

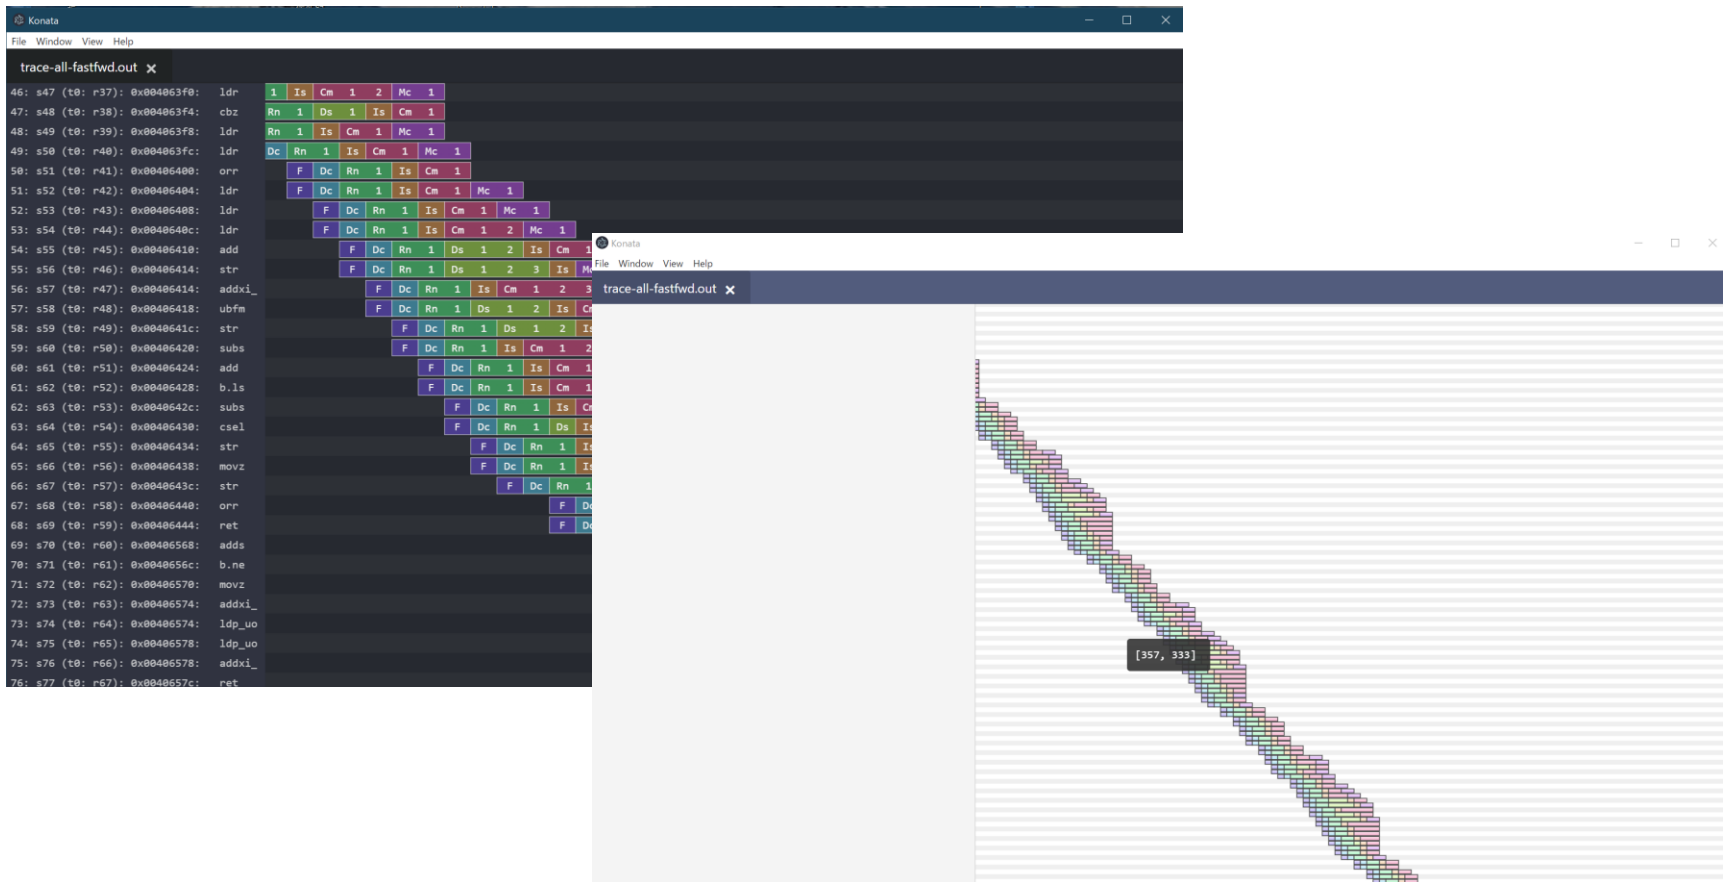
# Visualizing the out-of-order CPU model

---

Ryota Shioya  
Nagoya University

# Introduction

- This presentation introduces:
  - ◇ the visualization of the out-of-order CPU model in gem5



# Introduction

- Let's suppose
  - ◇ you come up with an excellent idea and
  - ◇ try to extend the CPU model in gem5 for adding your new method.
- You will probably tackle the following issues:
  - ◇ difficult bugs, especially performance related ones
  - ◇ a situation where your method cannot improve the performance as expected

