



HAL
open science

Performance Debugging through Microarchitectural Sensitivity and Causality Analysis

Alban Dutilleul, Hugo Pompougnac, Nicolas Derumigny, Gabriel Rodríguez, Valentin Trophime, Christophe Guillon, Fabrice Rastello

► **To cite this version:**

Alban Dutilleul, Hugo Pompougnac, Nicolas Derumigny, Gabriel Rodríguez, Valentin Trophime, et al.. Performance Debugging through Microarchitectural Sensitivity and Causality Analysis. INRIA. 2024. hal-04851704

HAL Id: hal-04851704

<https://inria.hal.science/hal-04851704v1>

Submitted on 20 Dec 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Performance Debugging through Microarchitectural Sensitivity and Causality Analysis

Alban Dutilleul
ENS Rennes
France
alban.dutilleul@ens-rennes.fr

Hugo Pompougnac
Univ. Grenoble Alpes, Inria, CNRS,
Grenoble INP, LIG, 38000 Grenoble,
France
hugo.pompougnac@inria.fr

Nicolas Derumigny
Télécom SudParis, France
derumigny.nicolas@telecom-
sudparis.eu

Gabriel Rodríguez
Universidade da Coruña
Spain
gabriel.rodriguez@udc.es

Valentin Trophime
Univ. Grenoble Alpes, Inria, CNRS,
Grenoble INP, LIG, 38000 Grenoble,
France
valentin.trophime@inria.fr

Christophe Guillon
Univ. Grenoble Alpes, Inria, CNRS,
Grenoble INP, LIG, 38000 Grenoble,
France
christophe.guillon@inria.fr

Fabrice Rastello
Univ. Grenoble Alpes, Inria, CNRS,
Grenoble INP, LIG, 38000 Grenoble,
France
fabrice.rastello@inria.fr

Abstract

Modern Out-of-Order (*OoO*) CPUs are complex systems with many components interleaved in non-trivial ways. Pinpointing performance bottlenecks and understanding the underlying causes of program performance issues are critical tasks to fully exploit the performance offered by hardware resources.

Current performance debugging approaches rely either on measuring resource utilization, in order to estimate which parts of a CPU induce performance limitations, or on code-based analysis deriving bottleneck information from capacity/throughput models. These approaches are limited by instrumental and methodological precision, present portability constraints across different microarchitectures, and often offer factual information about resource constraints, but not causal hints about how to solve them.

This paper presents a novel performance debugging and analysis tool that implements a resource-centric CPU model driven by dynamic binary instrumentation that is capable of detecting complex bottlenecks caused by an interplay of hardware and software factors. Bottlenecks are detected through **sensitivity-based analysis**, a sort of model parameterization that uses differential analysis to reveal constrained resources. It also implements a new technique we developed that we call **causality analysis**, that propagates constraints to pinpoint how each instruction contribute to the overall execution time.

To evaluate our analysis tool, we considered the set of high-performance computing kernels obtained by applying a wide range of transformations from the Polybench benchmark suite [38] and measured the precision on a few Intel CPU and Arm micro-architectures. We also took one of the

benchmarks (correlation) as an illustrative example to illustrate how our tool's bottleneck analysis can be used to optimize a code.

1 Introduction

In the last decades, high-performance computing has become one of the fundamental pillars of scientific innovation. The end of Dennard scaling in the early 21st century fundamentally shifted the architectural approach to performance improvements, which had been based on increasing clock frequencies with constant MOSFET power densities, and into a “dark silicon” [23] era in which systems are composed of a heterogeneous collection of computational devices with very different performance-power operating points.

While the basics of modern processor designs closely follow the ideas by Robert Tomasulo, debuted in the IBM 360/91 in 1966, CPU cores are nowadays more complex machines. In 2006, Intel launched its “Core” architecture, featuring a 65nm process node, a 4-uops issue width, a 96-entry ROB, and a 32-entry unified schedulers that feeds 6 superscalar ports. In 2023, the Meteor Lake architecture features both performance- and efficient-oriented processors, as well as an integrated GPU. The process node varies for different parts of the chip. The 7nm Redwood Cove P-processor has a 6-uops issue width, with three different instruction-decoding paths, a 512-entry ROB and 3 different scheduling units with a combined 215 entries that feed a 12-port superscalar architecture. Add to the mix Simultaneous Multi-Threading (SMT) support, and the result is an extraordinarily complex architecture, on which *performance debugging* becomes a daunting task.

A popular way to analyze the internal behavior of a CPU for a specific computational kernel is to rely on the performance-monitoring counters (PMCs) made available by the performance-monitoring unit (PMU) of modern CPUs. These model-specific registers provide a native way to measure the frequency of different low-level events in the CPU, offering a glimpse at the internal workings of this otherwise black box. The state-of-the-art *Top-down Microarchitecture Analysis* (TMA) method [46] defines a methodology in which low-level events are measured and used to compute aggregation formulas that pinpoint the areas of the processor that become performance bottlenecks during the execution of a given computational kernel. While TMA provides invaluable leads to guide optimization, it has limitations stemming both from precision, methodological, and portability concerns. *Binary code analysis* [2, 4, 8] is a complementary approach based on a performance model at the assembly level. Although it produces less precise estimations than PMC-based measurements, assembly models a wider range of metrics and greater flexibility in analysis, while not requiring any hardware support.

This work proposes two novel performance-debugging methods respectively called *causality* and *sensitivity analysis*. It is based on a coarse-level model-based simulation, implemented through instrumentation of a binary. This simulation is used to power: (1) a constraints propagation engine, that allows to pinpoint how each individual static instruction contributes to the overall execution time; (2) a differential analysis engine, that studies the performance effects of varying the capacities of different CPU resources, identifying both the CPU functional blocks that act as bottlenecks during the execution of a particular kernel, as well as the bottlenecked sections of code on which to focus the optimization efforts. The contributions of this work are as follows:

- A novel causality-based approach to performance debugging which provides information about which specific instruction contribute to the overall execution time.
- A novel sensitivity-based approach to performance debugging which provides information about the specific parts of the CPU which limit the performance of a given computational kernel.
- An end-to-end implementation of Gus, a sensitivity-based profiler for performance debugging.
- An extensive experimental validation over 1624 kernels on 6 different microarchitectures, demonstrating the accuracy of Gus compared to measurements.
- A case study illustrating the benefit of our bottleneck analyzer for performance debugging.

The paper is organized as follows. Section 2 gives a comprehensive description of existing performance analysis techniques and motivates our approach. Section 3 describes Gus. Section 4 reports our experimental results. Section 5 discussed our implementation choices. Section 6 concludes.

```

1 vmovaps (%rcx),%ymm1          1 vmovaps (%rcx),%ymm1
2 for                          2 vmovaps (%rax),%ymm2
3 vmovaps (%rax),%ymm2         3 for
4 vmadd231ps %ymm0,%ymm1,%ymm3 4 vmadd231ps %ymm0,%ymm1,%ymm3
5 vmadd231ps %ymm0,%ymm1,%ymm2 5 vmadd231ps %ymm0,%ymm1,%ymm2

```

(a) (b)

Figure 1. An example in pseudo-asm code of a kernel computing $ymm0 = ymm1 * ymm3 + ymm1 * ymm2$ iteratively. Integer operations and branches supporting pointer arithmetic and loop iteration have been removed for clarity. $ymm1$ and $ymm2$ are loaded from two non-overlapping memory arrays. $ymm3$ is constant. In version (b), the `vmovaps to $ymm2$` is hoisted out of the inner loop.

2 Performance debugging: state of the art and motivation

Consider the assembly kernel shown in Fig. 1(a), that yields a raw performance of 12.8 GFLOPS when executed on an Intel i9-12900k core with a theoretical peak performance of 153.6 GFLOPS¹. Figuring out where the *bottleneck* of this computation lies for the given microarchitecture is a remarkably complex task. Here, a bottleneck is defined as a microarchitectural resource or characteristic that becomes a fundamental cause of performance degradation due to the limitations it imposes to the flow of the computation. On the one hand, one could hypothesize that this kernel is memory-bound due to the `vmovaps` in line 3. However, depending on the access pattern and the prefetcher this load might be served from L1. If so, and depending on the arithmetic intensity of the kernel, the FP FMA units may become saturated and act as the computational bottleneck. This is the fundamental trade-off studied by the Roofline model [45], which predicts the expected performance of a kernel from its arithmetic intensity and the architectural parameters of the computer.

