

GEM-5 Technical Support Center: Troubleshooting Long Simulation Times

Author: BenchChem Technical Support Team. **Date:** December 2025

Compound of Interest

Compound Name: GEM-5

Cat. No.: B12410503

[Get Quote](#)

This guide provides troubleshooting steps and frequently asked questions to help researchers, scientists, and drug development professionals address long simulation times in their **GEM-5** experiments.

Frequently Asked Questions (FAQs)

Q1: My GEM-5 simulation is running very slowly. What are the common causes?

Several factors can contribute to slow **GEM-5** simulations. The most common culprits include:

- **High-Fidelity CPU Models:** Detailed CPU models like O3CPU provide cycle-accurate results but are computationally intensive.
- **Complex Memory System:** The Ruby memory model, while highly flexible and detailed, can be slower than the classic memory system.^{[1][2]}
- **Full System (FS) Mode:** Simulating a full operating system adds significant overhead compared to Syscall Emulation (SE) mode.^[3]
- **Large and Complex Workloads:** The nature of the application being simulated directly impacts the simulation time.
- **Host System Performance:** The CPU speed and cache size of the machine running **GEM-5** can be a bottleneck.^[4]

- Single-Threaded Execution: By default, **GEM-5** is a single-threaded simulator, which may not fully utilize modern multi-core processors.[\[5\]](#)

Q2: How can I speed up my simulation without sacrificing too much accuracy?

There are several techniques to accelerate **GEM-5** simulations, often involving a trade-off between speed and detail. Here are some effective strategies:

- Use Checkpoints: Run the simulation to a region of interest (ROI) and then create a checkpoint. Subsequent simulations can start directly from this checkpoint, skipping the often lengthy initialization and setup phases.[\[5\]](#)[\[6\]](#)
- Fast-Forwarding: Use a simpler, faster CPU model (e.g., AtomicSimpleCPU or TimingSimpleCPU) to quickly reach the ROI before switching to a more detailed model like O3CPU.[\[5\]](#)[\[6\]](#)
- KVM Acceleration: If the host and simulated machine share the same instruction set architecture (ISA), you can use the KVM-based CPU model (KvmCPU) for near-native execution speed to fast-forward to the ROI.[\[5\]](#)[\[7\]](#)[\[8\]](#)
- Sampling: Instead of simulating an entire workload, you can simulate representative portions. Techniques like SimPoints and LoopPoints help identify these representative simulation points.[\[2\]](#)[\[6\]](#)[\[9\]](#)
- Parallel Simulation: For multi-core system simulations, consider using parallel versions of **GEM-5** like parti-gem5, which can distribute the simulation across multiple host cores for significant speedups.[\[10\]](#)[\[11\]](#)

Q3: When should I use Full System (FS) mode versus Syscall Emulation (SE) mode?

The choice between FS and SE mode depends on the specific requirements of your experiment.

- Full System (FS) Mode: This mode simulates a complete system, including an operating system. It is necessary when the interaction between the workload and the OS is important

for the research. While more accurate, it is also significantly slower.[3][8][12]

- Syscall Emulation (SE) Mode: This mode is faster as it emulates system calls without booting a full OS. It is suitable for applications that do not have complex OS interactions.[3][8][12] It is often recommended to try SE mode first; if it works for your benchmark, it can save considerable time.[3]

Q4: How does the choice of memory system affect simulation speed?

GEM-5 offers two primary memory system models: Classic and Ruby.

- Classic Memory System: This model is generally faster and easier to configure. It's a good choice when you don't need to model a detailed, custom cache coherence protocol.[1][2]
- Ruby Memory System: Ruby provides a highly detailed and flexible memory hierarchy and is essential for accurately modeling various cache coherence protocols.[1][13] This detail comes at the cost of simulation speed.[1] Fast-forwarding is not supported when using the Ruby memory model.[2]

Troubleshooting Guides

Guide 1: My simulation is taking too long to boot the operating system.

Problem: The initial OS boot process in Full System mode is a major contributor to long simulation times.

Solution:

- Use KVM for Booting: If your host and guest systems have the same ISA (e.g., x86 on an x86 host), use the KvmCPU to boot the OS at near-native speed.[14]
- Create a Boot Checkpoint: Once the OS has booted and is idle, create a checkpoint. You can then restore from this checkpoint for all subsequent simulation runs, completely bypassing the boot process.