# Introduction

- You probably validate your *modified* gem5 as follows:
  - ◇ Check counters outputted by gem5
    - e.g., the number of LLC misses / branch mispredictions
    - These counters sometimes give us clues.
  - ◇ Check the behavior by using a debugger and step execution
- However, it is difficult to fix issues in the following situations:
  - ◇ You have no idea what causes it
  - ◇ You recognize some counters show that something is wrong, but you have no idea what happened

# Visualizing the pipeline behavior

- In such situations, pipeline visualization is very useful.
- In general, visualization is a powerful tool for investigating bugs or behavior.
  - ◇ If you have developed hardware with HDL such as Verilog, you may have used a waveform viewer.
    - In a waveform view, you can easily see signal transitions and relations between signals.
    - Such viewers may have helped you a lot.
- This is also true for gem5!



# Konata: a new GUI based viewer

- You can see the pipeline behavior as a map app.
- ◇ This presentation introduces Konata and best practices in gem5.

The image displays the Konata GUI, a tool for visualizing pipeline behavior. The main window shows a list of instructions and their corresponding pipeline stages, represented as a map. The instructions are listed on the left, and the pipeline stages are shown as colored blocks on the right. The stages are labeled with letters: I (Instruction), D (Decode), Rn (Register File), C (Cache), M (Memory), and W (Write Back). The map shows the flow of instructions through the pipeline stages over time.

A detailed view of a pipeline stage is shown in the bottom right. It displays the instruction being executed and the state of the pipeline registers. The instruction is `orr w8, w3r, w3`. The pipeline registers are shown as follows:

```
0x00406440: orr w8, w3r, w3
Fetched Tick: 5221000
Looking up IntRegClass arch reg 3, got phys reg 188 (IntRegClass)
Looking up IntRegClass arch reg 31, got phys reg 34 (IntRegClass)
Renaming arch reg 0 (IntRegClass) to physical reg 189 (189).
Access to int register 188, has data 0
Setting int register 189 to 0
Line: 31774
Serial ID: 224
Thread ID: 0
Retire ID: 186
```

# Outline

1. A brief explanation of how to use
2. Typical visualization examples
3. Use cases



# Preparation


1. Install: All you have to do is to download the package and unpack it.
  - ◇ <https://github.com/shioyadan/Konata/releases>
  - ◇ Windows/Linux/Mac packages are provided.
  - ◇ No additional runtime is not required.
2. Start the executable file such as Konata.exe

# How to Use


1. Generate a trace log from gem5 with the O3 CPU model
  - ◇ Execute gem5 with the following flags
  - ◇ `./build/ARM/gem5.opt --debug-flags=O3PipeView --debug-start=<first tick of interest> --debug-file=trace.out configs/example/se.py --cpu-type=detailed --caches -c <path to binary> -m <last cycle of interest>`
  - ◇ This example is from <http://www.m5sim.org/Visualization>
2. Load the generated "trace.out" to Konata
  - ◇ from the menu in the window or using drag&drop

# How to use

- After loading the file, contents like the following are shown.
  - ◇ Left side: instruction information such as a PC and mnemonic
  - ◇ Right side: the image of visualized pipeline behavior





# How to see the visualized image



- The clock cycle proceeds from left to right
  - ◇ F : Instruction fetch
  - ◇ Dc : Instruction Decode
  - ◇ Rn : Rename
  - ◇ Ds : Dispatch
  - ◇ IS : Issue
  - ◇ Cm : Completion of execution
    - The execution stage is not explicitly shown
    - The execution seems to be started at the first stage of "Cm" stages
  - ◇ (The end of Cm stages) : Retire

# Zoom in/out

- You can zoom in/out as follows:

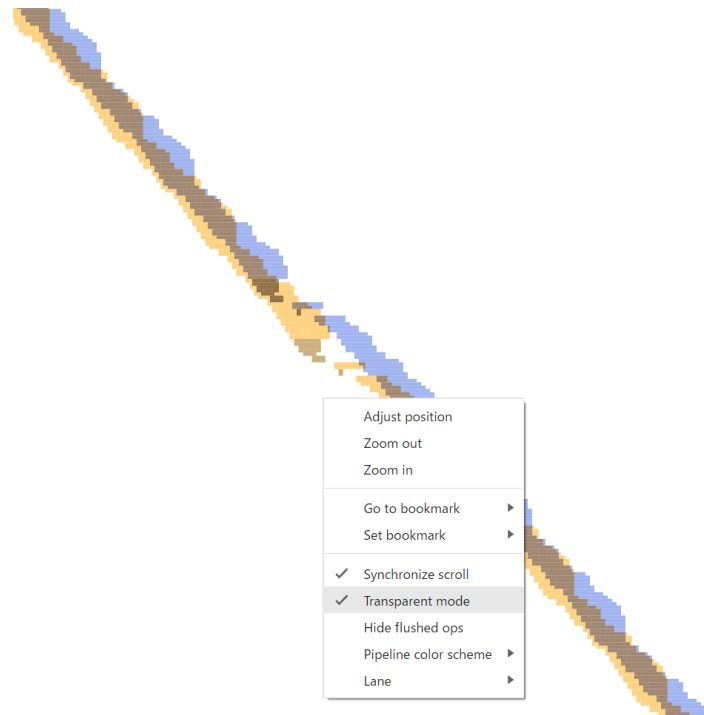


zoom-in

zoom-out

# Compare Two Pipelines

- Konata can show two pipelines overlapping.  
(Currently, it requires manual operations as follows:
  1. Load two files
  2. Right click -> "Transparent mode" & "Synchronize scroll"
  3. Right click -> "Pipeline color scheme" to change a pipeline color



