

Accelerate Cycle-Level Full-System Simulation of Multi-Core RISC-V Systems with Binary Translation

Xuan Guo
Gary.Guo@cl.cam.ac.uk
University of Cambridge
Cambridge, UK

Robert Mullins
Robert.Mullins@cl.cam.ac.uk
University of Cambridge
Cambridge, UK

ABSTRACT

It has always been challenging to balance the accuracy and performance of instruction set simulators (ISSs). Register-transfer level (RTL) simulators or systems such as gem5 [4] are used to execute programs in a cycle-accurate manner but are often prohibitively slow. In contrast, functional simulators such as QEMU [2] can run large benchmarks to completion in a reasonable time yet capture few performance metrics and fail to model complex interactions between multiple cores. This paper presents a novel multi-purpose simulator that exploits binary translation to offer fast cycle-level full-system simulations. Its functional simulation mode outperforms QEMU and, if desired, it is possible to switch between functional and timing modes at run-time. Cycle-level simulations of RISC-V multi-core processors are possible at more than 20 MIPS, a useful middle ground in terms of accuracy and performance with simulation speeds nearly 100 times those of more detailed cycle-accurate models.

1 INTRODUCTION

RISC-V is a free, open, and extensible ISA. With the ongoing ecosystem development of RISC-V and an increasing number of companies and institutions switching to RISC-V for both production and research, RISC-V has become the test bed instruction set of computer architecture research. A key tool when exploring new architectural trade-offs is the instruction-set simulator (ISS). Fast cycle-level simulation allows new ideas to be validated quickly at an appropriate level of abstraction and without the complexities of hardware development. In particular we focus on the challenge of simulating multi-core RISC-V systems.

The design of a processor can broadly be divided into the design of the core and memory subsystem. Characterising the performance of the core pipeline in isolation is often a simpler task to that of characterising the memory system. While smaller synthetic benchmarks are useful at the core level, larger more complex and longer running workloads are often needed to understand the memory system and the potential interactions between cores.

For example, the smaller synthetic MCU benchmark CoreMark [7] executes at a magnitude of 10^8 instructions per iteration, while SPEC2017 [8], a larger and more realistic benchmark running real-life applications, requires a magnitude of 10^{12} instructions for a single run [12]. SPEC takes from hours to days to run even on real

machines, and hardly possible for simulators to run to completion. It is therefore helpful to have a fast simulator.

Of course, fast simulation involves a trade-off between the fidelity of the model and the speed at which simulations can be completed. Unfortunately, we are currently forced to choose between slow cycle-accurate simulators or fast functional-only simulators such as QEMU. In particular, there is a lack of fast full-system simulators that can accurately model cache-coherent multi-core processors.

In this paper, we present the Rust RISC-V Virtual Machine (R2VM). R2VM is written in the increasingly popular high-level system programming language Rust [17]. R2VM is released ¹ under permissive MIT/Apache-2.0 licenses in the hope to encourage its adoption and expansion by the broader community. To our knowledge, this is the first binary translated simulator that supports cycle-level multi-core simulation. It can accurately model cache coherency protocols and shared caches. Cycle-level simulations are possible at more than 20 MIPS, while the performance of functional-only simulations can outperform QEMU and exceed 400 MIPS.

2 BACKGROUND

2.1 Instruction Set Simulators

ISSs can be classified as either execution-driven, emulation-driven or trace-driven [6]. We omit a detailed discussion of execution-driven simulators such as Cachegrind [13] that modify programs with binary instrumentation and execute them natively, because they require the host and the guest ISA to be identical and do not support full-system simulation. Emulation-driven simulators emulate the program execution, and gather performance metrics on-the-fly; in contrast, trace-driven simulators run emulation before-hand and gather traces from the program, e.g. branches or memory accesses, and later replay the trace against a specific model. Traces allow ideas to be evaluated quickly without the need to simulate in detail, but cannot easily capture effects that may alter the instructions that are executed, e.g. inter-core interactions or speculative execution [6]. Moreover, storage space required for traces grows linearly with the length of execution, making trace-driven simulators incapable of simulating large benchmarks. R2VM is an emulation-driven simulator and the remainder of this paper will focus exclusively on emulation-driven simulators.

