

getting started with GEM-5 for computer architecture research

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

Compound of Interest

Compound Name: GEM-5

Cat. No.: B12410503

[Get Quote](#)

An In-depth Technical Guide for Researchers

Getting Started with gem5 for Computer Architecture Research

This guide serves as a comprehensive introduction to the gem5 simulator, a powerful and modular platform for computer architecture research. It is designed for researchers and scientists who are new to gem5 and aims to provide a foundational understanding of its core concepts, setup, and basic operation.

Introduction to the gem5 Simulator

The gem5 simulator is a modular, open-source platform for computer system architecture research, covering everything from system-level architecture to processor microarchitecture.^[1]^[2] It is a discrete-event simulator, meaning it models the passage of time as a series of distinct events.^[3] gem5 is highly flexible, allowing researchers to configure, extend, or replace its components to suit their specific research needs.^[3] The simulator is primarily written in C++ and Python, with simulation configurations being handled by Python scripts.^[3]

Key features of gem5 include:

- **Multiple ISAs Support:** gem5 supports a variety of Instruction Set Architectures (ISAs), including x86, ARM, RISC-V, SPARC, and others, allowing for diverse and cross-architecture studies.^[2]^[4]^[5]

- **Interchangeable CPU Models:** It provides several CPU models with varying levels of detail, such as simple functional models for speed and detailed out-of-order models for accuracy.[\[2\]](#)[\[6\]](#)[\[7\]](#)
- **Detailed Memory System:** gem5 includes a flexible, event-driven memory system that can model complex, multi-level cache hierarchies and various DRAM controllers.[\[2\]](#)[\[8\]](#)
- **Multiple Simulation Modes:** gem5 can operate in two primary modes: Syscall Emulation (SE) and Full System (FS).[\[9\]](#)[\[10\]](#)

Simulation Modes: SE vs. FS

gem5 offers two main simulation modes that cater to different research needs.[\[9\]](#)

- **Syscall Emulation (SE) Mode:** This mode focuses on simulating the CPU and memory system for a single user-space application.[\[9\]](#)[\[10\]](#) It relies on the host operating system to handle system calls, which simplifies the simulation setup.[\[10\]](#)[\[11\]](#) SE mode is ideal for studies where the detailed interaction with the operating system is not critical.
- **Full System (FS) Mode:** In this mode, gem5 emulates a complete hardware system, allowing an unmodified operating system and its applications to run on the simulated hardware.[\[6\]](#)[\[9\]](#) This mode is akin to a virtual machine and is essential for research involving OS interactions, device drivers, and complex software stacks.[\[10\]](#)

Getting Started: Installation and Setup

This section provides a detailed protocol for downloading and compiling gem5 on a Unix-like operating system.

Experimental Protocol: gem5 Installation

Objective: To download the gem5 source code and compile the simulator binary.

Prerequisites: A Unix-like operating system (Linux is recommended) with necessary dependencies installed.[\[12\]](#)

Dependencies: Before compiling gem5, you need to install several packages. Key dependencies include:

- git: For cloning the source code repository.
- scon: The build system used by gem5.
- g++ or another C++ compiler.
- python-dev: Python development headers.
- swig: A software development tool that connects C/C++ programs with high-level programming languages.
- Other libraries such as zlib1g-dev and libprotobuf-dev.[13][14]

Procedure:

- Clone the gem5 Repository: Download the source code from the official gem5 GitHub repository. It is recommended to use the latest stable branch for research.[12]
- Compile gem5: Use scon to build the simulator. The build process can be time-consuming and memory-intensive.[12] The command specifies the target ISA and the desired optimization level. The -j flag specifies the number of parallel compilation jobs.[11]
 - As of gem5 v24.1, the ALL build includes all ISAs and Ruby protocols.[15]
 - For a machine with 8 cores, you might use -j 9.[16]
- Verification: Upon successful compilation, a gem5 binary will be created in the build/ALL/ directory (e.g., build/ALL/gem5.opt).[12]

The gem5 Architecture: Core Concepts

gem5's modularity is built upon a few fundamental concepts that are crucial for users to understand.