# Outline

1. A brief explanation of how to use
2. Typical visualization examples
3. Use cases

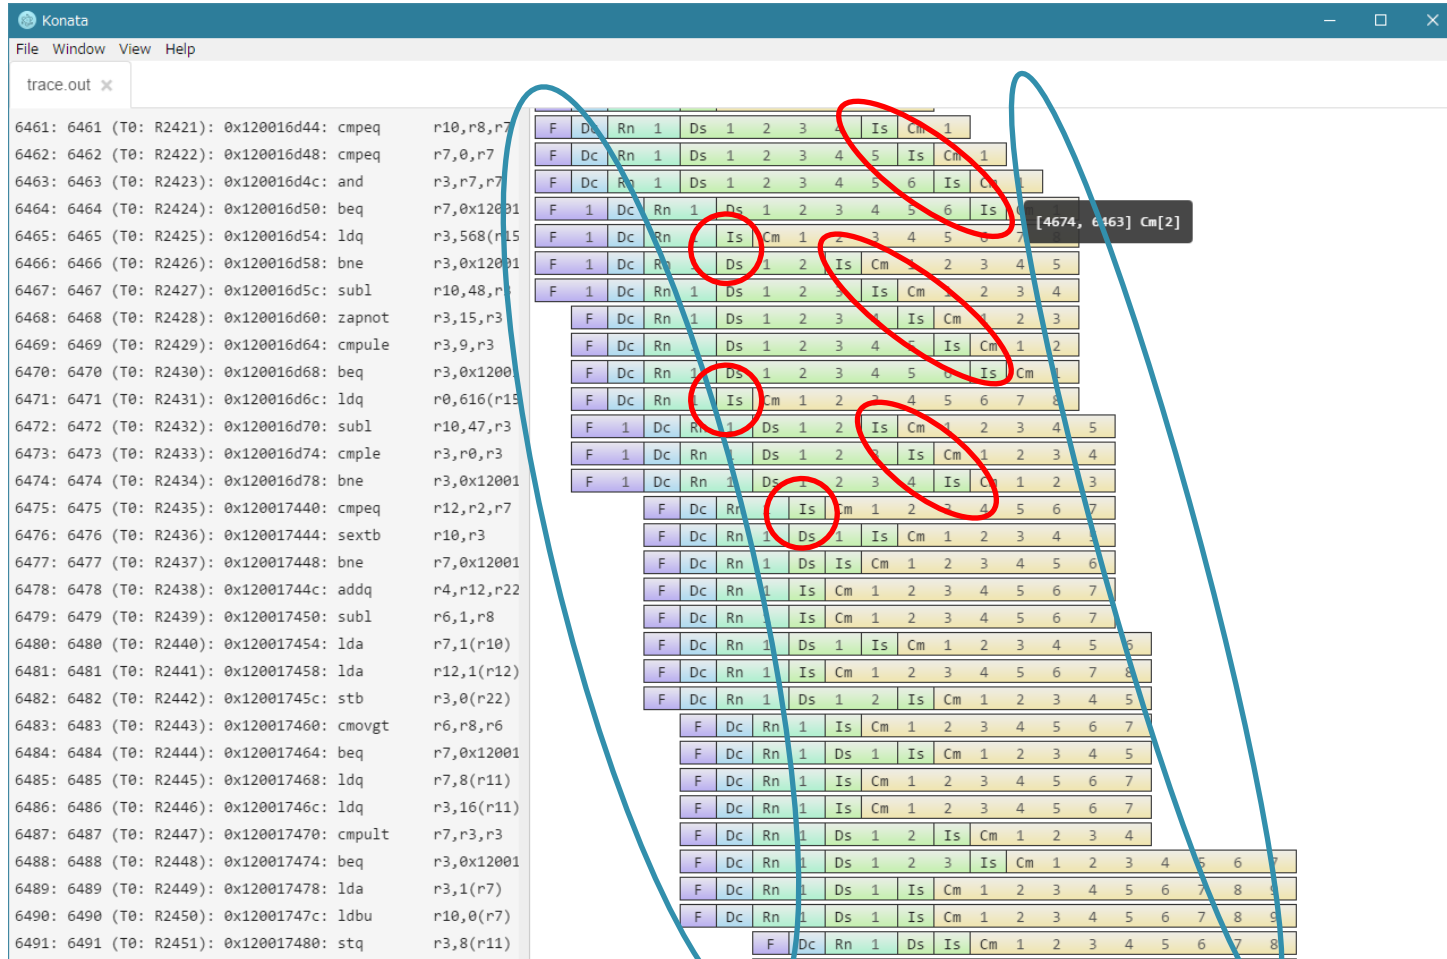
# Typical Visualization Examples

- Introduce how the following things are shown:
  1. Out-of-order execution
  2. Branch misprediction
  3. Cache miss
  4. Execution speed



