

System Simulation with gem5, SystemC and other Tools

Christian Menard (TU Dresden, Germany)

Matthias Jung (Fraunhofer IESE, Germany)

Detailed Description of Contents for this Talk



System Simulation with gem5 and SystemC

The Keystone for Full Interoperability

Christian Menard, Jeronimo Castrillon
Technische Universität Dresden
Dresden, Germany
Email: christian.menard@tu-dresden.de
jeronimo.castrillon@tu-dresden.de

Matthias Jung
Fraunhofer IESE
Kaiserslautern, Germany
Email: matthias.jung@iese.fraunhofer.de

Norbert Wehn
University of Kaiserslautern
Kaiserslautern, Germany
Email: wehn@cit.uni-kl.de

Abstract—SystemC TLM based virtual prototypes have become the main tool in industry and research for concurrent hardware and software development, as well as hardware design space exploration. However, there exists a lack of accurate, free, changeable and realistic SystemC models of modern CPUs. Therefore, many researchers use the cycle accurate open source system simulator gem5, which has been developed in parallel to the SystemC standard. In this paper we present a coupling of gem5 with SystemC that offers full interoperability between both simulation frameworks, and therefore enables a huge set of possibilities for system level design space exploration. Furthermore, we show that the coupling itself only induces a relatively small overhead to the total execution time of the simulation.

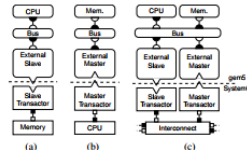


Figure 1: Possible scenarios for binding gem5 and SystemC.

I. INTRODUCTION

Today's companies have to deal with complex hardware architectures such as multi-cores, sophisticated interconnects and memory systems. *Virtual Prototypes* (VPs) are widely used to allow for early design space exploration and to decrease the *Time-to-Market* (TTM), costs, and efforts by developing software and hardware concurrently. They are high-speed, fully functional software models of physical hardware systems that can simulate the exact behavior of real hardware. With their help, complete *Multi-Processor System-on-Chips* (MPSoCs) can be simulated with reasonable simulation speed and visibility and controllability over the entire system. Case studies have shown that by employing VPs it is possible to deliver more competitive products up to six months earlier [1]. In recent years, the SystemC TLM2.0 IEEE1666 standard [2] has become the main developing tool for VPs in industry and research. It allows to quickly simulate HW and SW systems on different levels of abstraction in order to estimate and optimize the performance and power for different applications. In contrast to pin accurate models, *Transaction Level Modeling* (TLM) abstracts the communication mechanisms from the actual hardware. It encapsulates communication between interacting components in so-called transactions, which are transferred by function calls.

The industry offers several SystemC based CPU core models that provide a trade-off between simulation speed and accuracy. For instance, the *FastModels*, distributed by ARM Ltd. or the *QVP* [3] models use *just-in-time-compilation* for code execution and model communication using the TLM *loosely-timed* coding style. However, loosely-timed models do not reflect a realistic timing behavior and thus can mainly be used for software development and high level explorations. Cycle

accurate simulations can be performed with the commercially available Carbon models from ARM or the Aurix [4] TLM models from Infineon. However, these models are used as binary libraries, which makes them useless for micro-architectural research due to their inflexibility (i.e. they cannot be modified). Furthermore, they are slow and can, therefore, not be employed for fast design space exploration.

In contrast to industry, the academia lacks free, accurate and realistic SystemC models of modern CPUs. The most mature cycle accurate open source system simulator is the gem5 framework, which is a modular platform for computer-system architecture research [5]. It is not only widely used in academia, but also the industry employs gem5 for research. For instance, ARM and AMD use gem5 internally for design space exploration and actively contribute to the open source project. However, gem5 is not implemented in SystemC as its development started before the IEEE ratified the official SystemC and TLM standard in 2005 [2]. Since this time, both frameworks, gem5 and SystemC have evolved extensively and independent in parallel. Therefore, gem5 is incompatible to TLM models that exist in industry and academia.

In this paper, we present for the first time a comprehensive coupling between SystemC and gem5 that provides full interoperability. Through this coupling, any SystemC module that implements the TLM base protocol can be connected to any gem5 module, as shown in Figure 1. To the best of our knowledge, there exists no reference which describes syntax and semantics of both frameworks and how both simulation kernels can be coupled in order to enable full interoperability.

```
Terminal - README + (~/.Programming/gem5/util/tlm) - VIM
File Edit View Terminal Tabs Help

This directory contains a demo of a coupling between gem5 and SystemC-TLM. It
is based on the gem5-systemc implementation in utils/systemc. This Readme gives
an overall overview (I), describes the source files in this directory (II),
explains the build steps (III), shows how to run example simulations (IV-VI)
and lists known issues (VII).

I. Overview
=====

The sources in this directory provide three SystemC modules that manage the
SystemC/gem5 co-simulation: Gem5SimControl, Gem5MasterTransactor, and
Gem5SlaveTransactor. They also implement gem5's ExternalMaster::Port interface
(SCMasterPort) and ExternalSlave::Port interface (SCSlavePort).

**SCMasterPort** and **Gem5MasterTransactor** together form a TLM-to-gem5
bridge. SCMasterPort implements gem5's ExternalMaster::Port interface and forms
the gem5 end of the bridge. Gem5MasterTransactor is a SystemC module that
provides a target socket and represents the TLM side of the bridge. All TLM
requests sent to this target socket, are translated to gem5 requests and
forwarded to the gem5 world through the SCMasterPort. Then the gem5 world
handles the request and eventually issues a response. When the response arrives
at the SCMasterPort it gets translated back into a TLM response and forwarded
to the TLM world through target socket of the Gem5MasterTransactor.
SCMasterPort and Gem5MasterTransactor are bound to each other by configuring
them for the same port name.

**SCSlavePort** and **Gem5SlaveTransactor** together form a gem5-to-TLM bridge.
Gem5SlaveTransactor is a SystemC module that provides an initiator socket and
represents the TLM end of the bridge. SCSlavePort implements gem5's
ExternalSlave::Port interface and forms the gem5 side of the bridge. All gem5
requests sent to the SCSlavePort, are translated to TLM requests and forwarded
to the TLM world through the initiator socket of the Gem5SlaveTransactor. Then
the TLM world handles the request and eventually issues a response. When the
response arrives at the Gem5SlaveTransactor it gets translated back into a
gem5 response and forwarded to the gem5 world through the SCSlavePort. SCSlavePort
and Gem5SlaveTransactor are bound to each other by configuring them for the
same port name.

6,0-1 Top
```

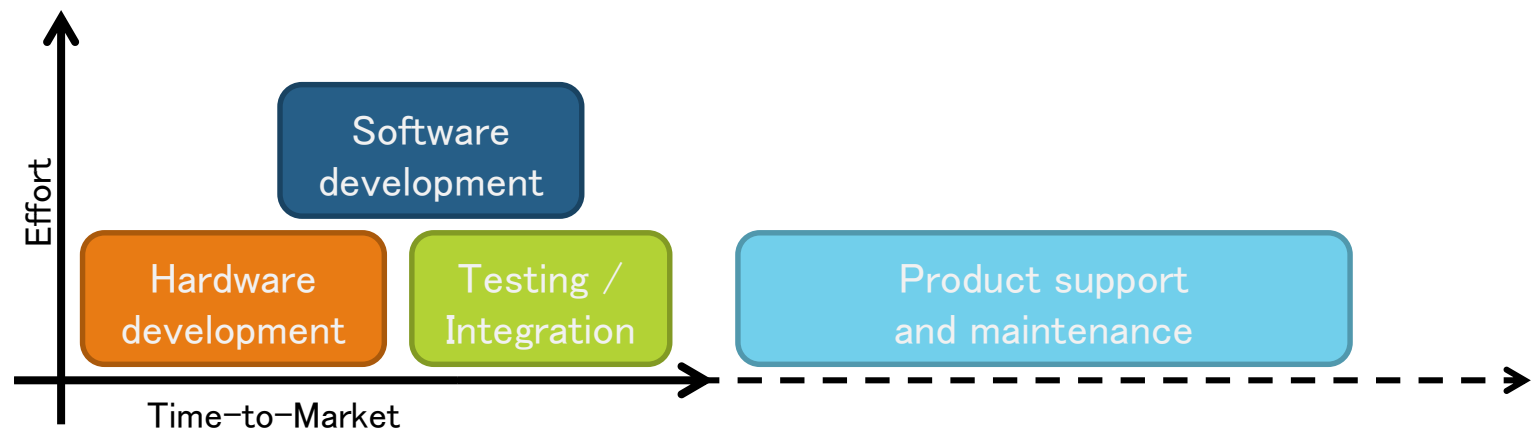
Paper: [Link](#) [1]

Sources: [/gem5/utlis/tlm/README](#)

Virtual Prototyping

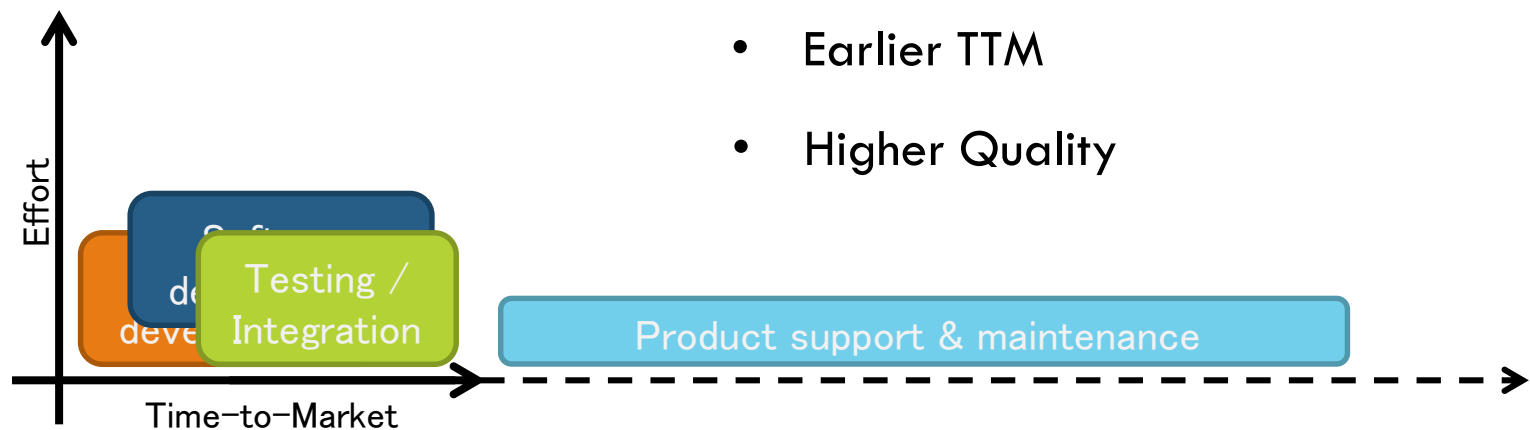
Functional software models of physical hardware:

- ❑ Visibility and controllability over the entire system
- ❑ Powerful debugging and analysis tools
- ❑ Reuse of components for future projects
- ❑ Fast Design Space Exploration (for HW engineers)
- ❑ Easy to exchange, worldwide
- ❑ Concurrent HW and SW development:

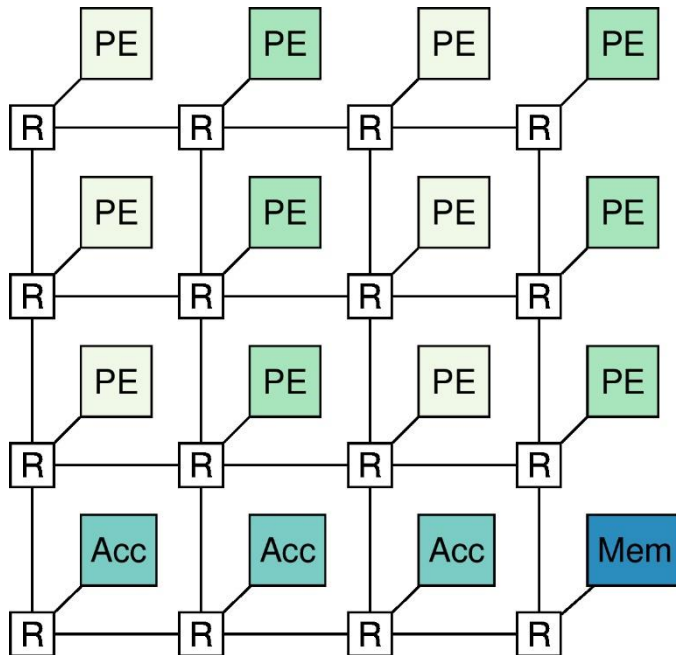


Functional software models of physical hardware:

- ❑ Visibility and controllability over the entire system
- ❑ Powerful debugging and analysis tools
- ❑ Reuse of components for future projects
- ❑ Fast Design Space Exploration (for HW engineers)
- ❑ Easy to exchange, worldwide
- ❑ Concurrent HW and SW development:

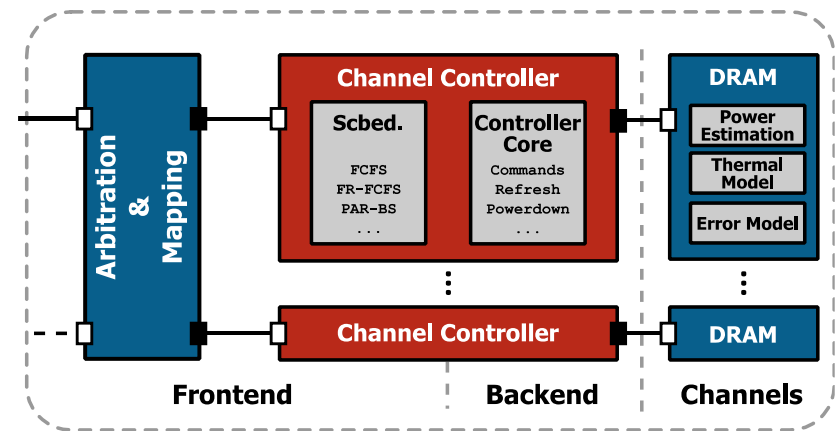


Simulation of (widely) heterogeneous systems

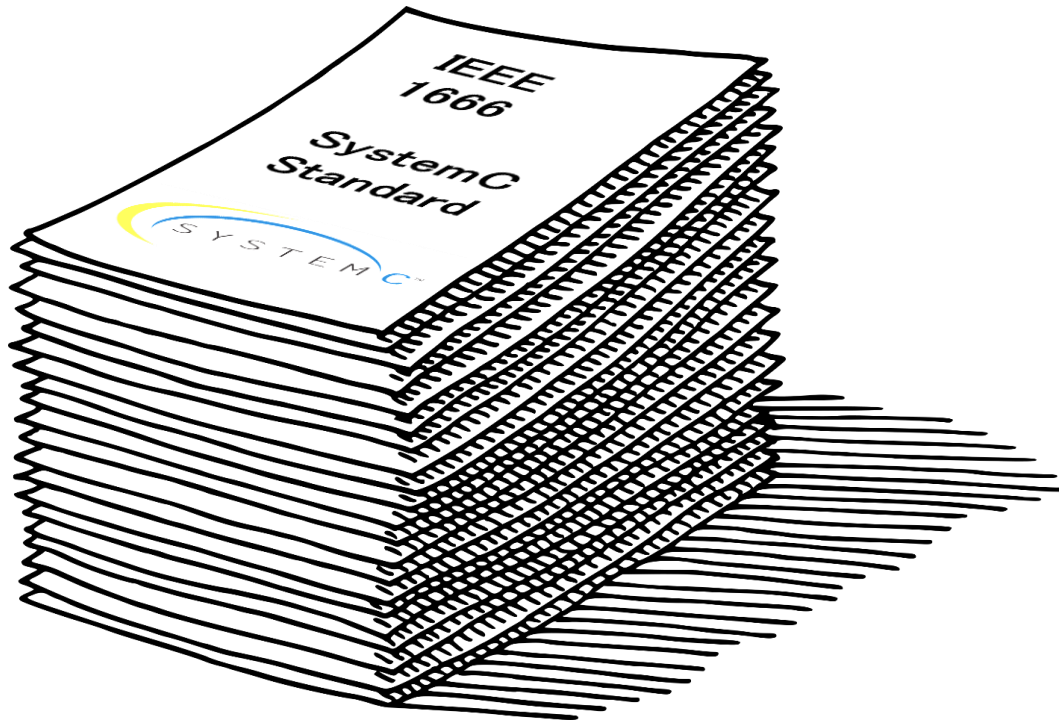


- Many different models of cores, accelerators, and communication infrastructure required.

Simulation of the Memory Subsystem



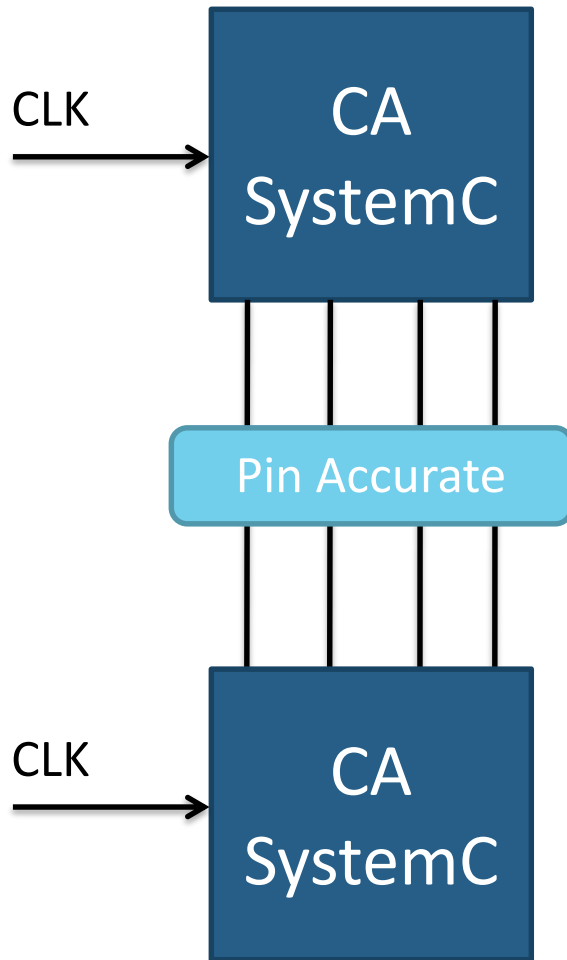
- Focuses on the memory subsystem, but detailed simulation of realistic workloads is required



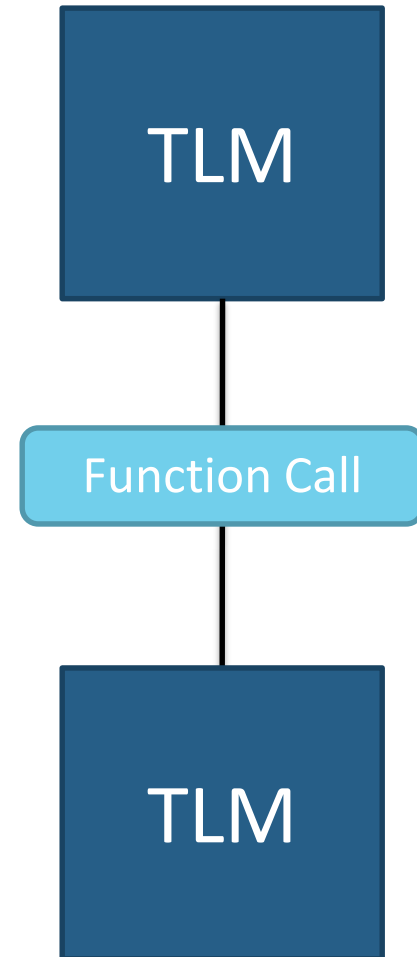
- ❑ Modeling language for HW and SW components
- ❑ Extends C++ to an event-driven simulation kernel
- ❑ Various levels of accuracy
- ❑ IEEE Standard, Maintained by Accellera
- ❑ 10-100x Faster than CA VHDL/Verilog Simulation

→ However, normal CA SystemC is not fast enough to, e.g., boot an OS.