SimObjects

The core of gem5's modular design is the SimObject.[9] Most simulated components, such as CPUs, caches, memory controllers, and buses, are implemented as SimObjects.[10][17] These


C++ objects are exported to Python, allowing researchers to instantiate and configure them within a Python script to define the simulated system's architecture.[9][17]

The gem5 Standard Library

To simplify the process of creating simulation scripts, gem5 provides a standard library (stdlib). [17][18] This library offers a collection of pre-defined, high-level components that can be easily combined to build a simulated system.[17][19] The philosophy behind the stdlib is analogous to building a real computer from off-the-shelf parts.[18] It abstracts away much of the low-level configuration, reducing boilerplate code and the potential for errors.[19]

The main components of the standard library are:

- Board: The backbone of the system where other components are connected.[19]
- Processor: Contains one or more CPU cores.[18]
- MemorySystem: Defines the main memory, such as a DDR3 or DDR4 system.[18]
- CacheHierarchy: Defines the components between the processor and main memory, such as L1 and L2 caches.[18]



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

gem5 Standard Library component relationships.

Your First Simulation: A "Hello World" Experiment

Running a simulation in gem5 involves executing the compiled binary with a Python configuration script as an argument.[\[12\]](#) This protocol details how to run a basic "Hello World" example in SE mode using the gem5 standard library.

Experimental Protocol: Running a "Hello World" Simulation

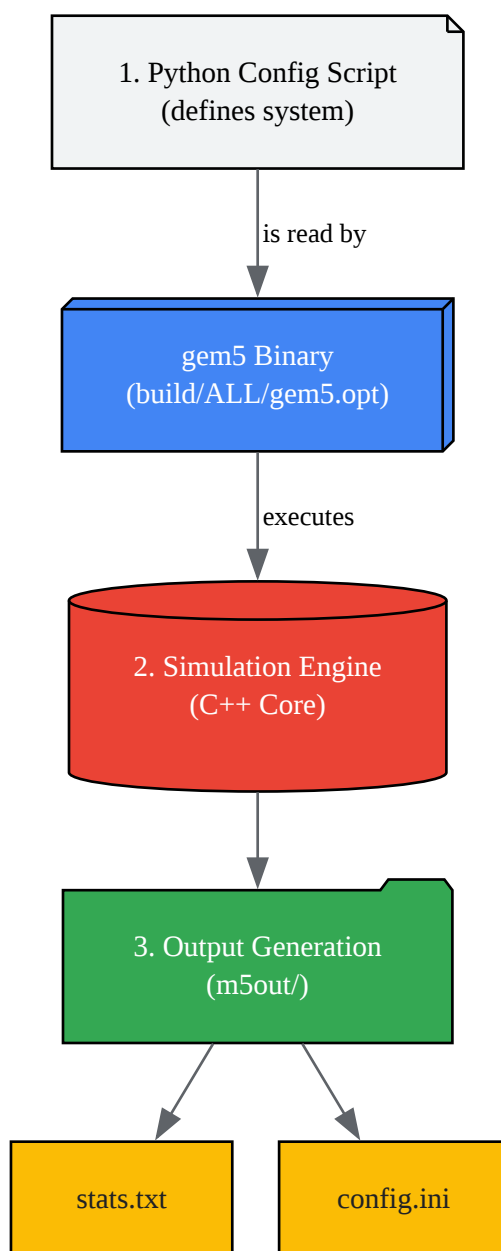
Objective: To run a pre-compiled "Hello World" binary on a simple simulated system in SE mode.

Procedure:

- **Create a Configuration Script:** Create a Python file (e.g., `hello.py`) to define the simulated system. This script will use components from the gem5 standard library.[\[15\]](#)
- **Run the Simulation:** Execute the gem5 binary with your Python script.
- **Observe the Output:** The simulation will run, and you should see "Hello world!" printed to your terminal, which is the output from the simulated binary.[\[12\]](#) An output directory named `m5out` will also be created.[\[20\]](#)

Simulation Workflow Diagram

The following diagram illustrates the high-level workflow of a gem5 simulation.



