

18-344 Recitation 7

10/31/2025

Outline

1. Logistics
2. Review
3. Lab 3


Logistics

- Homework 7 released, due November 5 (Wednesday) (in 5 days)
 - Covers lectures 15-16
- Lab 3
 - Code freeze: October 14 (Friday) (in two weeks)
 - Submit (push a commit to your repo) on GitHub Classroom
 - Report: November 17 (Monday) (in two weeks + three days)
 - Submit on Gradescope

Review

Pipeline *data* hazards

```
sub x6, x5, x4  
lw x16, 0xabc  
add x12, x6, x14
```



```
sub x8, x16, x4  
add x16, x6, x14  
lw x16, 0xabc
```



```
lw x6, 0xabc  
sub x6, x5, x4  
add x12, x6, x14
```



Read-After-Write (RAW)

Write-After-Read (WAR)

Write-After-Write (WAW)

There's always data hazards in programs, and we can't change that.

How to further exploit ILP?

Exploiting ILP using multiple issue

Common name	Issue structure	Hazard detection	Scheduling	Distinguishing characteristic	Examples
Superscalar (static)	Dynamic	Hardware	Static	In-order execution	Mostly in the embedded space: MIPS and ARM, including the Cortex-A53
Superscalar (dynamic)	Dynamic	Hardware	Dynamic	Some out-of-order execution, but no speculation	None at the present
Superscalar (speculative)	Dynamic	Hardware	Dynamic with speculation	Out-of-order execution with speculation	Intel Core i3, i5, i7; AMD Phenom; IBM Power 7
VLIW/LIW	Static	Primarily software	Static	All hazards determined and indicated by compiler	Most examples are in signal processing, such as the TI
EPIC	Primarily static	What we mean when we say “superscalar out-of-order” <i>in this class</i> .			

by the compiler

Exploiting ILP using multiple issue

Common	Issue	Hazard	Distinauishinga
--------	-------	--------	-----------------

Hardware-based speculation:


1. Dynamic branch prediction to choose which instructions to execute
2. Speculation to allow the execution of instructions before the control dependences are resolved
3. Dynamic scheduling to deal with the scheduling of different combinations of basic blocks

speculation


Superscalar (speculative)	Dynamic	Hardware	Dynamic with speculation	Out-of-order execution with speculation	Intel Core i3, i5, i7; AMD Phenom; IBM Power 7
VLIW/LIW	Static	Primarily software	Static	All hazards determined and indicated by compiler (often implicitly)	Most examples are in signal processing, such as the TI C6x
EPIC	Primarily static	Primarily software	Mostly static	All hazards determined and indicated explicitly by the compiler	Itanium

Register renaming

```
lw    x6, 0xabc  
      RAW  
sub   x5, x6, x4  
      WAR  
add   x6, x11, x14
```

The diagram shows three instructions in a vertical sequence. The first instruction is 'lw x6, 0xabc' where 'x6' is highlighted in yellow. The second instruction is 'sub x5, x6, x4' where 'x6' is highlighted in blue. A red arrow labeled 'RAW' points from the yellow 'x6' in the first instruction to the blue 'x6' in the second instruction. The third instruction is 'add x6, x11, x14' where 'x6' is highlighted in yellow. A green arrow labeled 'WAR' points from the yellow 'x6' in the third instruction to the blue 'x6' in the second instruction.

Instruction stream




Renaming table



Instruction stream
after renaming

Register renaming

```
lw  x6, 0xabc
    RAW
sub x5, x6, x4
    WAR
add x6, x11, x14
```

The diagram shows three instructions in a box. The first instruction is 'lw x6, 0xabc' where 'x6' is yellow. The second is 'sub x5, x6, x4' where 'x6' is blue. A red arrow labeled 'RAW' points from the yellow 'x6' in the first instruction to the blue 'x6' in the second. The third instruction is 'add x6, x11, x14' where 'x6' is yellow. A green arrow labeled 'WAR' points from the yellow 'x6' in the third instruction back to the blue 'x6' in the second instruction.

Instruction stream

```
lw:  x6 → r1
```


Renaming table

```
lw  r1, 0xabc
```

Instruction stream
after renaming

Register renaming

```
lw    x6, 0xabc  
      RAW  
sub   x5, x6, x4  
      WAR  
add   x6, x11, x14
```



Instruction stream

```
lw:   x6 → r1  
sub:  x5 → r2  
      x6 ← r1
```


Renaming table

```
lw    r1, 0xabc  
sub   r2, r1, x4
```

Instruction stream
after renaming

Register renaming

```
lw  x6, 0xabc
sub x5, x6, x4
add x6, x11, x14
```



A red arrow labeled 'RAW' points from the 'x6' in the first instruction to the 'x6' in the second instruction. A green arrow labeled 'WAR' points from the 'x6' in the third instruction back to the 'x6' in the second instruction.