Transaction Level Modeling (TLM)

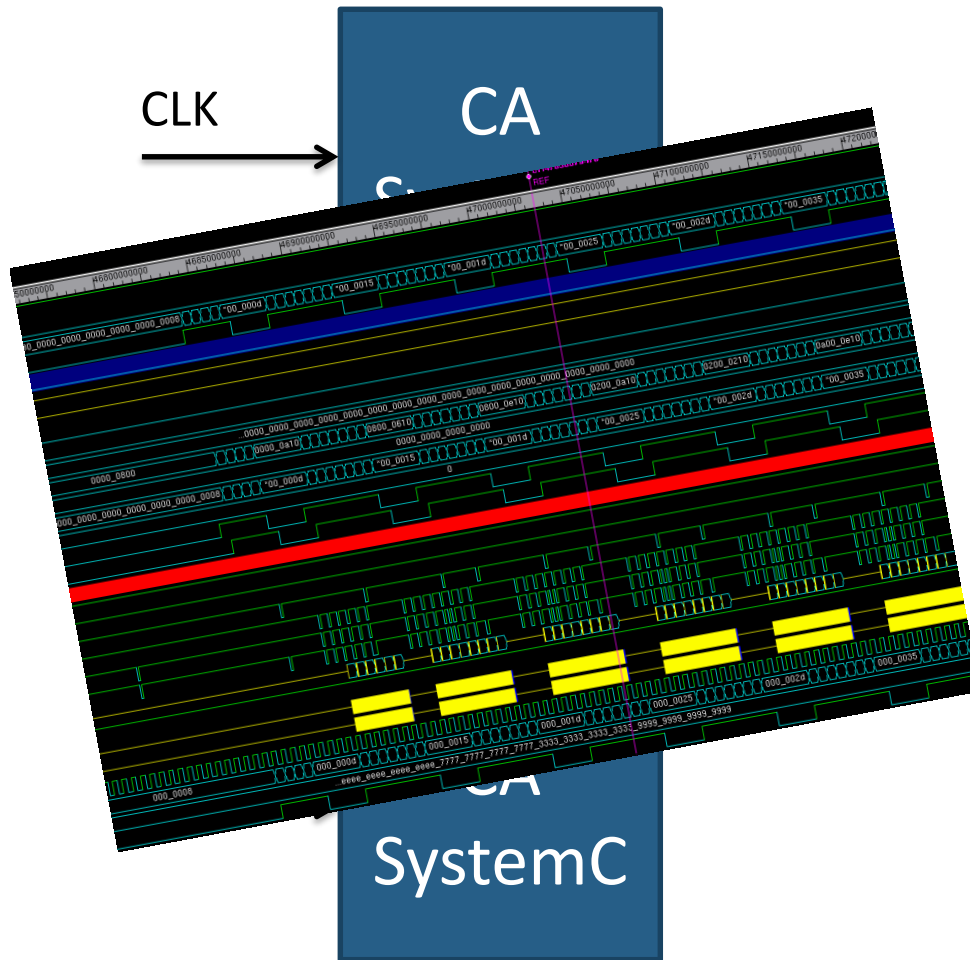


Simulate each pin separately

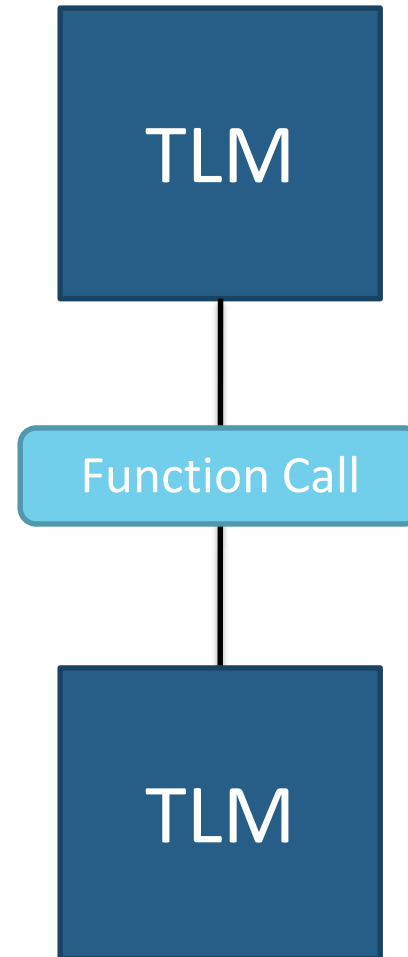


Simulate transactions up to 10,000x Faster

Transaction Level Modeling (TLM)

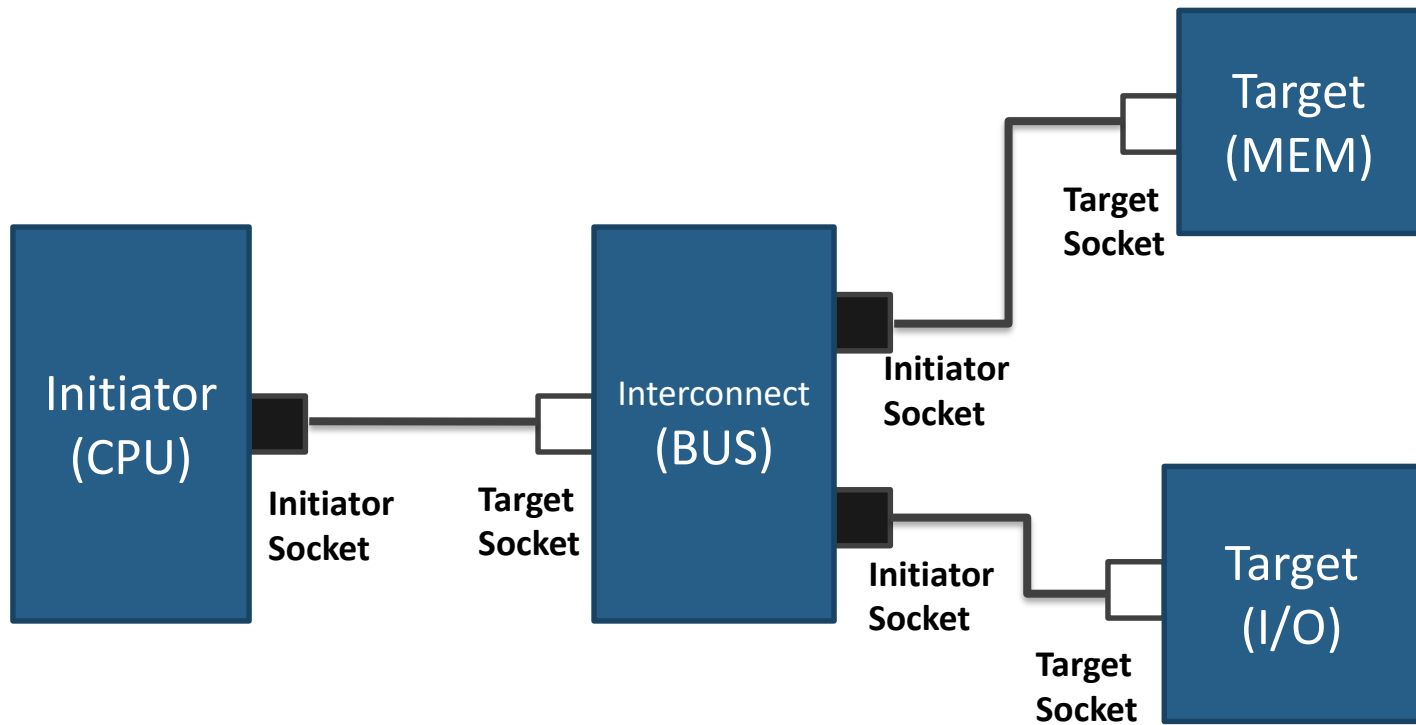


Simulate each pin separately



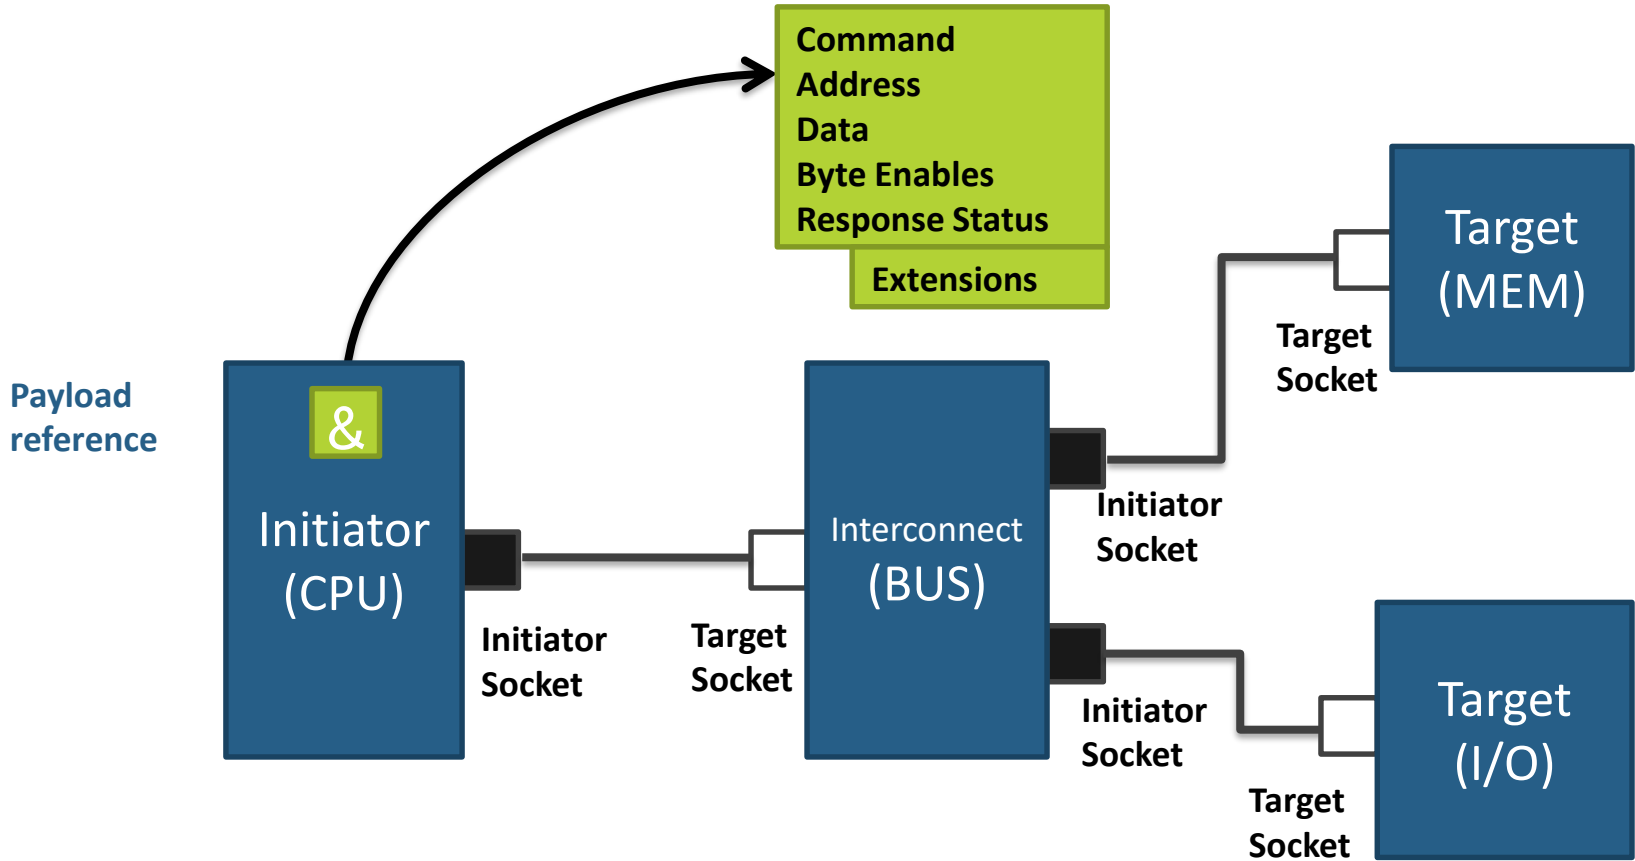
Simulate transactions up to 10,000x Faster

Generic Payload



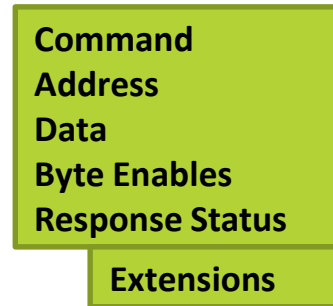
Generic Payload

Generic payload object

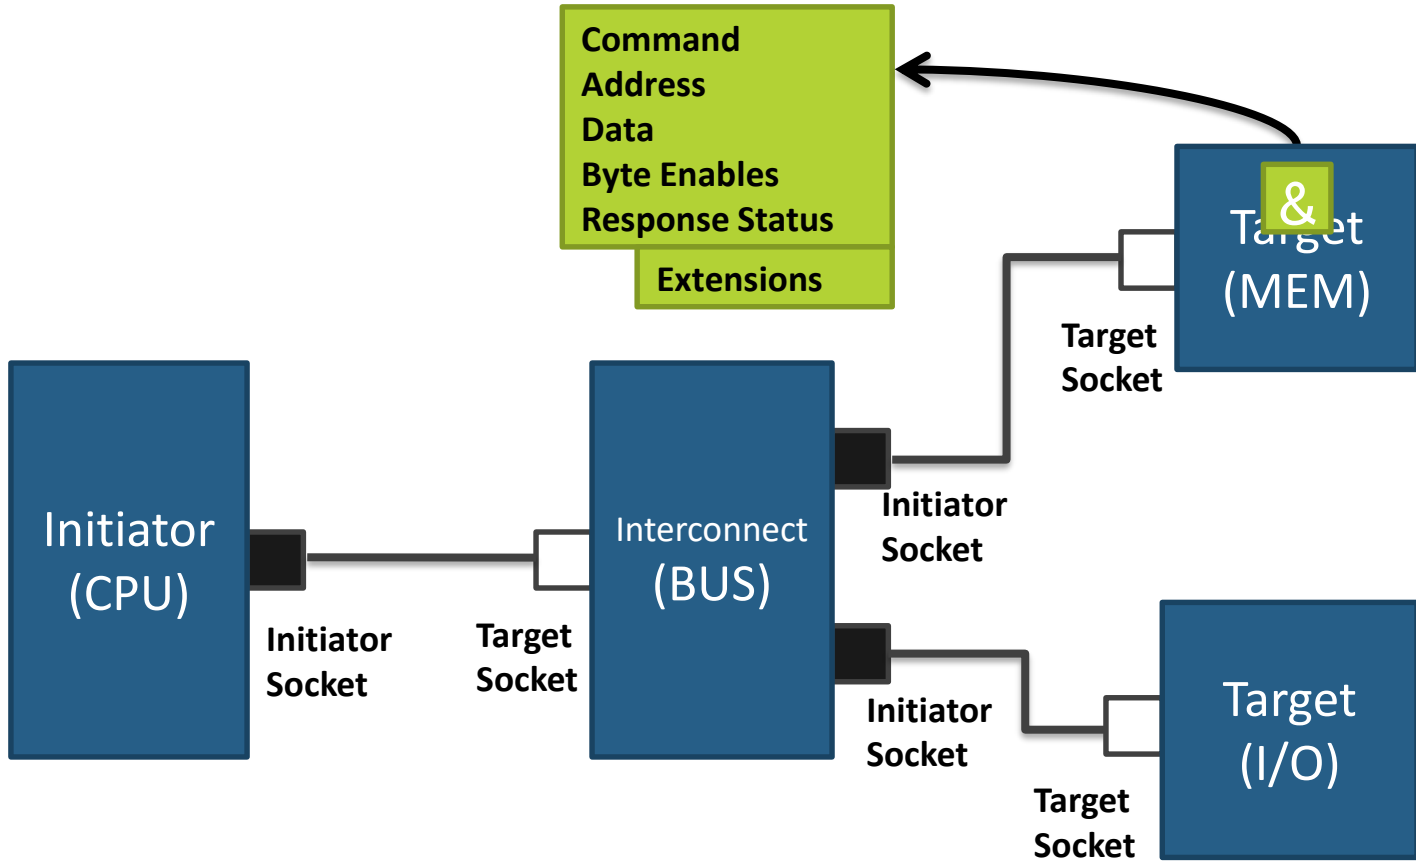


Generic Payload

Generic payload object



Payload reference



TLM Use Cases

SW Application
Development

SW Performance
Analysis

Architecture
Analysis

Hardware
Verification

TLM 2.0 Coding Style (*Just Guidelines*)

Loosely-timed

Single-phase, blocking API
`debug_transport`, `b_transport`

`nb_transport`

Multi-phase, non-blocking API

Approximately -timed

TLM Mechanisms (*Definitive API for enabling Interoperability*)

Blocking
transport

DMI

Quantum

Sockets

Generic
payload

Extensions

Phases

Non-
blocking
transport

TLM is widely used in Industry:

- ❑ The market of virtual platform tools:
 - ❑ Synopsys - Platform Architect
 - ❑ Cadence - Virtual System Platform
 - ❑ Mentor Graphics - Vista Virtual prototyping
 - ❑ Imperas - OpenVP
 - ❑ ASTC - VLAB Works

- ❑ Virtual Platform Core Models:
 - ❑ ARM (Fastmodels):
 - only LT models based on JIT, non-free, library
 - ❑ ARM Carbon (Former Carbon Design Systems):
 - Cycle Accurate (CA) Models in TLM Wrapper, non-free, library
 - ❑ Imperas / OVP:
 - only LT, Free

→ An accurate, free available and changeable core model is needed

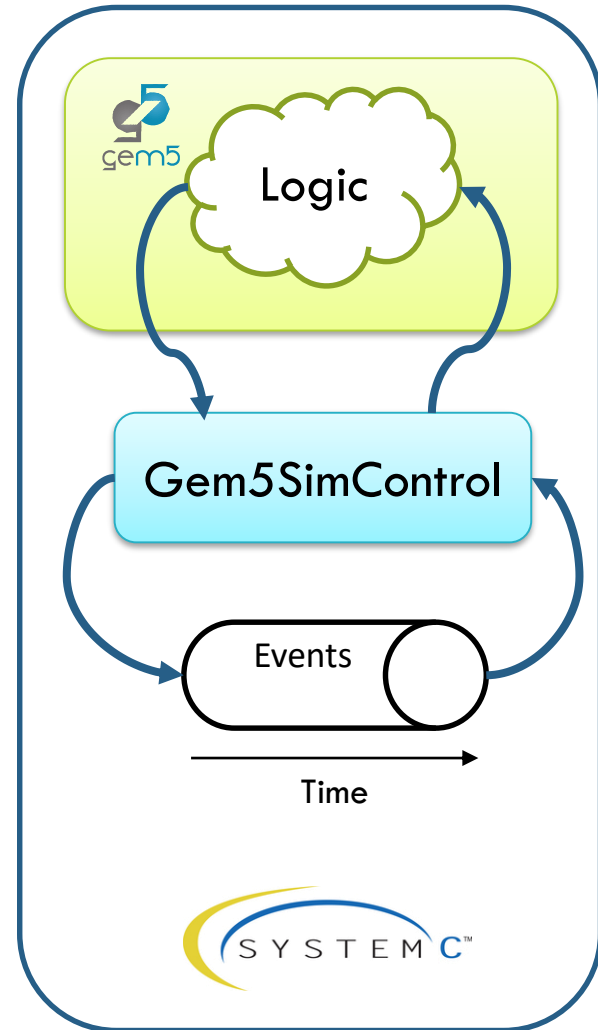
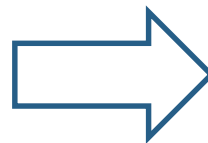
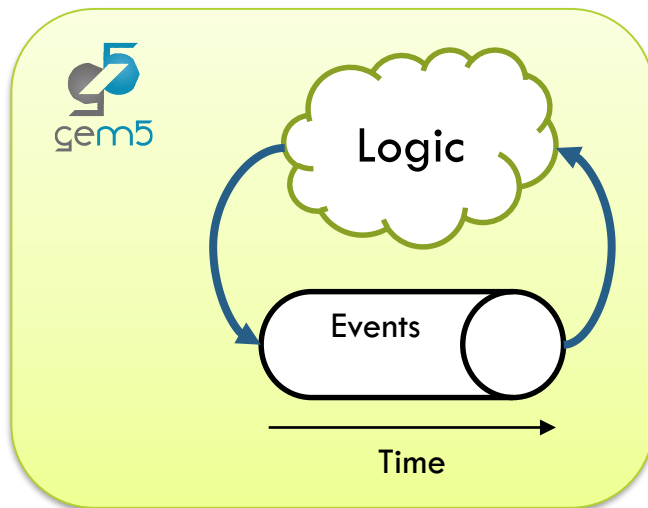


Coupling gem5 with SystemC

Coupling gem5 with SystemC

gem5 supports a SystemC coupling:

- ❑ Gem5 is build as a C++ library.
- ❑ It is linked into a SystemC simulation.
- ❑ A SystemC object implements the gem5 event queue.
- ❑ **How can we communicate with other SystemC modules?**



Timing

- ❑ The most detailed access: queuing delay + resource contention
- ❑ Similar to the TLM `nb_transport` interface.