Experimental Protocol for Creating a Boot Checkpoint:

- Configure your **GEM-5** simulation script to use the KvmCPU.
- Run the simulation to allow the operating system to fully boot.
- Once the OS is at the login prompt or idle state, insert an m5 exit command into your script to terminate the simulation at this point.
- Before the exit, include a command to create a checkpoint.
- For subsequent runs, modify your script to restore from this checkpoint and switch to a more detailed CPU model for the workload execution.

Guide 2: How do I identify the performance bottlenecks in my **GEM-5** simulation itself?

Problem: It's unclear which part of the **GEM-5** configuration is causing the primary slowdown.

Solution: Profiling the **GEM-5** simulation can reveal performance bottlenecks. A recent study highlighted that the host machine's L1 cache size significantly impacts simulation speed.[\[4\]](#)[\[15\]](#)
[\[16\]](#)

Methodology for Performance Profiling:

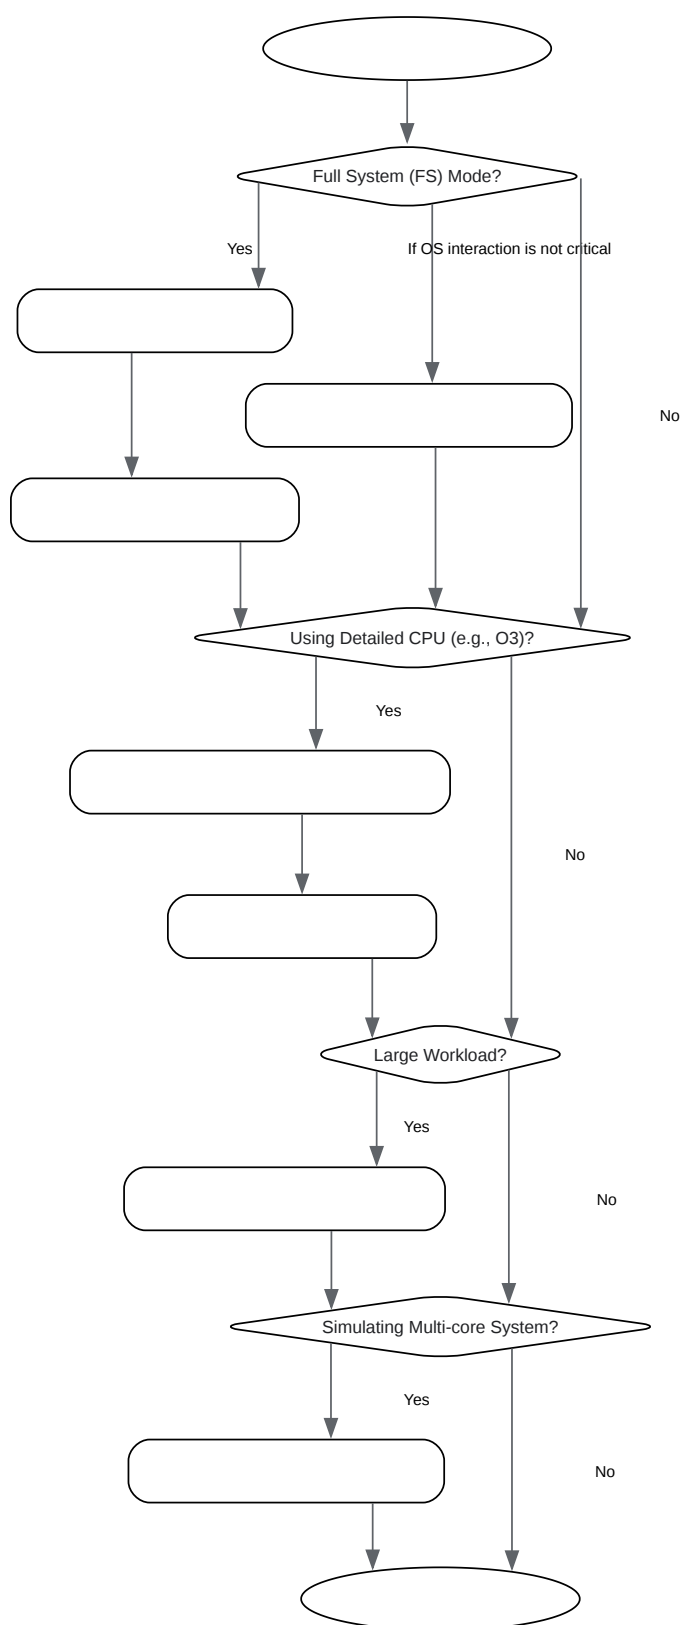
- Vary Host Machine Configurations: If possible, run the same **GEM-5** simulation on different host machines with varying CPU architectures and cache sizes to observe the performance impact. For instance, a MacBook Pro with an M1 chip has been shown to complete simulations 1.7x to 3.02x faster than a server with Xeon Gold CPUs.[\[4\]](#)
- Analyze **GEM-5** Statistics: Use **GEM-5**'s built-in statistics to understand the behavior of the simulated system. High cache miss rates or other performance counters can indicate areas for optimization in the simulated architecture.
- Compile **GEM-5** with Optimization Flags: Ensure you are using an optimized build of **GEM-5** (e.g., gem5.opt or gem5.fast). Compiling with the -O3 flag can provide a modest speedup.[\[4\]](#)
[\[12\]](#)