Roofline analysis models a simplified, ideal version of the processor, useful to assess whether a computation achieves peak performance. However, when that does not happen, it does not provide any information as to *why*. In a real processor, different parts can become a performance bottleneck, depending on the computation itself, its implementation, the microarchitectural details, and even the execution context (e.g., load by other tasks, operational temperature, etc.). For example, if the calculation of the addresses accessed by the `vmovaps` in line 3 is complex enough, it could saturate the integer ALUs and become a bottleneck itself.

Fig. 2 shows a simplified, high-level representation of a pipelined out-of-order CPU core. During its lifecycle, each μop passes through four distinct states: it is *issued* by the front-end, *dispatched* by the allocator, *executed* by the functional units during its stay in the Re-Order Buffer (ROB),

¹This processor features 3xFP-FMA AVX-2 units running at a base frequency of 3.2 GHz. These measurements refer to single-precision FLOPS.

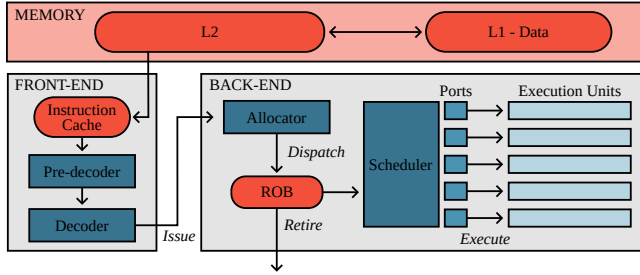


Figure 2. Simplified view of a pipelined OoO CPU core.

and eventually *retired* once its effects have been committed. During this process, any of the following parts of the CPU may incur performance penalties, and ultimately become a performance bottleneck:

Front-end (F-E) This is the part of the processor in charge of fetching and decoding instructions from memory. It includes the Branch Prediction Unit (BPU), which guides the control flow by setting the value of the Program Counter (PC), potentially speculating on branch outcomes. Instructions are either read from the Instruction Cache (IC) and decoded into micro-operations (μ ops), or fetched from a μ op-cache. If the F-E is not capable of decoding μ ops at the rate that the processor executes them, it can become a bottleneck.

Re-Order Buffer (ROB) The F-E issues μ ops for execution to the Back-end (B-E) through a set of *pipeline slots*². Upon arrival, the Allocator assigns μ ops to an execution port [3], and places them in the ROB, a circular buffer which allows for in-order termination with out-of-order execution, enabling simple recovery from branch mispredictions. When a μ op is *retired*, its architectural effects are committed and it is removed from the ROB. When all ROB slots become full the processor is stalled, i.e., the issue of μ ops from the F-E to the B-E is stopped, until a slot becomes free.

Scheduler/Reservation Station (RS) Upon execution, μ ops are held in a RS, which includes a copy of their available input data, until its dependencies are satisfied and its assigned execution port becomes free. When these conditions are met, the μ op is effectively executed. When all the RS become full the processor must be stalled.

Execution ports and execution units Each port is linked to execution units responsible for a class of operations: integer ALUs, divider, memory accesses, vector computations, etc.. A port can process a maximum of one μ op per cycle. A lack of execution ports of a given type may cause a backlog, possibly leading to the ROB and/or RS becoming saturated and processor stalls.

²In this context, this term refers to each of the multiple paths between the F-E and the B-E. In each cycle, the front-end sends either zero or one μ ops to the back-end for execution through each of these paths. The number of pipeline slots is a microarchitectural parameter.

$$\begin{aligned} \text{Retiring} &= \frac{\text{SlotsRetired}}{\text{TotalSlots}} \\ \text{Front End Bound} &= \frac{\text{FetchBubbles}}{\text{TotalSlots}} \\ \text{Bad Spec} &= \frac{\text{SlotsIssued} - \text{SlotsRetired} + \text{RecoveryBubbles}}{\text{TotalSlots}} \\ \text{Back End Bound} &= 1 - (\text{Front End Bound} + \text{Bad Spec} + \text{Retiring}) \end{aligned}$$

Figure 3. The four original formulas of TMA L1. On recent Intel microarchitectures, these are directly provided by ad-hoc PMC events.

Memory instructions Memory loads and stores are placed into Load and Store Buffers, respectively. The latency of Load-Store units is determined by the level of the memory hierarchy into which the retrieved data is located. A high memory latency may cause μ ops depending on loaded data to backlog, possibly inducing saturation of the ROB and/or RS and processor stalls.

2.1 Performance debugging and optimization

A state-of-the-art technique for characterizing and debugging the performance of a computational kernel is the Top-down Microarchitecture Analysis (TMA) method [46], developed by Intel engineers and actively supported as part of Perfmon Metrics [17]. Conceptually, TMA studies what happens to each pipeline slot in each cycle of the execution of an application. If I_w is the architectural issue width, T_c is the execution time of the application, and F is the average frequency, TMA categorizes each of the $I_w \cdot T_c \cdot F$ pipeline slots into one of four categories:

- **Retired:** this slot issued a μ op that was eventually retired. This indicates useful computation³.
- **Bad speculation:** this slot issued a μ op that was eventually *not* retired, i.e., it was purged from the ROB without being committed. This is indicative of misspeculation.
- **F-E bound:** this slot remained unused due to no decoded μ ops being available. This is indicative of F-E boundedness.
- **B-E bound:** this slot remained unused due to the ROB/RS being full. This is indicative of B-E boundedness.

In an ideal situation, all the pipeline slots belong to the Retired category. In reality, the distribution of the slots among the four categories indicates how efficiently the CPU is being used, and which of its parts become bottlenecks. TMA defines a hierarchy of subcategories, further refining this top-level classification to give more insight about specific microarchitectural bottlenecks. For instance, a slot into the back-end bound category can be further classified into *core bound* and *memory bound* categories, and in turn a memory bound slot can be attributed to the latency of any of the levels of the memory hierarchy.

³But not necessarily *efficient* computation. For instance, a non-vectorized computation might have a large fraction of its pipeline slots in the *Retired* category, yet be far from peak performance.

The TMA method has been implemented both in Intel’s own VTune [28] tool, as well as on the open-source perf [1] tool. Both use sampling to inspect PMCs that are used to classify pipeline slots into the categories mentioned above. Historically, the low-level counters shown in Fig. 3 were used to calculate the relevant metrics, as described by A. Yasin [46]. On recent generations of Intel CPUs, these metrics can be obtained through dedicated PMC events.

While TMA provides useful information about which parts of the microarchitecture act as bottlenecks for a given computation, it does not provide information about *which parts of the code* trigger the bottlenecks, or why. Tools like perf-annotate [19] help correlate the obtained metrics to specific lines of code, but additional (human) analysis is required. Furthermore, this technique is subject to the inherent inaccuracy of sampling techniques [44, 47].

In addition to sampling inaccuracies, TMA is also subject to methodological ones. Consider again the code in Fig. 1(a). When run on an Intel Core i9-12900k, the code achieves the aforementioned performance of 12.8 GFLOPS, and TMA indicates that the code is fundamentally B-E bound (87.5%), and as memory-bound (44.1%) as it is core-bound (43.4%). Digging deeper into the TMA hierarchy, it finds that while the core is stressed by the utilization of the execution ports, the memory boundedness stems from pressure on the L1 cache. Given that on-peak FP computations are typically core-bound, the user, following TMA’s guidance, decides to focus on the memory-bound component and optimize for register locality. The computation is rescheduled as shown in Fig. 1(b), on which the same ymm2 value can be reused for the entirety of the inner loop.

After the optimization step, the code is re-run, but the same 12.8 GFLOPS performance is obtained. In this case, TMA indicates that the code is almost completely core-bound (91.4%), due to the utilization of the execution ports in the B-E (71.4%). In fact, both these computations are neither memory-bound nor bound by port capacity. The actual bottleneck lies on the dependence chain imposed by the reduction on ymm0 performed by the vfmadd instructions, which limits the throughput to a single FMA at a time. The execution is therefore limited by the *latency* of the FMA unit (4 cycles in this microarchitecture), yielding a throughput of a single vfmadd, i.e., 16 FLOPS, each 4 cycles at 3.2 GHz, that is, exactly 12.8 GFLOPS.