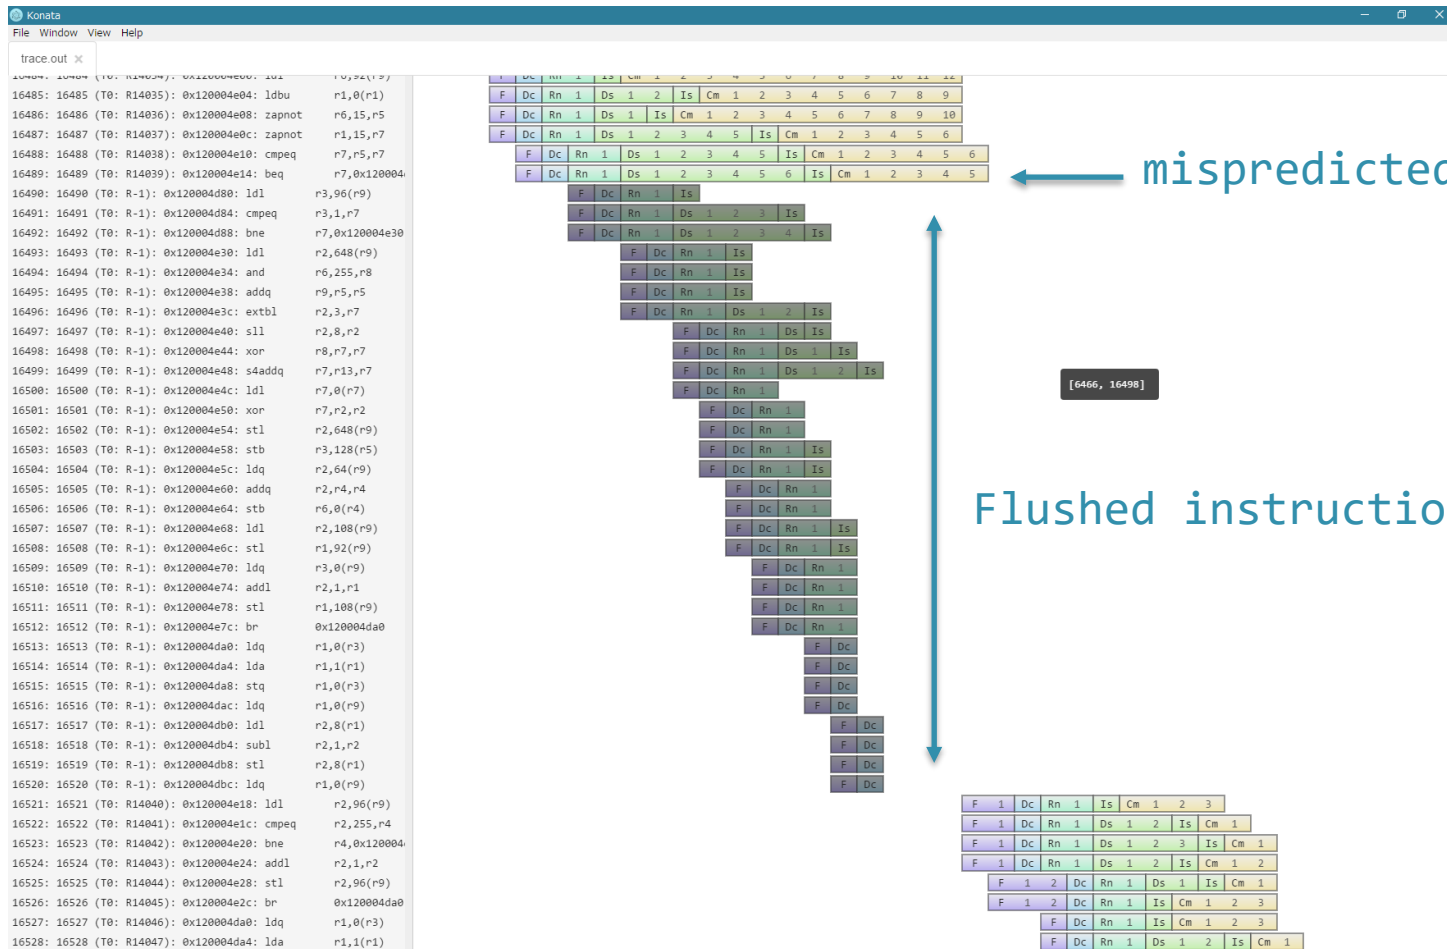
# Example: Out-of-order Execution

- Fetch and retirement, marked with the blue circles, are performed in-order
- ◇ Instruction issue, marked with the red circles, is performed out-of-order