Simulators can also be categorised by their levels of abstraction. One category of simulators is functional simulators. Functional simulators simulate the effects of instructions without taking microarchitectural details into account. Because less information is

This paper is published under the Creative Commons Attribution 4.0 International (CC-BY 4.0) license.

CARRV 2020, May 29, 2020, Virtual Workshop
© 2020 Copyright held by the owner/author(s).

¹Available at <https://github.com/nbdd0121/r2vm>.

needed, aggressive optimisations can be performed, and the performance is usually several magnitudes faster than timing simulators. QEMU falls into this category. It should be noted though that while QEMU itself is a purely functional simulator, it can be modified to collect metadata for off-line or on-line cache simulation [18].

The other category is timing simulators. Out of timing simulators, RTL simulators can model processor microarchitectures very precisely, but the difficulty in implementing a feature in RTL simulator is not much different from implementing it in hardware directly. RTL simulators are also poor in performance, usually run at a magnitude of kIPS [16].

At a higher level, there are cycle-level microarchitectural simulators. These are able to omit RTL implementation details to improve performance while retaining a detailed microarchitectural model. An popular example is the gem5 simulator running with In-Order or O3 mode [4]. For faster performance, we can give up some extra microarchitectural details and predict the number of cycles taken for each non-memory instruction instead of computing them in real-time, and in the extreme case, assume all non-memory operation only takes 1 cycle to execute as gem5’s “timing simple” CPU model assumes. This approach is no longer cycle-accurate, but this cycle-approximate model is often adequate to perform cache and memory simulations.

2.2 Binary Translation

Binary translation is a technique that accelerates instruction set architecture (ISA) simulation or program instrumentation [11]. An interpreter will fetch, decode and execute the instruction pointed by the current program counter (PC) one-by-one, while binary translation will, either ahead of time (static binary translation) or in the runtime, i.e. when the block of code is first executed (dynamic binary translation (DBT)), translate one or more basic blocks from the simulated ISA to the host’s native code, cache the result, and use the translation result next time the same block is executed.

QEMU uses binary translation for cross-ISA simulation or when there is no hardware virtualisation support [2]. Böhm et al. proposed a method to introduce binary translation to single-core timing simulation in 2010 [5].

2.3 Multi-core Simulation

Extending single core simulators to handle multiple cores is complicated by the performance implications of the ways in which cores may interact. As cores share caches and memory, simulations of individual cores cannot simply be run independently. For example, accurate modelling of cache coherence, atomic memory operations and inter-processor interrupts (IPIs) must be considered.

Böhm et al.’s modified ARCSim simulator [5] can model single-core processors with high accuracy and reasonable performance; however, Almer et al.’s extension to Böhm et al.’s work [1] that essentially runs multiple copies of the single-core simulator in parallel threads to provide multi-core support is limited in its fidelity. The author comments “detailed behaviour of the shared second-level cache, processor interconnect and external memory of the simulated multi-core platform” cannot be modelled accurately. QEMU is able to exploit multiple cores to emulate a multi-core guest but

provides only a functional simulation mode and supports no timing or modelling of the memory system.

An accurate model of cache coherence and the memory hierarchy requires that multiple cores are simulated in lockstep (or in a way that guarantees equivalent results). Simulators that forego this are unable to properly simulate race conditions and shared resources. Existing cycle-level simulators such as gem5 achieve lockstep by iterating through all simulated cores each cycle. This causes a significant performance drop. Spike (or riscv-isa-sim), on the other hand, switches the active core to simulate less frequently. Its default compilation option only switches the core every 1000 cycles, making it impossible to model race conditions where all cores are trying to acquire a lock simultaneously. No existing binary translated simulators can model multi-core interaction in lockstep, and therefore none of these can model cache coherence or shared second-level cache properly.

3 IMPLEMENTATION

3.1 Overview

The high-level control flow of R2VM, as shown in Figure 1, is similar to other binary translators. When an instruction at a particular PC is to be executed, the code cache is looked up and the cached translated binary is executed directly if found; otherwise, the binary translator is invoked and an entire basic block is fetched, decoded, and translated. We have used a variety of techniques to improve the binary translator performance that are often found in other binary translators, such as block chaining [2].

