

The gem5 Tutorial @ISCA 2024



Plan for the day

- ▶ Intro
 - ▶ Break
- ▶ Using gem5, standard library, and gem5-resources
- ▶ **[Alen Sabu]** Elfies and gem5
 - ▶ Lunch
- ▶ Extending gem5
 - ▶ Break
- ▶ **[Matt S., Matt P., & Vishnu]** Running ML workloads and gem5's GPU model

Introduction

In this section we'll talk about gem5's history, the purpose and uses of computer architecture simulation, some nomenclature, and gem5's software architecture



Outline

- ▶ What is gem5 and a bit of history
- ▶ My perspective on architecture simulation
- ▶ gem5's (software) architecture
- ▶ Getting started using gem5



Created at Michigan by students of Steve Reinhardt, principally Nate Binkert.

“A tool for simulating systems”



Two Views of M5

1. A framework for event-driven simulation
 - Events, objects, statistics, configuration
 2. A collection of predefined object models
 - CPUs, caches, busses, devices, etc.
-
- This tutorial focuses on #2
 - You may find #1 useful even if #2 is not





Created at Michigan by students of Steve Reinhardt, principally Nate Binkert.

“A tool for simulating systems”



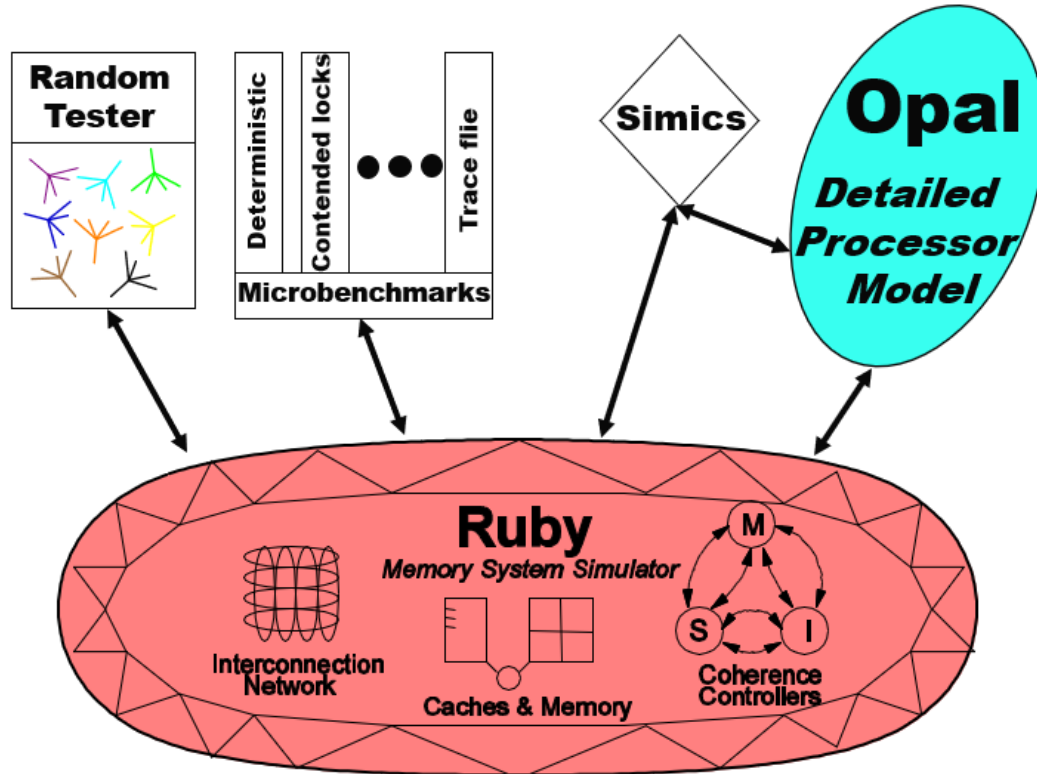
Multifacet GEMS

General Execution-driven Multiprocessor Simulator

Created at Wisconsin by students of Mark Hill and David Wood.

Detailed memory system

GEMS From 50,000 Feet





Created at Michigan by students of Steve Reinhardt, principally Nate Binkert.

“A tool for simulating systems”



Multifacet GEMS

General Execution-driven Multiprocessor Simulator

Created at Wisconsin by students of Mark Hill and David Wood.

Detailed memory system

What is gem5?