This TMA behavior stems from the fact that some bottlenecks cannot be *causally* established by the usage of PMCs alone. Without any notion of dependencies among instructions and no knowledge of the code, TMA heuristically assumes that stalls caused by B-E pressure in the presence of an elevated Load Queue utilization must be counted towards the memory bound component, and is therefore misguided by this code.

	p0	p1	p2	p3
0	vfmadd _{0,0}		vmovaps ₀	vmovaps ₁
1			vmovaps ₂	
2				vmovaps ₃
3			vmovaps ₄	
4		vfmadd _{0,1}	vmovaps ₆	vmovaps ₅
5			vmovaps ₈	vmovaps ₇
6				
7				vmovaps ₉
8	vfmadd _{1,0}		vmovaps ₁₀	
9				vmovaps ₁₁
10			vmovaps ₁₂	vmovaps ₁₃
11			vmovaps ₁₄	
12		vfmadd _{1,1}		vmovaps ₁₅
⋮	⋮	⋮	⋮	⋮

Figure 4. Port occupancy over time during the execution of version (b) of the inner loop from Fig. 1. Subindices indicate the iteration to which each instruction corresponds and, in the case of vfmadds, whether it is the first or the second FMA in the loop.

2.2 Gus’s approach: Sensitivity and causality analysis

The PMC-based analysis presented in the previous section relies on correlating high-pressured resources with stalls in order to identify performance bottlenecks. However, modern out-of-order architectures feature a high degree of instruction-level parallelism and complex resource interdependencies, and the fact that a resource is stalled while one it depends on is saturated does not necessarily mean that this latest has become a performance bottleneck.

To illustrate this notion, Fig. 4 shows the port pressure during the execution of the first few cycles of the inner loop of the code in Fig 1(a). Ports p2 and p3, used for dispatching the vmovaps memory instructions, are under heavy pressure. TMA assumes correlation between this pressure with the stalls that lead to performance degradation. PMCs provide no way to realize that the dispatched vmovaps are out-of-order with respect to the vfmadds. As such, TMA cannot discover that the vfmadds are stalling due to dependencies with other vfmadds, and not due to RAW hazards with the memory loads. The determination that any back-end stall in the presence of unresolved memory loads must be counted towards memory-boundedness is a heuristic one. However, as shown in Fig. 4, it would be possible to introduce up to 7 additional memory loads in the inner loop before p2 and p3 became bottlenecks, decreasing the throughput from 8 to 9 cycles per iteration.⁴

⁴Note that the difference in the number of vmovaps dispatched per cycle (one or two) depends on code alignment and the width of instruction cache ports.

In this paper, a sensitivity-based approach is proposed. It consists in measuring the performance variations associated to the acceleration of different microarchitectural resources (e.g., port capacity, instruction latency, ROB size, etc.). This approach immediately reveals which factors directly affect performance. When applied to the codes in Fig. 1, the sensitivity-based analysis determines that the performance is linearly correlated with the instruction latency.

To find which static instruction is involved in the bottleneck, this paper also propose a causality-based approach. It consists of propagating resources (and dynamic instructions) that constrain the availability (resp. retiring time) of resources (resp. instructions) that depend on them. In this specific example, it points to the chain of dependencies that constitutes the critical path of the code: $\text{vmadd}_{i,0} \rightarrow \text{vmadd}_{i,1} \rightarrow \dots$. Those approaches are covered in detail in Sec. 3. Identifying this bottleneck without sensitivity would require a detailed analysis of the execution, the code, and the dependencies between the different μops .

Being based on code instrumentation, Gus has yet another advantage over TMA: portability. Due to the very specific PMC events needed to implement the full hierarchy described by the TMA methodology, it needs a remarkably high level of support by vendors. As such, TMA is very well-supported on Intel’s processors, but much more limited on AMD and ARM products, on which it is limited to a few hierarchical levels of detail.

2.3 Other performance analysis techniques

Although Intel’s TMA has become one of the most popular performance analysis techniques, there are other approaches that are covered in the following paragraphs.

2.3.1 Static binary analysis All the tools described here take a basic-block in isolation as an input and statically analyze it as if it were the body of a loop.

Intel *iaca* [2] was the code analyzer distributed by Intel until 2019. When applied to a basic block such as the one in Fig. 1, it delivers a report including a throughput estimation, a detailed ports pressure analysis and a bottleneck analysis which outputs the most probable bottleneck cause. In *iaca*, bottleneck analysis essentially consists in looking for the source of probable *stalls*. It first looks for a traffic jam in the back-end of the CPU and, if it finds one, considers it the bottleneck. Otherwise, it computes the throughput of the basic block when ignoring latencies. If a significant difference with the nominal throughput prediction is found, the basic block is latency-bound. Finally, if the bottleneck is neither the back-end nor instruction latency, then the front-end is marked as bottleneck. The *llvm-mca* [4] code analyzer is based on a similar approach. Its bottleneck analysis engine exploits the underlying performance model used for throughput estimation. It provides indicators of back-end resource saturation

and of the influence of dependence-induced latencies on performance.

facile and *uica* use automatically generated performance models, and differ simply in the algorithm used to make the basic block throughput estimation, which influences the bottleneck analysis. On the one hand, *uica* [8] associates the throughput of a basic block with the timestamp at which its last instruction was retired (after possibly many iterations). A resource is considered to be a bottleneck if its throughput exceeds a certain hard-coded proportion of the total system throughput. On the other hand, *facile* [9] considers that the basic block throughput is the throughput of the slowest CPU resource. With this approach, the latter is by definition the (only) performance-limiting bottleneck.

Both Intel *iaca*, *llvm-mca*, *facile*, and *uica* are meant to provide bottleneck analysis. Other tools in this category are limited to throughput estimation. It is the case of neural-network based ones such as *granite* [43] and *ithemal* [35]. Others only partially model the microarchitecture. For instance, *osaca* [32] assumes that a basic block is never front-end bound.

All the static tools listed above are known for suffering from a lack of context [40]: by ignoring actual values of variables when entering the loop, and by not simulating the evolution of scalar variables they are enabling to model the flow of data through cache and memory leading to critical inaccuracies, even for L1-resident computing kernels.

2.3.2 Simulation-based analysis A possible approach to estimating performance is the use of simulators which give insight into what happens inside the CPU during the execution of an application, without the need to actually measure events in real hardware. This allows to estimate the performance of full or isolated subsystems without actually building them, and for this reason it is the predominant approach in microarchitectural exploration works. *gem5* [33] is an architectural simulator capable of simulating out-of-order execution. More details about how *gem5* relates to Gus are given in Section 3. In general, the use of cycle-level simulators for performance debugging is not a popular approach, due to two fundamental limitations. The first one is methodological: building a cycle-level model of all the events taking place inside a CPU is a very different problem from actually understanding its performance. Similar to the limitations described for TMA, it is necessary to build causal models that relate instruction latency, port capacity, dependencies, and code. The second one is accuracy: architectural simulators are useful tools for differential analysis, i.e., discovering the relative differences between two different designs, but they are not particularly accurate at predicting the actual execution time of a kernel on a given hardware. To overcome these limitations, the current work turns to binary code instrumentation, on which a coarse-grained architecture is simulated

in order to discover first-order performance effects. This approach offers a much more accurate performance prediction, as detailed in Sec. 4.

3 Gus, a sensitivity-oriented code analyzer

We introduce in the following Gus, a microarchitectural simulator for performance debugging. The design of Gus is based around three fundamental goals:

Portability The design of Gus decouples the functional part of the execution from the architectural aspects that drive performance. In the front-end of the tool, an emulator (QEMU) captures the functional behavior of the binary to be analyzed. It then uses binary instrumentation to feed an abstract simulator that models the performance of and interactions between the different architectural blocks. Decoupling the functional and architectural aspects allows using abstract performance models, including those developed for compilers or analyzers such as `llvm-mca` or `uica`. This enables the integration of performance models generated automatically through micro-benchmarking, ensuring portability, which is crucial for precision. These performance models can be improved orthogonally to the development of the tool, benefiting precision and simulation performance.

This design contrasts with traditional architectural simulators such as Gem5, which integrate both aspects and simulate executions by modeling the behavior of each architectural block in an extremely detailed way. This immensely complicates the implementation of new architectures, often leading to “simulation by proxy” approaches which introduce inaccuracies.