Instruction stream

```
lw:  x6 → r1
sub: x5 → r2
      x6 ← r1
add: x6 → r3
```

Renaming table

```
lw  r1, 0xabc
sub r2, r1, x4
add r3, x11, x14
```

Instruction stream
after renaming

Register renaming

```
lw  x6, 0xabc
sub x5, x6, x4
add x6, x11, x14
```

RAW (red arrow from x6 to x6)

WAR (green arrow from x6 to x6)

Instruction stream

```
lw:  x6 → r1
sub: x5 → r2
      x6 ← r1
add: x6 → r3
```

Renaming table

```
lw  r1, 0xabc
sub r2, r1, x4
add r3, x11, x14
```

RAW (red arrow from r1 to r1)

WAR (green dashed arrow from r1 to r3)

Instruction stream
after renaming

Gets rid of the false dependence.

Lab 3


Overview

- A VM simulator, instead of a pintool
 - Virtual-to-physical mapping using hierarchical page table
 - Translation Lookaside Buffer
 - Page fault handler

What's provided, and what's not

- Provided by us, can make small changes if you want
 - `pte-util.h`, `pte.{h,cpp}`
 - Simple declarations of PTE types for you to use
 - `vm.h`, `vm-util.h`
 - VM class and various helpful declarations for you to use
 - `Vm-provided.cpp`
 - Contains implementation of VM methods that you do not need to implement (e.g. allocating and replacing physical pages)
 - `vm_trace/*`
 - A pintool we provide that you can use to collect a trace of memory accesses for testing
- Provided, but you'll need to implement/modify
 - `ptab.{h,cpp}`
 - Skeleton code and declarations of page table operations that you'll need to implement
 - `tlb.{h,cpp}`
 - Skeleton code and declarations of a TLB that you'll need to implement
 - `vm-student.cpp`
 - This is where you'll implement VM methods (e.g. `map`, `translate`) using the page table and TLB
 - `vm-test.cpp`
 - A driver for the VM simulator that runs on a given trace. Skeleton code with a simple, hard-coded trace is provided. You'll need to modify it so it works with traces generated by `vm_trace`.


General structure



Task 1: Implement a hierarchical page table

VPN: Virtual Page Number
PT: Page Table
PTB: Page Table Base
PTE: Page Table Entry
PTO: Page Table Offset

48-bit virtual addresses and 4KB pages => 512 entries per page table.



Task 1: Implement a hierarchical page table

`vm-student.cpp`: implement three functions (using your PT implementation):

- `vmMap(vaddr, size)`
 - Update Page Table to map `size` bytes at address `vaddr`; Could span multiple pages; create Page Tables or PTEs here if not exist
- `vmTranslate(vaddr)`
 - TODOs in the handout; Return the translated physical address
- `vmPageFaultHandler(*pte)`
 - Handle Page Faults (page not residing in memory) here; if there exist free physical pages in memory, allocate them with `bumpAllocate()` function; if you run out of free pages, replace an existing page with the `replacePage()` function. These functions return the PPN to store in your PTE.
 - `bumpAllocate()` and `replacePage()` are provided to you; declarations in `vm.h`, implementations in `vm-provided.cpp`.

Task 2: Implement a TLB


`tlb.{h,cpp}`: Implement a TLB structure, with your choice of organization (size, ways, replacement). You will also need two functions: `lookup()` and `update()`.

- `lookup(vaddr, &PPN)`
 - Attempt to assign cached-PPN to the argument `&PPN`, and return true for hits.
- `update(vaddr, new_PPN)`
 - Update the TLB entries with this new mapping.
 - This may evict an existing entry, your TLB organization should account for replacement.

Testing

Sample mmap log:

```
PC= ... : mmap(vaddr = ... len = ... <zip> ) ret = ...
```



Modify `vm_test.cpp` to parse* the traces generated by `vm_trace`.
`vm_trace` logs mmaps and optionally memory accesses.

Reporting the results

- Page Table implementation should report `#page_faults`, `#num_accesses` and `#tlb_hits`.
- You will justify your TLB organization with a ***quantitative*** analysis of these three parameters.
 - Show a plot of TLB misses vs TLB configurations.
 - Reason about the reduction in Page Table walks due to your TLB.
- Your code should also be robust to ***handle exceptions***: page faults and segmentation faults (unmapped accesses).
- Turn in all code, writeup, and test traces!
 - Include any scripts/tests you used as well.