Michigan m5 + Wisconsin GEMS = gem5

“The gem5 simulator is a modular platform for computer-system architecture research, encompassing system-level architecture as well as processor microarchitecture.”

Lowe-Power et al. **The gem5 Simulator: Version 20.0+**. ArXiv Preprint ArXiv:2007.03152, 2021.
<https://doi.org/10.48550/arXiv.2007.03152>

Nathan Binkert, Bradford Beckmann, Gabriel Black, Steven K. Reinhardt, Ali Saidi, Arkaprava Basu, Joel Hestness, Derek R. Hower, Tushar Krishna, Somayeh Sardashti, Rathijit Sen, Korey Sewell, Muhammad Shoaib, Nilay Vaish, Mark D. Hill, and David A. Wood. 2011. **The gem5 simulator**. *SIGARCH Comput. Archit. News* 39, 2 (August 2011), 1-7.
DOI=<http://dx.doi.org/10.1145/2024716.2024718>

► gem5-20+: A new era in CA simulation




Version 20.0

Abdul Mutaal	Bagus	Curtis Dunham	Gabe Black	Jakub Jermar	Koan-Sin Tan	Martinasso	Nicolas Zea	Riken Gohil	Srikant	Uri Wiener
Ahmad	Hanindhito	Dam Sunwoo	Gabe Loh	James Clarkson	Korey Sewell	Maximilian Stein	Nikos Nikoleris	Rizwana Begum	Bharadwaj	Victor Garcia
Adrian Herrera	Benjamin Nash	Dan Gibson	Gabor Doza	Jan-Peter	Krishnendra	Maximilien	Nils Asmussen	Robert Kovacsics	Stan Czerniawski	Vilas Sridharan
Adrien Pesle	Bertrand	Daniel Carvalho	Gedare Bloom	Larsson	Nathella	Breughe	Nuwan Jayasena	Robert Scheffel	Stanislaw	Vince Weaver
Adrià Armejach	Marquis	Daniel Johnson	Gene WU	Jason Lowe-	Lena Olson	Michael Adler	Ola Jeppsson	Rohit Kurup	Czerniawski	Vincentius
Akash Bagdia	Binh Pham	Daniel Sanchez	Gene Wu	Power	Lisa Hsu	Michael LeBeane	Omar Naji	Ron Dreslinski	Stephan	Robby
Alec Roelke	Bjoern A. Zeeb	David Guillen-	Geoffrey Blake	Javier Bueno	Lluç Alvarez	Michael	Pablo Prieto	Ruben	Diestelhorst	Wade Walker
Alexandru Dutu	Blake Hechtman	Fandos	Georg Kotheimer	Hedo	Lluís Vilanova	Levenhagen	Palle Lyckegaard	Ayrapetyan	Stephen Hines	Weiping Liao
Ali Jafri	Bobby R. Bruce	David Hashe	Giacomo	Javier Cano-Cano	Mahyar Samani	Michiel Van Tol	Pau Cabre	Rune Holm	Steve Raasch	Wendy Elsasser
Ali Saidi	Yori Mingarov	David Oehmke	Gabrielli	Javier Setoain	Malek Musleh	Raquel Serrano	Paul Rosenfeld	Richard H. ...	Stephan ...	William Wang
Amin Farmahini	Beck	Derek ...	Giacomo ...	Jaroslav ...	Mark ...	Patrick ...	Paul ...	Ryan ...	Shu ...	Willy Wolff
Anders Handler	Benjamin ...	Daniel ...	Travis ...	Wesley ...	Yuan ...	Zhen ...	John ...	Polina ...	Sandipan ...	Swapnil ...
Andrea Mondelli	Brandon ...	Djordje	Hamid Reza	Jieming Yin	Marco Balboni	Mingyuan	Polina Dudnik	Sandipan Das	Swapnil Haria	Ma
Andrea Pellegrini	Brandon Potter	Djordje	Hamid Reza	Jieming Yin	Marco Balboni	Mingyuan	Polina Dudnik	Sandipan Das	Swapnil Haria	Ma
Andreas Hansson	Brian Grayson	Kovacevic	Khaleghzadeh	Jing Qu	Marco Elver	Mitch Hayenga	Polydoros	Santi Galan	Taeho Kgil	...
Andreas	Cagdas Dirik	Dongxue Zhang	Hanhwi Jang	Jiuyue Ma	Marjan Fariborz	Mohammad	Petrakis	Sascha Bischoff	Tao Zhang	...
Sandberg	Chander	Doğukan	Hoa Nguyen	Joe Gross	Matt DeVuyst	Alian	Pouya Fotouhi	Sean McGoogan	Thomas Grass	...
Andrew Bardsley	Sudanathi	Korkmaztürk	Hongil Yoon	Joel Hestness	Matt Evans	Monir	Prakash	Sean Wilson	Thomas Mück	...
Andrew Lukefahr	Chen Zou	Dylan Johnson	Hsuan Hsu	John Alsop	Matt Horsnell	Mozumder	Ramrakhiani	Sergei Trofimov	Tim Harris	...
Andrew Schultz	Chris Adeniyi-	Earl Ou	Hussein	John	Matt Poremba	Moyang Wang	Pritha Ghoshal	Severin	Timothy Hayes	...
Andriani	Jones	Edmund Grimley	Elnawawy	Kalamatianos	Matt Sinclair	Mrinmoy Ghosh	Radhika Jagtap	Wischmann	Timothy M.	...
Mappoura	Chris Emmons	Evans	Ian Jiang	Jordi Vaquero	Matteo	Nathan Binkert	Rahul Thakur	Shawn Rosti	Jones	...
Ani Udipi	Christian Menard	Emilio Castillo	IanJiangICT	Jose Marinho	Andreozzi	Nathanael	Reiley Jeapaul	Sherif Elhabbal	Tom Jablin	...
Anis Peysieux	Christoph Pfister	Erfan Azarkhish	Ilias Vougioukas	Jui-min Lee	Matteo M. Fusi	Premillieu	Rekai Gonzalez-	Siddhesh	Tommaso	...
Anouk Van Laer	Christopher	Eric Van	Isaac Richter	Kanishk Sugand	Matthew	Nayan	Alberquilla	Poyarekar	Marinelli	...
Arthur Perais	Tornng	Hensbergen	Isaac Sánchez	Karthik Sangaiah	Poremba	Deshmukh	Rene de Jong	Somayeh	Tony Gutierrez	...
Ashkan Tousi	Chuan Zhu	Erik Hallnor	Barrera	Ke Meng	Matthias Hille	Neha Agarwal	Ricardo Alves	Sardashti	Trivikram Reddy	...
Austin Harris	Chun-Chen Hsu	Erik Tomusk	Ivan Pizarro	Kevin Brodsky	Matthias Jung	Nicholas Lindsay	Richard D. Strong	Sooraj Puthoor	Tuan Ta	...
Avishai Tvila	Ciro Santilli	Faissal Sleiman	Jack Whitham	Kevin Lim	Maurice Becker	Nicolas	Richard Strong	Sophiane Senni	Tushar Krishna	...
Ayaz Akram	Clint Smullen	Fernando Endo	Jairo Balart	Khalique	Maxime	Derumigny	Rico Amslinger	Soumyaroop Roy	Umesh Bhaskar	...