Sensitivity As previously mentioned, focusing solely on resource usage rates is insufficient on its own, and can lead to incorrect bottleneck characterization. Additionally, bottlenecks like latency issues are hard to detect through resource usage alone. In Gus, the key criterion for identifying a resource as a bottleneck is sensitivity analysis: adjusting its capacity (throughput or latency) to verify its effect on the performance of the system.

An example of sensitivity-based analysis can be borrowed from the HPC realm, on which users often characterize the bottleneck of a computation by lowering the frequency of the CPU through DVFS while keeping the frequency of the memory constant. If this does not slow down the kernel’s overall execution, it is indicative of memory boundedness. Generalizing this approach using a simulator is another feature of Gus. To enable such analysis, the simulator must allow for parameterizing resource capacity. Unlike the detailed behavioral simulation in Gem5, our abstract model-based design is particularly suited for this purpose. *Given this ability of changing the capacity of each individual resource in our simulator, the implementation of sensitivity in Gus is pretty simple: for each abstract resource, the simulator is run with increased*

capacity (e.g., decreasing the instruction latencies to pinpoint latency-bound kernels, or increasing the L2 cache bandwidth to pinpoint L2-bound kernels) and the predicted execution time is compared.

Causality While identifying the set of instructions that stress execution ports is straightforward, it is more challenging to do the same for other bottlenecks like the ROB or instruction latencies in ILP-bound kernels. The analysis we developed in Gus to highlight bottleneck instructions is called *causality analysis*. Its basic idea is to use a “tainting” mechanism during simulation to trace the instructions or resources that constrain the execution time of others. It is important to note that such propagation would not be possible with an event-based simulator e.g. relying on a SystemC [11] architectural description. Instead of this standard technique, we designed our simulator based on constraints propagation. This technique does not provide the state of each resource at any given time. Instead, it processes the stream of decoded instructions and maintains the earliest (optimistic) availability time for each resource (including hardware resources, operands, instructions, etc.). When a resource is needed, it checks the resources it depends on and propagates time constraints (usually taking the maximum). *This approach is enabled by using conjunctive resource mapping, as described in [20], rather than the standard port mapping used in [5, 39]. This property of the simulator allowed us to develop our tainting-based causality analysis, which we will describe below.*

The remainder of this section covers the design and implementation details of Gus. First, we detail the underlying performance model of Gus based on abstract resources. Then, we explain how sensitivity-based analysis is used in Gus. Finally, we demonstrate the usefulness of sensitivity and how it overcomes the shortcomings of previous bottlenecks detection techniques.

3.1 An abstract resource-centric performance model and causality analyzer

Algorithm 1 shows in detail the coarse-grained simulation approach implemented by Gus. We first ignore tainting. The algorithm takes as input a stream (T) of instructions provided by QEMU. For the current instruction `i`, it updates the state of resources used during its different execution stages. Before `i` can be dispatched, the dispatch queue (IDQ) must have an empty slot. The earliest time at which this can be done is represented as `dispatch.tavail`. This value is updated in line 21 by constraining it to be greater than the time at which the latest instruction in the IDQ can be popped, that is, its retiring time. The resource itself that decodes and dispatches the instruction is represented in this pseudocode as `frontend`. Here we assume it has an inverse throughput of `frontend.inversethroughput`. The earliest time at which this abstract resource can be used is represented as

Algorithm 1: Core algorithm of Gus

```

1 function ConstrainBy(this, c)
2   if this.tavail = c.tavail then
3     this.taint ← this.taint ∪ c.taint
4   else if this.tavail < c.tavail then
5     this.tavail ← c.tavail
6     this.taint ← c.taint
7 function SetBy(this, c)
8   this.tavail ← c.tavail
9   this.taint ← c.taint
10 function UsedBy(this, i, tmin)
11  if this.tavail < tmin then
12    this.taint ← {i}
13    this.tavail ← tmin
14  else
15    this.taint ← this.taint ∪ {i}
16  this.tavail ← this.tavail + this.inversethroughput
17 function Simulate()
18  foreach i ∈ T do
19    taintqueue.push(i)
20    // IDQ/retiring
21    iretired ← dispatchqueue.pop()
22    dispatch.ConstrainBy(iretired)
23    // Front-end
24    frontend.ConstrainBy(dispatch)
25    frontend.UsedBy(i, tmin = dispatch.tavail)
26    // IDQ/dispatch
27    dispatch.ConstrainBy(frontend)
28    dispatchqueue.push(i)
29    i.SetBy(dispatch)
30    // Cache
31    foreach memory line loc accessed by i do
32      foreach cache level Li hit by access of loc do
33        Li.UsedBy(i, tmin = i.dispatch)
34    // Dependencies
35    foreach loc ∈ i.reads do
36      i.ConstrainBy(loc)
37    // Back-end
38    foreach res ∈ i.resources do
39      i.ConstrainBy(res)
40      res.UsedBy(i, tmin = i.dispatch)
41    // Execution of i
42    i.tstart ← i.tavail
43    i.tend ← i.tstart + i.latency
44    i.tavail ← i.tend
45    foreach loc ∈ i.writes do
46      // Assume register renaming works perfectly
47      if loc ≡ Register or loc.tavail < i.tend then
48        loc.SetBy(i)
49    // Critical path
50    i' ← taintqueue.pop()
51    if i' ∈ dispatch.taint then
52      taint[i'.pc] ← taint[i'.pc] + 1
53    i.taint ← i.taint ∪ {i}

```

frontend.t_{avail}. Just as for any abstract resource, in our simulator this value can only increase. When increased by its inverse throughput in line 16 it means it is used during a time inversethroughput. When increased (line 13) because t_{min} is greater than frontend.t_{avail} (line 11) it means it is idle between frontend.t_{avail} and t_{min}. It must be noted that this does not reflect actual usage but provides an approximation of a lower bound on the available timestamp of a resource.

Only once i has been decoded can it be dispatched. This constraint is expressed on line 24. Instruction i is pushed on the dispatch queue and the earliest time at which it can be

executed is set to this timestamp. We continue to follow the lifetime of i by traversing all the locations of its memory operands. A simple cache simulator tells at which cache level the corresponding cache line loc of the operand data is located. L_i (with i the cache level) is an abstract resource that models the throughput between L_i and the previous cache level. Line 29 states that earliest time at which loc can be made available in L1 is not lesser than the earliest availability of L_i. The limited bandwidth between L_i and the previous cache level is represented as L_i.inversethroughput. Similarly to any other abstract resource, its usage is modeled by incrementing L_i.t_{avail} by L_i.inversethroughput. Also, a stall in the dispatch queue will create a bubble in this resource equal to t_{min} - L_i.t_{avail}.

Once the cache related resource states are updated, the constraints imposed by read operands on the time at which i can be executed are propagated to i.t_{avail}. A subtlety is that here loc is a word while before it was a line. The shadow memory which is used to store this information contains both information at the granularity of bytes and information at the granularity of lines. The one at the granularity of bytes is both constrained by the cache behavior, but also by the earliest time at which the last instruction that modified it could be executed.

Similarly to previously considered abstract resources, the t_{avail} of the resources res that represent the compute units (combined ports) are taken into account to update i.t_{avail} (line 34), and inversally the res.t_{avail} is increased accordingly (line 35). The rest of the code reflects the next stages where i.latency is the expected execution latency (obtained using microbenchmarks) of i.

The tainting does nothing else than propagating a set of instructions from the constraining resource (c in line 1) to the constrained instruction (lines 32 and 34) or resource (lines 21, 22, 24, 29): if the constraining resource c creates a bubble (line 4), then the tainting of the constrained resource (this in line 1) is cleared and set to the tainting set of c (line 6). Here, the semantic of this.taint corresponds to the set of all dynamic instructions that may impact the t_{avail} of this one. In case of equality (line 2) both the original set of this.taint and the set of c.taint are possible contributors (line 3). Obviously the tainting does not keep all past dynamic instructions. Instead, a queue is used. The pop of i' in line 42 starts only once taintqueue is sufficiently large (its length is set to twice the ROB size by default). As such, i' corresponds to a sufficiently old instruction. For this instruction, we consider it to contribute to the overall execution time of the program if it delayed the dispatch of i, which is reflected as soon as it belongs to the tainting set dispatch.taint.