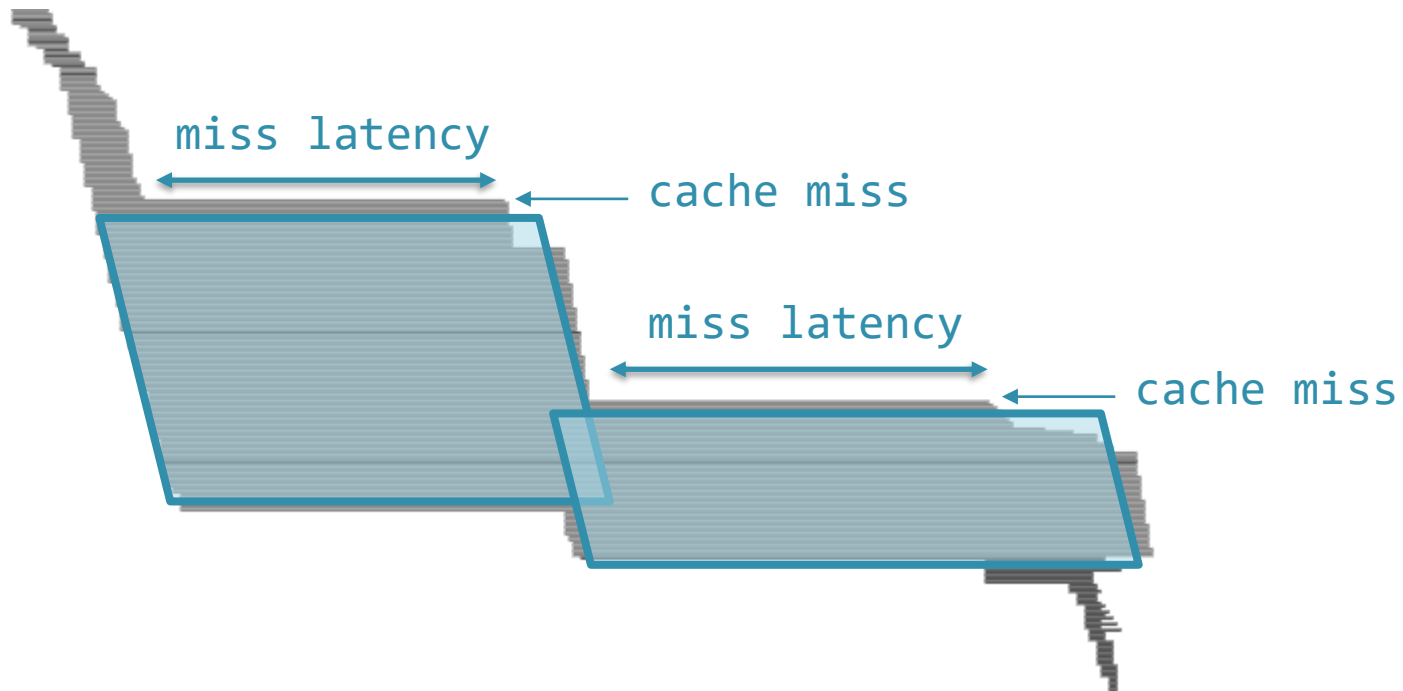
# Example: Branch Misprediction

◆ Flushed instructions are shown as dark ones



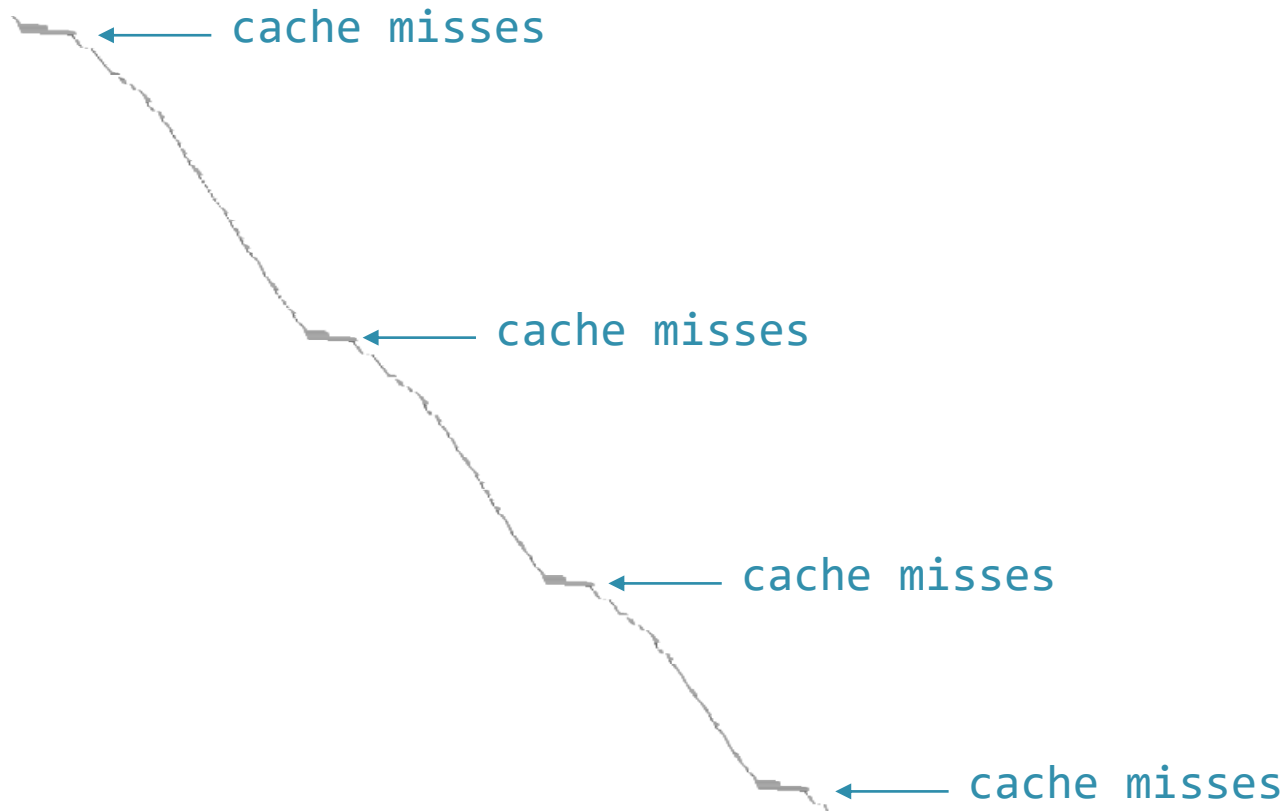
# Example: Cache Misses

- A cache miss is typically shown as a diamond-like shape when the image is zoomed out as follows