As full-system simulation is supported, we have to deal with the case that a 4-byte uncompressed instruction spans two pages. We handle this by creating a stub that reads the 2 bytes that lie on the second page each time the stub is executed, and patches the generated code if 2 bytes read are different from that of initial translation.

Cota et al. [9] suggests sharing a code cache between multiple cores to promote code reuse and boost performance. In contrast, we provide each hardware thread its own code cache. This allows different code to be generated for each core, e.g. in the case of heterogeneous cores. This also lessens the synchronisation requirements when modifying the code cache, simplify the implementation.

3.2 Pipeline Simulation

The main difference between our simulator’s flow and existing ones such as QEMU is that we introduce “pipeline model”s, which comprises several hooks. Hooks can process relevant instructions and generate essential microarchitectural simulation code if necessary. The hooks can also indicate the number of cycles it would take for the instruction to complete. It should be noted that this is only for the execution pipeline, while memory systems and cache are in a separate component.

For simple models, such as gem5’s “timing simple” model where each instruction takes 1 cycle to execute, implementation is straightforward as shown in Listing 1.

We have also implemented and validated an in-order pipeline model that accurately models a classic 5-stage pipeline with a static branch predictor. Our implementation captures pipeline hazards, such as data hazards caused by load-use dependency and stalls due

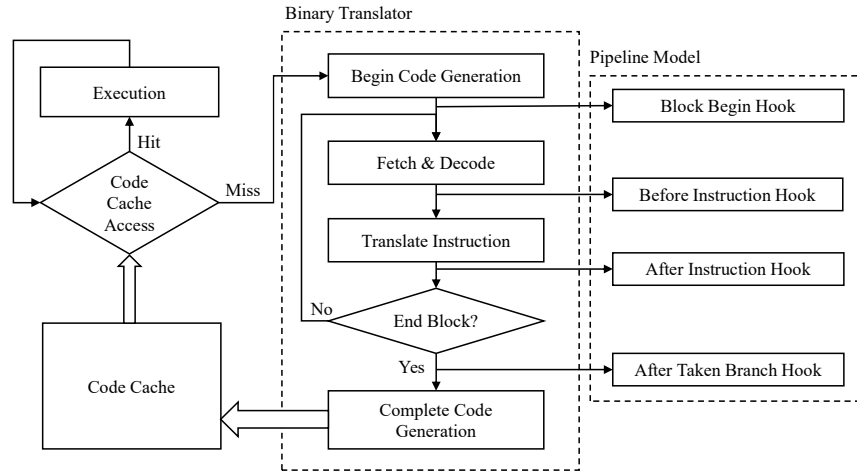


Figure 1: Control flow overview of the simulator

```

1  #[derive(Default)]
2  pub struct SimpleModel;
3
4  impl PipelineModel for SimpleModel {
5      fn after_instruction(&mut self, compiler: &mut
6  ↪ DbtCompiler, _op: &Op, _compressed: bool) {
7          compiler.insert_cycle_count(1);
8      }
9
10     fn after_taken_branch(&mut self, compiler: &mut
11 ↪ DbtCompiler, _op: &Op, _compressed: bool) {
12         compiler.insert_cycle_count(1);
13     }
14 }

```

Listing 1: Timing simple model implementation

to a branch/jump into a misaligned 4-byte instruction. Unlike Böhm et al’s simulator [5] which needs to call a “pipeline” function after each instruction, our implementation models pipeline behaviours during DBT code generation and reflects them as number of cycles taken, therefore requires no explicit code to be executed in runtime.

More complex processors may need either to make an estimation of pipeline states (and sacrifice some accuracy) or generate custom assembly in the hooks to maintain these states during execution (and sacrifice some performance).

3.3 Multi-core Simulation

The techniques we described in the previous section works well for single-core systems. But as described in the background section, running them in parallel or switch between them in a coarse-grained manner has a huge impact on simulation accuracy of multi-threaded programs. The ideal scheduling granularity is, therefore, a cycle, i.e. having all simulated cores run in lockstep. This is, however, difficult to achieve for binary translators.

We experimented the idea of using thread barriers to synchronise multiple threads each simulating a single core. It turns out we could

only synchronise 1 million times per second even after careful optimisation at the assembly level.