Atomic

- ❑ Accesses are a faster than detailed access
- ❑ Used for **fast forwarding** and **warming up caches**
- ❑ Similar to the TLM `b_transport` interface
- ❑ Not good for performance simulation

Functional

- ❑ Similar to `transport_dbg` e.g. loading binaries, avoiding deadlocks in multi-level cache coherent networks

Converting between TLM and gem5

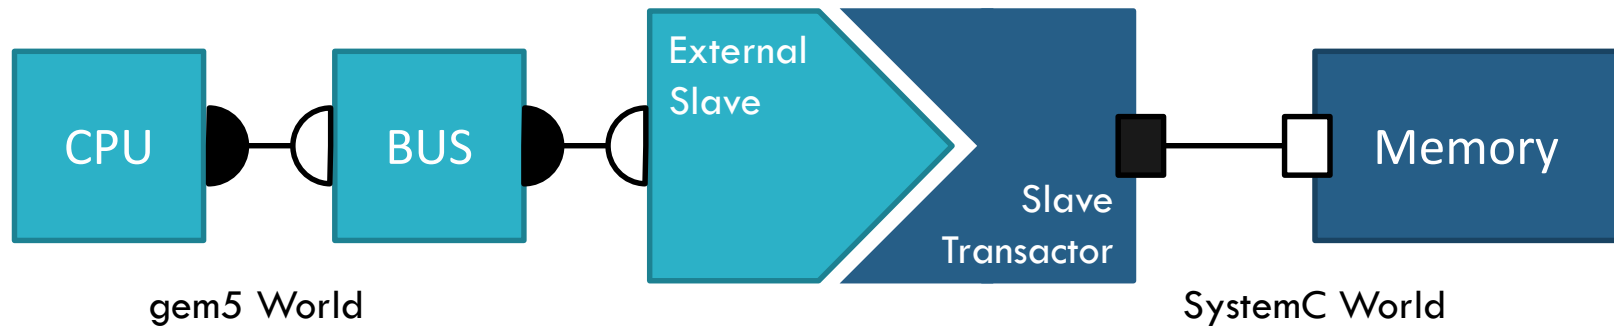


`recvFunctional(...)` → `transport_dbg(...)`
`recvAtomic(...)` → `b_transport(...)`
`recvTimingReq(...)` → `nb_transport(...)`

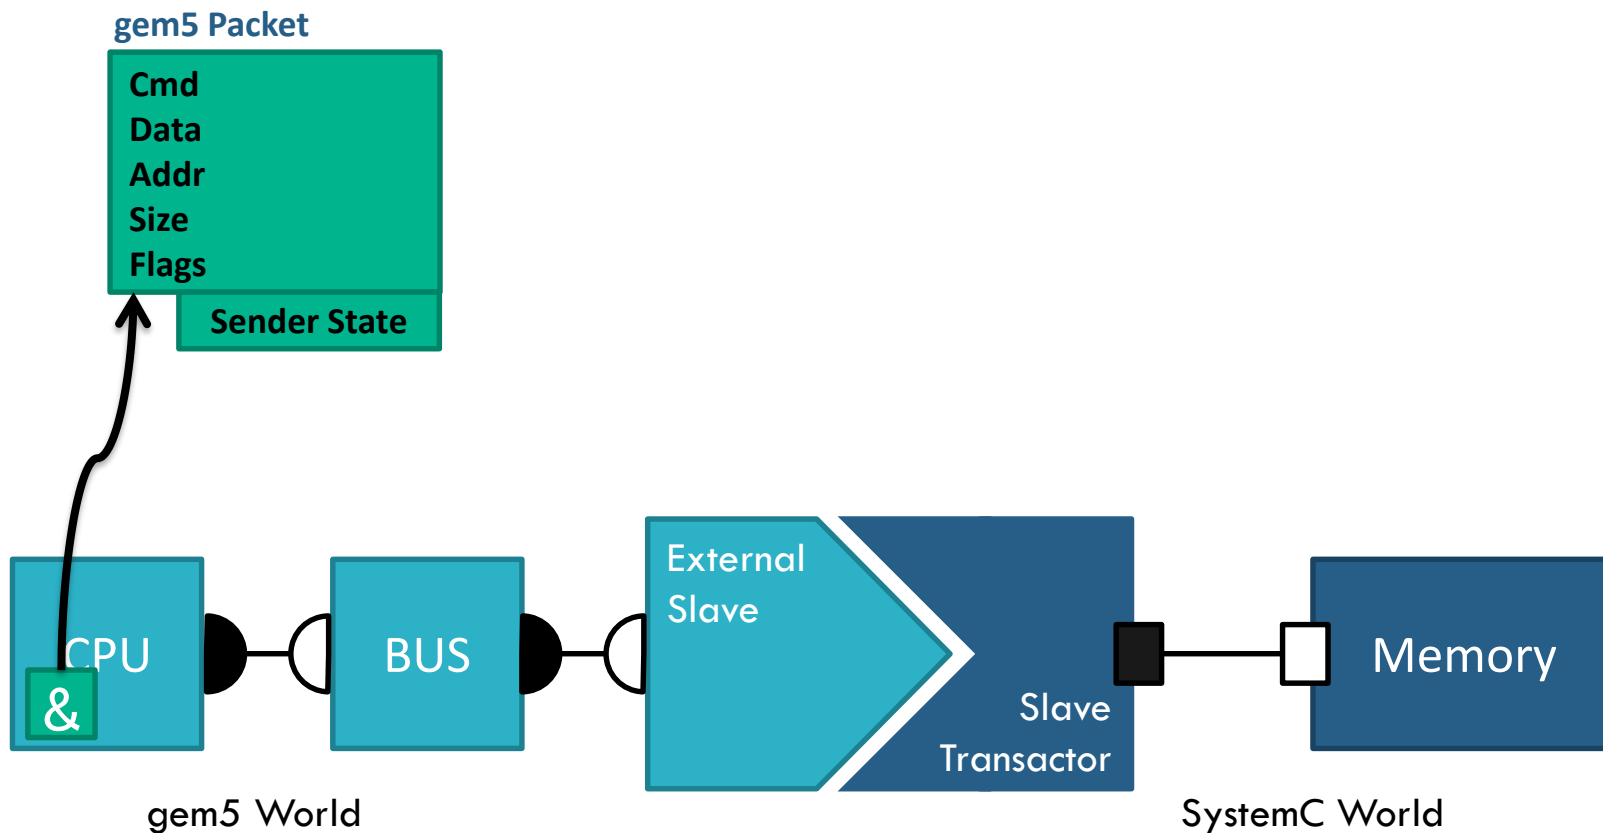


`transport_dbg(...)` → `recvFunctional(...)`
`b_transport(...)` → `recvAtomic(...)`
`nb_transport(...)` → `recvTimingReq(...)`