3.1.1 Implementation details

Front-end The front-end is modeled by two resources representing 1) decoding unit whose behavior is described

by uiCA [8] and 2) the μ op cache tied to a resource of fixed throughput.

Execution ports and functional units The classical formalism to describe the throughput and sharing of back-end resources is a port mapping, a tripartite graph which describes how instructions decompose into μ ops and which functional units μ ops can execute on.

Gus uses a simpler two-level representation [20] instead, called a resource mapping, where instructions are directly associated with a list of abstract resources. To account for μ op decomposition, a resource can appear in this list multiple times. Aside from avoiding the need to explicitly model the port scheduling algorithm in the back-end and thus allows to achieve a lower runtime overhead, this representation allows a simple propagation of constraints which constitutes the foundation of causality analysis.

Latencies The latencies used for execution units in our experiments are extracted from UOPS.INFO [6] and used to build a resource mapping. For bottleneck analysis, we present latencies induced by dependencies as a separate resource. The same approach can be found, for instance, in iaca. The latter helps guide optimization by highlighting the need to discover more parallelism, rather than the need to relieve the load on a hardware component.

Dispatch Queue The ROB capacity is modeled by the use of a dispatch queue, i.e., a finite-sized buffer of instructions that are in-flight (issued but not retired).

The dispatch queue is bounded by the number of slots it has. As such, if the queue becomes full, no more instructions can be issued until one is retired.

Caches Gus uses a fork [25] of the Dinero IV [22] simulator to determine hits and misses in the different levels of the cache hierarchy, enriched with access latencies and a next-line prefetch heuristic. The PLRU (Pseudo-LRU) replacement policy from this fork is used for all levels of the hierarchy. However, cache replacement policies implemented in commercial processors are a complex topic [7], and an exhaustive modelization is out of the scope of Gus.

3.2 A sensitivity-oriented code analyzer

Sensitivity analysis allows pinpointing the source of a performance bottleneck precisely across instructions. We briefly describe here how it compares with other related works and how it is implemented and used in Gus.

Sensitivity analysis [27, 31], also called differential profiling [34], works by executing a program multiple times, each time varying the usage or capacity of one or more resources. Bottlenecks are then identified by observing how these changes impact the overall performance, i.e. how sensitive performance is to each resource.

Moreover, sensitivity analysis is an automatic approach that does not require as much expertise as existing performance debugging tools [16, 26]. This approach requires careful consideration of different trade-offs in order to be practical.

Simulation granularity Gus employs a coarse-grained simulation approach known as ‘Instruction-window centric simulation’ [15, 24]. The fundamental concept of this simulator category [14, 41] is to represent only components that are critical for performance, thereby enhancing simulation speed. As such Gus’s model assigns a fixed, per-instruction cost to each simulated resource, driven by reverse-engineered microarchitectural features.

This is computationally less intensive than the usual approach in cycle-level simulators such as gem5 [33]. The latter embeds a single model to simulate OoO execution, the O3CPU model. While this model is designed to be configurable to accurately reflect the microarchitectural layout, some components *must* be defined for it to be usable: fetch and decode logic, structures for issuing instructions such as the issue queue, register renaming, or pipeline stages.

Some modes (FSA [42] or functional mode) are designed to *swiftly* establish a consistent checkpoint of Gem5’s internal structures at a predetermined point in the execution process (usually post-boot), foregoing any simulation. A switch to a detailed model, based on the O3CPU, is then done for any kind of simulation. Similarly, Gus can also restrict its instrumentation to a subset of the program, leaving the other code generated by QEMU to run as is. While this reduces the computational cost of the simulation by a significant margin, the bottleneck of running the region of interest itself still remains. The coarse-grained approach implemented by Gus has proved to be a good compromise between simulation speed and precision.

Sensitivity analysis in the broader world The idea of varying microarchitectural features is not new in itself either, as it has been used by hardware designers for Microarchitecture Design Space Exploration [10, 29] to guide microarchitecture parameter tuning to explore the trade-offs amongst performance, power, and area (PPA).

DECAN [31] is a dynamic performance analysis tool based on the MAQAO [21] binary analysis and instrumentation framework. DECAN finds bottlenecks by sensitivity analysis based on binary rewriting; it removes or modifies instructions in a kernel and checks by how much each transformation affects the overall performance via performance counters.

Coz [18] relies on *causal profiling*, which consists in sensitivity on code lines. Each one is successively ‘accelerated’ to quantify its impact on performance, by slowing down the surrounding code (in practice, the latter is embedded in threads that are paused when necessary).

ASM	L2	p23	LAT/PORTS
mov rdx, qword ptr [rsp - 0x10]	0%	10%	2/p23
vmovsd xmm0, qword ptr [rdx + rax]	0%	10%	4/p23
vaddsd xmm0, xmm0, qword ptr [rdx + rax + 8]	0%	10%	4/p016 p01 p015 p0156 p23
vaddsd xmm0, xmm0, qword ptr [rdx + rax + 0x10]	0%	10%	4/p016 p01 p015 p0156 p23
vmulsd xmm0, xmm0, xmm1	0%	0%	4/p016 p01 p015 p0156
mov rdx, qword ptr [rsp - 0x18]	0%	10%	2/p23
vmovsd qword ptr [rdx + rax + 8], xmm0	1%	0%	4/p4
mov rdx, qword ptr [rsp - 0x18]	0%	10%	2/p23
vmovsd xmm0, qword ptr [rdx + rax - 8]	0%	10%	4/p23
vaddsd xmm0, xmm0, qword ptr [rdx + rax]	0%	10%	4/p016 p01 p015 p0156 p23
vaddsd xmm0, xmm0, qword ptr [rdx + rax + 8]	0%	10%	4/p016 p01 p015 p0156 p23
vmulsd xmm0, xmm0, xmm1	0%	0%	4/p016 p01 p015 p0156
mov rdx, qword ptr [rsp - 0x10]	0%	10%	2/p23
vmovsd qword ptr [rdx + rax], xmm0	0%	0%	4/p4
add rax, 0x18	0%	0%	1/
cmp rax, rcx	0%	0%	1/p0156

Table 1. Fine grain report analysis of resource usage by instruction by Gus on a Jacobi kernel (Skylake architecture). Only a subset of resources are depicted. The cells colored in orange indicate the bottleneck port and the corresponding instructions.

The SAAKE system [27] is conceptually closer to Gus in its implementation of sensitivity analysis. It uses a fast symbolic execution engine that estimates the runtime of GPU programs to drive sensitivity analysis for finding bottlenecks.

Sensitivity analysis in Gus In addition to execution ports, resources which are subject to sensitivity analysis in Gus are: the instruction latencies, the front-end throughput, the size of the ROB, the throughput of the retire buffer, and the communication bandwidths between the different memory levels.

Each resource has a *capacity*, a finite quantity whose variation may impact the execution time of a program. These may be quantities immediately linked to a hardware component, such as the size of the ROB, or more abstract in nature, such as instruction latency. We represent each of these quantities by a real number. The sensitivity analysis implemented in Gus consists in varying the n resources represented in the model during successive execution time estimates performed on the same program p . At each iteration, the capacity c_r of a resource r is successively weighted by real numbers w_0, \dots, w_m to discover a w that minimizes the estimation function f_p . In this framework, the other resource capacities and the input program p can be seen as constants. Thus, f_p is a function from real numbers to real numbers, associating an estimate of duration with c_r . The speed-up s_r obtained by weighting c_r by w is therefore calculated as follows:

$$s_{w,r} = \frac{f_p(c_r)}{f_p(w * c_r)} - 1$$

Resources whose variations result in a speedup are bottleneck resources, and should be the focus of optimization efforts.

Gus instruction-level report One can also look at the resource usage per instruction to pinpoint it in the kernel. Table 1 contains the instruction-level report generated by Gus from a Jacobi stencil. Sensitivity analysis shows that

the runtime of this kernel is not affected by increasing the throughput of any other resource aside than ports p2 and p3, but the report is more specific. It shows that multiple instructions (and which) use the resource p23 by loading from memory. This level of precision is useful for performance debugging, as it allows a level of granularity that is not usually possible through the measurement of PMC events [46].