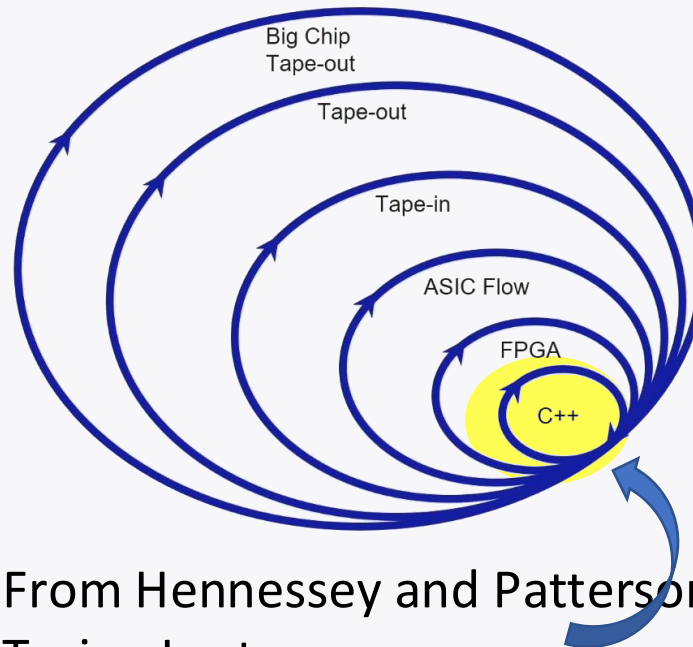
Your name here!



gem5's goals



Agile Hardware Dev. Methodology



From Hennessey and Patterson
Turing Lecture

gem5's goals

- ▶ Anyone (including non-architect) can download and use gem5
- ▶ Used for cross-stack research:
 - Change kernel, change runtime, change hardware, all in concert
 - Run full ML stacks, full AR/VR stacks... other emerging apps

We're close... just a lot of rough edges! You can help!