3.3.1 Lockstep Simulation. The approach we use takes inspiration from fibers, sometimes also referred to as coroutines or green threads. Fibers are cooperatively scheduled by the user-space application, and they voluntarily “yield” to other fibers, in contrast to traditional threads which are preemptively scheduled by the operating system and are generally heavy-weight constructs. Fibers are often used in I/O heavy, highly concurrent workloads, such as network programming, but this time we borrowed it to our simulator.

In our implementation, we create one fiber for each hardware thread simulated, plus a fiber for the event loop. Each time the pipeline model instructs the DBT to wait for a few cycles, we will generate a number of yields. Listing 2 shows an example of generated code under timing simple model.

```

1  mov    rax, qword [rbp+0x78] ; \
2  mov    qword [rbp+0x70], rax ; | add a4, zero, a5
3  call   fiber_yield_raw     ; /
4  mov    eax, dword [rbp+0x78] ; \
5  add    eax, -0x1           ; |
6  cdqe                      ; | addiw a5, a5, -1
7  mov    qword [rbp+0x78], rax ; |
8  call   fiber_yield_raw     ; /
9  mov    eax, dword [rbp+0x70] ; \
10 imul  eax, dword [rbp+0x50] ; |
11 cdqe                      ; | mulw a0, a4, a0
12 mov    qword [rbp+0x50], rax ; |
13 call   fiber_yield_raw     ; /

```

Listing 2: Example of generated code with yield calls. RBP points to the array of RISC-V registers.

Different from normal fiber implementation, we engineered the fiber’s memory layout to look like Figure 2 to suit the need of a simulator. Each fiber is allocated with a 2M memory aligned to 2M boundary, and the stack for running code under the fiber is contained within this memory range. The alignment requirement allows the fiber’s start address to be recovered from the stack pointer

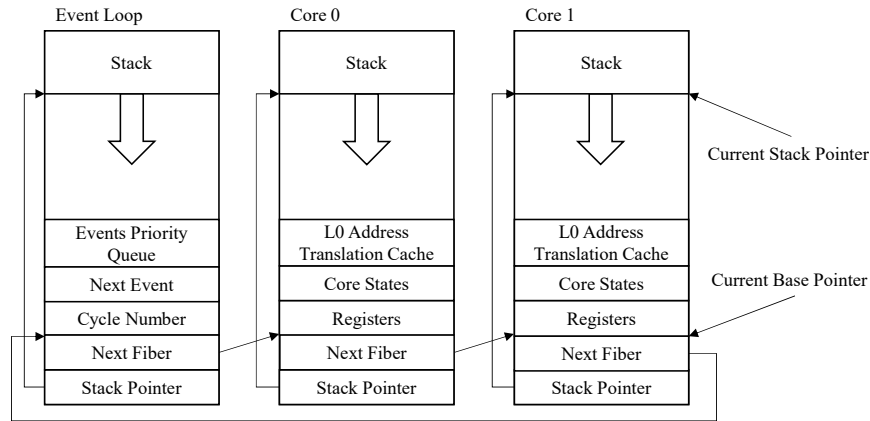


Figure 2: Memory layout of fibers

by simply masking out least significant 21 bits. The base pointer points to the end of fixed fiber structures rather than the beginning, so that positive offsets from the base pointer can be used freely by the DBT-ed code, while the negative offsets are used for fiber management.

The ABI of the host platform for DBT-ed code is not respected; we rather specify all registers other than the base pointer and stack pointer to be volatile, or caller-saved. By doing so, `fiber_yield_raw` does not need to bear the cost of saving any registers. To yield in non-DBT-ed code, we can alternatively push ABI-specified callee-saved registers into the stack and switch.

This careful design makes fiber switching lightning fast; the `fiber_yield_raw` function is as simple as 4 instructions on AMD64, shown in Listing 3.

```

1 fiber_yield_raw:
2     mov [rbp - 32], rsp ; Save current stack pointer
3     mov rbp, [rbp - 16] ; Move to next fiber
4     mov rsp, [rbp - 32] ; Restore stack pointer
5     ret

```

Listing 3: Implementation of the fiber yielding code

3.3.2 Synchronisation Points. Simply yielding a few cycles after every executed instruction will severely limit performance and in many cases will be unnecessary. In practice, we only need to synchronise at points where the execution pipeline can produce visible side-effects to other cores and/or the rest of the system, or where the rest of the system’s behaviour would affect the running pipeline.

