

# debugging and verifying custom memory models in GEM-5

**Author:** BenchChem Technical Support Team. **Date:** March 2026

## Compound of Interest

Compound Name: GEM-5  
Cat. No.: B12410503

[Get Quote](#)

## GEM-5 Custom Memory Model: Technical Support Center

This technical support center provides troubleshooting guides and frequently asked questions (FAQs) to assist researchers, scientists, and drug development professionals in debugging and verifying custom memory models within the **GEM-5** simulator.

### Frequently Asked Questions (FAQs)

Q1: What are the first steps I should take when my custom memory model is not behaving as expected?

A1: Start by enabling **GEM-5's** powerful debug tracing capabilities. Use the `--debug-flags` command-line option with relevant flags to get detailed execution traces. For memory-specific issues, the `DRAM` and `MemoryAccess` flags are a good starting point. If you are using the Ruby memory system, `ProtocolTrace` is invaluable for observing the coherence protocol transitions.

[\[1\]](#)[\[2\]](#)[\[3\]](#)

Q2: My simulation is terminating with a "fatal" error. How can I pinpoint the cause?

A2: A "fatal" error in **GEM-5** typically indicates a configuration issue or a critical state violation that the simulator cannot recover from.[4] The error message itself is the first clue, as it often points to the C++ file and line number where the error was triggered.[4] Common causes include unconnected ports in your memory system configuration or invalid parameter values being passed to your memory model.[4] Carefully review your Python configuration scripts and the C++ implementation of your custom model.

Q3: What is the difference between the "Classic" and "Ruby" memory systems in **GEM-5**, and how does this affect debugging?

A3: The "Classic" memory system is a simpler, faster model primarily focused on basic memory hierarchy simulation.[5][6] Debugging here often involves flags like Cache and Bus. Ruby, on the other hand, is a highly detailed and flexible memory system simulator designed for modeling complex cache coherence protocols.[5][6][7] Debugging custom models in Ruby requires a deeper understanding of its components (Sequencers, Controllers, SLICC) and utilizing Ruby-specific debug flags such as ProtocolTrace, RubyNetwork, and RubyGenerated.  
[3]

Q4: How can I verify the functional correctness and performance of my custom memory model?

A4: Verification should be a multi-step process.

- Unit Testing: Develop targeted tests that exercise specific functionalities of your model in isolation.
- Synthetic Traffic Generation: Use **GEM-5**'s traffic generators, like PyTrafficGen, to create controlled memory access patterns (e.g., sequential, random) and measure key performance metrics like bandwidth and latency under specific loads.[8][9]
- Comparative Analysis: Compare the performance of your model against established models in **GEM-5** or other validated simulators like DRAMSim3.[8][9][10] This helps in identifying discrepancies in timing and behavior.
- Random Testing: For coherence protocols developed in Ruby, leverage the Ruby random tester to issue semi-random requests and check for data correctness and protocol deadlocks.[3]

Q5: My simulation is running extremely slowly after integrating my custom memory model. What are the likely causes?

A5: Performance degradation can stem from several sources. Excessive use of DPRINTF statements can significantly slow down the simulation, so ensure they are only enabled when actively debugging.<sup>[1]</sup><sup>[11]</sup> Inefficient implementation of your memory model's C++ code, particularly in frequently accessed functions, can be a bottleneck. Additionally, complex Ruby protocols with many transient states and messages can inherently have a higher simulation overhead. Profile your simulation to identify the components consuming the most time.

## Troubleshooting Guides

### Issue 1: Simulation Hangs or Deadlocks with a Custom Ruby Protocol

Symptoms: The simulation time stops advancing, but the **GEM-5** process does not terminate. This is a classic sign of a deadlock in the memory system.

Troubleshooting Steps:

- **Enable Protocol Tracing:** The most critical tool for debugging deadlocks is the ProtocolTrace debug flag.<sup>[3]</sup> Rerun the simulation with `--debug-flags=ProtocolTrace`. This will generate a detailed log of every state transition in every controller.<sup>[3]</sup>
- **Analyze the Trace:** Examine the end of the trace file to identify the last few transitions that occurred. Look for requests that were sent but never received a response, or controllers that are stuck waiting for a particular event that never happens.
- **Visualize the Deadlock:** Use the protocol trace to manually diagram the sequence of events leading to the hang. This often reveals a circular dependency where multiple controllers are waiting on each other.
- **Check SLICC State Machine Logic:** Review your SLICC (.sm) files for logical errors in your state transitions. Ensure that for every possible event in a given state, there is a defined transition or a deliberate stall. Pay close attention to resource allocation and deallocation (e.g., network buffers, transient block entries).

- Use the Ruby Random Tester: The random tester is designed to uncover corner-case bugs that can lead to deadlocks by issuing concurrent read and write requests to the same cache block from different controllers.[3]

## Issue 2: Data Corruption or Incorrect Values Read from Memory

Symptoms: The simulated program produces incorrect results, or there are explicit data mismatches reported by the simulator.