The gem5 community

- ▶ 100s of contributors & 1000s(?) of users
- ▶ Aim to meet the needs of
 - Academic research (most of you all!)
 - Industry research and development
 - Classroom use
- ▶ Code of conduct (see repo)
- ▶ **I want to see the community grow!**

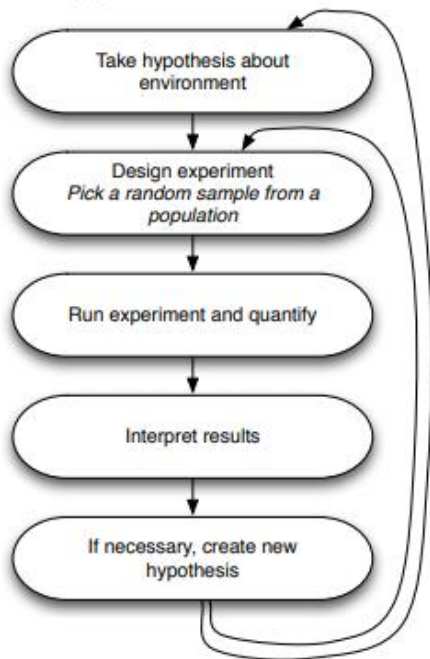


My views on
simulation

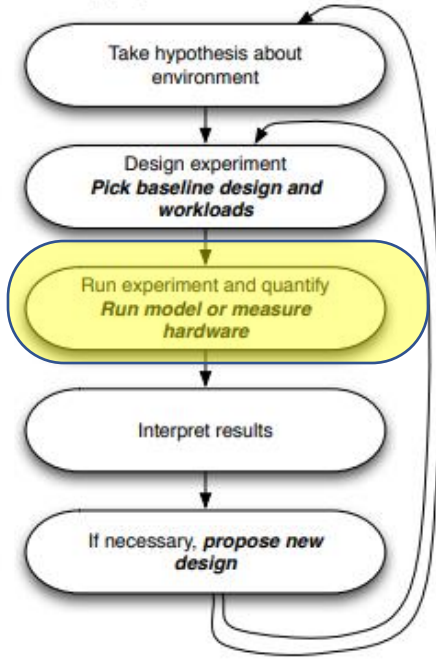


Computer systems research/engineering

(a) Scientific research



(b) Systems research



From Computer Architecture Performance Evaluation Methods by Lieven Eeckhout



Computer architecture simulation!

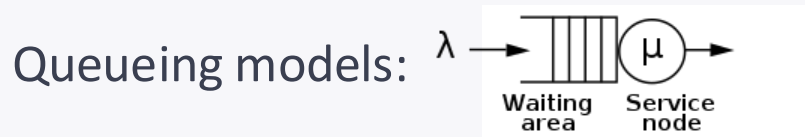
Why simulation

- ▶ Need a tool to evaluate systems that don't exist (yet)
Performance, power, energy, etc.
- ▶ Very costly to actually make the hardware
- ▶ Computer systems are complex with many interdependent parts
Not easy to be accurate without the full system
- ▶ Simulation can be parameterized
Design-space exploration
Sensitivity analysis

Alternatives to cycle-level simulation

Analytic models

Amdahl's Law:
$$S_{\text{latency}}(s) = \frac{1}{(1 - p) + \frac{p}{s}}$$



Kinds of simulation

- ▶ Functional simulation
- ▶ Instrumentation-based
- ▶ Trace-based
- ▶ Execution-driven
- ▶ Full system

Kinds of simulation

- ▶ Functional simulation

Executes programs correctly. Usually no timing information
Used to validate correctness of compilers, etc.
RISC-V Spike, QEMU, gem5 “atomic” mode

- ▶ Instrumentation

Often binary translation. Runs on actual hardware with callbacks
Like trace-based. Not flexible to new ISA. Some things opaque
PIN, NVBit

