006, pp. 95–104. doi:10.1145/1152154.1152172
- [27] S. Karandikar, H. Mao, D. Kim, D. Biancolin, A. Amid, D. Lee, N. Pemberton, E. Amaro, C. Schmidt, A. Chopra, Q. Huang, K. Kovacs, B. Nikolic, R. Katz, J. Bachrach, and K. Asanović, “FireSim: FPGA-accelerated cycle-exact scale-out system simulation in the public cloud,” in *Proceedings of the 45th Annual International Symposium on Computer Architecture*, ser. ISCA ’18. Piscataway, NJ, USA: IEEE Press, Jun. 2018, pp. 29–42. doi:10.1109/ISCA.2018.00014
- [28] S. V. Kodakara, J. Kim, D. J. Lilja, W.-C. Hsu, and P.-C. Yew, “Analysis of statistical sampling in microarchitecture simulation: metric, methodology and program characterization,” in *2007 IEEE 10th International Symposium on Workload Characterization*, Sep. 2007, pp. 139–148. doi:10.1109/IISWC.2007.4362190
- [29] C. Lattner and V. Adve, “LLVM: a compilation framework for lifelong program analysis & transformation,” in *Proceedings of the International Symposium on Code Generation and Optimization: Feedback-Directed and Runtime Optimization*, ser. CGO ’04. USA: IEEE Computer Society, Mar. 2004, p. 75. doi:10.1109/CGO.2004.1281665

- [30] P. Lavin, J. Young, and R. Vuduc, “Multifidelity memory system simulation in SST,” in *Proceedings of the International Symposium on Memory Systems*, ser. MEMSYS ’23. New York, NY, USA: Association for Computing Machinery, Oct. 2024. doi:10.1145/3631882.3631890
- [31] A. Limaye and T. Adegbiya, “A workload characterization of the SPEC CPU2017 benchmark suite,” in *2018 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*, Apr. 2018, pp. 149–158. doi:10.1109/ISPASS.2018.00028
- [32] C. Liu, A. Sabu, A. Chaudhari, Q. Kang, and T. E. Carlson, “Pac-Sim: simulation of multi-threaded workloads using intelligent, live sampling,” Oct. 2023. [Online]. Available: <https://arxiv.org/abs/2310.17089>
- [33] J. Lowe-Power, A. M. Ahmad, A. Akram, M. Alian, R. Amslinger, M. Andreozzi, A. Armejach, N. Asmussen, S. Bharadwaj, G. Black, G. Bloom, B. R. Bruce, D. R. Carvalho, J. Castrillón, L. Chen, N. Derumigny, S. Diestelhorst, W. Elsasser, M. Fariborz, A. F. Farahani, P. Fotouhi, R. Gambord, J. Gandhi, D. Gope, T. Grass, B. Hanindhito, A. Hansson, S. Haria, A. Harris, T. Hayes, A. Herrera, M. Horsnell, S. A. R. Jafri, R. Jagtap, H. Jang, R. Jeyapaul, T. M. Jones, M. Jung, S. Kannoth, H. Khaleghzadeh, Y. Kodama, T. Krishna, T. Marinelli, C. Menard, A. Mondelli, T. Mück, O. Naji, K. Nathella, H. Nguyen, N. Nikoleris, L. E. Olson, M. S. Orr, B. Pham, P. Prieto, T. Reddy, A. Roelke, M. Samani, A. Sandberg, J. Setoain, B. Shingarov, M. D. Sinclair, T. Ta, R. Thakur, G. Travaglini, M. Upton, N. Vaish, I. Vougioukas, Z. Wang, N. Wehn, C. Weis, D. A. Wood, H. Yoon, and É. F. Zulian, “The gem5 simulator: version 20.0+,” *CoRR*, vol. abs/2007.03152, Jul. 2020. [Online]. Available: <https://arxiv.org/abs/2007.03152>
- [34] P. Messina, “The Exascale Computing Project,” *Computing in Science and Engineering*, vol. 19, no. 3, pp. 63–67, May 2017. doi:10.1109/MCSE.2017.57
- [35] H. Patil, R. Cohn, M. Charney, R. Kapoor, A. Sun, and A. Karunanidhi, “Pin-pointing representative portions of large Intel® Itanium® programs with dynamic instrumentation,” in *Proceedings of the 37th Annual IEEE/ACM International Symposium on Microarchitecture*, ser. MICRO 37. USA: IEEE Computer Society, Dec. 2004, pp. 81–92. doi:10.1109/MICRO.2004.28
- [36] E. Perelman, J. Lau, H. Patil, A. Jaleel, G. Hamerly, and B. Calder, “Cross binary simulation points,” in *2007 IEEE International Symposium on Performance Analysis of Systems & Software*, Apr. 2007, pp. 179–189. doi:10.1109/ISPASS.2007.363748
- [37] Z. Qiu, “The LoopPoint methodology in the gem5 simulator,” 2023, presentation slides, gem5 Workshop at International Symposium on Computer Architecture (ISCA 2023). [Online]. Available: <https://www.gem5.org/assets/files/workshop-isca-2023/slides/looppoint-in-the-gem5-simulator.pdf>