Quantitative Data Summary

Parameter	Impact on Simulation Speed	Key Findings
Host L1 Cache Size	Significant	Increasing the L1 data and instruction cache size from 8KB to 32KB on a RISC-V core improved GEM-5's simulation speed by 31% to 61%. [4] [15] [17]
Parallel Simulation (part-gem5)	High	Achieved speedups of up to 42.7x when simulating a 120-core ARM MPSoC on a 64-core x86-64 host system. [10] [11]
Host CPU Architecture	High	An Apple M1 chip can be 1.7x to 3.02x faster for GEM-5 simulations compared to an Intel Xeon Gold 6242R CPU. [4]
Host CPU Frequency	Linear	Reducing the CPU frequency of the host machine from 3.1GHz to 1.2GHz increased the simulation time by 2.67x. [4]

Visualizations

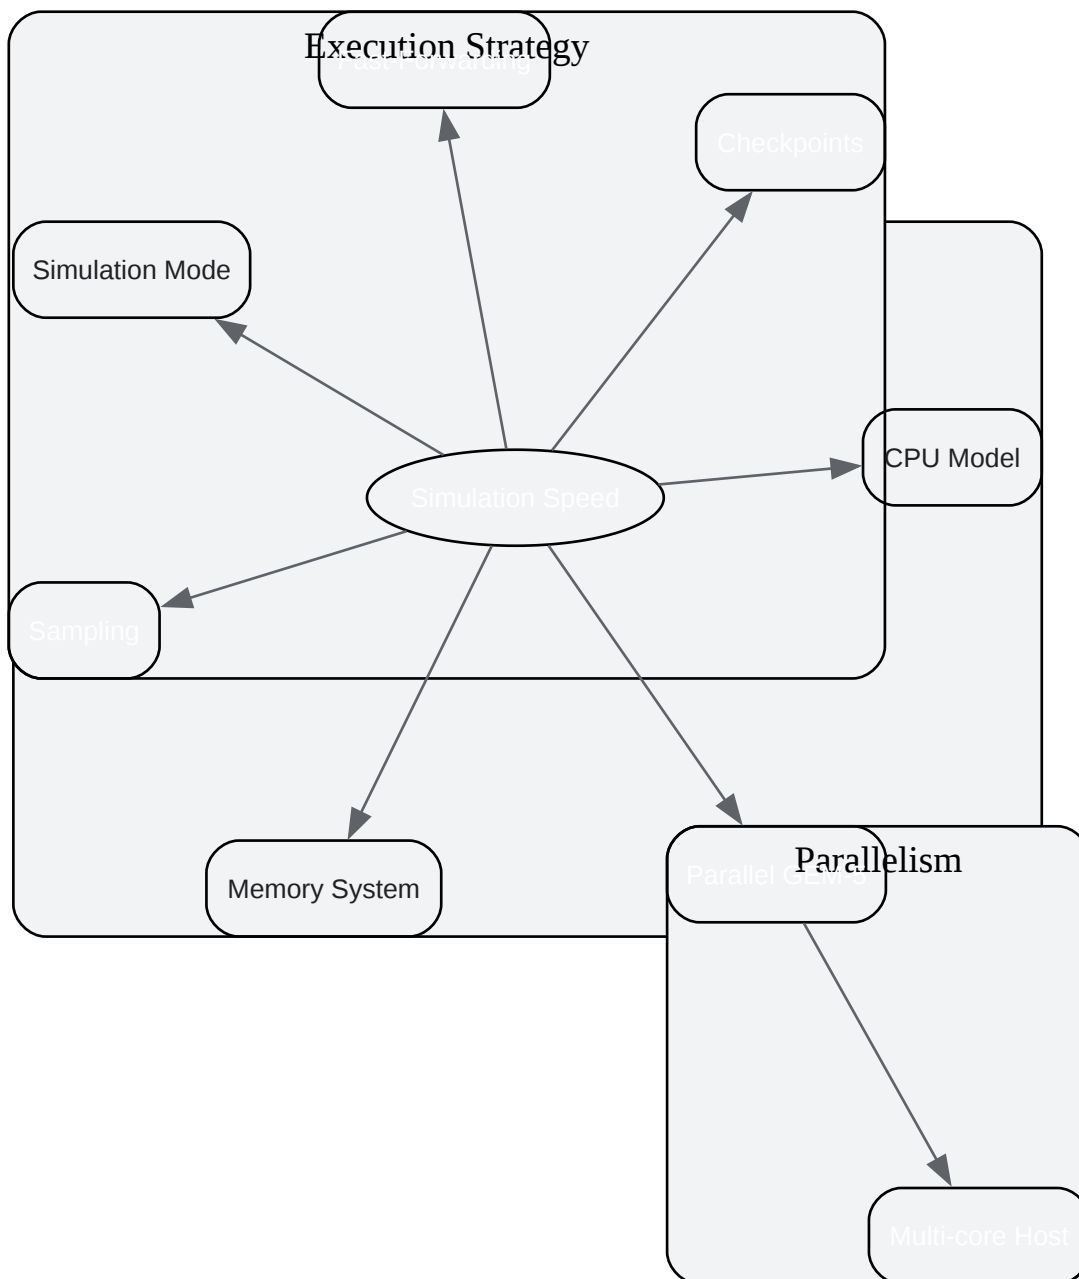
Workflow for Accelerating GEM-5 Simulations