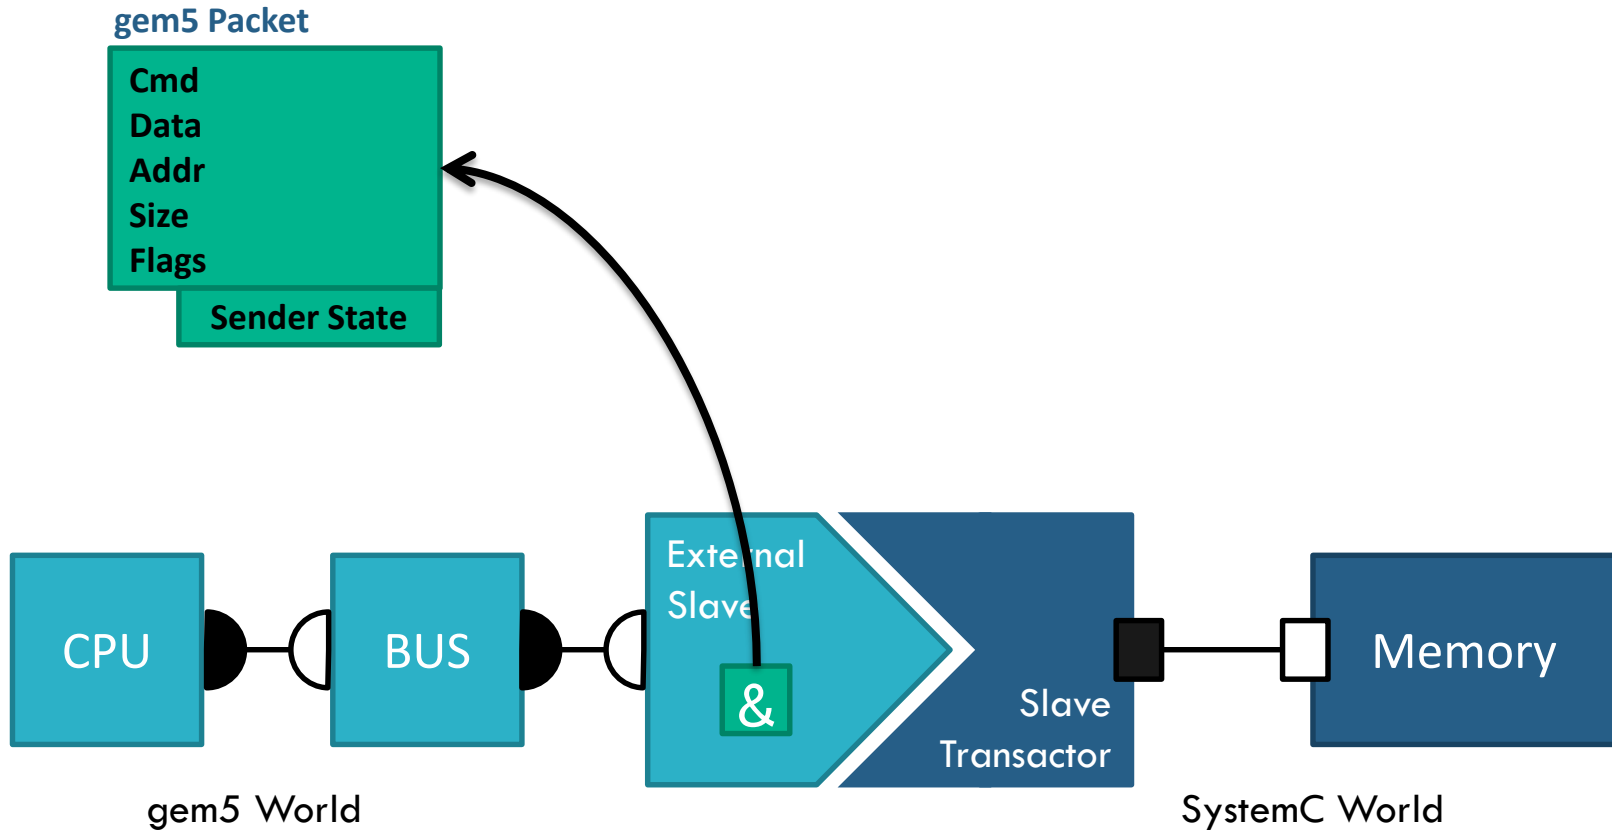
Transaction Explained



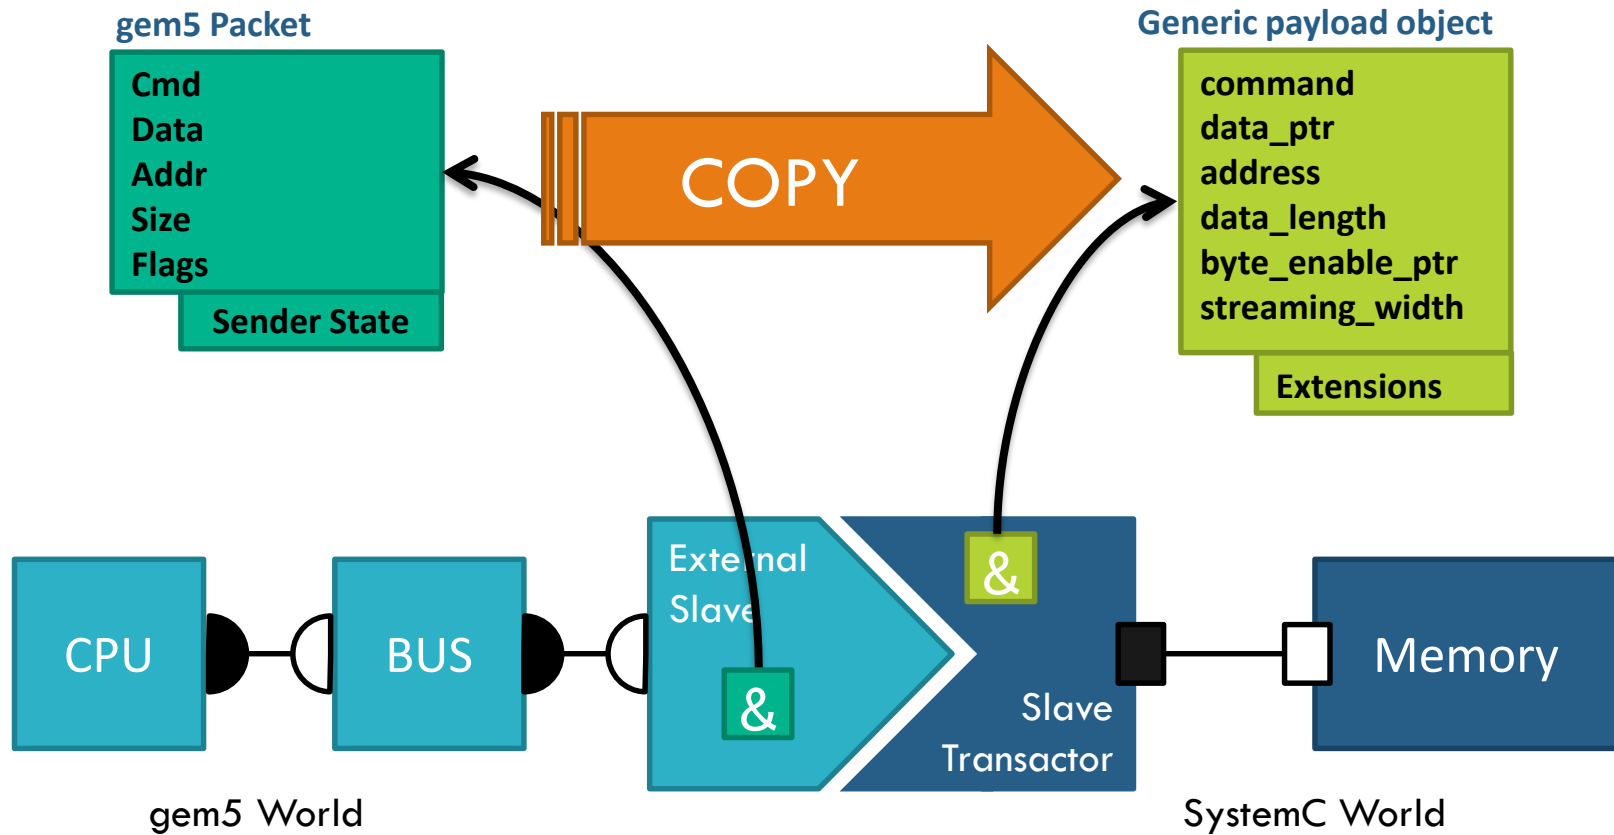
Transaction Explained



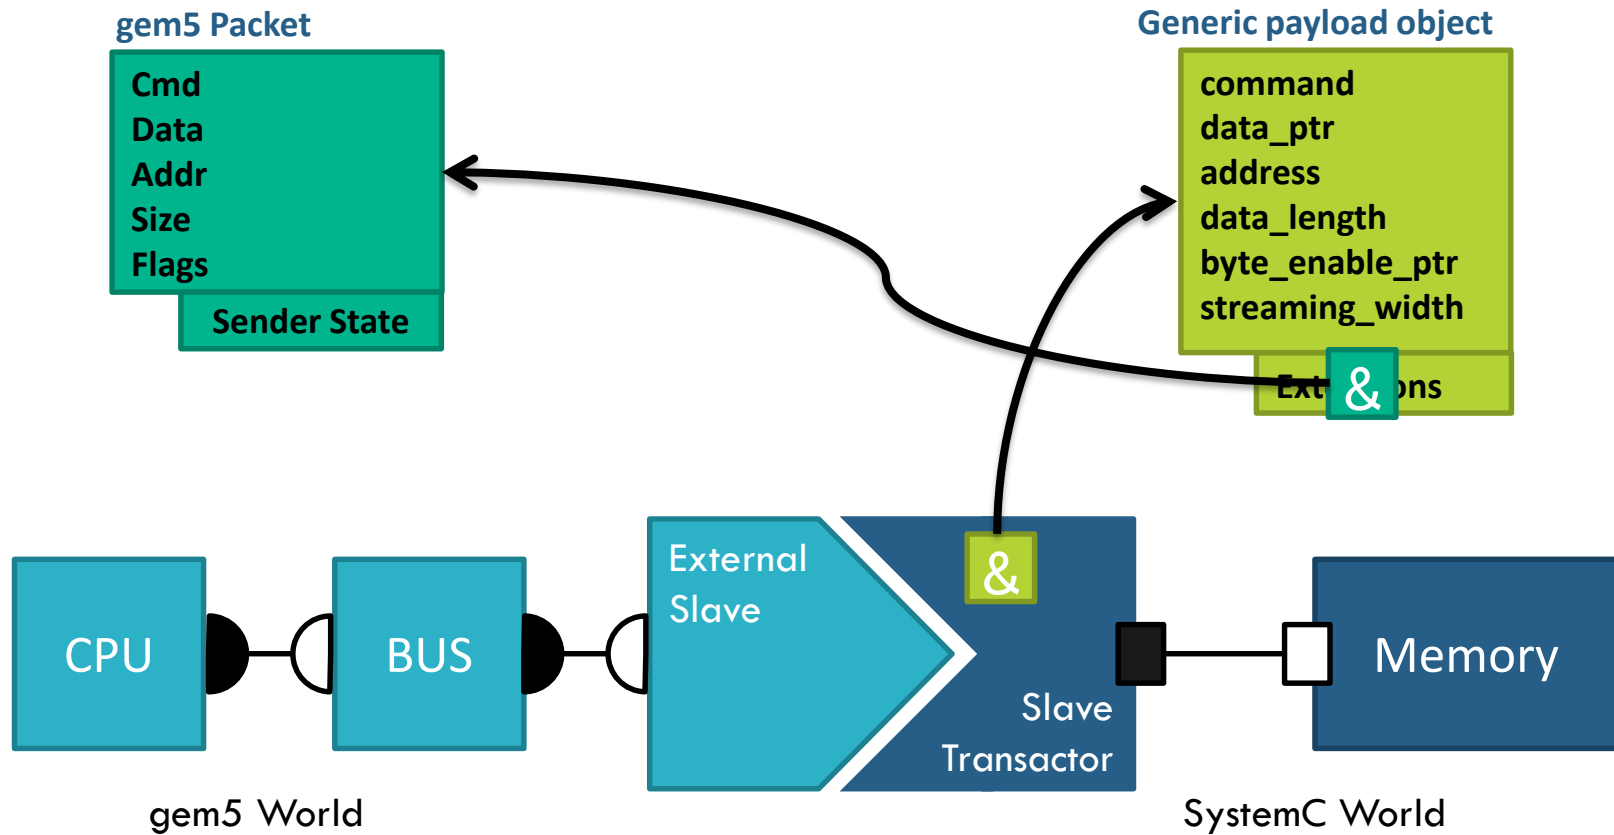
Transaction Explained



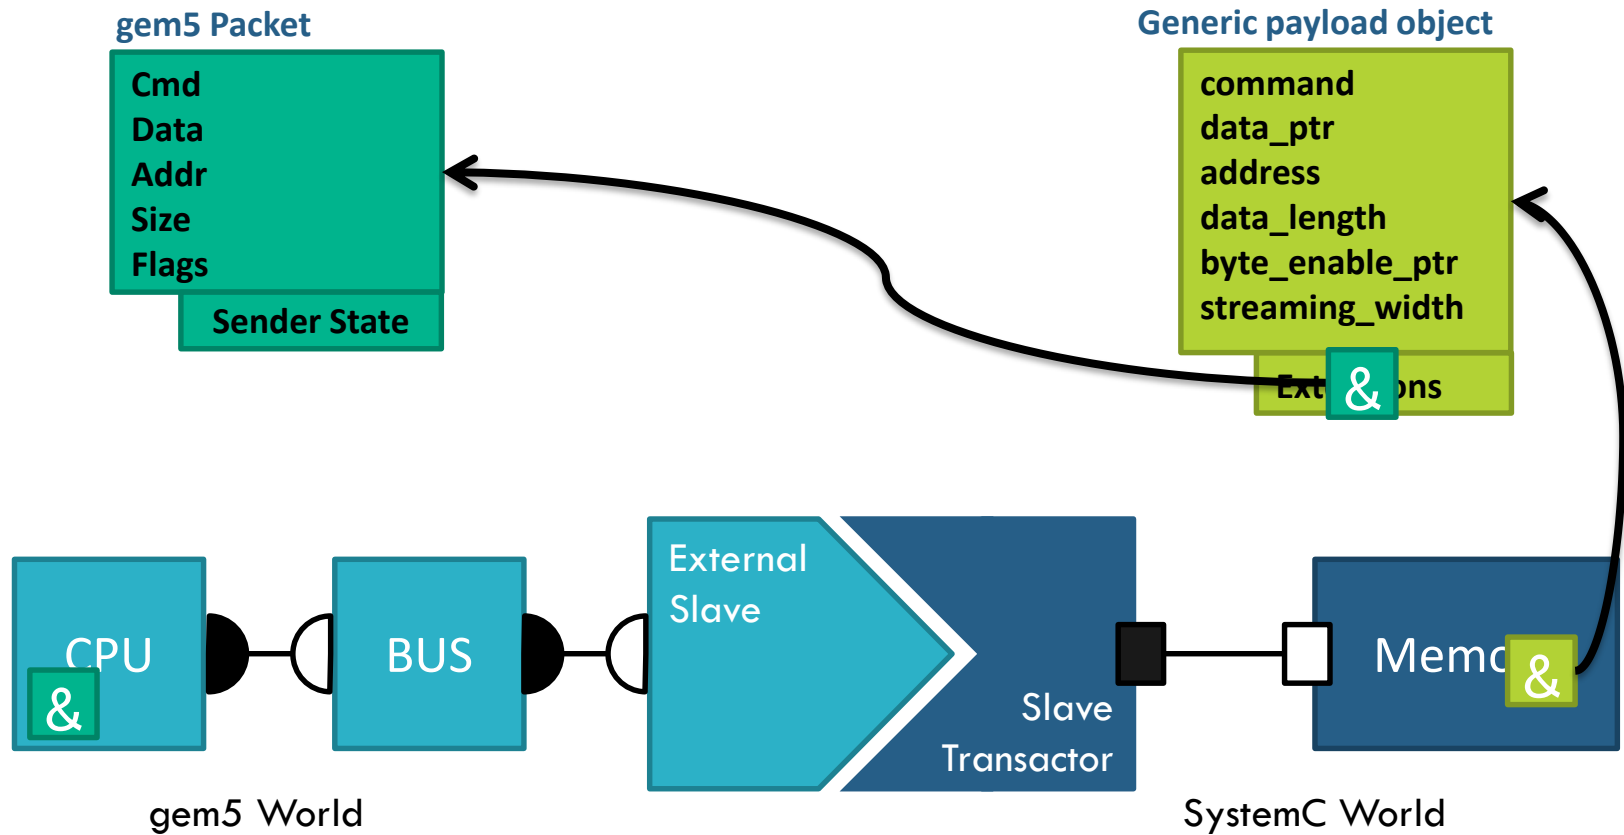
Transaction Explained



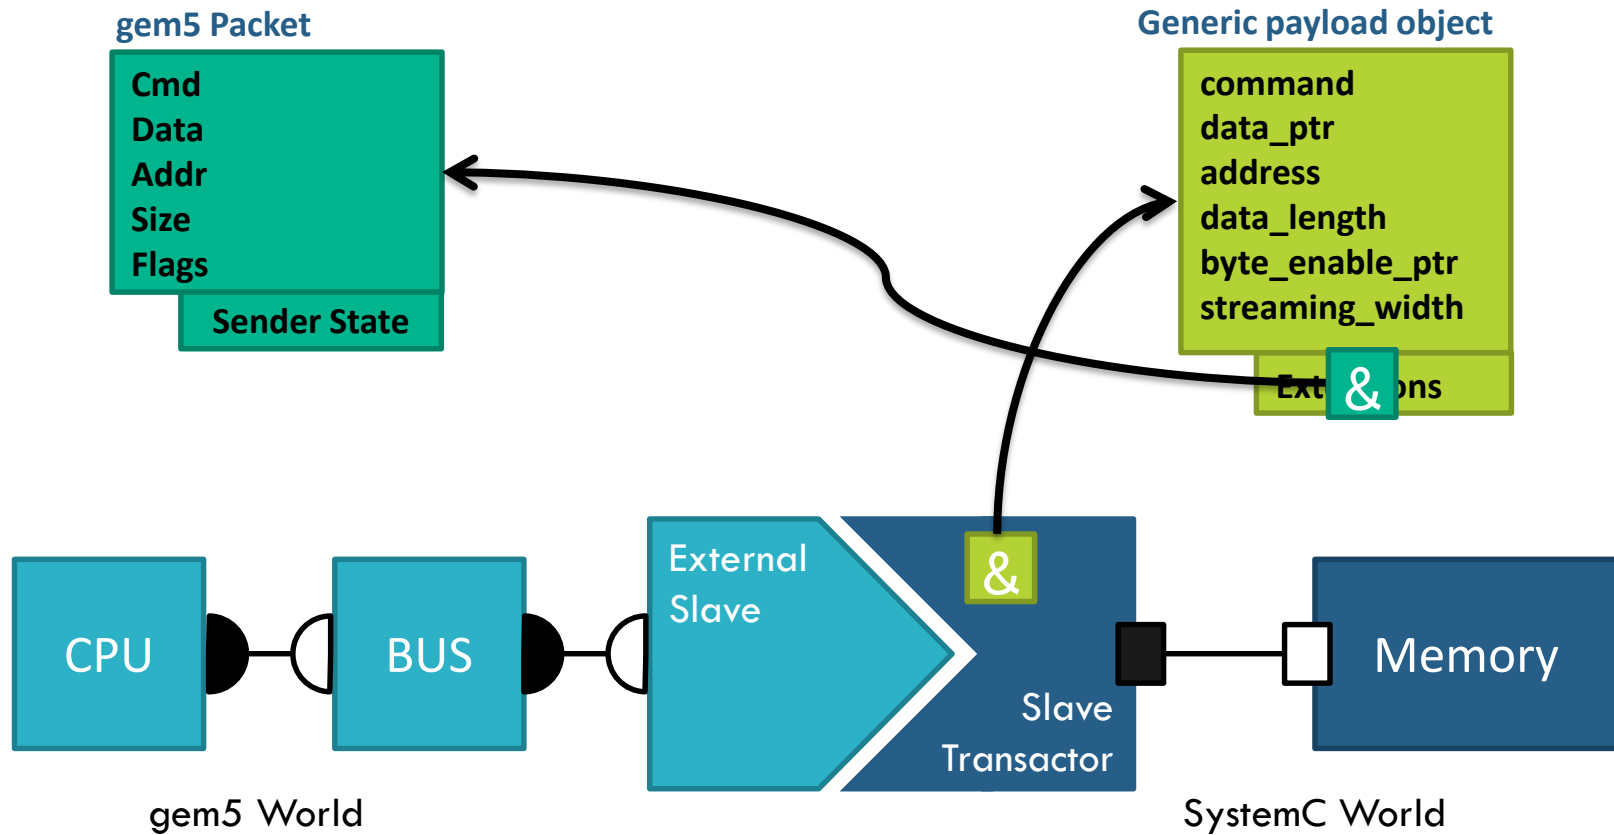
Transaction Explained



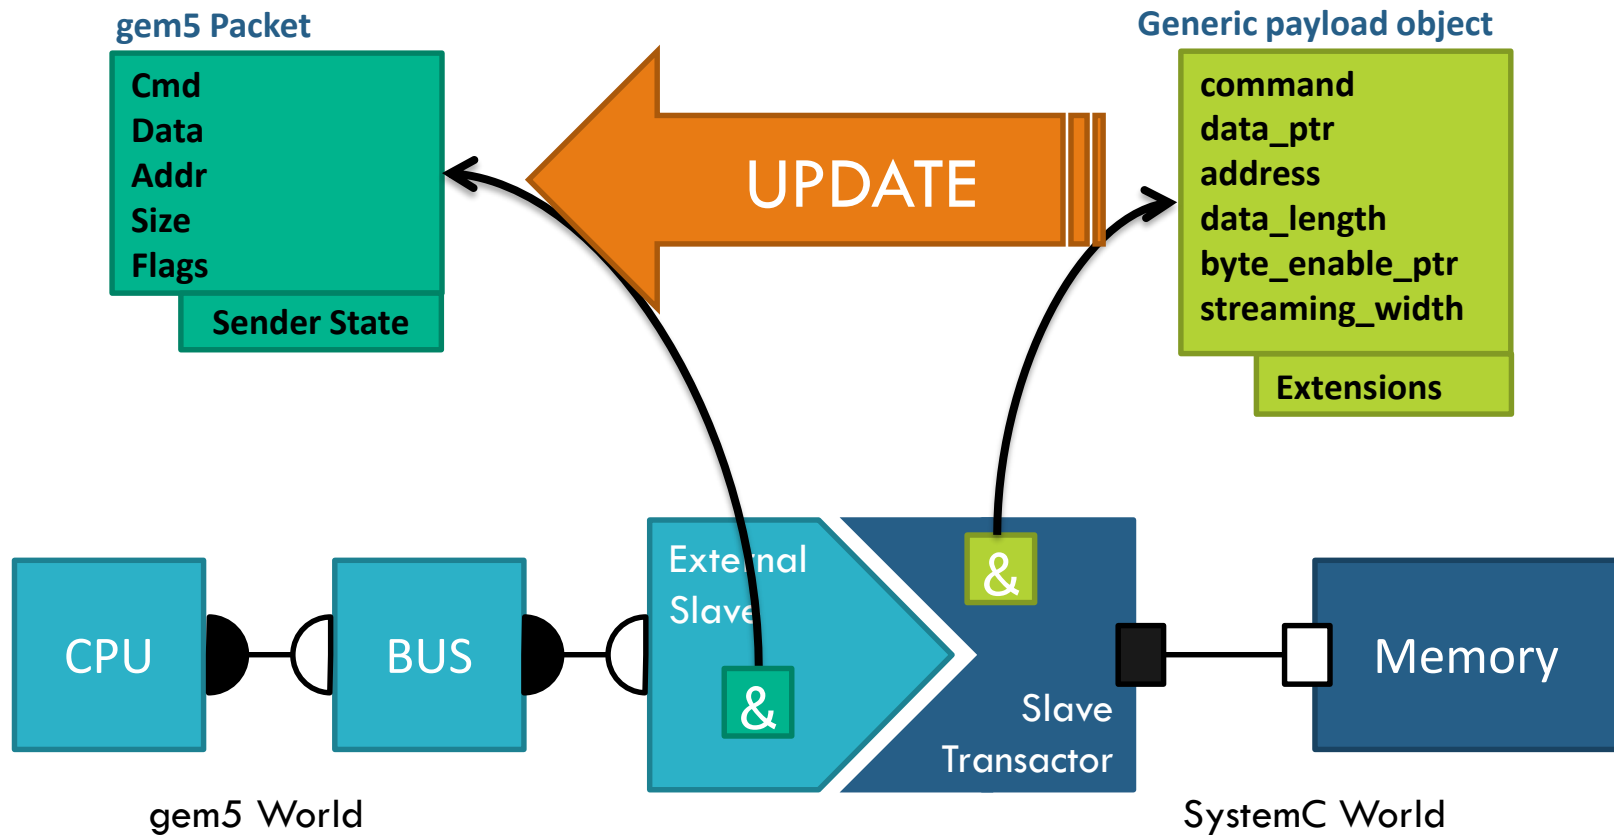
Transaction Explained



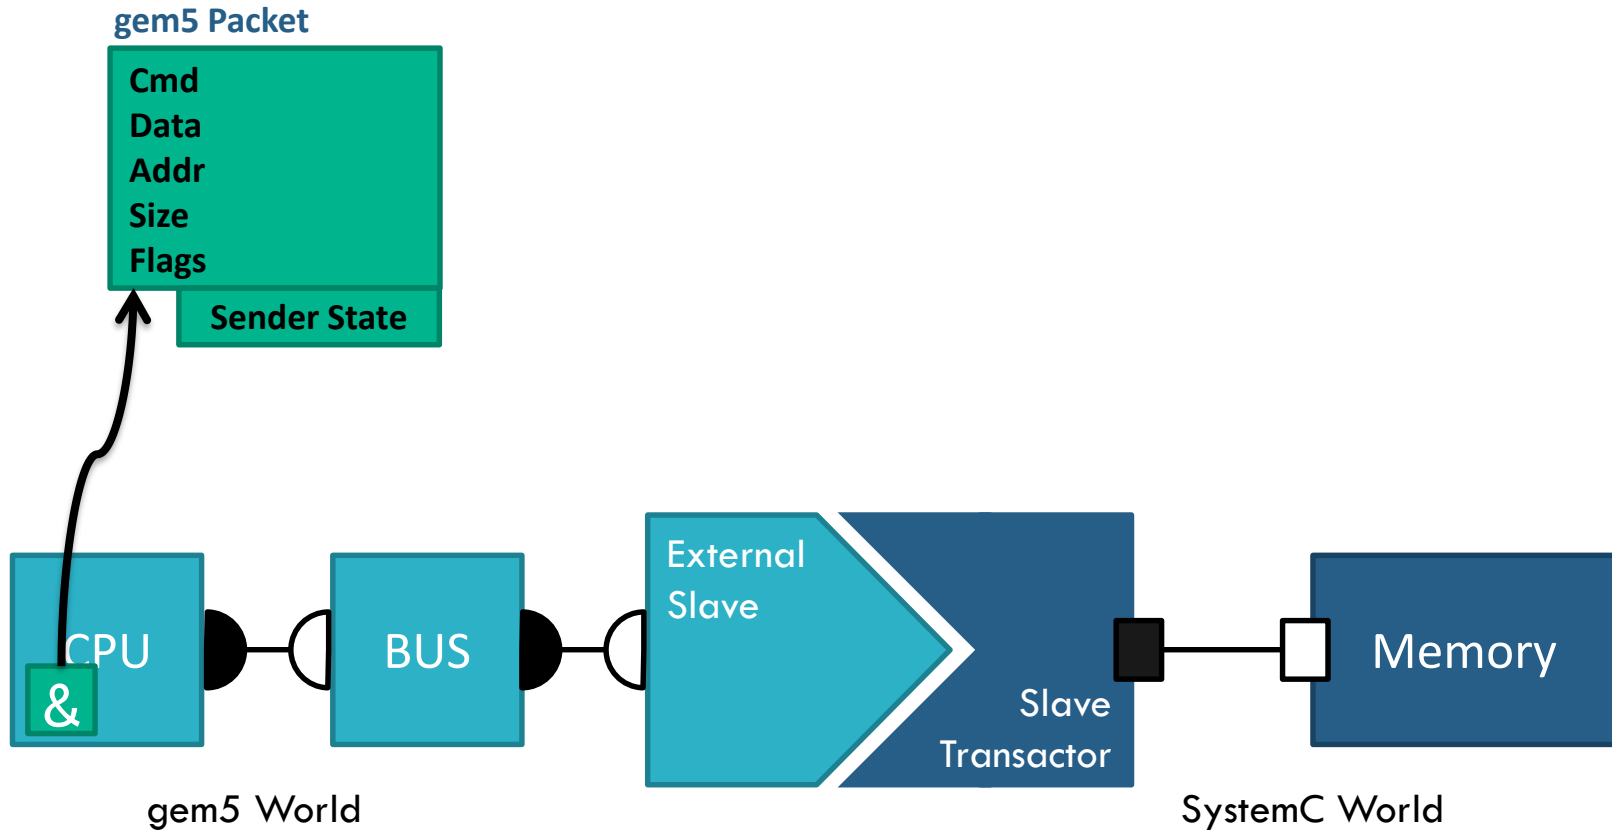
Transaction Explained



Transaction Explained



Transaction Explained



How to get Started?