We observe that there are three ways that a pipeline interacts with another:

- An memory operation is performed.
- An control register operation is performed. This includes read/write to performance monitor registers, or control registers related to the memory system.
- An interrupt happens.

For the first two types of interaction, we insert a synchronisation point before and after they are executed. For the third case (interrupts), because it is generally difficult to interrupt an DBT-ed code

mid-way, we choose to check for interrupts only at the end of basic blocks. We believe that this decision will not affect the accuracy of our simulation due to the inherent entropy of I/O operations.

The positioning of yielding that lies in between two synchronisation points, therefore, would have no visible side-effects and cannot be distinguished. Our implementation postpones all yielding until the next synchronisation point. We tweaked our yield implementation as shown in Listing 3 slightly to allow multi-cycle yield, and it demonstrates around 10% performance gain compared to naive yielding.

3.4 Memory Simulation

Previous sections described how we simulate each core’s processing pipeline and how we achieve simulation in lockstep. The techniques we described and implemented speeds up pipeline simulation, but the speedup could be very limited when all memory accesses are still simulated. Moreover, the instruction cache and translation-lookaside buffers (TLBs) would also need to be simulated for accuracy.

3.4.1 L0 Data Cache. For memory operations, each running core has its own “L0 data cache”. When a core needs to read from or write to a memory address, it first checks if it is in the L0 data cache. If it hits, then memory access is performed entirely within DBT-ed code, bypassing the memory model entirely.

As a result, the memory model will not intercept all memory accesses. It is therefore important to control what could be in the L0 data cache. We maintain a property that if an access hits the L0 data cache, then it must be a cache hit would the memory access reach the memory model. We speed up TLB simulation with a similar approach in our previous work [10].

In our previous TLB simulation work, the property mandates an invariant that all entries in L0 TLB are in the L1 data TLB. The invariant kept in R2VM is that all L0 data cache entries are contained both in L1 data TLB and L1 data cache. Therefore, as shown in Figure 3, when entries are evicted from either the simulated TLB or cache model, corresponding entries need to be flushed from the L0 data cache for the inclusiveness property.

We carefully engineered the memory layout of L0 data cache entries for maximum efficiency. The L0 data cache is direct-mapped, with each entry representing a cache line. Each entry has a memory

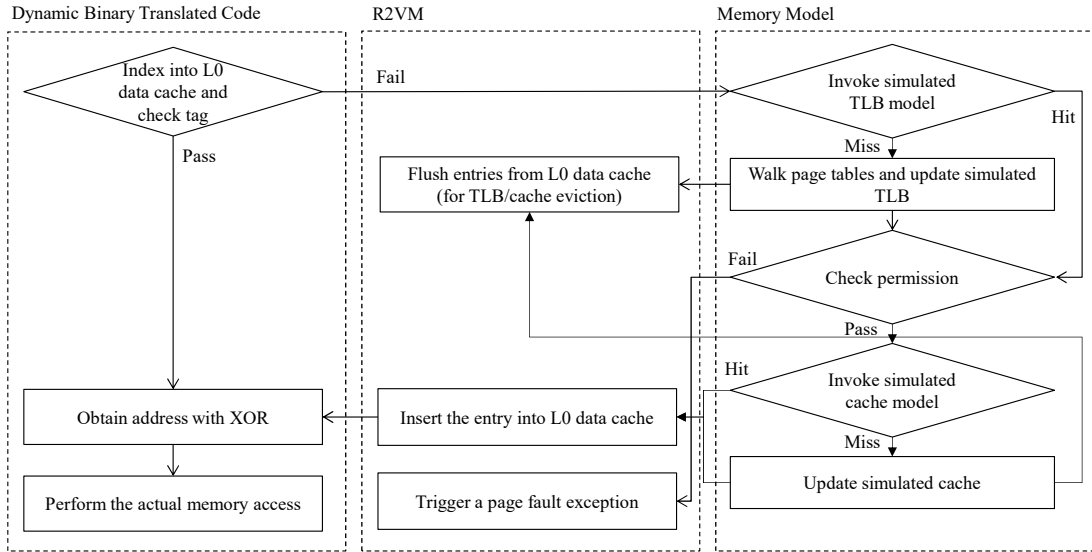


Figure 3: Control flow for memory accesses




Figure 4: Memory layout of a tag entry in L0 data cache

layout like Figure 4. It does not store actual memory contents; it rather stores a translation from the virtual tag to a physical address. In a sense, it is more like a TLB with cache-line granularity than a cache. We pack the XOR-ed value of guest physical address and corresponding guest virtual address, plus a bit indicating if the cache line is read-only to a single machine word.