[Click to download full resolution via product page](#)

Caption: A decision workflow for troubleshooting and accelerating slow **GEM-5** simulations.

Logical Relationship of GEM-5 Speed Optimization Techniques



[Click to download full resolution via product page](#)

Need Custom Synthesis?

BenchChem offers custom synthesis for rare earth carbides and specific isotopic labeling.

Email: info@benchchem.com or [Request Quote Online](#).

References

- 1. General Memory System - gem5 [old.gem5.org]
- 2. gem5 simulation time is high - Stack Overflow [stackoverflow.com]
- 3. When to use full system FS vs syscall emulation SE with userland programs in gem5? - Stack Overflow [stackoverflow.com]
- 4. ws.engr.illinois.edu [ws.engr.illinois.edu]
- 5. [gem5-users] Running gem5 simulation faster on multiple host CPU? [gem5-users.gem5.narkive.com]
- 6. m.youtube.com [m.youtube.com]
- 7. gem5: Setting Up and Using KVM on your machine [gem5.org]
- 8. arxiv.org [arxiv.org]
- 9. youtube.com [youtube.com]
- 10. parti-gem5: gem5's Timing Mode Parallelised [arxiv.org]
- 11. [2308.09445] parti-gem5: gem5's Timing Mode Parallelised [arxiv.org]
- 12. developer.arm.com [developer.arm.com]
- 13. gem5: Introduction to Ruby [gem5.org]
- 14. gem5: X86 Full-System Tutorial [gem5.org]
- 15. Optimizing gem5 Simulator Performance: Profiling Insights and Userspace Networking Enhancements | I2S | Institute for Information Sciences [i2s-research.ku.edu]
- 16. ieeexplore.ieee.org [ieeexplore.ieee.org]
- 17. Optimizing gem5 Simulator Performance: Profiling Insights and Userspace Networking Enhancements | Electrical Engineering and Computer Science [eecs.ku.edu]
- To cite this document: BenchChem. [GEM-5 Technical Support Center: Troubleshooting Long Simulation Times]. BenchChem, [2025]. [Online PDF]. Available at: [https://www.benchchem.com/product/b12410503#what-to-do-when-a-gem-5-simulation-is-taking-too-long]

Disclaimer & Data Validity:

The information provided in this document is for Research Use Only (RUO) and is strictly not intended for diagnostic or therapeutic procedures. While BenchChem strives to provide

accurate protocols, we make no warranties, express or implied, regarding the fitness of this product for every specific experimental setup.

Technical Support: The protocols provided are for reference purposes. Unsure if this reagent suits your experiment? [[Contact our Ph.D. Support Team for a compatibility check](#)]

Need Industrial/Bulk Grade? [Request Custom Synthesis Quote](#)

BenchChem

Our mission is to be the trusted global source of essential and advanced chemicals, empowering scientists and researchers to drive progress in science and industry.

Contact

Address: 3281 E Guasti Rd
Ontario, CA 91761, United States
Phone: (601) 213-4426
Email: info@benchchem.com