3.3 Using Gus for manual code optimization

In this subsection, we illustrate how sensitivity analysis can be used to drive manual code optimization. For this purpose we analyze the correlation kernel, which is one of the evaluation benchmarks discussed below (in section 4). Its most costly loops nest is shown in Fig. 5. The different steps are reported in Table 2. The compiler (gcc-13 in this example) does a very poor job of vectorizing the original version, leading to a very low performance. We thus start by vectorizing the code by hand, leading to version **v1**, which is still not very efficient. On this version, both TMA and Gus pinpoint L2 bandwidth as a bottleneck. We apply a register tiling transformation to improve data locality. Considering the target microarchitecture (Skylake-X), we choose a 2x16 register tile size. With AVX2, this executes 4x SIMD Fuse-Multiply-Adds (FMAs) per iteration, 2 iterations running in parallel. This version **v2** reaches 33.1% of the peak performance.

On this new version, Gus highlights the retire buffer as the bottleneck. The detailed report shows that the stores (vmovapd) cause this behavior as they take 2 slots in the retire buffer. We hoist the load/store instructions outside of the innermost loop to reduce the overall number of instructions, thus the pressure on the retire buffer. The obtained version **v3** reaches 46.8% of the peak performance. Gus pinpoints now the L2 cache accesses as the bottleneck. We tile the axis of the matrix multiplication in order to improve L1-reuse. This optimization (leading to version **v4**) unlocks 81.2% of the peak performance. Gus pinpoints the retire buffer by showing too many vmovapd suggesting a more squared shaped register tile. Forgetting that our architecture contained only 16 vector registers, we had inadvertently changed the register tile to 4x16. Sensitivity analysis on **v5** outputs instruction latency as a bottleneck, introduced by dependency chains on spilled registers; which lead to our final version **v6**. There, the bottleneck shifts to ports p0 and p1 that process FMAs. As these are exactly the core operations we want to execute, we stop here the optimization process.

Overall we improve the performance by 662.4%, achieving 82.8% of the peak performance. As summarized in Table 2, optimization is guided at every stage by Gus’s sensitivity analysis and instruction-level report. TMA provides little guidance from step 3 onwards. In particular in step 5, while TMA points to backend-boundedness, which is not straightforward to detect spilling: only sensitivity analysis on instruction latency assures that dependency chains are the bottleneck.

Version	Transformations	Peak perf.	TAM (perf)	Bottleneck (Gus): ressource, instructions
v0	None	4.5%		
v1	<i>Vectorized</i>	12.5%	Backend(L2)	L2
v2	Vectorized, <i>Register tiling 2x16</i>	33.1%	Retiring	Retire Buffer
v3	Vectorized, Register tiling 2x16, <i>Hoisting</i>	46.8%	Backend (ports usage)	L2, vbroadcastsd, vmovaddpd (loads)
v4	Vectorized, Register tiling 2x16, Hoisting, <i>Tiling</i>	81.2%	Backend (ports usage)	Retire buffer, includ. vmovaddpd
v5	Vectorized, Register tiling 4x16, Hoisting, Tiling	52.6%	Backend (ports usage)	Latency, includ. vmovaddps[rsp] (spill)
v6	Vectorized, Register tiling 3x12, Hoisting, Tiling	82.8%	Backend (ports usage)	P01

Table 2. Optimization process using Gus on the most costly loops extracted from the benchmark correlation.

```

1 for (int j = 0; j < _PB_M; j++) {
2   for (int i = j; i < _PB_M; i++) {
3     for (int k = 0; k < _PB_N; k++) {
4       corr[i][j] += data[k][i] * data[k][j];
5     }
6     corr[j][i] = corr[i][j];
7   }
8 }

```

Figure 5. Untransformed correlation kernel.

4 Experiments

We first demonstrate the simulation capabilities of Gus by comparing it against a cycle-level simulator over a set of numerical kernels.

We then validate the portability and accuracy of Gus against measurements on several microarchitectures.

4.1 Benchmarking framework

Our benchmarking framework relies on the benchmark suite PolyBench [38], composed of 30 numerical computation kernels. We combine C loop nest optimizers – Pluto [12] and PoCC [37] – and gcc to apply several transformations (tiling, loop fusion, unroll and jam, etc.) to PolyBench. We evaluate Gus on the resulting variants.

The measurements have mostly been performed using the Performance Application Programming Interface library (PAPI) [13]. On the system where the latter is unavailable (the ARM-based one), we used the Linux perf API. In both cases, we instrumented the benchmark in order to profile a single section of code (a function call).

To carry out the measurements, we checked that three conditions were met on the target machines:

- Simultaneous Multi-Threading must be deactivated, as it tends to increase pressure on the front-end, resulting in overestimated hardware counters.
- Processor frequency must be set to the manufacturer’s base frequency, as frequency variations change memory-boundness.
- The benchmark should use huge pages to be less sensitive to TLB misses.

4.2 Validation of accuracy and speed against a cycle-level simulator

We present in Figure 6 a comparison of the precision and simulation speed of Gus against a state-of-the-art cycle-level simulator, gem5 [33] (v24.0.0.1) on a Skylake microarchitecture⁵. We use the configuration with the highest level of optimization, fast for gem5, with the Deriv03CPU CPU model and a verbatim description of the processor based on reverse-engineered characteristics. Since gem5 does not support more recent SIMD extensions than SSE, we limit our gcc-13 compiler to generate SSE code. Over the 410 benchmarks generated by our benchmark framework, 147 had an optimization which did not affect the binary code and were discarded. An arbitrary timeout of an hour was set for both simulators, that was reached by 35 benchmarks for gem5 and none for Gus.

We achieve an order-of-magnitude faster simulation over gem5, with a mean of 1.23 million instructions per second (MIPS) compared to 0.11 MIPS. We also surpass gem5 in accuracy, with a MAPE (Mean absolute percentage error) of 14.56% compared to 87.28%, but also in terms of Kendall tau [30] ($\tau_K \in [0, 1]$), which measures the fraction of preserved pairwise ordering).

We would like to outline that we focused our implementation efforts for Gus on the relevance of the output and not its speed yet. In particular, the size of the shadow memory which is the main source of slowdown could be easily reduced.


The low accuracy achieved by gem5 can be explained through some of its design characteristics:

- Macro-operations are decomposed into micro-operations in gem5. However, the decomposition is imprecise, as it is defined for the whole x86 ISA, and not for the actual microarchitecture. We observe that while Gus achieves a MAPE of 2.44% on the number of retired micro-operations, gem5 presents a MAPE of 52.63%.
- No μ op-cache is modeled in gem5. This results in a pessimistic front-end model.

⁵Cascade Lake generation, Intel Xeon Gold 6230R CPU.

Simu	#benchs	MIPS mean	Cycles MAPE	Cycles τ_K	Retired μ ops MAPE
Gus	263	1.23	14.56%	0.92	2.44%
gem5	228	0.11	87.28%	0.84	52.63%

(a)



(b)

Figure 6. Statistical comparison (a) and corresponding error distribution (b) of accuracy and speed of Gus against the cycle-level simulator gem5.

4.3 Validation of accuracy against measurements

We validate the accuracy of Gus against measurements on several microarchitectures, as shown in Table 3. This demonstrates its portability and modularity (especially for its performance model, which is fed in this case by `uops.info` tables [5] for our Intel micro-architectures and by PALMED [20] table for our Cortex-A72). This approach achieves satisfying results on different architectures.

Microarchitecture	Linux	Release	CPU model
Maia	6.1	2015	ARM Cortex-A72
Sandy Bridge	6.11	2011	Intel Core i5-2520m
Skylake X	6.11	2015	Intel Core i9-7980XE
Skylake SP	5.10	2019	Intel Xeon Gold 6230R
Ice Lake	6.11	2020	Intel Core i7-1065G7
Alder Lake	6.11	2021	Intel Core i9-12900k

Table 3. Microarchitectures used for validation of Gus.

We generate numerous sets of benchmarks for each microarchitecture, with SIMD extensions (AVX2 for x86, NEON for ARM) with a wide range of optimizations as described in section 4.

We would like to stress that our performance model relies solely on automatic resource characterization [6, 20], which is an orthogonal active research topic.