Kinds of simulation

- ▶ Trace-based simulation
 - Generate addresses/events and re-execute
 - Can be fast (no need to do functional simulation). Reuse traces
 - If execution depends on timing, this will not work!
 - “Specialized” simulators for single aspect (e.g., just cache hit/miss)
- ▶ Execution-driven
 - Functional and timing simulation is combined
 - gem5 and many others
 - gem5 is “execute in execute” or “timing directed”

Full system simulation

- ▶ Components modeled with enough fidelity to run mostly unmodified apps
- ▶ Often “Bare metal” simulation
- ▶ All of the program is functionally emulated by the simulator
- ▶ Often means running the OS in the simulator, not faking it

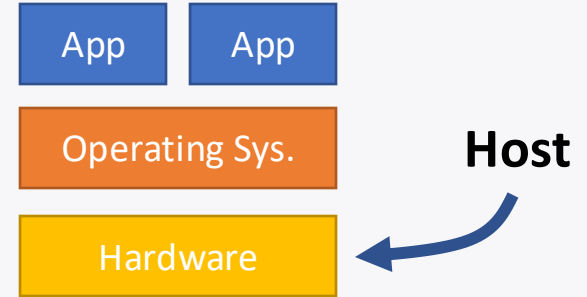
- ▶ “Full system” simulators are often combine functional and execution-based

Nomenclature


Host: the actual hardware you're using
Running things directly on the hardware:
Native execution

Guest: Code running on top of "fake"
hardware
OS in virtual machine is guest OS
Running "on top of" hypervisor
Hypervisor is emulating hardware

Your system



Virtual machines



Nomenclature

Host: the actual hardware you're using

Simulator: Runs on the host

Exposes hardware to the guest

Guest: Code running on *simulated* hardware

OS running on gem5 is guest OS

gem5 is simulating hardware

Simulator's code: Runs natively

executes/emulates the guest code

Guest's code: (or benchmark, workload, etc.)

Runs on gem5, not on the host.



Your system



Host



Simulation

Workload



Guest



Host



Nomenclature

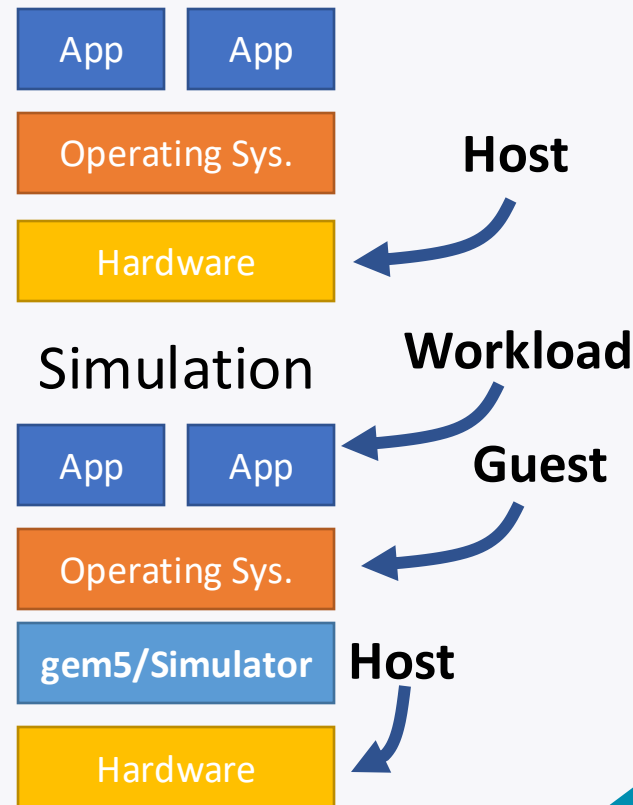
Host: the actual hardware you're using

Simulator: Runs on the host
Exposes hardware to the guest

Simulator's performance:
Time to run the simulation on host
Wallclock time as you perceive it

Simulated performance:
Time predicted by the simulator
Time for guest code to run on simulator