Troubleshooting Steps:

- Enable Network and Data Tracing: Use the `RubyNetwork` debug flag to inspect the contents of messages being passed through the interconnection network.[3] This allows you to see the data being written to and read from memory. The `Exec` flag can be used to trace the instructions and the data they operate on at the CPU level.[2]
- Verify Port Connections: In your Python configuration script, ensure that all memory object ports (master and slave) are correctly connected.[12] An unconnected port can lead to requests being dropped or not being responded to.
- Debug with GDB: For deep inspection, run **GEM-5** within GDB. You can set breakpoints in your custom memory model's C++ code to inspect the state of memory packets (Packet objects) and internal data structures at specific points in time. Use the `--debug-break` option to stop the simulation at a specific tick before the corruption is expected to occur.[2][13]
- Check Memory Address Mapping: Verify that the address ranges in your memory controllers and other memory objects are configured correctly and do not have unintended overlaps or gaps.

## Experimental Protocols

### Protocol 1: Memory Bandwidth and Latency Verification

This protocol details a methodology for evaluating the performance of a custom DRAM controller model using a synthetic traffic generator.

- System Configuration:
  - CPU: TrafficGen (synthetic traffic generator).
  - Memory System: Your custom memory controller connected to a simple, single-level cache hierarchy. This isolates the DRAM controller's performance.[\[9\]](#)
  - Reference Model: A standard **GEM-5** memory model (e.g., DDR4\_2400\_8x8) for baseline comparison.[\[14\]](#)
- Traffic Generation:
  - Configure PyTrafficGen to generate a stream of random memory requests.[\[9\]](#)
  - Sweep the demand bandwidth from a low value (e.g., 1 GB/s) to a value exceeding the theoretical maximum of your memory model.
  - For each bandwidth point, run the simulation for a fixed number of requests (e.g., 1 million).
- Data Collection:
  - From the **GEM-5** statistics output (stats.txt), record the simulated memory bandwidth (system.mem\_ctrl.bw\_total::total) and average memory latency (system.mem\_ctrl.read\_average\_latency).
- Analysis:
  - Plot the measured bandwidth and latency as a function of the demand bandwidth for both your custom model and the reference model.
  - Compare the saturation points and latency curves to validate the performance characteristics of your model.



## FULL PROTOCOL TRUNCATED

To view exact molar ratios, purification steps, and HRP optimization data, please view the interactive version.

[Unlock Full Protocol on Website](#)

Note: The data in this table is illustrative and will vary based on the specific memory models and system configuration.

## Mandatory Visualizations

### Debugging Workflow for Custom Memory Models



## FULL PROTOCOL TRUNCATED

To view exact molar ratios, purification steps, and HRP optimization data, please view the interactive version.

[Unlock Full Protocol on Website](#)

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

Caption: A logical workflow for debugging custom memory models in **GEM-5**.

## GEM-5 Ruby Memory System Component Interaction



## FULL PROTOCOL TRUNCATED

To view exact molar ratios, purification steps, and HRP optimization data, please view the interactive version.

[Unlock Full Protocol on Website](#)

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

Caption: High-level interaction of components in the **GEM-5** Ruby memory system.

### *Need Custom Synthesis?*

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

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

## References

- [1. gem5: Debugging gem5 \[gem5.org\]](#)
- [2. gem5: Trace-based Debugging \[gem5.org\]](#)
- [3. gem5: Debugging SLICC Protocols \[gem5.org\]](#)
- [4. gem5: Common errors within gem5 \[gem5.org\]](#)
- [5. scribd.com \[scribd.com\]](#)
- [6. m.youtube.com \[m.youtube.com\]](#)
- [7. gem5: Introduction to Ruby \[gem5.org\]](#)
- [8. escholarship.org \[escholarship.org\]](#)
- [9. arch.cs.ucdavis.edu \[arch.cs.ucdavis.edu\]](#)
- [10. Methodologies for Evaluating Memory Models in gem5 \[escholarship.org\]](#)
- [11. Debugging gem5 — gem5 Tutorial 0.1 documentation \[courses.grainger.illinois.edu\]](#)

- [12. gem5: Memory system \[gem5.org\]](#)
- [13. gem5: Debugger-based Debugging \[gem5.org\]](#)
- [14. A Tutorial on the Gem5 Memory Model | Nitish Srivastava \[nitish2112.github.io\]](#)
- To cite this document: BenchChem. [debugging and verifying custom memory models in GEM-5]. BenchChem, [2026]. [Online PDF]. Available at: [\[https://www.benchchem.com/product/b12410503#debugging-and-verifying-custom-memory-models-in-gem-5\]](https://www.benchchem.com/product/b12410503#debugging-and-verifying-custom-memory-models-in-gem-5)

---

### 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?

**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](mailto:info@benchchem.com)

[Contact our Ph.D. Support Team for a compatibility check](#)