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

High-level gem5 simulation workflow.

Analyzing the Output

After a simulation completes, gem5 generates an output directory, typically named m5out, which contains detailed information about the simulation run.^[20]

The key files in the m5out directory are:

- `config.ini` / `config.json`: These files contain a complete record of every `SimObject` created for the simulation and all of its parameters, including those set by default.[\[20\]](#) This is crucial for ensuring reproducibility.
- `stats.txt`: This file contains a dump of all the statistics collected during the simulation.[\[20\]](#) Statistics are registered by `SimObjects` and provide detailed insights into the behavior and performance of the simulated system.

Data Presentation: Key Simulation Statistics

The `stats.txt` file provides a wealth of quantitative data. Below is a table summarizing some of the most important high-level statistics.

Statistic Name	Description	Example Use Case
<code>sim_seconds</code>	The total simulated time that has passed. [20]	Calculating simulated performance.
<code>sim_insts</code>	The total number of instructions committed by the CPU(s). [20]	Measuring workload progress.
<code>host_inst_rate</code>	The simulation speed in terms of host instructions per second. [20]	Assessing the performance of the simulator itself.
<code>system.cpu.ipc</code>	Instructions Per Cycle for the CPU.	Core performance analysis.
<code>system.cpu.dcache.miss_rate</code>	The miss rate of the L1 data cache.	Memory system performance analysis.
<code>system.mem_ctrl.bw_total</code>	Total memory bandwidth utilized.	Analyzing memory system bottlenecks.

Building a System: CPU, Cache, and Memory

To conduct meaningful research, you will need to move beyond the simplest configurations and build systems with more detailed components, such as caches. The modular nature of `gem5`

and its standard library makes this straightforward.

Experimental Protocol: Simulating a System with Caches

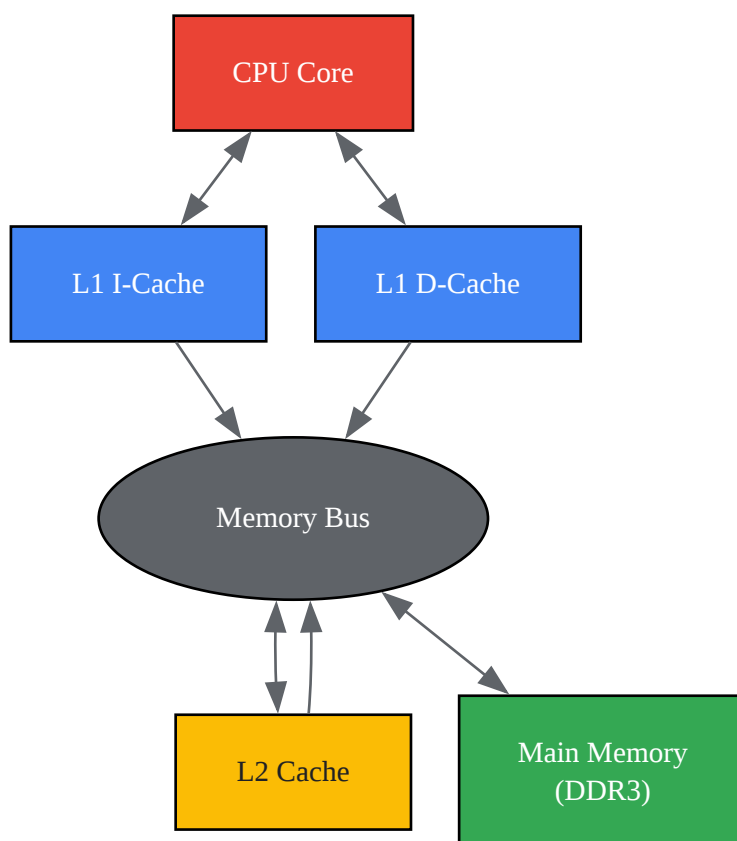
Objective: To configure and simulate a simple system consisting of a CPU, L1 instruction and data caches, an L2 cache, and a memory bus.