Your system



Tradeoffs

Development time: time to make the simulator/models

Evaluation time: wallclock time to run the simulator

Accuracy: How close is the simulator to *real* hardware

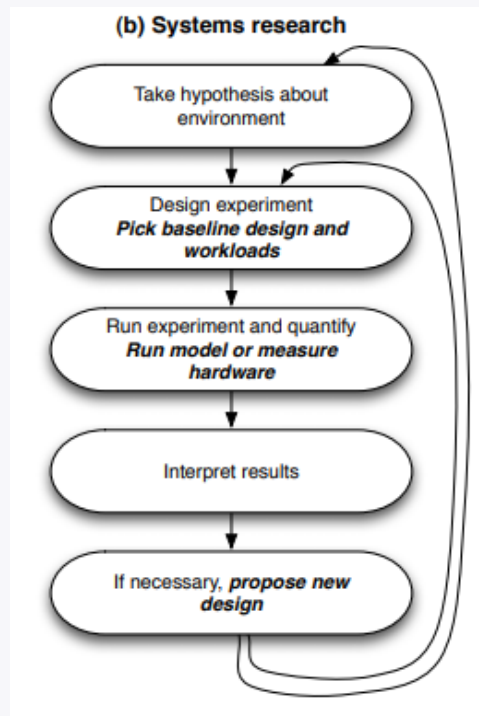
Coverage: How broadly can the simulator be used?

	Development time	Evaluation time	Accuracy	Coverage
functional simulation	excellent	good	poor	poor
instrumentation	excellent	very good	poor	poor
specialized cache and predictor simulation	good	good	good	limited
full trace-driven simulation	poor	poor	very good	excellent
full execution-driven simulation	very poor	very poor	excellent	excellent

<https://www.morganclaypool.com/doi/abs/10.2200/S00273ED1V01Y201006CAC010>

What “level” should we simulate?

- ▶ Ask yourself: What fidelity is required for this question?
Example: New register file design
- ▶ Often, the answer is a mix.
- ▶ gem5 is well suited for this mix
Models with different fidelity
Drop-in replacements for each other
- ▶ “Cycle level” vs “cycle accurate”



RTL simulation

- ▶ RTL: Register transfer level/logic
 - The “model” is the hardware design
 - You specify every wire and every register
 - Close to the actual ASIC
- ▶ This is “cycle accurate” as it should be the same in the model and in an ASIC
- ▶ Very high fidelity, but at the cost of configurability
 - Need the entire design
 - More difficult to combine functional and timing

Cycle-level simulation

- ▶ Models the system cycle-by-cycle
- ▶ Often “event-driven” (we’ll see this soon)
- ▶ Can be quite accurate
 - Not the exact same cycle-by-cycle as the ASIC, but similar timing
- ▶ Easily parameterizeable
 - No need for a full hardware design
- ▶ Faster than cycle-accurate
 - Can “cheat” and functionally emulate some things



gem5's architecture



C++

Models → gem5 has 100s
of models

Cache

↳ parameters: size

Core

↳ parameters
width

↳ #
stages

DRAM

↳ parameters

tRCD

tRAS

tCAS

⋮

gem5 architecture: SimObject

- ▶ **Model**
C++ code in src/
- ▶ **Parameters**
Python code in src/
In SimObject declaration file
- ▶ **Instance or configuration**
A particular choice for the parameters
In standard library, your extensions, or python runscript

Some nomenclature

You can **extend** a model, to model new things
You would want to *inherit* from the object in C++

```
class O3CPU : public BaseCPU  
{
```


You can **specialize** a model with specific parameters
You would want to *inherit* from the object in python

```
class i7CPU(O3CPU):  
    issue_width = 10
```