How to get Started?

❑ Study the Examples in [/gem5/utils/tlm/](https://github.com/gem5/gem5/tree/master/utils/tlm/)

❑ Slave Example:



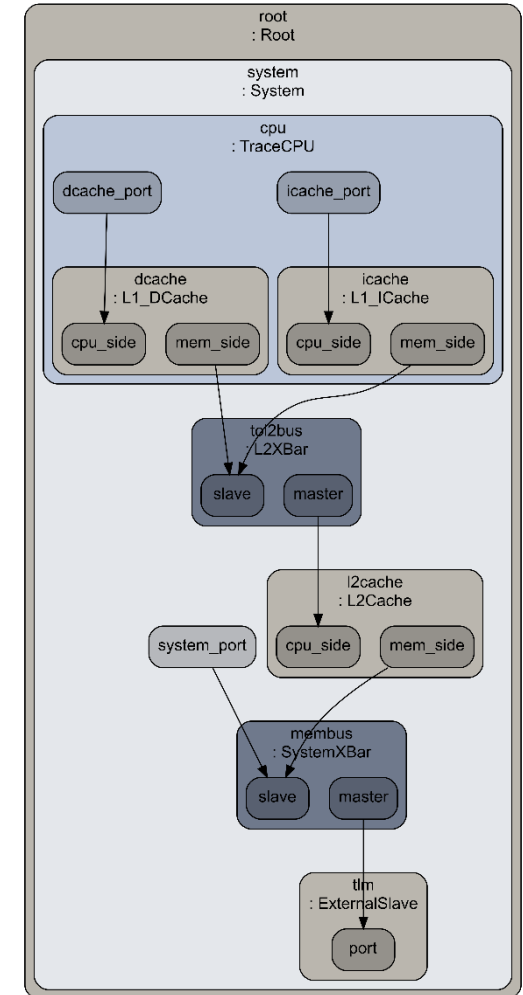
❑ Master Example:



❑ Elastic Trace Example [5] (see left)

❑ Full System Example:

```
../../../../build/ARM/gem5.opt ../../configs/example/fs.py \
--tlm-memory=transactor --cpu-type=TimingSimpleCPU --num-cpu=1 \
--mem-type=SimpleMemory --mem-size=512MB --mem-channels=1 --caches \
--l2cache --machine-type=VExpress_EMM \
--dtb-filename=vexpress.aarch32.11_20131205.0-gem5.1cpu.dtb \
--kernel=vmlinux.aarch32.11_20131205.0-gem5 \
--disk-image=linux-aarch32-ael.img
```



1. Compile gem5 normally:

```
scons build/ARM/gem5.opt
```
2. Compile gem5 as a library:

```
scons --with-cxx-config --without-python --without-tcmalloc \  
build/ARM/libgem5_opt.so
```
3. Include the gem5 modules `Gem5SimControl` and `Gem5SlaveTransactor` and/or `Gem5MasterTransactor` in your SystemC project and connect them to your SystemC models. Be sure to pass an individual port name to the constructor of each transactor.
4. Compile your project and link against the gem5 library.
5. Run normal gem5 with a custom python script or `fs.py` with `--tlm-memory=<port-name>` to generate `m5out/config.ini`. Be sure to set the `tlm_data` attribute of the External Masters/Slaves to the port name of the corresponding SystemC transactor.
6. Run your SystemC project and pass the `m5out/config.ini` file to your `Gem5SimControl` object.

Hands On: A Memory Module in SystemC



```
struct Target: public sc_module {
    // TLM interface socket:
    tlm_utils::simple_target_socket<Target> socket;

    // Storage
    unsigned char *mem;

    // Constructor
    Target(sc_core::sc_module_name name, /* ... */);
    SC_HAS_PROCESS(Target);

    // TLM interface functions
    virtual void b_transport(tlm::tlm_generic_payload& trans,
                            sc_time& delay);
    virtual unsigned int transport_dbg(tlm::tlm_generic_payload& trans);
    virtual tlm::tlm_sync_enum nb_transport_fw(
        tlm::tlm_generic_payload& trans,
        tlm::tlm_phase& phase,
        sc_time& delay);

    // ...
};
```

→ util/tlm/examples/slave_port/sc_target.hh

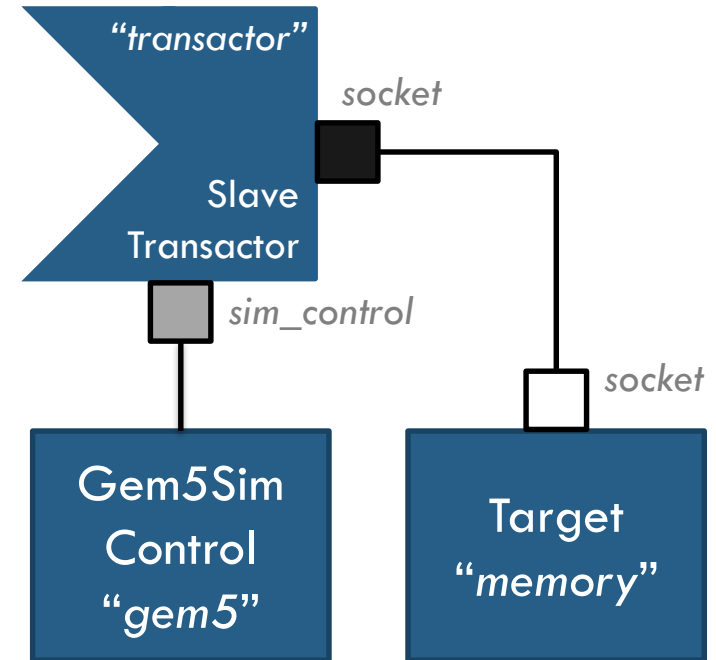
Hands On: Connect the Memory to gem5

```
int sc_main(int argc, char **argv)
{
    // Instantiate all modules
    Gem5SystemC::Gem5SimControl
        sim_control("gem5", /* config ... */);
    Gem5SystemC::Gem5SlaveTransactor
        transactor("transactor", "transactor");
    Target memory("memory", /* config ... */);

    // Bind modules
    memory.socket.bind(transactor.socket);
    transactor.sim_control.bind(sim_control);

    // Start simulation
    sc_core::sc_start();

    return EXIT_SUCCESS;
}
```



→ [util/tlm/examples/slave_port/main.cc](#)

Hands On: Configure gem5

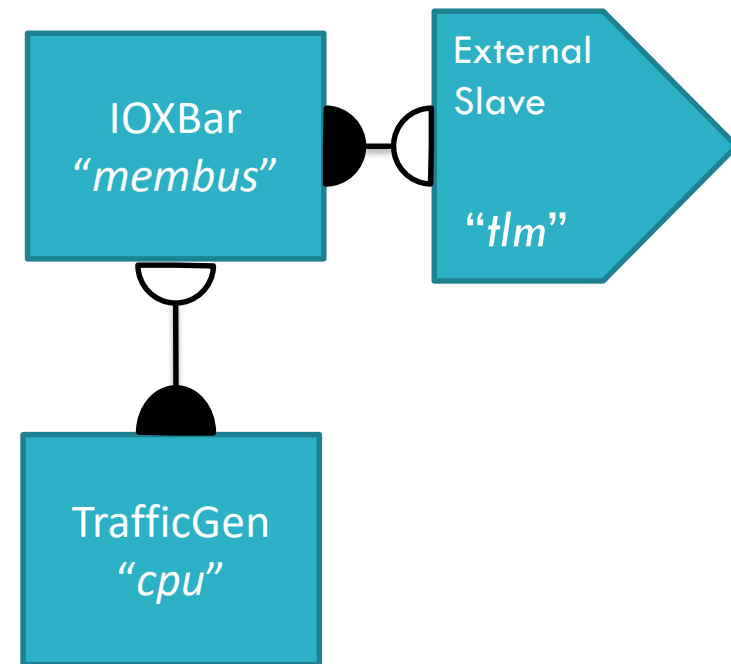


```
# Create a system with a Crossbar and a
TrafficGenerator
system = System()
system.membus = IOXBar(width = 16)
# This must be instantiated, even if not needed
system.physmem = SimpleMemory()
system.cpu = TrafficGen(config_file = "tgen.cfg")
system.clk_domain = SrcClockDomain(clock = '1.5GHz',
    voltage_domain = VoltageDomain(voltage = '1V'))

# Create an external TLM port:
system.tlm = ExternalSlave()
system.tlm.addr_ranges = [AddrRange('512MB')]
system.tlm.port_type = "tlm_slave"
system.tlm.port_data = "transactor"

# Route the connections:
system.cpu.port = system.membus.slave
system.system_port = system.membus.slave
system.membus.master = system.tlm.port

# Start the simulation:
root = Root(full_system = False, system = system)
root.system.mem_mode = 'timing'
m5.instantiate()
m5.simulate()
```



Hands On: Run the Simulation

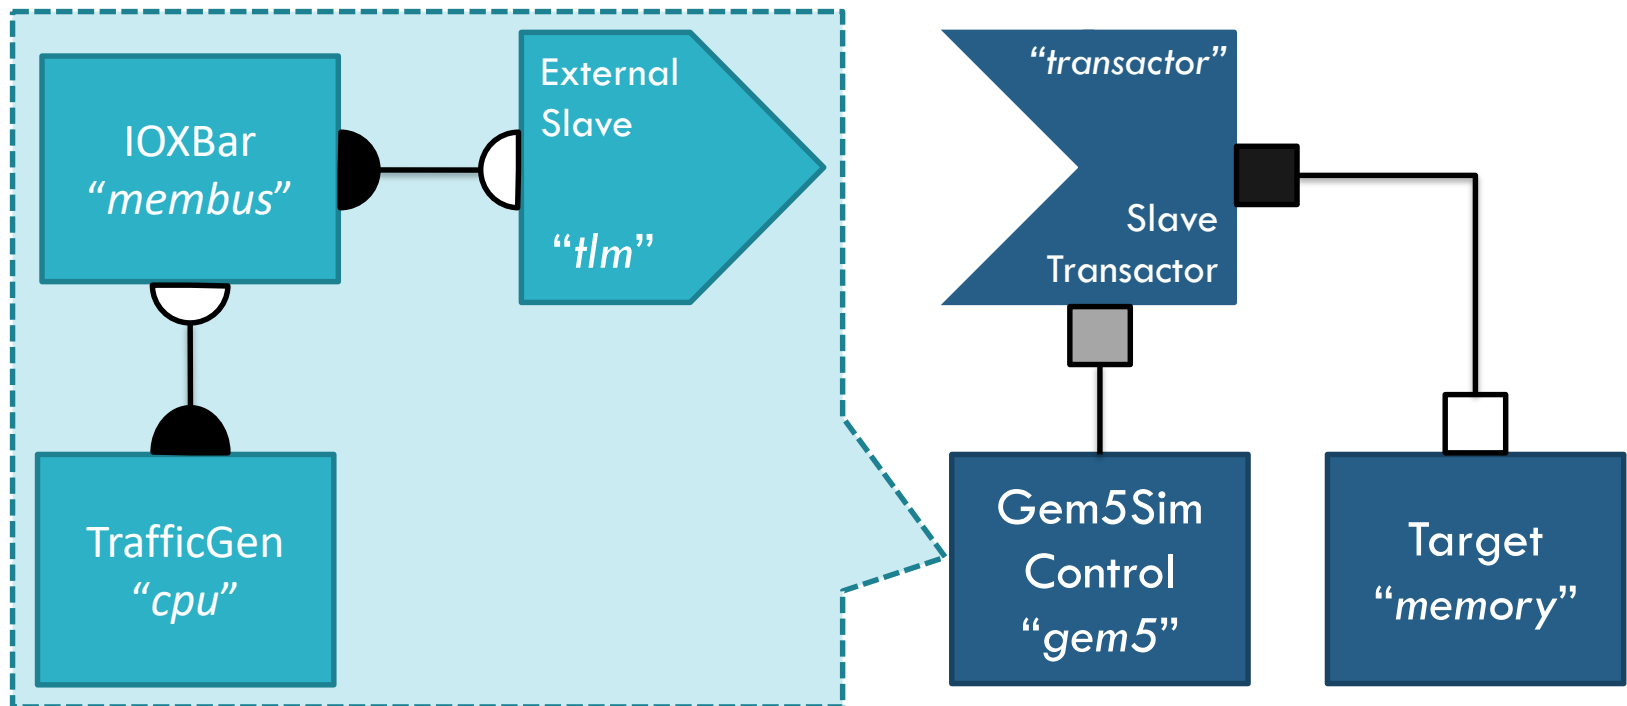
1. Build the example: `$ cd util/tlm && scons`

2. Create a gem5 config.ini file:

```
$ ../../build/ARM/gem5.opt conf/tlm_slave.py
```

3. Run the simulation:

```
$ build/examples/slave_port/gem5.sc m5out/config.ini
```



```
$ build/examples/slave_port/gem5.sc m5out/config.ini -e 200000 -d TrafficGen
[...]
```

0 s (=) : sc_main Start of Simulation
info: Entering event queue @ 0. Starting simulation...

5 ns (=) : system.cpu LinearGen::getNextPacket: r to addr 0, size 4
5 ns (=) : system.cpu Next event scheduled at 10000

10 ns (=) : system.cpu LinearGen::getNextPacket: w to addr 4, size 4
15 ns (=) : system.cpu Received retry
15 ns (=) : system.cpu LinearGen::getNextPacket: r to addr 8, size 4
16675 ps (=) : system.cpu Received retry
75 ns (=) : system.cpu Received retry
75 ns (=) : system.cpu LinearGen::getNextPacket: r to addr c, size 4
76038 ps (=) : system.cpu Received retry
135 ns (=) : system.cpu Received retry
135 ns (=) : system.cpu LinearGen::getNextPacket: r to addr 10, size 4
136068 ps (=) : system.cpu Received retry
195 ns (=) : system.cpu Received retry
195 ns (=) : system.cpu LinearGen::getNextPacket: w to addr 14, size 4
196098 ps (=) : system.cpu Received retry
Exit at tick 200000, cause: simulate() limit reached

➔ The binary accepts various options:

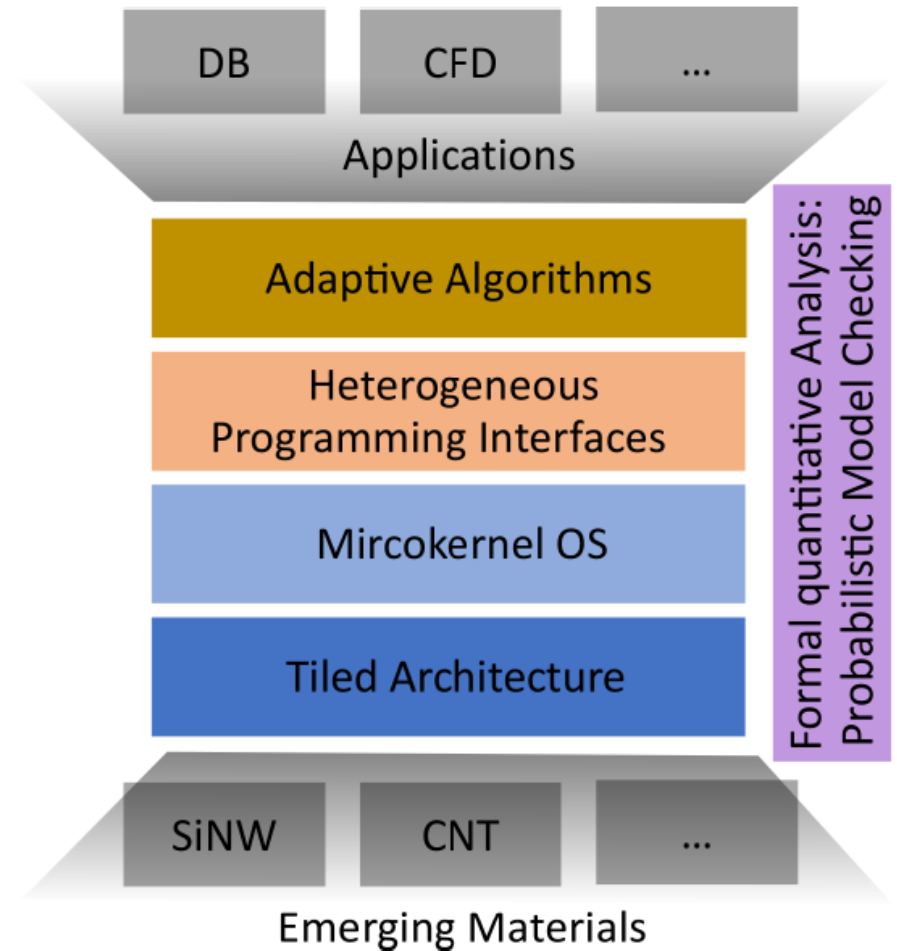
- -e end of simulation at tick
- -d set a gem5 debug flag