For each memory access, the L0 data cache is indexed into using the virtual tag. For read access, we compare if $T \gg 1$ is equal to $vtag$. For write access, we compare if $vtag \ll 1$ is equal to T . If the check passes, the requested virtual address is XOR-ed with A to produce the address to access directly within DBT-ed code. If the check fails, the cold path is executed and the memory model is invoked. The memory model will simulate both TLB and data cache, and either triggers a page fault or inserts an entry into the L0 data cache.

The existence of L0 data cache promises the performance of R2VM’s fast-path, because it requires only 3 memory operations for each memory operation simulated. In the default configuration, because the memory model does not intercept all memory accesses, replacement policies such as least-recently used (LRU) cannot be used for the simulated TLB and cache. Generally, we believe this is an acceptable accuracy loss to trade vastly better simulation performance. If LRU-like policies are really needed, the L0 data cache could be bypassed and the memory model be invoked for each memory access, in sacrifice of performance.

3.4.2 L0 Instruction Cache. R2VM also simulates instruction TLB and caches similar to the data cache. Each core also has its own L0 instruction cache, with a simpler entry layout because read/write permission needs not to be distinguished. To keep the overhead of simulating instruction cache down, instead of accessing it each time

an instruction is executed, we instead do it only when a basic block begins, or when the instruction being translated is in a different cache line compared to the previous instruction. For a cache line size of 64 bytes, this means that we only need to generate a single L0 instruction cache access for every 16-32 instructions.

We also creatively use the L0 instruction cache to optimise jumps across pages. Traditionally, because the page mapping might change and therefore the actual target of jump instruction might change, DBTs have to conservatively not link these blocks together. We instead check the L0 instruction cache (which we would need to check anyway when next block begins) and see if the target is the same as the cached target. If so, the cached target is used and the control does not go back to the main loop.

3.4.3 Cache Coherency. Our design for the memory system inherently supports the use of cache coherency. Whenever the cache coherency protocol requires an invalidation, it can be flushed from the L0 data cache of the target core. Because all simulated cores execute in lockstep, and there are synchronisation points before all memory accesses, the effect of the invalidation will be visible before the next memory access.

3.5 Runtime Reconfiguration

R2VM is capable of doing user-level simulation, supervisor-level simulation and machine-level simulation. For user-level simulation, Linux syscalls are emulated, and for supervisor-level, supervisor binary interface (SBI) calls are emulated.

In many cases, we want to gather cache statistics with the behaviour of operating system (OS) taken into account, but we do not want to count the OS booting and workload preparation steps before the region of interest, and do not want to pay for the performance overhead of detailed models for these portions. The design of R2VM takes this into account, and both pipeline and memory models can be switched dynamically in the runtime. The switching is controlled by writing a special control and status register (CSR) in the vendor-specific CSR range.

R2VM supports pipeline model switching by simply flushing the code cache for translated binary, and let the DBT engine to use the new model’s hooks for code generation. Moreover, since as mentioned previously in Section 3.1, each core has its own code cache for DBT-ed code, we allow the pipeline models to be specified per core rather than at once.

The memory model is switched in the runtime by flushing the L0 data cache and the instruction cache. The cache line size is also a runtime-configurable property. For example, if both TLB and cache are simulated, the cache line size can be set to 64 bytes. If only TLB is simulated, the cache line size can be set to 4096 bytes, turning L0 data cache effectively into an L0 data TLB.

If the memory model permits, R2VM can also switch between lockstep execution and parallel execution like other binary translators during the runtime. Parallel execution is enabled on the “atomic” memory model. When paired with the “atomic” pipeline model this behaves functionally equivalent to QEMU and gem5’s atomic model which permits fast-forwarding of aforementioned booting and preparation steps.