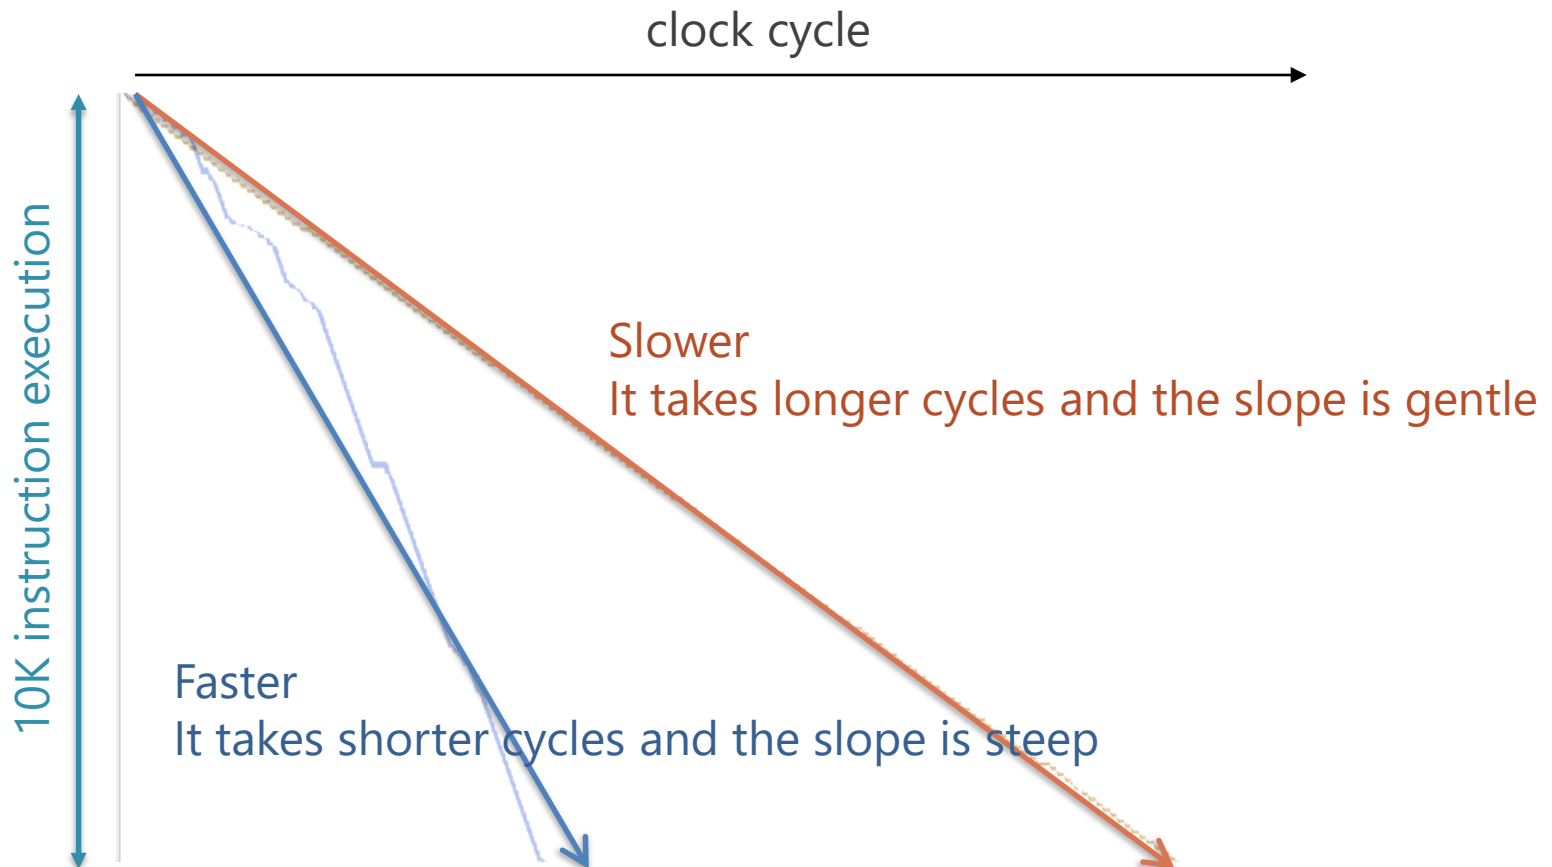
# Example: Cache Misses

- As it is zoomed out more, the pipeline is typically shown as follows
  - ◇ This is the pipeline behavior of MCF in SPEC CPU 2006.
  - ◇ This figure shows the performance is degraded by the cache misses.