Microarchitecture	#benchmarks	MAPE
Maia	289	32.13%
Skylake SP	267	18.60%
Skylake X	267	31.64%
Ice Lake	267	24.11%
Sandy Bridge	267	38.99%
Alder Lake	267	25.01%

Table 4. Accuracy of Gus on different microarchitectures.

4.4 Consistency of sensitivity analysis

The experimental set-ups where Gus obtains a low MAPE compared to the measurements indicate that it models well the impact of the program’s instruction mix on the concerned microarchitectures. Thus, we also checked that the sensitivity gave consistent results. This means that, for a benchmark B and its optimized version V (*i.e.*, V has a smaller execution time), then the bottlenecks of B discovered by the sensitivity analysis should appear equally or less stressed in V than in B. By alleviating the pressure on these bottlenecks, we have indeed improved the predicted execution time. All the pairs benchmark/variant of our dataset verify this property.

5 Discussion

In this section, we discuss the design choices of Gus and why they are relevant for performance debugging, and how we differentiate from typical computer architecture simulators.

5.1 Trade-offs in the design of Gus

Over the years, many microarchitectural simulators [14, 33, 36, 41] have been proposed for computer-system architecture research, each with its own set of trade-offs. Ideally one would like a one-size-fits-all simulator that is both fast and accurate, with a wide applicability to different workloads and fast development time. However, simulators were designed with different goals in mind (e.g. parallelism, memory hierarchy, etc.), and thus focus on a peculiar dimension of the design space.

Especially, one aspect that is often overlooked with simulators is development cost – describing a detailed configuration for a microarchitecture down to the μ op level can be a daunting task. In contrast, we propose what we call a *model-based* approach, where we feed a reverse-engineered performance model to a coarse-grained core model to find performance bottlenecks. This approach covers less of the design space than a full-fledged simulator, as it sacrifices coverage and some accuracy, but enables finding performance bottlenecks in a reasonable amount of time and effort.

We summarize briefly the major trade-offs of simulation approaches in Table 5 to position Gus among them. We differentiate between cycle-level simulators (e.g. gem5 [33]), cycle-approximate simulators (e.g. Sniper [14], ZSim [41]), and model-based simulators (e.g. Gus).

Type	Dev	Accuracy	Speed	Coverage
Cycle-level	Slow	High	Low	High
Cycle-approximate	Slow	Medium	High	High
Model-based	Fast	Medium	High	Low

Table 5. Comparison of different approaches to microarchitectural simulation.

6 Conclusion

Identifying performance bottlenecks in programs that need to make the most of the architecture is an increasingly critical task. We present micro-architectural mechanisms on which program performance depends, as well as the metrics and analysis methods that have been designed to guide optimization. We propose to generalize bottleneck analysis by means of microarchitectural sensitivity and constraint causality. Sensitivity analysis aims to automatically discover which resources influence a program’s overall performance by successively accelerating each of them and observing the overall speedup generated by this variation. Causality analysis aims to propagate timing constraints through the simulation process to pinpoint how each instruction contribute to the overall execution time. We present Gus, a dynamic code analyzer that implements sensitivity and causality analysis. We evaluate both its performance model and bottleneck analysis algorithm on a set of benchmarks generated from the PolyBench benchmark suite, and highlight its strengths and limitations compared with existing methods based on hardware counters or performance models. Gus’s performance model achieves state-of-the-art accuracy in throughput estimation over our experimental harness. As for our bottleneck analysis, it allows us to enhance the results of existing approaches, especially where saturation is not sufficient to identify the one that limits parallelism.

References

- [1] Performance analysis tools for linux. <https://man7.org/linux/man-pages/man1/perf.1.html>.
- [2] Intel architecture code analyzer user’s guide. <https://www.intel.com/content/dam/develop/external/us/en/documents/intel-architecture-code-analyzer-3-0-users-guide-157552.pdf>, 2017.
- [3] Intel 64 and ia-32 architectures optimization reference manual, 2023. <https://cdrdv2-public.intel.com/671488/248966-046A-software-optimization-manual.pdf>.
- [4] Llm machine code analyzer. <https://llvm.org/docs/CommandGuide/llvm-mca.html>, 2023.
- [5] Andreas Abel and Jan Reineke. uops.info: Characterizing latency, throughput, and port usage of instructions on intel microarchitectures. In *ASPLOS, ASPLOS ’19*, pages 673–686, New York, NY, USA, 2019. ACM.
- [6] Andreas Abel and Jan Reineke. uops.info: Characterizing latency, throughput, and port usage of instructions on intel microarchitectures. In *ASPLOS, ASPLOS ’19*, pages 673–686, New York, NY, USA, 2019. ACM.
- [7] Andreas Abel and Jan Reineke. nanobench: A low-overhead tool for running microbenchmarks on x86 systems. In *2020 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*, August 2020.
- [8] Andreas Abel and Jan Reineke. Uica: Accurate throughput prediction of basic blocks on recent intel microarchitectures. In *Proceedings of the 36th ACM International Conference on Supercomputing, ICS ’22*, New York, NY, USA, 2022. Association for Computing Machinery.
- [9] Andreas Abel, Shrey Sharma, and Jan Reineke. Facile: Fast, accurate, and interpretable basic-block throughput prediction. In *2023 IEEE International Symposium on Workload Characterization (IISWC)*, pages 87–99. IEEE Computer Society, 10 2023.
- [10] Chen Bai, Jiayi Huang, Xuechao Wei, Yuzhe Ma, Sicheng Li, Hongzhong Zheng, Bei Yu, and Yuan Xie. Archexplorer: Microarchitecture exploration via bottleneck analysis. In *Proceedings of the 56th Annual IEEE/ACM International Symposium on Microarchitecture, MICRO ’23*, page 268–282, New York, NY, USA, 2023. Association for Computing Machinery.
- [11] David C. Black, Jack Donovan, Bill Bunton, and Anna Keist. *SystemC: From the Ground Up, Second Edition*. Springer Publishing Company, Incorporated, 2nd edition, 2009.
- [12] Uday Bondhugula, J. Ramanujam, and P. Sadayappan. Pluto: A practical and fully automatic polyhedral parallelizer and locality optimizer. Technical Report OSU-CISRC-10/07-TR70, The Ohio State University, October 2007.
- [13] S. Browne, J. Dongarra, N. Garner, G. Ho, and P. Mucci. A portable programming interface for performance evaluation on modern processors. *Int. J. High Perform. Comput. Appl.*, 14(3):189–204, August 2000.
- [14] Trevor E. Carlson, Wim Heirman, and Lieven 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*, pages 1–12, 2011.
- [15] Trevor E. Carlson, Wim Heirman, Stijn Eyerman, Ibrahim Hur, and Lieven Eeckhout. An evaluation of high-level mechanistic core models. *ACM Trans. Archit. Code Optim.*, 11(3), aug 2014.
- [16] Andres S. Charif-Rubial, Emmanuel Oseret, José Noudouhouenou, William Jalby, and Ghislain Lartigue. Cqa: A code quality analyzer tool at binary level. In *2014 21st International Conference on High Performance Computing (HiPC)*, pages 1–10, 2014.
- [17] Intel Corporation. Perfmon metrics. <https://github.com/intel/perfmon>.
- [18] Charlie Curtsinger and Emery D. Berger. Coz: finding code that counts with causal profiling. In *Proceedings of the 25th Symposium on Operating Systems Principles, SOSP ’15*, page 184–197, New York, NY, USA, 2015. Association for Computing Machinery.
- [19] Arnaldo Carvalho de Melo and Red Hat. The new linux ’perf’ tools. 2010.
- [20] Nicolas Derumigny, Théophile Bastian, Fabian Gruber, Guillaume Iooss, Christophe Guillon, Louis-Noël Pouchet, and Fabrice Rastello. Palmed: Throughput characterization for superscalar architectures. In *Proceedings of the 20th IEEE/ACM International Symposium on Code Generation and Optimization, CGO ’22*, page 106–117. IEEE Press, 2022.
- [21] Lamia Djoudi, Denis Barthou, Patrick Carribault, Christophe Lemuet, Jean-Thomas Acquaviva, and William Jalby. Maqao: Modular assembler quality analyzer and optimizer for itanium 2. In *Workshop on Explicitly Parallel Instruction Computing Techniques*, Santa Jose, California, March 2005.
- [22] Jan Edler and Mark D. Hill. Dinero iv trace-driven uniprocessor cache simulator. <https://pages.cs.wisc.edu/markhill/DineroIV/>.
- [23] Hadi Esmailzadeh, Emily Blem, Renée St. Amant, Karthikeyan Sankaralingam, and Doug Burger. Dark silicon and the end of multicore scaling. In *2011 38th Annual International Symposium on Computer Architecture (ISCA)*, pages 365–376, 2011.
- [24] Davy Genbrugge, Stijn Eyerman, and Lieven Eeckhout. Interval simulation: Raising the level of abstraction in architectural simulation.