Usecases

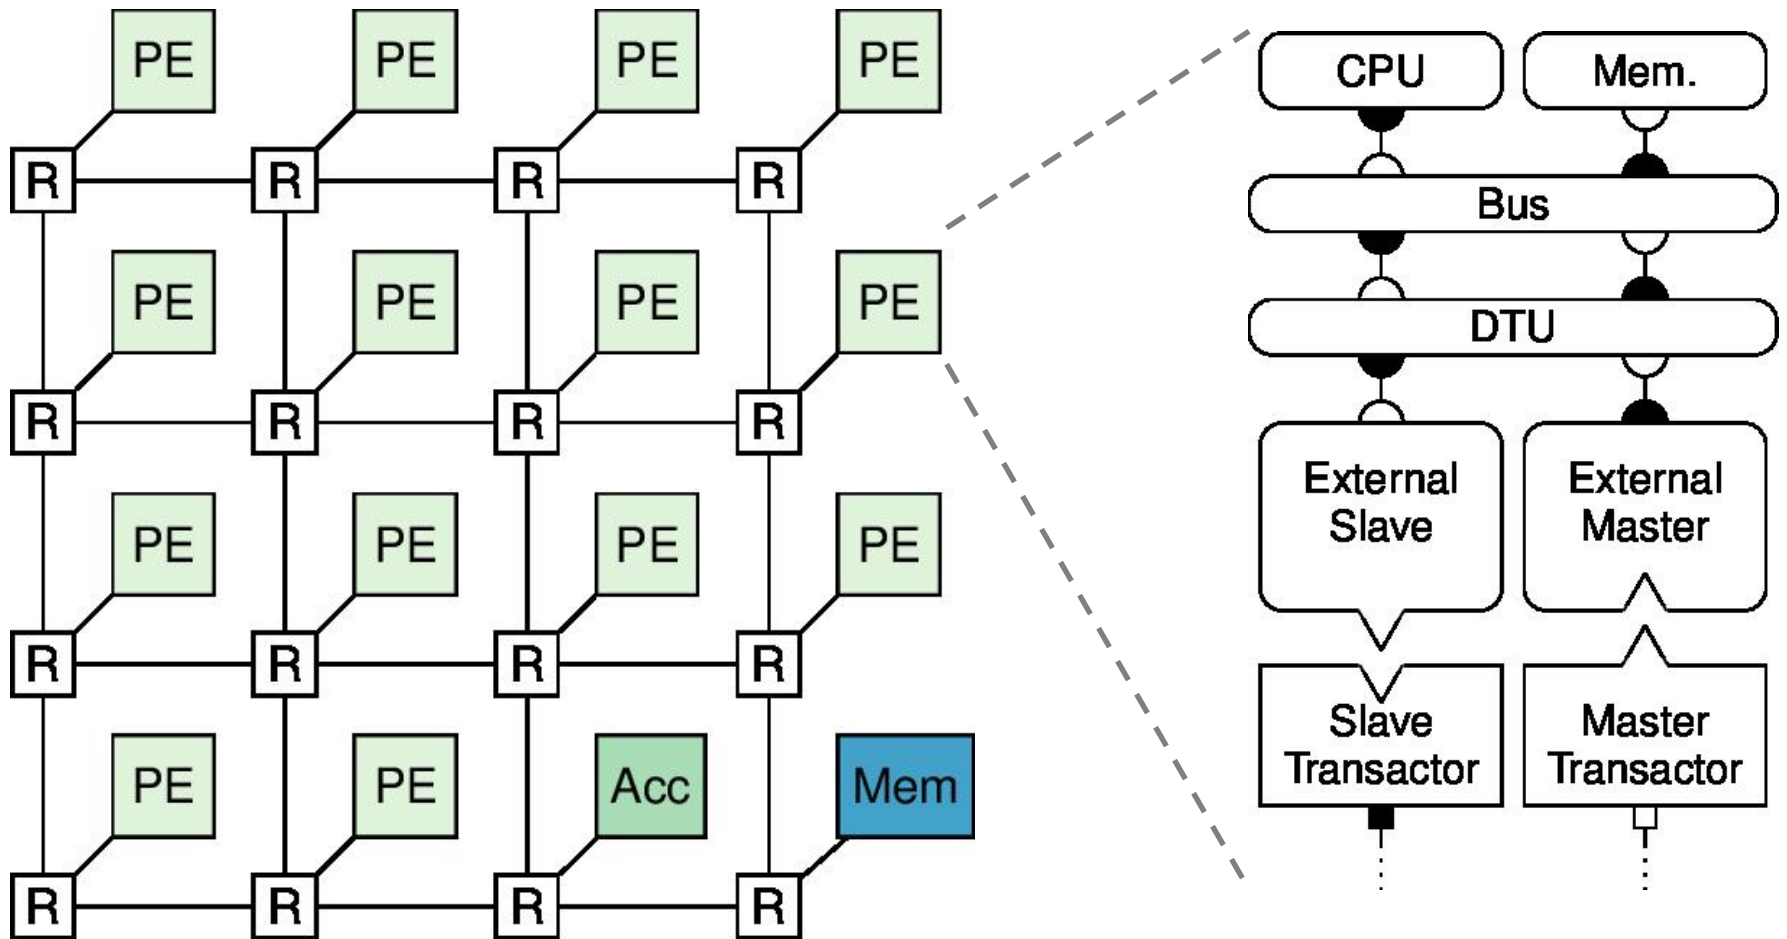
A programming stack for wildly heterogeneous systems including:

- ❑ dataflow programming models
- ❑ dataflow compiler
- ❑ adaptive runtime systems
- ❑ capability-based OS
- ❑ model checker

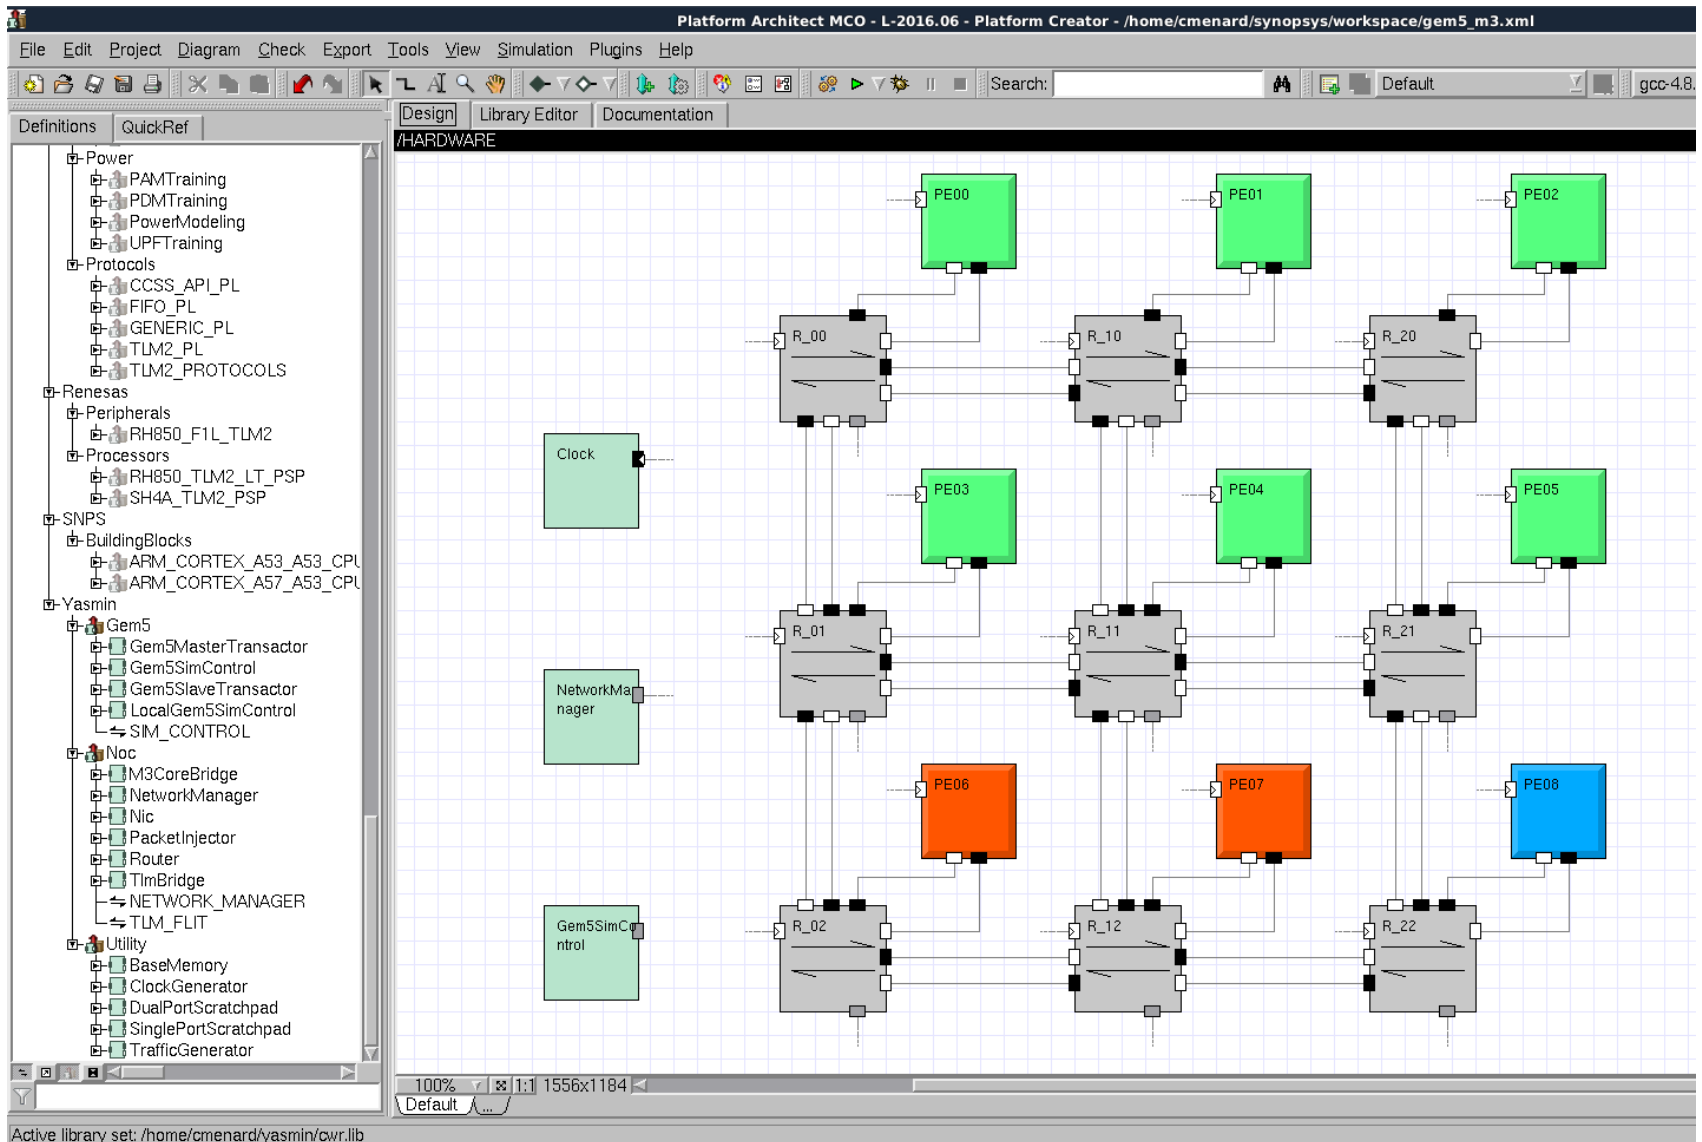
➔ A flexible simulation platform is required to try new designs and technologies.



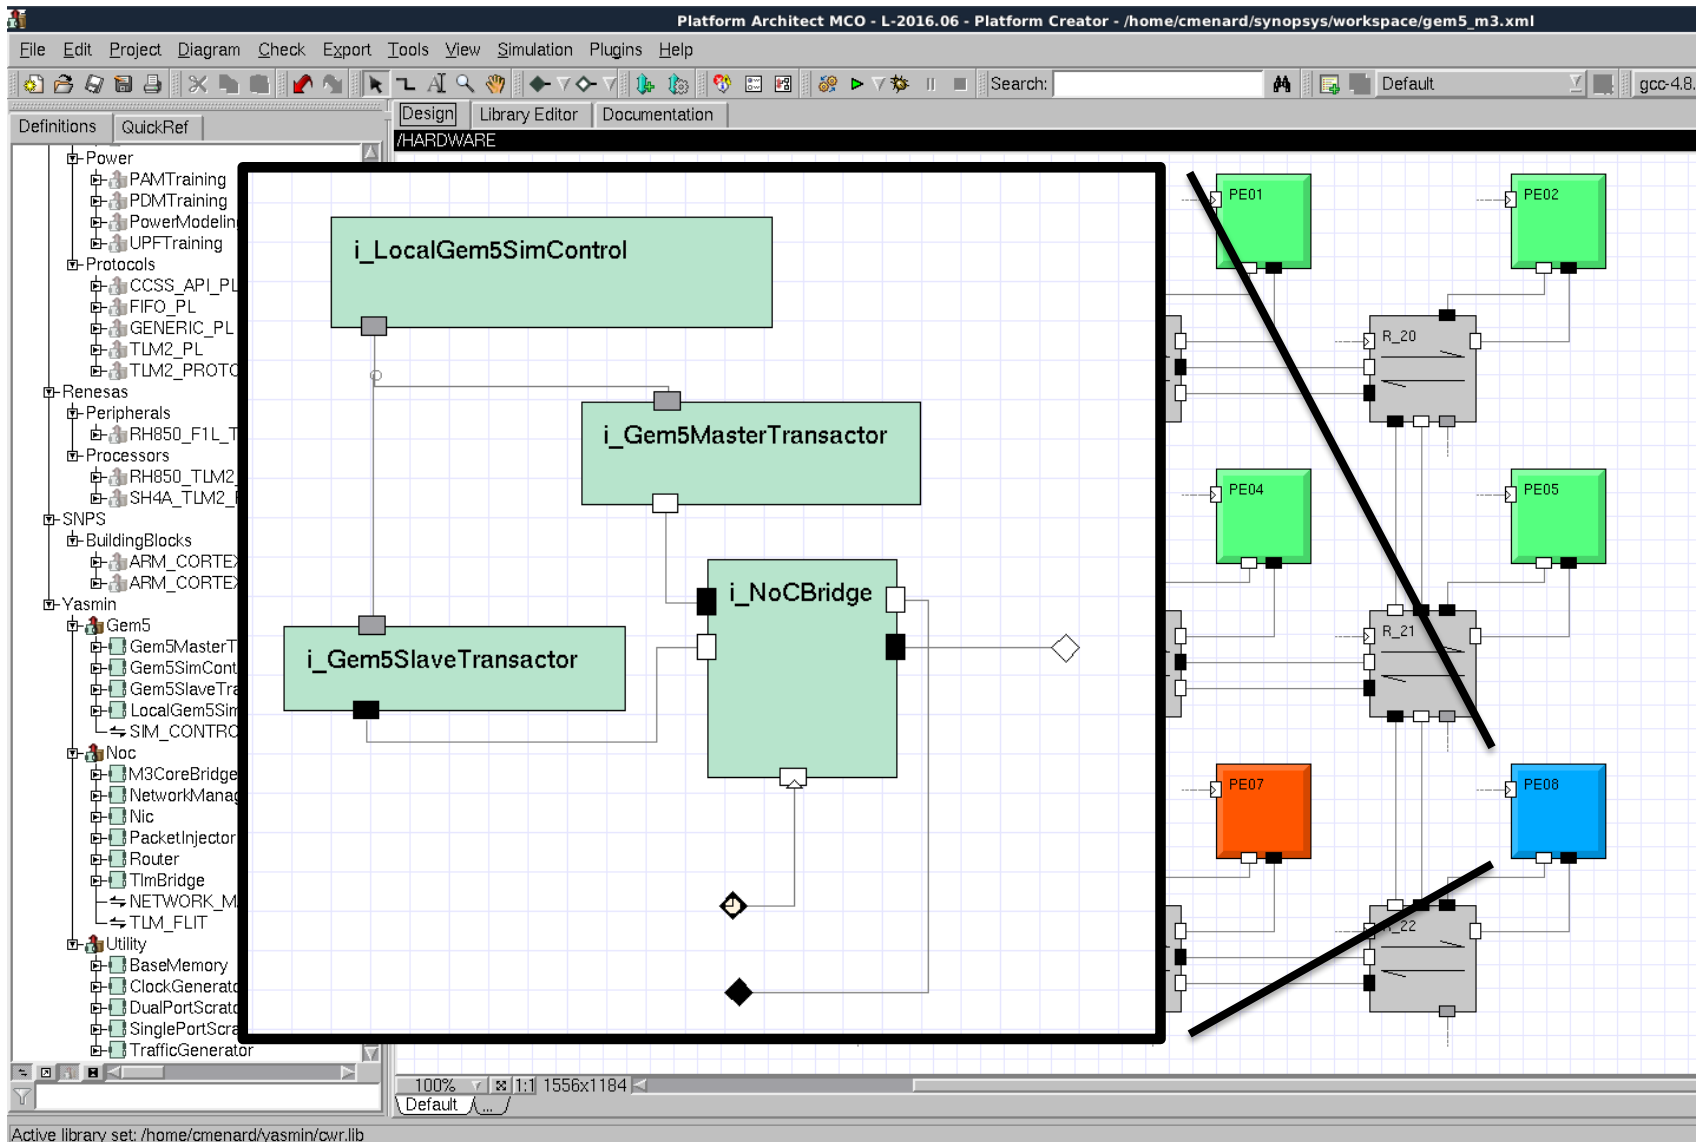
Building Heterogeneous MPSoCs: gem5 as a Tile