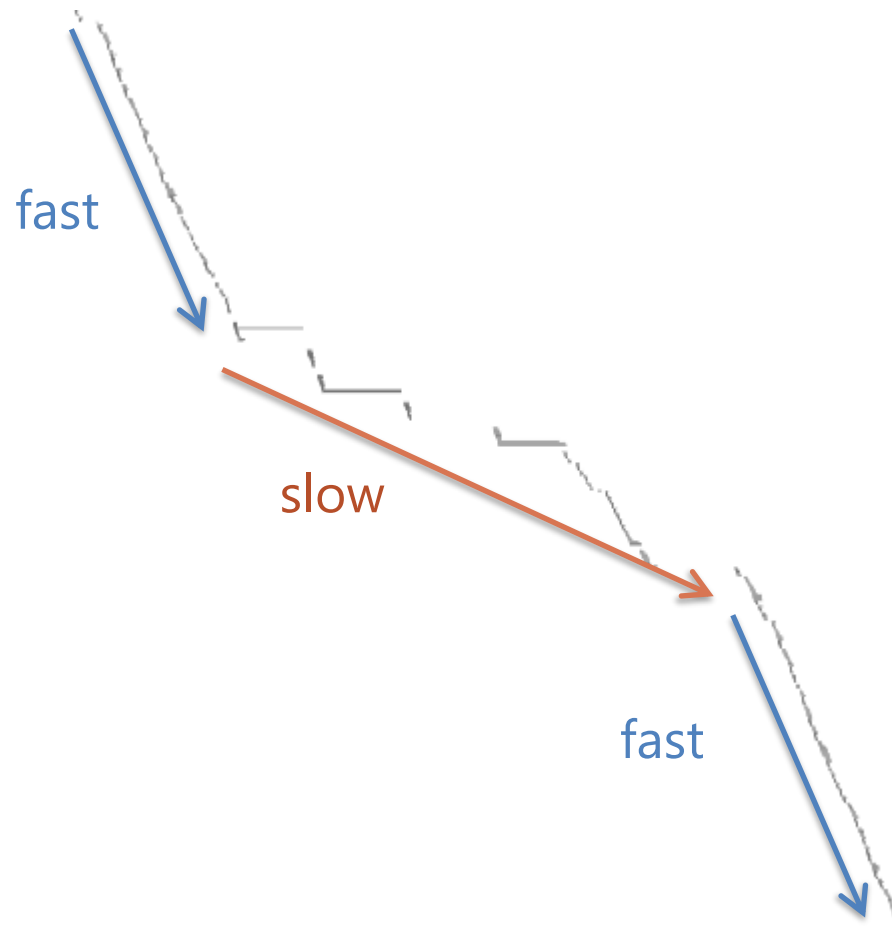
# The slope of a pipeline shape roughly represents the execution speed (IPC)

- ◇ The following two pipelines show the execution of the same 10K instructions



# The slope of a pipeline shape roughly represents the execution speed (IPC)

- You can see the transition in the execution speed for each part of the program as follows



# Outline

1. A brief explanation of how to use
2. Typical visualization examples
3. Use cases
  1. Grasping the pipeline behavior
  2. Comparing pipelines

# Grasping the pipeline behavior

- The pipeline visualization makes it easy to grasp the pipeline behavior.
  - ◇ Explain this by using some examples
- Let's suppose you newly add speculative execution with branch prediction.
  - ◇ (Of course, gem5 already has this feature
  - ◇ Something wrong happens in recovery from mispredictions.



# Investigating with a log


```
15 O3PipeView:fetch:25069250:0x120016358:0:177:addq      r12,r12,r1
16 O3PipeView:decode:25069500
17 O3PipeView:rename:25069750
18 O3PipeView:dispatch:25070250
19 O3PipeView:issue:25071500
20 O3PipeView:complete:25071750
21 O3PipeView:retire:0:store:0
22 O3PipeView:fetch:25069250:0x12001635c:0:178:ldq      r3,104(r22)
23 O3PipeView:decode:25069500
24 O3PipeView:rename:25069750
25 O3PipeView:dispatch:25070250
26 O3PipeView:issue:25071250
27 O3PipeView:complete:25071500
28 O3PipeView:retire:0:store:0
29 O3PipeView:fetch:25069500:0x120016360:0:179:addq      r3,r1,r3
30 O3PipeView:decode:25069750
31 O3PipeView:rename:25070000
32 O3PipeView:dispatch:25070500
33 O3PipeView:issue:25072000
34 O3PipeView:complete:25072250
35 O3PipeView:retire:25072500:store:0
36 O3PipeView:fetch:25069500:0x120016364:0:180:ldwu     r1,0(r3)
37 O3PipeView:decode:25069750
38 O3PipeView:rename:25070000
39 O3PipeView:dispatch:25070500
40 O3PipeView:issue:25072250
41 O3PipeView:complete:0
42 O3PipeView:retire:0:store:0
43 O3PipeView:fetch:25069500:0x120016368:0:181:srl      r1,13,r1
44 O3PipeView:decode:25069750
45 O3PipeView:rename:25070000
46 O3PipeView:dispatch:25070500
47 O3PipeView:issue:0
48 O3PipeView:complete:0
49 O3PipeView:retire:0:store:0
```

■ For investigating your implementation, you probably:

- ◇ Check custom logs or your "printf" outputs such as the left example
- ◇ It records when/what instructions are flushed.