4 EVALUATION

As described in Section 3, R2VM offers a range of pipeline models and memory models to select from, and allows switching between them mid-simulation. Each model shows different trade-offs. The list of pre-implemented pipeline and memory models can be found in Table 1 and Table 2.

Name	Description
Atomic	Cycle count not tracked
Simple	Each non-memory instruction takes one cycle
InOrder	Models a simple 5-stage in-order scalar pipeline

Table 1: List of pre-implemented pipeline models

Name	Description
Atomic	Memory accesses not tracked
TLB	TLB hit rate collected; cache not simulated
Cache	Cache hit rate collected; TLB and cache coherency not modelled; parallel execution allowed
MESI	A directory-based MESI cache coherency protocol with a shared L2. Lockstep execution required.

Table 2: List of pre-implemented memory models

4.1 Accuracy and Validation

For pipeline models, we validated the accuracy of the in-order model against an actual RTL implementation of a RISC-V core using CoreMark [7]. CoreMark is particularly helpful for this validation, as CoreMark’s working set is small enough to fit into caches and therefore the memory system of the RTL implementation would not affect the benchmark result. In our run, the RTL implementation reports 2.10 CoreMark/MHz where the in-order model, when paired with the atomic memory model, reports 2.09 CoreMark/MHz. The difference is less than 1%. The “simple” model is simply validated by checking that all cores have their MCYCLE and MINSTRET CSR equal.

For memory models, we used a few micro-benchmarks to cover the use case for each model. For TLB and cache simulation, we

used a single-core micro-benchmark that is similar to the MemLat tool from the 7-zip LZMA benchmark [14]. For the MESI cache-coherency model, we used a micro-benchmark to simulate a scenario where two cores are heavily contending over a shared spin-lock. The memory model under test is used together with the validated in-order pipeline model, and we compare the number of cycles taken to execute a benchmark in R2VM and in RTL simulation. The error is around for the 10% for the cache coherency model and lower for non-coherent models. Though not as accurate as the pipeline model, we believe at this accuracy the simulation can provide representative-enough metrics for exploring design decisions.

4.2 Performance

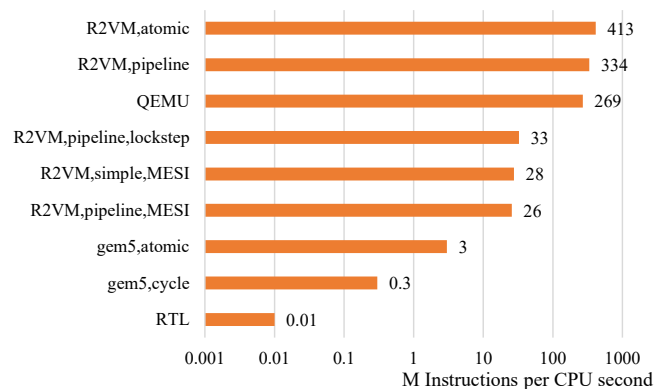


Figure 5: Performance comparison between models and other simulators

We evaluated the performance of R2VM against QEMU using the deduplication workload from PARSEC [3] on 4 cores to test the integer performance of the simulator (as both R2VM and QEMU interprets floating-point operations). The kIPS numbers of the gem5 simulator are from Saidi et al.’s presentation [15].

As shown in Figure 5, the techniques we use lead to superb performance. When caches are not simulated and therefore cores can run in parallel threads, R2VM runs at >300 MIPS per core, even outperforming QEMU. Lockstep execution brings down performance by 10x to ~30 MIPS (for 4 simulated cores in a single-threaded), but this is still significantly faster than gem5.

Thanks to our pipeline model design which moves most simulation to DBT compilation time rather than runtime, and to our memory model design which offloads most memory accesses by using L0 caches, simulating pipelines and cache coherency protocols did not add a significant overhead themselves, compared to the overhead of lockstep execution.

5 CONCLUSION

We have introduced R2VM, a multi-purpose binary translating simulator that is able to simulate multi-core RISC-V systems at the cycle-level at high-speed. This is done by leveraging the use of fibers to support fast lockstep execution. Overall, optimisations made R2VM possible to achieve functional simulation performance that exceeds that of QEMU and cycle-level simulation nearly 100 times faster than gem5.