- [38] Z. Qiu, M. Samani, and J. Lowe-Power, “Nugget: portable program snippets,” in *2026 IEEE International Symposium on High-Performance Computer Architecture (HPCA)*, Feb. 2026. doi:10.1109/HPCA68181.2026.11408606
- [39] R. Rahman, D. Owen, and J. You, “Tracking large-scale AI models,” 2024, accessed: 2024-10-18. [Online]. Available: <https://epochai.org/blog/tracking-large-scale-ai-models>
- [40] A. Sabu, H. Patil, W. Heirman, and T. E. Carlson, “LoopPoint: checkpoint-driven sampled simulation for multi-threaded applications,” in *2022 IEEE International Symposium on High-Performance Computer Architecture (HPCA)*, Apr. 2022, pp. 604–618. doi:10.1109/HPCA53966.2022.00051
- [41] A. Sabu, Z. Qiu, H. Patil, C. Liu, W. Heirman, J. Lowe-Power, and T. E. Carlson, “Accelerating the simulation of parallel workloads using loop-bounded checkpoints,” *ACM Trans. Archit. Code Optim.*, Feb. 2026, just Accepted. doi:10.1145/3799430. [Online]. Available: <https://doi.org/10.1145/3799430>
- [42] A. Sandberg, N. Nikoleris, T. E. Carlson, E. Hagersten, S. Kaxiras, and D. Black-Schaffer, “Full speed ahead: detailed architectural simulation at near-native speed,” in *2015 IEEE International Symposium on Workload Characterization*, Oct. 2015, pp. 183–192. doi:10.1109/IISWC.2015.29
- [43] T. Sherwood, E. Perelman, G. Hamerly, and B. Calder, “Automatically characterizing large scale program behavior,” in *Proceedings of the 10th International Conference on Architectural Support for Programming Languages and Operating Systems*, ser. ASPLOS X. New York, NY, USA: Association for Computing Machinery, Oct. 2002, pp. 45–57. doi:10.1145/605397.605403
- [44] Standard Performance Evaluation Corporation (SPEC), “SPEC CPU2000 benchmark suite,” <https://www.spec.org/cpu2000/>, 2000, [Online; accessed 21-Oct-2025].
- [45] Standard Performance Evaluation Corporation (SPEC), “SPEC CPU2006 benchmark suite,” <https://www.spec.org/cpu2006/>, 2006, [Online; accessed 21-Oct-2025].
- [46] Standard Performance Evaluation Corporation (SPEC), “SPEC CPUv8 benchmark suite,” <https://www.spec.org/cpuv8/>, 2025.
- [47] Vertical Research Group, “microbench: extremely simple microbenchmarks,” 2018, MIT License. Accessed 2025-10-21. [Online]. Available: <https://github.com/VerticalResearchGroup/microbench>
- [48] Y. Wang, G. M. Stocks, W. A. Shelton, D. M. C. Nicholson, W. M. Temmerman, and Z. Szotek, “Order-N multiple scattering approach to electronic structure calculations,” *Phys. Rev. Lett.*, vol. 75, pp. 2867–2870, Oct. 1995. doi:10.1103/PhysRevLett.75.2867

- [49] V. M. Weaver and S. A. McKee, “Using dynamic binary instrumentation to generate multi-platform SimPoints: methodology and accuracy,” in *High Performance Embedded Architectures and Compilers*, P. Stenström, M. Dubois, M. Katevenis, R. Gupta, and T. Ungerer, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, Jan. 2008, pp. 305–319. doi:10.1007/978-3-540-77560-7\_21
- [50] T. F. Wenisch, R. E. Wunderlich, M. Ferdman, A. Ailamaki, B. Falsafi, and J. C. Hoe, “SimFlex: statistical sampling of computer system simulation,” *IEEE Micro*, vol. 26, no. 4, pp. 18–31, 2006. doi:10.1109/MM.2006.79
- [51] Q. Wu, S. Flolid, S. Song, J. Deng, and L. K. John, “Invited paper for the hot workloads special session hot regions in SPEC CPU2017,” in *2018 IEEE International Symposium on Workload Characterization (IISWC)*, Oct. 2018, pp. 71–77. doi:10.1109/IISWC.2018.8573479
- [52] R. E. Wunderlich, T. F. Wenisch, B. Falsafi, and J. C. Hoe, “SMARTS: accelerating microarchitecture simulation via rigorous statistical sampling,” in *30th Annual International Symposium on Computer Architecture, 2003. Proceedings.*, Jun. 2003, pp. 84–95. doi:10.1109/ISCA.2003.1206991
- [53] J. Xue, T. Xiong, L. Chao, and R. Xue, “SimPoint+: more stable, accurate and efficient program analysis,” in *Euro-Par 2025: Parallel Processing*, W. E. Nagel, D. Goehring, and P. C. Diniz, Eds. Cham: Springer Nature Switzerland, 2026, pp. 3–17. doi:10.1007/978-3-031-99857-7\_1
- [54] J. J. Yi, S. V. Kodakara, R. Sendag, D. J. Lilja, and D. M. Hawkins, “Characterizing and comparing prevailing simulation techniques,” in *11th International Symposium on High-Performance Computer Architecture*, Feb. 2005, pp. 266–277. doi:10.1109/HPCA.2005.8