■ It's very difficult to detect which point is incorrect from such text logs.

# Investigating with visualization




- By visualizing it, you can easily notice the incorrect point.
  - ◇ There is the light instruction (not flushed) between the dark flushed instructions.
- Although this is an artificial example,
  - ◇ visualization gives us a lot of hints intuitively!

# Another example: memory level parallelism

- One of my friends tackled a topic related to memory level parallelism.
  - ◇ In short, his method improves the performance by performing multiple memory accesses in parallel.
- He enlarged the size of the OoO scheduling window so that more memory accesses are performed in parallel.
  - ◇ But, the performance is not improved...

# Another example: memory level parallelism

- He realized a curious behavior from the following zoomed-out image,
  - ◇ because the shape is unnatural
- He realized that the pipeline was flushed on a cache miss
  - ◇ In this sequence, memory accesses should be performed in parallel



# The cause of the flush

- He examined the flushed instruction in detail and found the cause.
- This was because he used Alpha ISA.
  - ◇ In Alpha architecture, TLB miss causes a trap and the pipeline is flushed.
  - ◇ On a cache miss, a TLB miss often occurs.
  - ◇ So, memory accesses cannot be performed in parallel.
- It is not easily noticed simply by observing the counters in gem5.
  - ◇ The shape or pattern of visualized pipelines often tell us hints.

# Outline


1. A brief explanation of how to use
2. Typical visualization examples
3. Use cases
  1. Grasping the pipeline behavior
  2. Comparing pipelines

# Comparing pipelines

- Let's suppose
  - ◇ your new method seems to work correctly,
  - ◇ but it does not improve the performance as you expected.
- Konata can compare two pipelines!
  - ◇ It is useful when investigating the above situation.

# Example of comparing


- Konata can show two pipelines overlapping.
  - ◇ Blue shows a baseline processor pipeline.
  - ◇ Orange shows a pipeline with a proposed method.





# Example of comparing

- The orange one (proposed) is basically faster than the blue one (baseline)
- ◇ The insn. fetch and retirement of the orange one are clearly fast.



# Example of comparing

- In the zoomed-in image,
  - ◇ in some places, the fetching of the orange is unreasonably delayed.
  - ◇ This was caused by a bug, and this bug degraded the performance.



# Comparing pipelines

- Visual comparison is very effective for analysis when adding new features to gem5.
  - ◇ If the performance is not improved as expected, something is delayed.
  - ◇ You can detect such parts by visual comparison.
    - It is easy to see which part is different.

# Conclusion


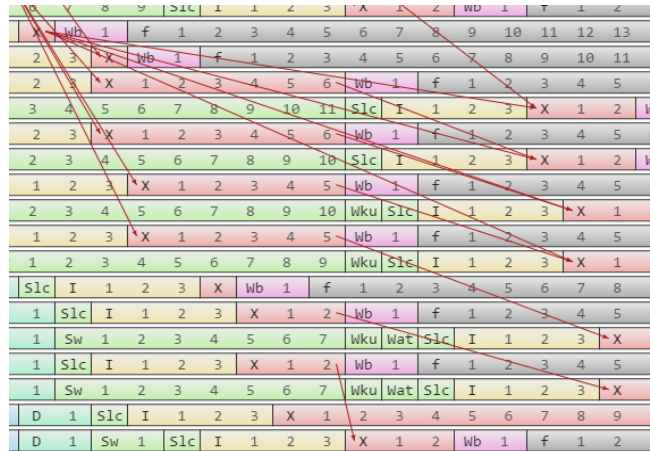
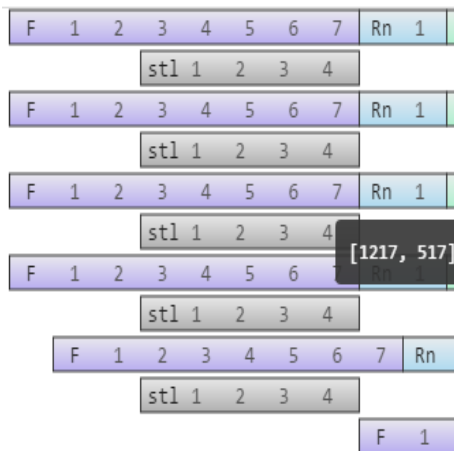
- It is generally difficult to investigate the cause of a bug related to the performance.
  - ◇ Especially, when you have no idea what happened.
- In such cases, visualization is very useful.
  - ◇ This presentation introduced a new pipeline viewer Konata and best practice in gem5
- Please try it!
  - ◇ <https://github.com/shioyadan/Konata/releases>

# Appendix

# Future work

(In 2022, the features 2. and 3. have been already implemented.  
You can use them with "--debug-flags=O3PipeView,O3CPUAll."

- ◇ Add new features:
  1. Showing concurrent events such as pipeline stall etc.
  2. Dependency arrows between instructions
  3. Custom messages for each stage
- ◇ Konata has already had these features.
  - Because Konata was developed for other simulators.
  - I have a plan to add support for these features to gem5.



1. Multiple stage lanes

2. Dependency arrow

3. Custom log message

# The slope of a pipeline shape roughly represents the execution speed (IPC)

- ◇ It is not accurate because flushed instructions are also shown.
- ◇ If you want to compare accurately, use "Hide flushed ops" option from the right click menu