- In *HPCA - 16 2010 The Sixteenth International Symposium on High-Performance Computer Architecture*, pages 1–12, 2010.
- [25] Christophe Guillon. Dinerio iv with plru replacement policy support. <https://github.com/atos-tools/dinerioIV>.
- [26] Julian Hammer, Jan Eitzinger, Georg Hager, and Gerhard Wellein. Kerncraft: A tool for analytic performance modeling of loop kernels. In Christoph Niethammer, José Gracia, Tobias Hilbrich, Andreas Knüpfer, Michael M. Resch, and Wolfgang E. Nagel, editors, *Tools for High Performance Computing 2016*, pages 1–22, Cham, 2017. Springer International Publishing.
- [27] Changwan Hong, Aravind Sukumaran-Rajam, Jinsung Kim, Prashant Singh Rawat, Sriram Krishnamoorthy, Louis-Noël Pouchet, Fabrice Rastello, and P. Sadayappan. Gpu code optimization using abstract kernel emulation and sensitivity analysis. In *Proceedings of the 39th ACM SIGPLAN Conference on Programming Language Design and Implementation*, PLDI 2018, page 736–751, New York, NY, USA, 2018. Association for Computing Machinery.
- [28] Intel. Vtune profiler. <https://www.intel.com/content/www/us/en/developer/tools/oneapi/vtune-profiler.html>, 2011.
- [29] Engin İpek, Sally A. McKee, Rich Caruana, Bronis R. de Supinski, and Martin Schulz. Efficiently exploring architectural design spaces via predictive modeling. *SIGOPS Oper. Syst. Rev.*, 40(5):195–206, oct 2006.
- [30] Maurice G Kendall. A new measure of rank correlation. *Biometrika*, 30(1/2):81–93, 1938.
- [31] Souad Koliaï, Zakaria Bendifallah, Mathieu Tribalat, Cédric Valensi, Jean-Thomas Acquaviva, and William Jalby. Quantifying performance bottleneck cost through differential analysis. In *Proceedings of the 27th International ACM Conference on International Conference on Supercomputing*, ICS '13, page 263–272, New York, NY, USA, 2013. Association for Computing Machinery.
- [32] Jan Laukemann and Georg Hager. Core-level performance engineering with the open-source architecture code analyzer (osaca) and the compiler explorer. In *Companion of the 2023 ACM/SPEC International Conference on Performance Engineering*, ICPE '23 Companion, page 127–131, New York, NY, USA, 2023. Association for Computing Machinery.
- [33] Jason Lowe-Power, Abdul Mutaal Ahmad, Ayaz Akram, Mohammad Alian, Rico Amslinger, Matteo Andreozzi, Adrià Armejach, Nils Asmussen, Brad Beckmann, Srikant Bharadwaj, Gabe Black, Gedare Bloom, Bobby R. Bruce, Daniel Rodrigues Carvalho, Jeronimo Casttrillon, Lizhong Chen, Nicolas Derumigny, Stephan Diestelhorst, Wendy Elsasser, Carlos Escuin, Marjan Fariborz, Amin Farmahini-Farahani, Pouya Fotouhi, Ryan Gambord, Jayneel Gandhi, Dibakar Gope, Thomas Grass, Anthony Gutierrez, Bagus Hanindhito, Andreas Hansson, Swapnil Haria, Austin Harris, Timothy Hayes, Adrian Herrera, Matthew Horsnell, Syed Ali Raza Jafri, Radhika Jagtap, Hanhwi Jang, Reiley Jeyapaul, Timothy M. Jones, Matthias Jung, Subash Kan-noth, Hamidreza Khaleghzadeh, Yuetsu Kodama, Tushar Krishna, Tommaso Marinelli, Christian Menard, Andrea Mondelli, Miquel Moreto, Tiago Mück, Omar Naji, Krishnendra Nathella, Hoa Nguyen, Nikos Nikoleris, Lena E. Olson, Marc Orr, Binh Pham, Pablo Prieto, Trivikram Reddy, Alec Roelke, Mahyar Samani, Andreas Sandberg, Javier Setoain, Boris Shingarov, Matthew D. Sinclair, Tuan Ta, Rahul Thakur, Giacomo Travaglini, Michael Upton, Nilay Vaish, Ilias Vougioukas, William Wang, Zhengrong Wang, Norbert Wehn, Christian Weis, David A. Wood, Hongil Yoon, and Éder F. Zulian. The gem5 simulator: Version 20.0+, 2020.
- [34] Paul E. McKenney. Differential profiling. In *Proceedings of the 3rd International Workshop on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems*, MASCOTS '95, page 237–241, USA, 1995. IEEE Computer Society.
- [35] Charith Mendis, Saman P. Amarasinghe, and Michael Carbin. Ithemal: Accurate, portable and fast basic block throughput estimation using deep neural networks. *CoRR*, abs/1808.07412, 2018.
- [36] Avadh Patel, Furat Afram, Shunfei Chen, and Kanad Ghose. Marss: A full system simulator for multicore x86 cpus. In *2011 48th ACM/EDAC/IEEE Design Automation Conference (DAC)*, pages 1050–1055, 2011.
- [37] PoCC, the polyhedral compiler collection. <https://www.cs.colostate.edu/~pouchet/software/poccc/>.
- [38] Louis-Noël Pouchet and Tomofumi Yuki. PolyBench/C: The polyhedral benchmark suite, version 4.2, 2016. <http://polybench.sf.net>.
- [39] Fabian Ritter and Sebastian Hack. Pmevo: portable inference of port mappings for out-of-order processors by evolutionary optimization. In *Proceedings of the 41st ACM SIGPLAN Conference on Programming Language Design and Implementation*, PLDI 2020, page 608–622, New York, NY, USA, 2020. Association for Computing Machinery.
- [40] Fabian Ritter and Sebastian Hack. Anica: Analyzing inconsistencies in microarchitectural code analyzers. *Proc. ACM Program. Lang.*, 6(OOP-SLA2), oct 2022.
- [41] Daniel Sanchez and Christos Kozyrakis. Zsim: Fast and accurate microarchitectural simulation of thousand-core systems. In *Proceedings of the 40th Annual International Symposium on Computer Architecture*, ISCA '13, page 475–486, New York, NY, USA, 2013. Association for Computing Machinery.
- [42] Andreas Sandberg, Nikos Nikoleris, Trevor E. Carlson, Erik Hagersten, Stefanos Kaxiras, and David Black-Schaffer. Full speed ahead: Detailed architectural simulation at near-native speed. In *2015 IEEE International Symposium on Workload Characterization*, pages 183–192, 2015.
- [43] O. Sykora, P. Phothilimthana, C. Mendis, and A. Yazdanbakhsh. Granite: A graph neural network model for basic block throughput estimation. In *2022 IEEE International Symposium on Workload Characterization (IISWC)*, pages 14–26, Los Alamitos, CA, USA, nov 2022. IEEE Computer Society.
- [44] Vincent M Weaver, Dan Terpstra, and Shirley Moore. Non-determinism and overcount on modern hardware performance counter implementations. In *2013 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*, pages 215–224. IEEE, 2013.
- [45] Samuel Williams, Andrew Waterman, and David Patterson. Roofline: an insightful visual performance model for multicore architectures. *Commun. ACM*, 52(4):65–76, apr 2009.
- [46] Ahmad Yasin. A top-down method for performance analysis and counters architecture. In *2014 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*, pages 35–44, 2014.
- [47] Jifei Yi, Benchao Dong, Mingkai Dong, and Haibo Chen. On the precision of precise event based sampling. In *Proceedings of the 11th ACM SIGOPS Asia-Pacific Workshop on Systems*, APSys '20, page 98–105, New York, NY, USA, 2020. Association for Computing Machinery.