REFERENCES

- [1] Oscar Almer, Igor Böhm, Tobias Edler Von Koch, Björn Franke, Stephen Kyle, Volker Seeker, Christopher Thompson, and Nigel Topham. 2011. Scalable multi-core simulation using parallel dynamic binary translation. In *2011 International Conference on Embedded Computer Systems: Architectures, Modeling and Simulation*. IEEE, 190–199.
- [2] Fabrice Bellard. 2005. QEMU, a fast and portable dynamic translator.. In *USENIX Annual Technical Conference, FREENIX Track*, Vol. 41. 46.
- [3] Christian Biemia, Sanjeev Kumar, Jaswinder Pal Singh, and Kai Li. 2008. The PARSEC benchmark suite: Characterization and architectural implications. In *Proceedings of the 17th international conference on Parallel architectures and compilation techniques*. ACM, 72–81.
- [4] Nathan Binkert, Bradford Beckmann, Gabriel Black, Steven K Reinhardt, Ali Saidi, Arkaprava Basu, Joel Hestness, Derek R Hower, Tushar Krishna, Somayeh Sardashti, et al. 2011. The gem5 simulator. *ACM SIGARCH computer architecture news* 39, 2 (2011), 1–7.
- [5] Igor Böhm, Björn Franke, and Nigel Topham. 2010. Cycle-accurate performance modelling in an ultra-fast just-in-time dynamic binary translation instruction set simulator. In *2010 International Conference on Embedded Computer Systems: Architectures, Modeling and Simulation*. IEEE, 1–10.
- [6] Hadi Brais, Rajshekar Kalayappan, and Preeti Ranjan Panda. 2020. A Survey of Cache Simulators. *ACM Computing Surveys (CSUR)* 53, 1 (2020), 1–32.
- [7] The Embedded Microprocessor Benchmark Consortium. 2020. CoreMark. <https://www.eembc.org/coremark/>. Accessed: 2020-04-14.
- [8] The Standard Performance Evaluation Corporation. 2017. SPEC CPU® 2017. <https://www.spec.org/cpu2017/>. Accessed: 2020-04-14.
- [9] Emilio G Cota and Luca P Carloni. 2019. Cross-ISA machine instrumentation using fast and scalable dynamic binary translation. In *Proceedings of the 15th ACM SIGPLAN/SIGOPS International Conference on Virtual Execution Environments*. 74–87.
- [10] Xuan Guo and Robert Mullins. 2019. Fast TLB Simulation for RISC-V Systems. In *Third Workshop on Computer Architecture Research with RISC-V*.
- [11] Kim Hazelwood. 2011. Dynamic binary modification: Tools, techniques, and applications. *Synthesis Lectures on Computer Architecture* 6, 2 (2011), 1–81.
- [12] Ankur Limaye and Tosiron Adegbiya. 2018. A workload characterization of the SPEC CPU2017 benchmark suite. In *2018 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*. IEEE, 149–158.
- [13] Nicholas Nethercote and Julian Seward. 2007. Valgrind: a framework for heavy-weight dynamic binary instrumentation. *ACM Sigplan notices* 42, 6 (2007), 89–100.
- [14] Igor Pavlov. 2019. 7-Zip LZMA Benchmark. <https://www.7-cpu.com/>. Accessed: 2020-04-14.
- [15] Ali Saidi and Andreas Sandberg. [n. d.]. gem5 Virtual Machine Acceleration. http://www.m5sim.org/wiki/images/c/c3/2012_12_gem5_workshop_kvm.pdf.
- [16] Tuan Ta, Lin Cheng, and Christopher Batten. 2018. Simulating Multi-Core RISC-V Systems in gem5. In *Workshop on Computer Architecture Research with RISC-V*.
- [17] The Rust Team. 2020. Rust Programming Language. <https://www.rust-lang.org/>. Accessed: 2020-04-14.
- [18] Tran Van Dung, Ittetsu Taniguchi, and Hiroyuki Tomiyama. 2014. Cache simulation for instruction set simulator QEMU. In *2014 IEEE 12th International Conference on Dependable, Autonomic and Secure Computing*. IEEE, 441–446.