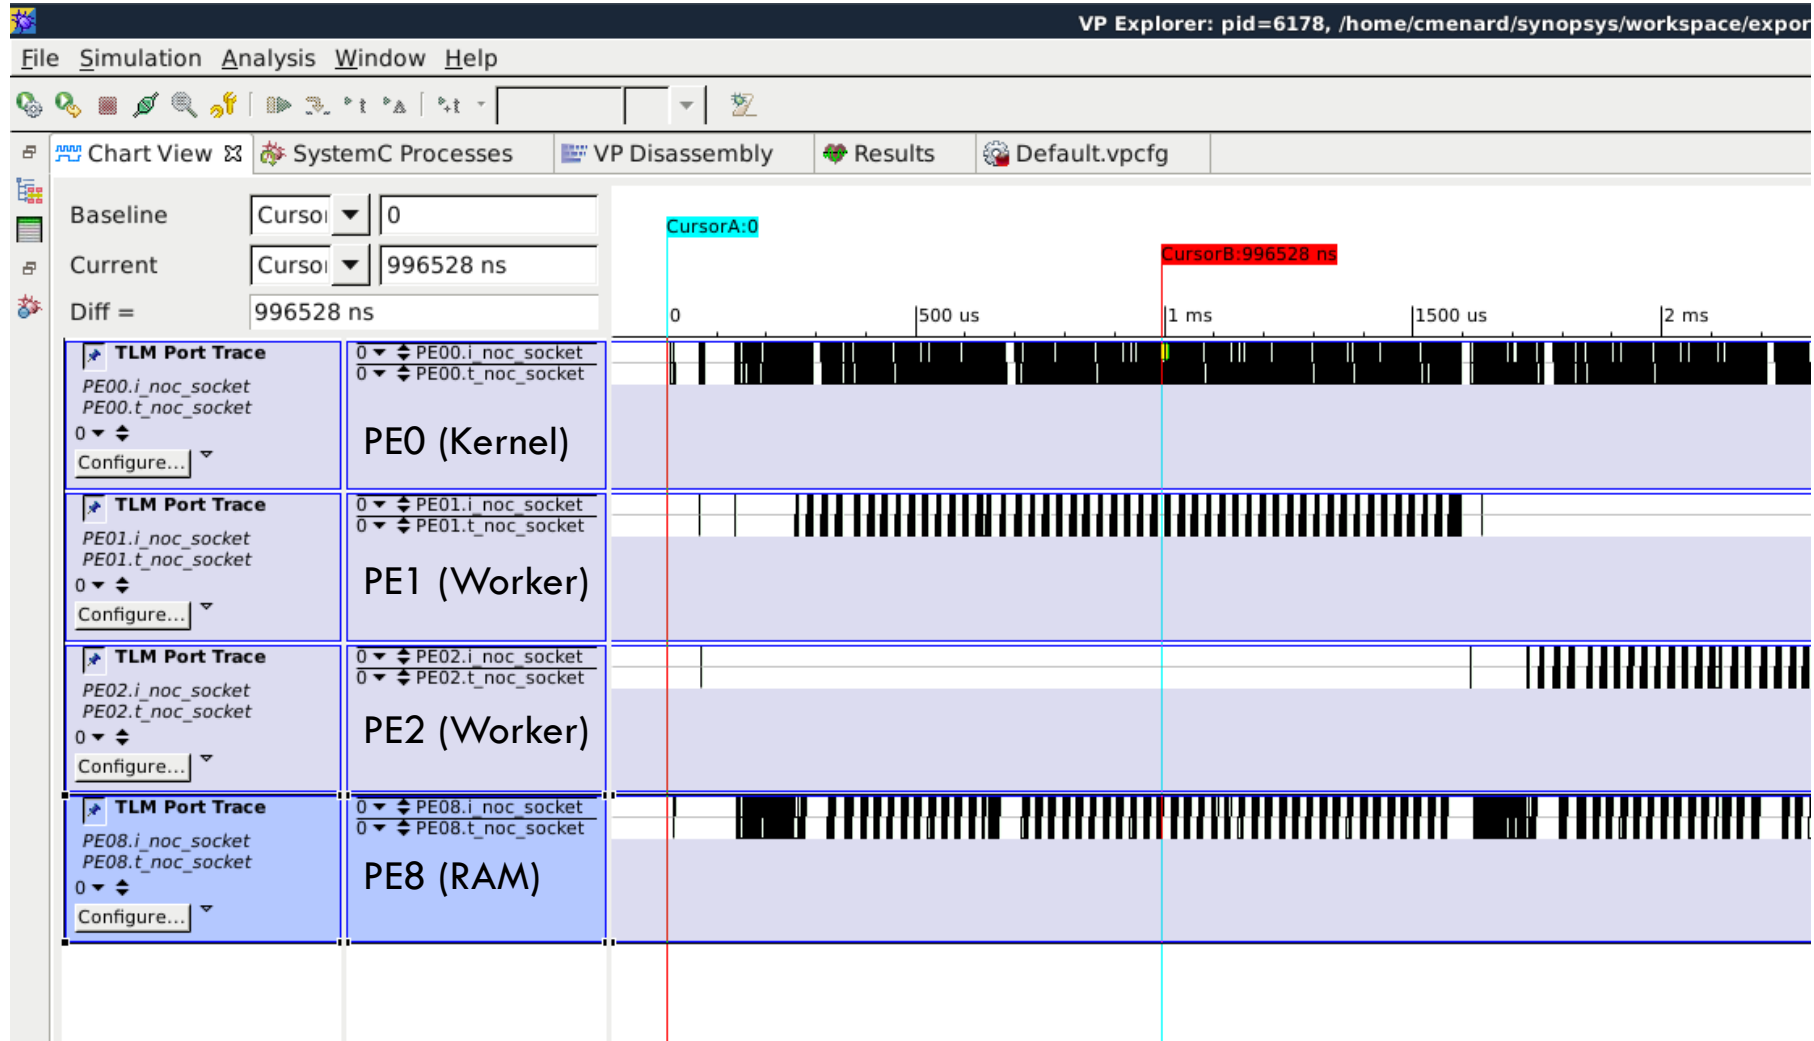
gem5 in Synopsys Platform Architect



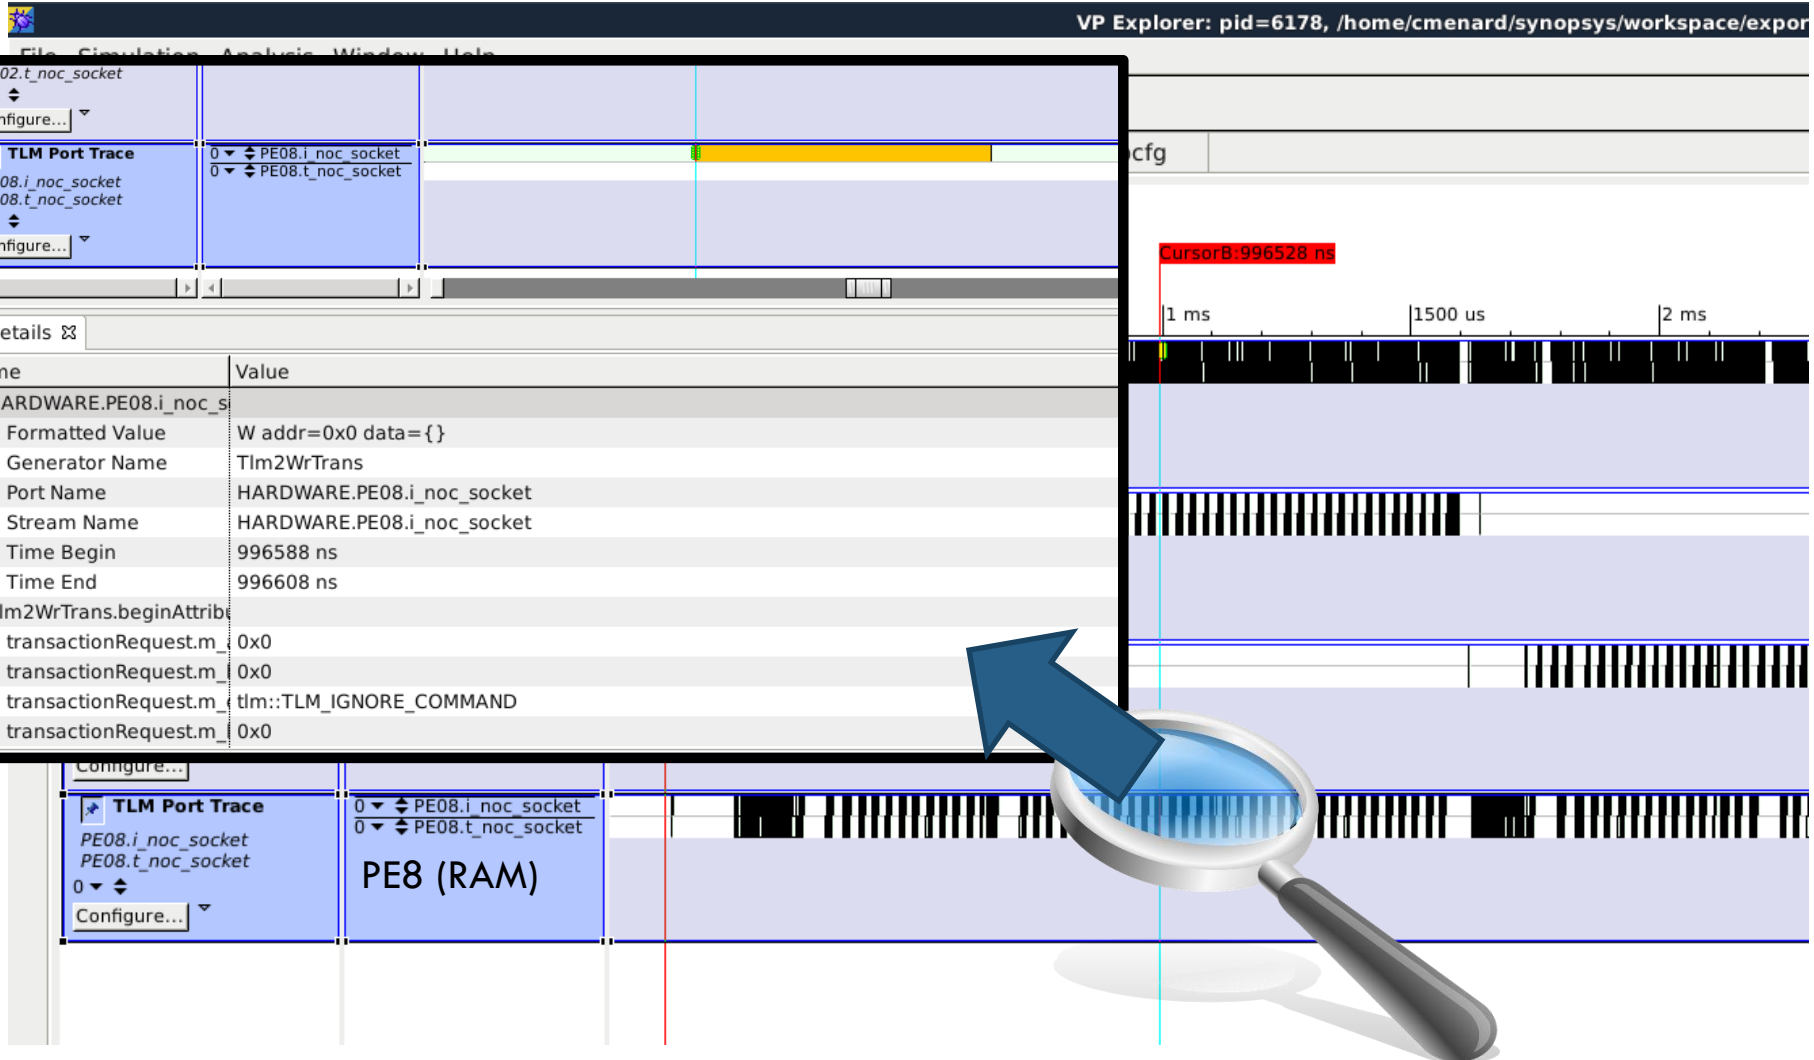
gem5 in Synopsys Platform Architect



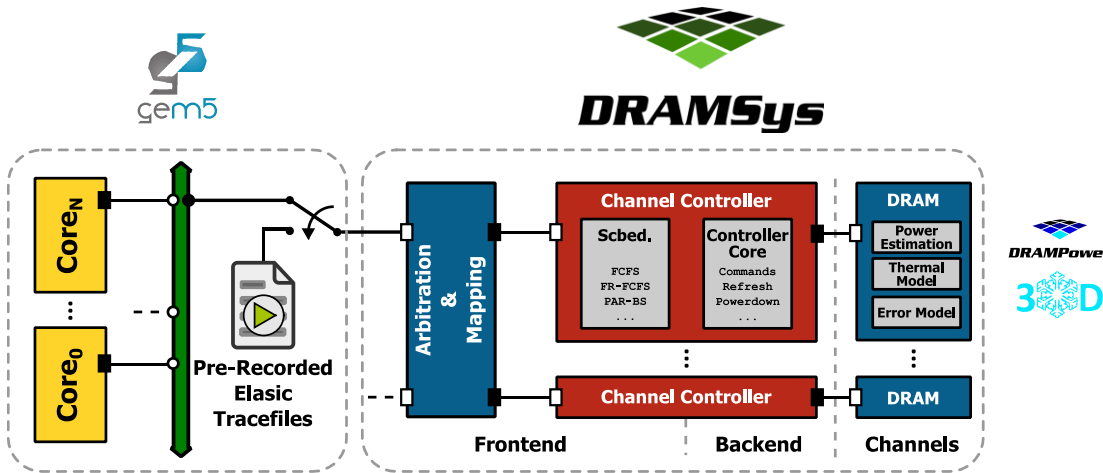
gem5 in Synopsys Platform Architect: Trace Analysis



gem5 in Synopsys Platform Architect: Trace Analysis



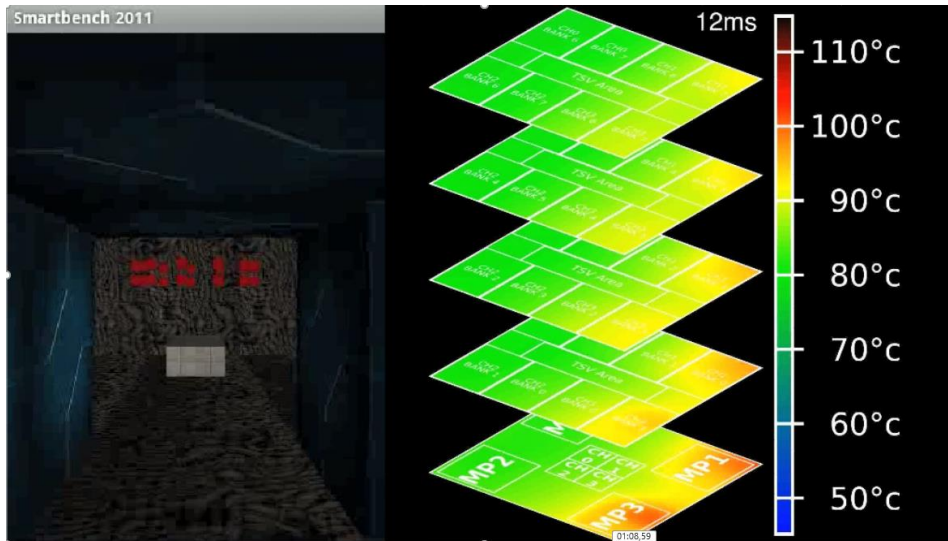
Coupling gem5 with DRAMSys [3], [4]



DRAMSys is a design space exploration framework for DRAM and memory controller

It includes:

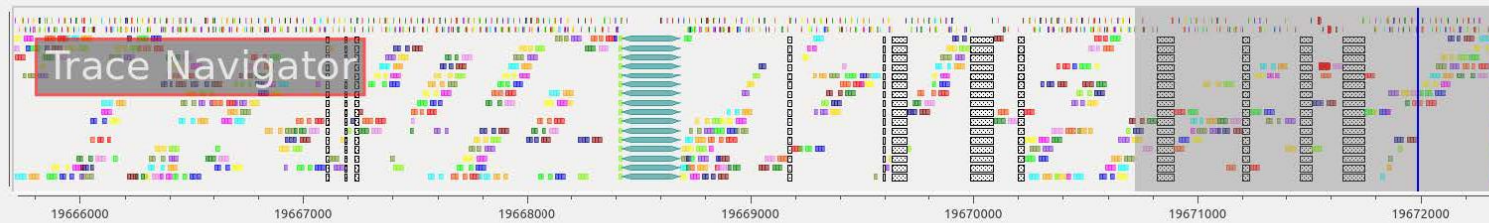
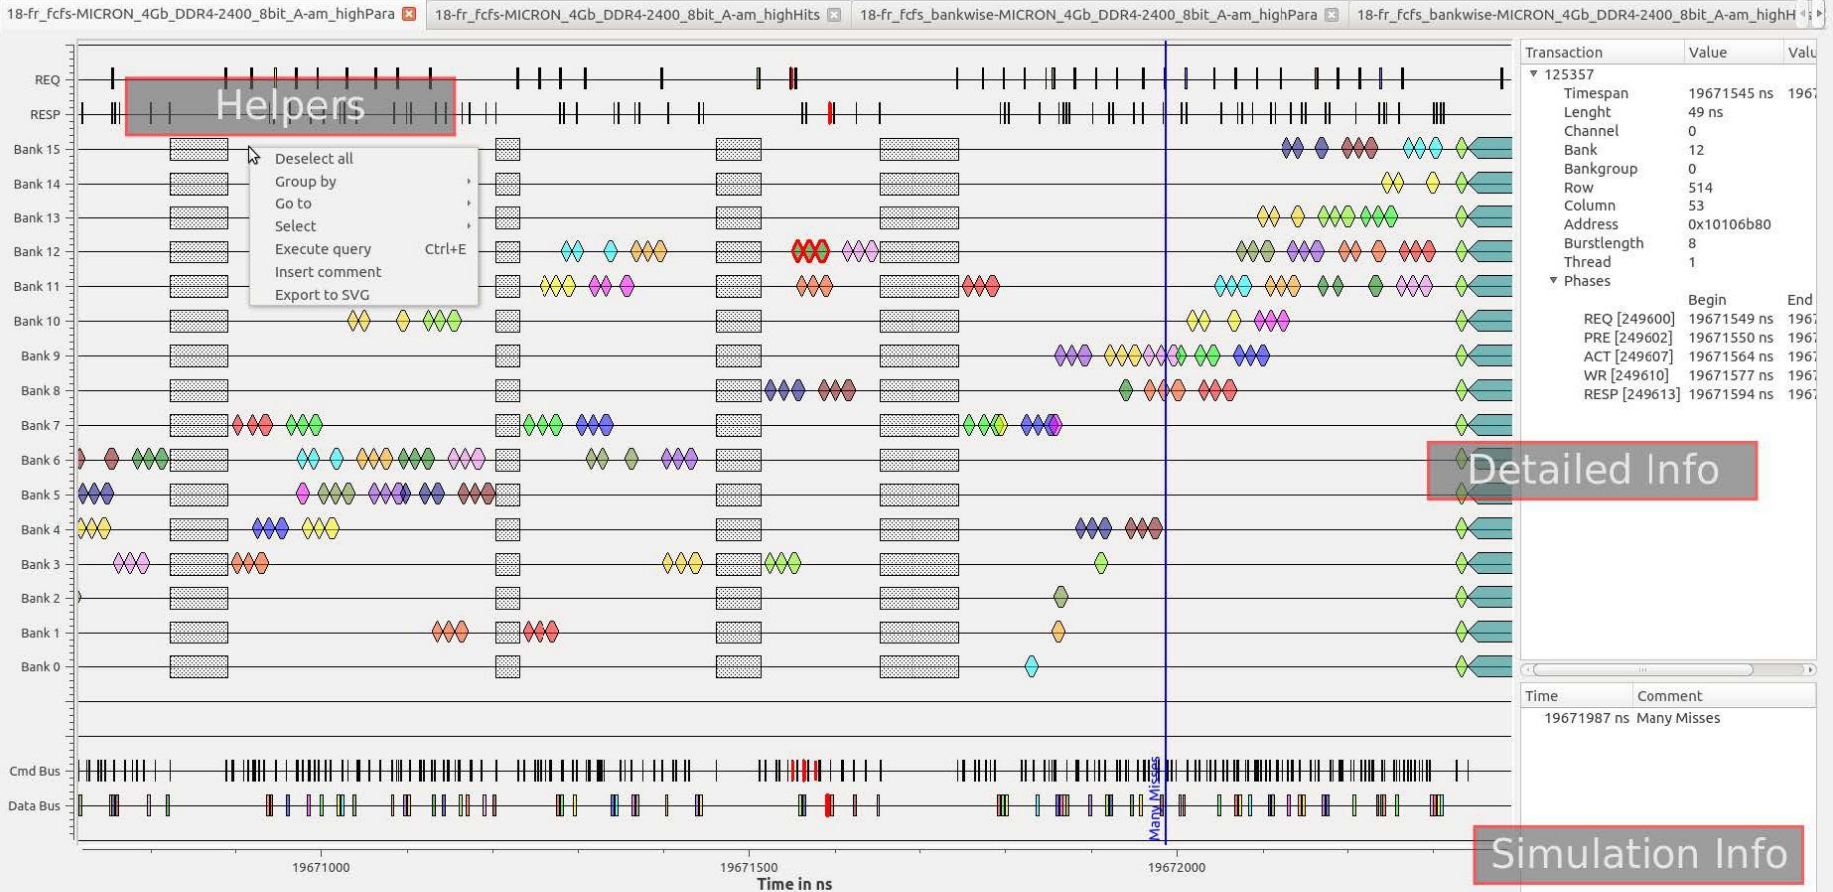
- Power model
- Thermal model
- Retention error model.