gem5 architecture: Simulating

gem5 is a **discrete event simulator**

Event Queue



- 1) Event at head dequeued
- 2) Event executed
- 3) More events queued

gem5 architecture: Simulating

gem5 is a **discrete event simulator**

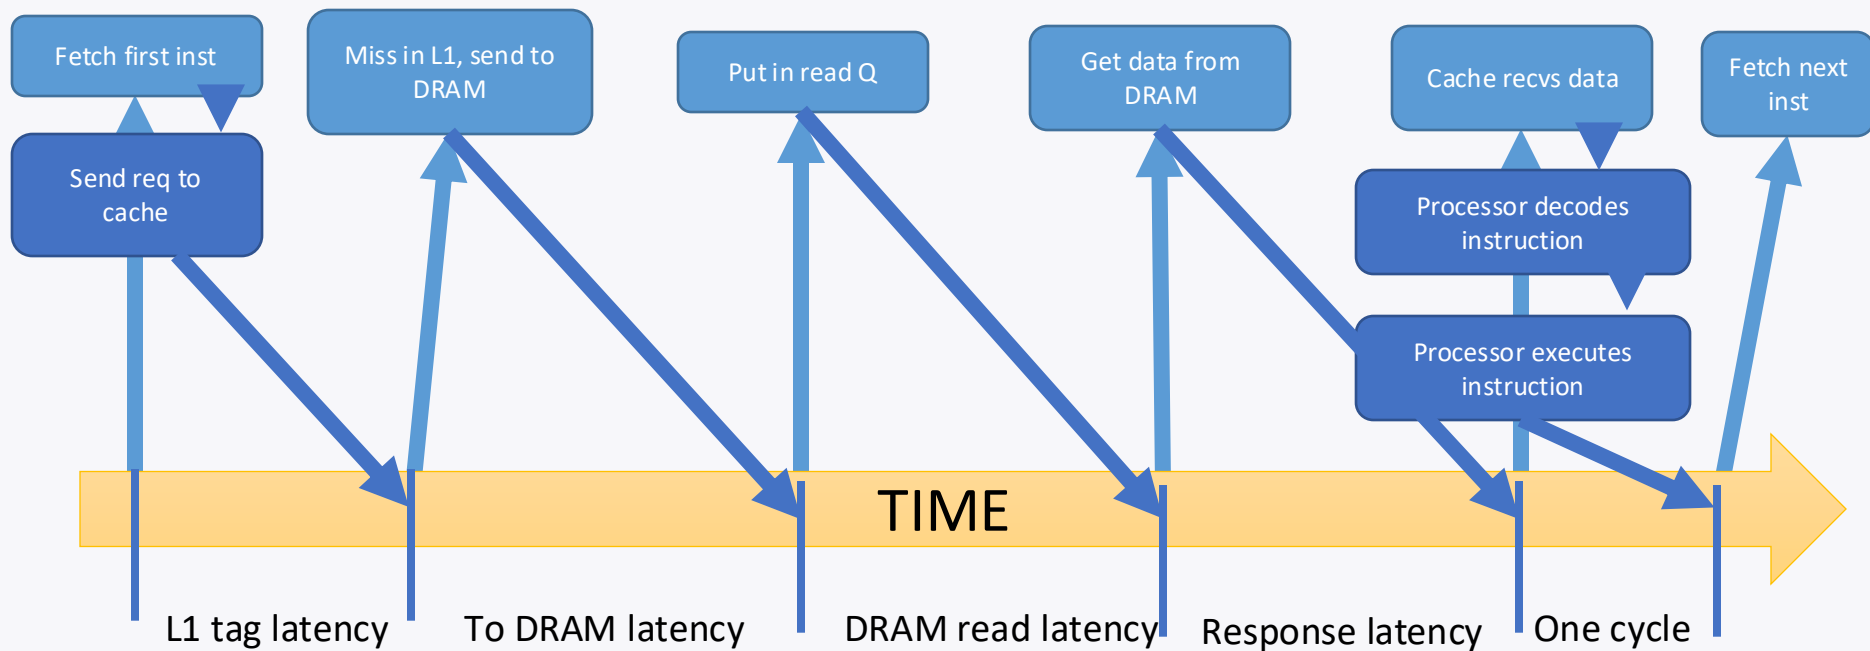
Event Queue



- 1) Event at head dequeued
- 2) Event executed
- 3) More events queued

All SimObjects can enqueue events to the event queue

Discrete event simulation example




Discrete event simulation

- ▶ "Time" needs a unit
In gem5, we use a unit called "Tick"
- ▶ Need to convert a simulation "tick" to user-understandable time
E.g., seconds
- ▶ This is the global simulation tick rate
Usually this is 1 ps per tick or 10^{12} ticks per second

gem5's Main abstractions

- ▶ **ISA vs CPU model & Memory request *abstractions***
- ▶ **Ports** allow you to send requests and receive responses
 - ▶ **Ports** are unidirectional (two types, request/response)
 - ▶ Anything* with a *Request* port can be connected to any *Response* port



gem5's Main abstractions

- ▶ **ISA vs CPU model & Memory requests *abstractions***
- ▶ **CPU model *abstracted* from ISA**
 - ▶ **Any*** CPU model can be used with any* ISA
 - ▶ ISA provides *StaticInst* with *execute()*, *initiateAccess()*, etc.
 - ▶ If you implement the interface, you can add new CPU model or new ISA