Procedure:

- **Modify the Configuration Script:** Start with the "Hello World" script and replace the NoCache hierarchy with a classic cache hierarchy like PrivateL1PrivateL2CacheHierarchy.
- **Run and Analyze:** Run the simulation as before. The new configuration will be reflected in m5out/config.ini. You can now analyze cache-related statistics (e.g., miss rates, hit latency) in m5out/stats.txt to understand the performance of your memory system.

System Component Connection Diagram

The following diagram illustrates the logical connections in the simple cached system you configured.



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

Logical connections in a simple cached system.

Summary of Quantitative Data

For quick reference, the following tables summarize key quantitative and categorical data about gem5's capabilities.

Table 1: Supported Instruction Set Architectures (ISAs)

ISA	Support Status
X86	Supports 64-bit extensions; can boot unmodified Linux kernels.[5]
ARM	Supports ARMv8-A profile (AArch32 and AArch64); can boot Linux.[5]
RISC-V	Support for privileged ISA spec is a work in progress.[2]
SPARC	Models a single core of an UltraSPARC T1; can boot Solaris.[5]
MIPS	Supported.[2]
Alpha	Models a DEC Tsunami system; can boot Linux 2.4/2.6 and FreeBSD.[5]
POWER	Limited to syscall emulation mode based on POWER ISA v3.0B.[5]

Table 2: gem5 CPU Models

CPU Model	Type	Key Characteristics
AtomicSimpleCPU	Functional	Uses atomic memory accesses for speed; not cycle-accurate for memory.[21]
TimingSimpleCPU	In-Order	Uses timing-based memory accesses; stalls on cache misses.[21]
O3CPU	Out-of-Order	A detailed, cycle-accurate model of an out-of-order processor.[7]
MinorCPU	In-Order	A more realistic in-order CPU model with a fixed pipeline.[6]
KVMCPU	KVM-based	Uses virtualization to accelerate simulation, especially for non-interesting code regions.[2]

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. gem5: The gem5 simulator system [gem5.org]
- 2. gem5: About [gem5.org]
- 3. gem5: Learning gem5 [gem5.org]
- 4. Vayavya Labs Pvt. Ltd. - Introducing gem5 : An Open-Source Computer Architecture Simulator [vayavyalabs.com]
- 5. gem5: Architecture Support [gem5.org]
- 6. developer.arm.com [developer.arm.com]

- 7. gem5: gem5's CPU models [gem5.org]
- 8. gem5: gem5_memory_syste [gem5.org]
- 9. gem5: Creating a simple configuration script [gem5.org]
- 10. Creating a simple configuration script — gem5 Tutorial 0.1 documentation [courses.grainger.illinois.edu]
- 11. cs.pomona.edu [cs.pomona.edu]
- 12. gem5: Getting Started with gem5 [gem5.org]
- 13. eshita_omar.gitbooks.io [eshita_omar.gitbooks.io]
- 14. Gem5 and GS Gem5-Validate Tutorial [web-archive.southampton.ac.uk]
- 15. gem5: Hello World Tutorial [gem5.org]
- 16. Instruction Set Assignmnet: T · GitHub [gist.github.com]
- 17. gem5: Creating a simple configuration script [courses.grainger.illinois.edu]
- 18. gem5: Standard Library Overview [gem5.org]
- 19. scribd.com [scribd.com]
- 20. gem5: Understanding gem5 statistics and output [gem5.org]
- 21. gem5: Simple CPU Models [gem5.org]
- To cite this document: BenchChem. [getting started with GEM-5 for computer architecture research]. BenchChem, [2025]. [Online PDF]. Available at: [https://www.benchchem.com/product/b12410503#getting-started-with-gem-5-for-computer-architecture-research]

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