Linux boot (without thermal model) using the DRAMSys model → slowdown of 1.9×

This slowdown mostly comes from detailed DRAMSys model

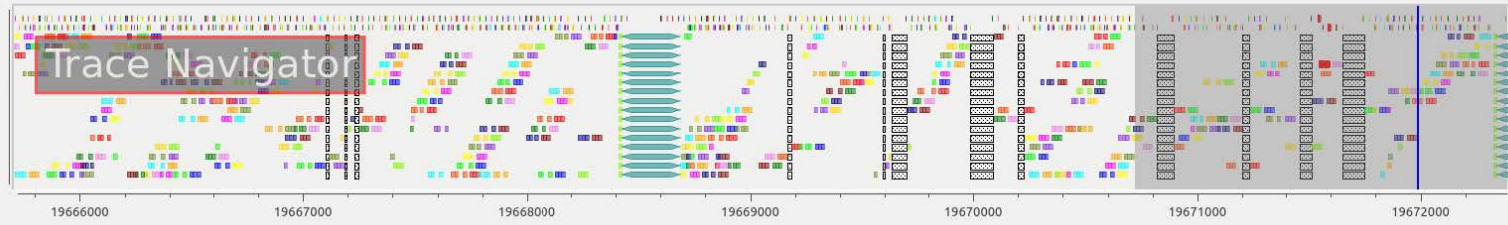
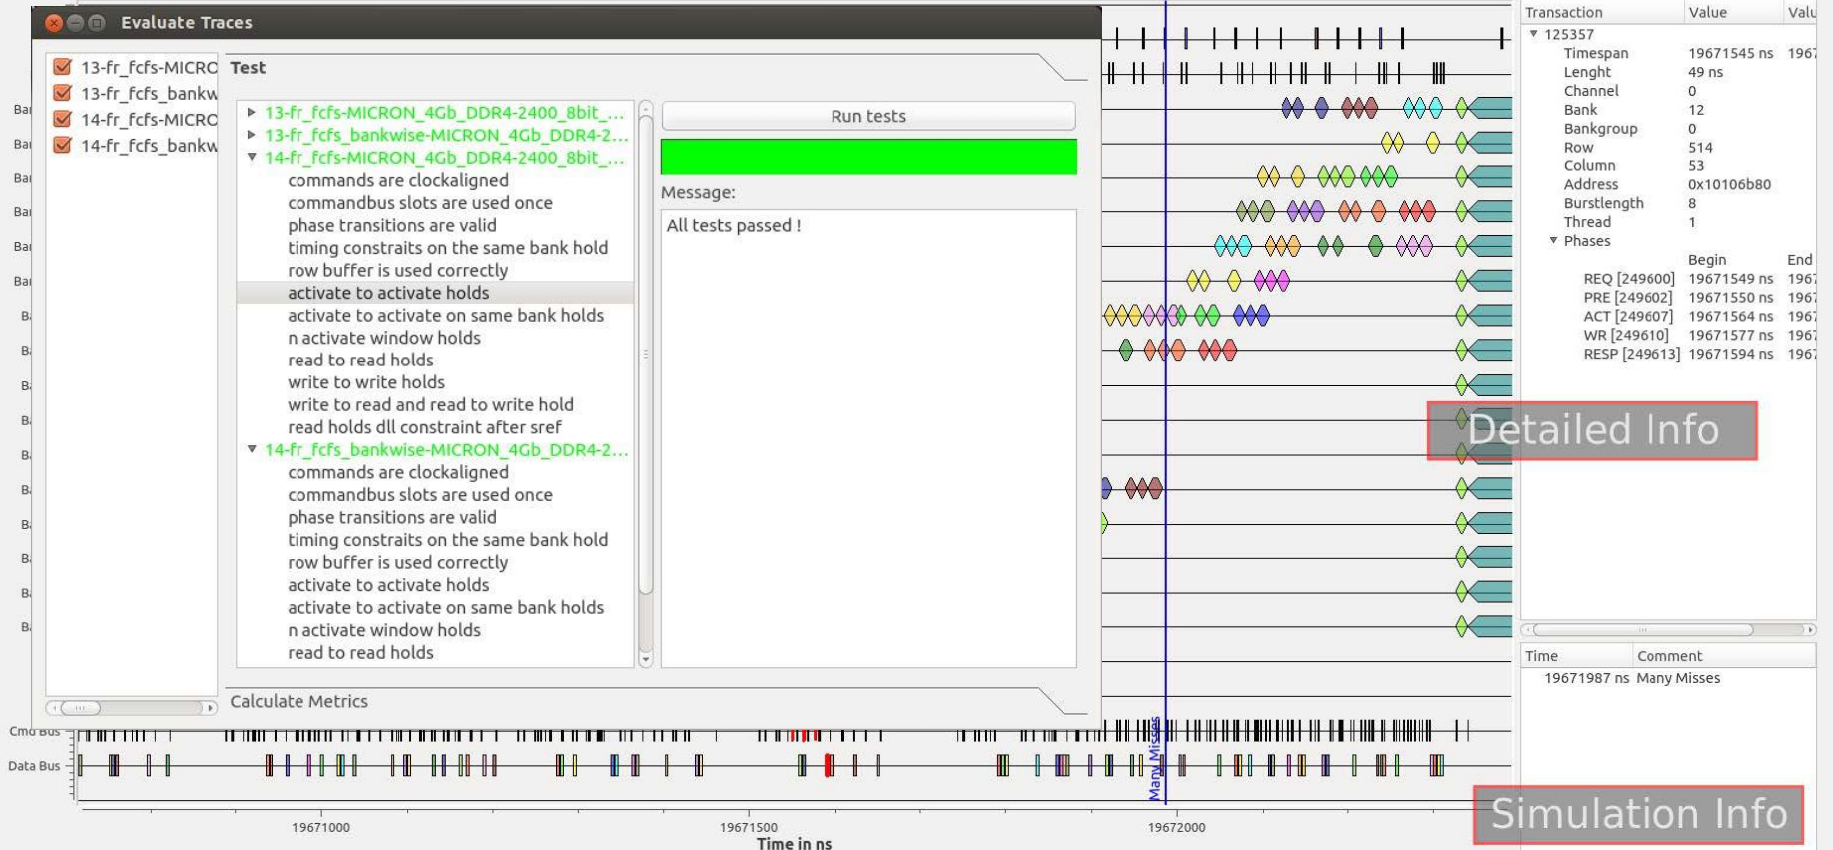
Coupling gem5 with DRAMSys Continued



Coupling gem5 with DRAMSys Continued



18-fr_fcfs-MICRON_4Gb_DDR4-2400_8bit_A-am_highPara 18-fr_fcfs-MICRON_4Gb_DDR4-2400_8bit_A-am_highHits 18-fr_fcfs_bankwise-MICRON_4Gb_DDR4-2400_8bit_A-am_highPara 18-fr_fcfs_bankwise-MICRON_4Gb_DDR4-2400_8bit_A-am_highH



Coupling gem5 with DRAMSys Continued



The screenshot displays the DRAMSys Evaluate Traces interface. The main window shows a list of tests on the left, a central message area, and a transaction timeline on the right. A secondary window in the foreground shows detailed metrics for three configurations.

Test Results:

- 13-fr_fcfs-MICRON_4Gb_DDR4-2400_8bit_A-am_highPara: commands are clockaligned, commandbus slots are used once, phase transitions are valid, timing constraints on the same bank hold, row buffer is used correctly, activate to activate holds, activate to activate on same bank holds, n activate window holds, read to read holds, write to write holds, write to read and read to write hold, read holds dll constraint after sref.
- 13-fr_fcfs_bankwise-MICRON_4Gb_DDR4-2400_8bit_A-am_highHits: (Message: All tests passed!)
- 14-fr_fcfs-MICRON_4Gb_DDR4-2400_8bit_A-am_highPara: (Message: All tests passed!)
- 14-fr_fcfs_bankwise-MICRON_4Gb_DDR4-2400_8bit_A-am_highHits: (Message: All tests passed!)

Transaction Details (Transaction 125357):

Transaction	Value	Value
Timespan	19671545 ns	19671545 ns
Lenght	49 ns	49 ns
Channel	0	0
Bank	12	12
Bankgroup	0	0
Row	514	514
Column	53	53
Address	0x10106b80	0x10106b80
Burstlength	8	8
Thread	1	1
Phases		
Begin		End
REQ [249600]	19671549 ns	19671549 ns
PRE [249602]	19671550 ns	19671550 ns
ACT [249607]	19671564 ns	19671564 ns

Calculate Metrics Results:

- 13-fr_fcfs-MICRON_4Gb_DDR4-2400_8bit_A-am_highPara**
 - average response latency in ns: 84.3
 - median response latency in ns: 37.6
 - number of activates: 16454
 - number of precharges: 108546
 - accesses per activate: 2.1
 - Active time (%): 6.79003
 - Time in PDNA (%): 27.334
 - Time in PDNP (%): 23.5452
 - Time in SREF (%): 42.3307
- 13-fr_fcfs_bankwise-MICRON_4Gb_DDR4-2400_8bit_A-am_highPara**
 - average response latency in ns: 126.1
 - median response latency in ns: 37.5
 - number of activates: 15623
 - number of precharges: 15614
 - accesses per activate: 2.2
 - Active time (%): 1.00676
 - Time in PDNA (%): 4.56337
 - Time in PDNP (%): 3.38575
 - Time in SREF (%): 91.0441
- 14-fr_fcfs-MICRON_4Gb_DDR4-2400_8bit_A-am_highPara**
 - average response latency in ns: 307.7
 - median response latency in ns: 304.9

Coupling gem5 with DRAMSys Continued



18-fr_fcfs-MICRON_4Gb_DDR4-2400_8bit_A-am_highPara 18-fr_fcfs-MICRON_4Gb_DDR4-2400_8bit_A-am_highHits 18-fr_fcfs_bankwise-MICRON_4Gb_DDR4-2400_8bit_A-am_highPara 18-fr_fcfs_bankwise-MICRON_4Gb_DDR4-2400_8bit_A-am_highH

Evaluate Traces

- 13-fr_fcfs-MICRO
- 13-fr_fcfs_bankw
- 14-fr_fcfs-MICRO
- 14-fr_fcfs_bankw

Test

- ▶ 13-fr_fcfs-MICRON_4Gb_DDR4-2400_8bit_...
- ▶ 13-fr_fcfs_bankwise-MICRON_4Gb_DDR4-2...
- ▼ 14-fr_fcfs-MICRON_4Gb_DDR4-2400_8bit_...
- com
- com
- phas
- timin
- row
- activ
- activ
- n act
- read
- write
- read

▼ 14-fr_fc...

- com
- com
- phas
- timin
- row
- activ
- activ
- n act
- read

Calculate M...

CmU bus

Data Bus

1967

Evaluate Traces

Run tests

CPU Bandwidth in Gbps

Bandwidth [Gbit/s]

Time [ns]

- CPU + GPU (SMS-SJF)
- CPU + GPU (STVQ)

GPU Bandwidth in Gbps

Bandwidth [Gbit/s]

Time [ns]

- CPU + GPU (SMS-SJF)
- CPU + GPU (STVQ)

Transaction	Value	Valu
125357		
Timespan	19671545 ns	1967
Lenght	49 ns	
Channel	0	
Bank	12	
Bankgroup	0	
Row	514	
Column	53	
Address	0x10106b80	
Burstlength	8	
Thread	1	
Phases		
Begin		End
REQ [249600]	19671549 ns	1967
PRE [249602]	19671550 ns	1967
ACT [249607]	19671564 ns	1967

Calculate metrics Export to CSV

n_highPara

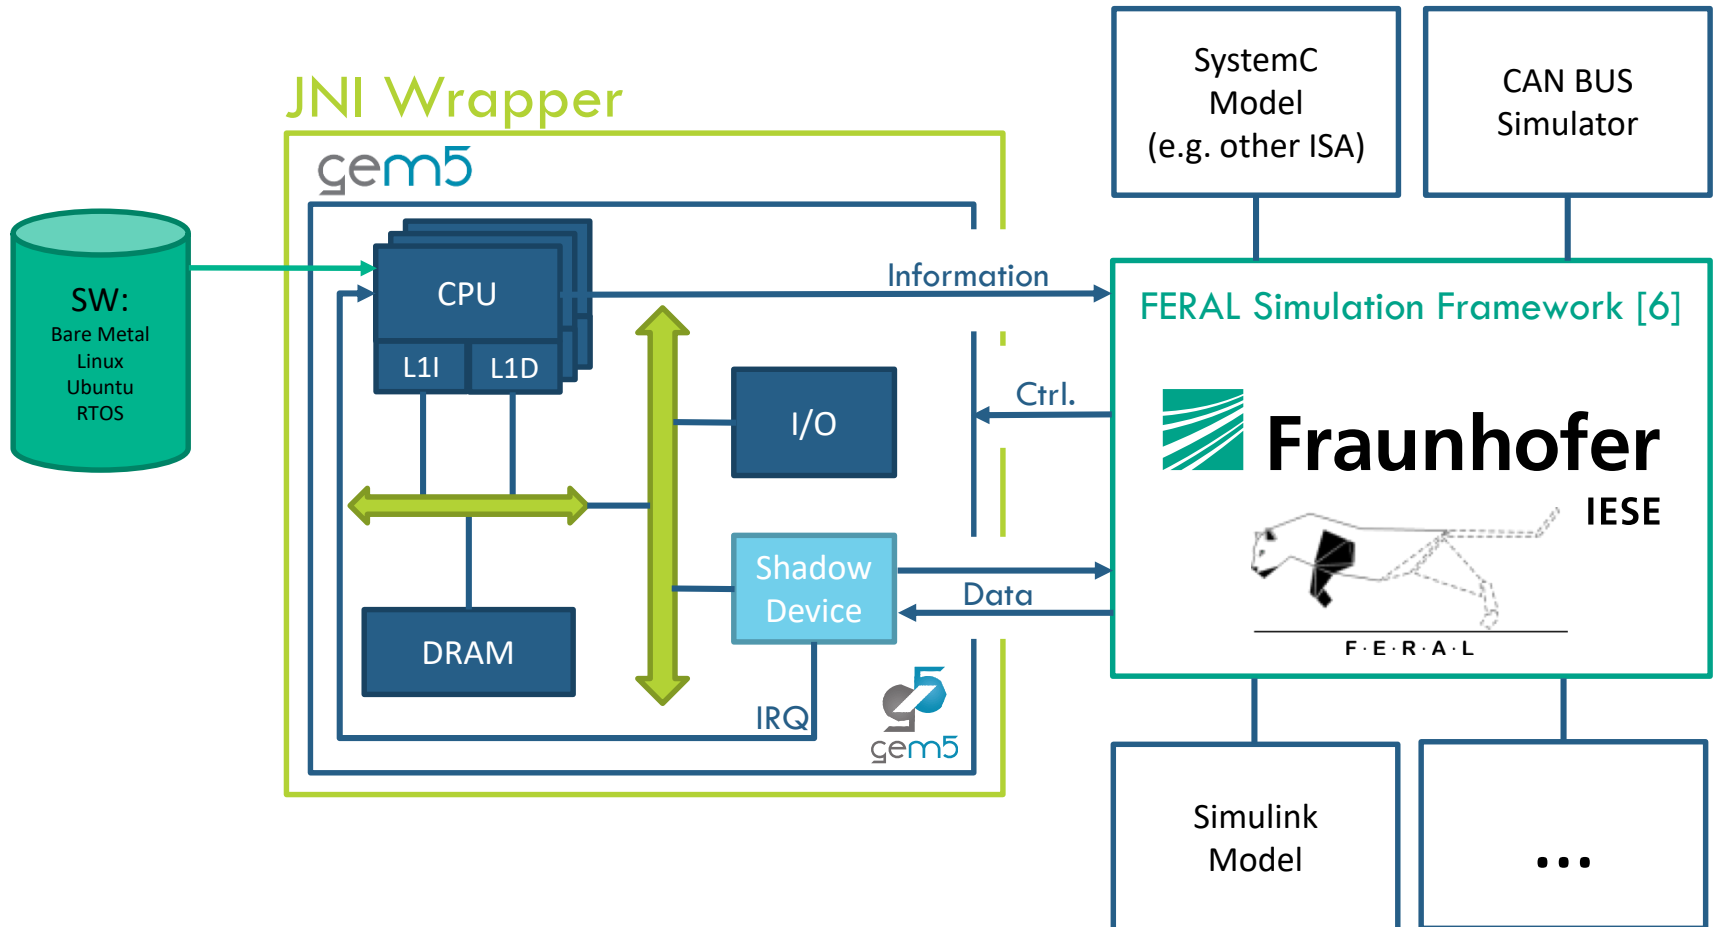
Time in PDNA (%): 4.30337
Time in PDNP (%): 3.38575
Time in SREF (%): 91.0441

▼ 14-fr_fcfs-MICRON_4Gb_DDR4-2400_8bit_A-am_highPara
average response latency in ns: 307.7
median response latency in ns: 304.9

Trace Navigator

19666000 19667000 19668000 19669000 19670000 19671000 19672000

Connecting gem5 to Non-SystemC Simulators



- ❑ Coupling of different Simulators with different models of computation (e.g. Simulink with gem5 or SystemC)
- ❑ gem5 compiled as C++ library and wrapped in JNI wrapper

- ❑ Development of software concepts
 - ❑ Simulation of systems of systems
 - ❑ Combining different levels of abstraction

- ❑ Software Testing:
 - ❑ Normal software testing instruments source code or binary → Intrusive
 - ❑ gem5 is instrumented instead
 - ❑ Supervised testing of concurrent software
 - ❑ Fault Injection
 - ❑ Coverage

- ❑ Full interoperability between gem5 and SystemC
 - ❑ Fully compliant to the SystemC standard
 - ❑ It is part of the gem5 repository!
- Next steps:
- replace the entire simulation kernel and communication System by SystemC/TLM (?)
 - Remove step for .ini generation (?)
 - Suggestions? We are open!

Thank you!

- [1] Christian Menard, Matthias Jung, Jeronimo Castrillon, Norbert Wehn, "**System Simulation with gem5 and SystemC: The Keystone for Full Interoperability**", *Proceedings of the IEEE International Conference on Embedded Computer Systems Architectures Modeling and Simulation (SAMOS)*, Jul 2017.
- [2] Marcus Völp, Sascha Klüppelholz, Jeronimo Castrillon, Hermann Härtig, et al. "**The Orchestration Stack: The Impossible Task of Designing Software for Unknown Future Post-CMOS Hardware**", *Proceedings of the 1st International Workshop on Post-Moore's Era Supercomputing (PMES), Co-located with The International Conference for High Performance Computing, Networking, Storage and Analysis (SC16)*, Salt Lake City, USA, Nov 2016
- [3] Matthias Jung, Christian Weis, Norbert Wehn, Karthik Chandrasekar, "**TLM Modelling of 3D Stacked Wide I/O DRAM Subsystems, A Virtual Platform for Memory Controller Design Space Exploration**" *Rapid Simulation and Performance Evaluation: Methods and Tools (RAPIDO)*, January, 2013, Berlin.
- [4] Matthias Jung, Christian Weis, Norbert Wehn, "**DRAMSys: A flexible DRAM Subsystem Design Space Exploration Framework**", *IPSI Transactions on System LSI Design Methodology (T-SLDM)*, October, 2015.
- [5] Radhika Jagtap, Stefan Diestelhorst, Andreas Hansson, Matthias Jung, Norbert Wehn, "**Exploring System Performance using Elastic Traces: Fast, Accurate and Portable**", *IEEE International Conference on Embedded Computer Systems Architectures Modeling and Simulation (SAMOS)*, July, 2016, Samos Island, Greece.
- [6] Thomas Kuhn, Thomas Forster, Tobias Braun, Reinhard Gotzhein, "**FERAL - Framework for Simulator Coupling on Requirements and Architecture Level**", *IEEE Computer Society: Eleventh ACM/IEEE International Conference on Formal Methods and Models for Codesign, MEMOCODE 2013*. Proceedings : 8-20 Oct. 2013, Portland, OR