

This lecture is about how virtual memory is managed in a process and a computer system.



So far, we have only considered three-stage pipeline architecture.

Consider the maximum clock frequency such a pipelined CPU can achieve.

Each stage involves two time constraints: 1) propagation delay of the logic  $t_p$ , 2) the setup time of the clocked circuit (i.e. register or D-FF).

A D-FF only works correctly if the input data is STABLE some time BEFORE the active clock edge. The minimum time that data MUST BE stable before the clock edge is known as setup time  $t_s$  (shown in BLUE).

Data must also stay STABLE some time AFTER the active clock edge. This is known as hold time  $t_h$  (shown in GREEN).

If data changes within the setup and hold time window, the output of the D-FF is unpredictable. This will cause a crash in the computer if it happens.

Therefore, the minimum period between two active clock edges is the worstcase delay in the logic =  $\max(t_{p_1}, t_{p_2}, t_{p_3}, t_{p_4}, t_{p_5}) + t_s$ . Hold time places not part in this consideration.

In other words, to increase the clock speed of a processor, we can attempt to reduce the worst-case propagation delays between pipeline registers. One way to achieve this is add more pipeline stages.



This idea leads the technique known as deep pipelining. Here logic stages are splitted into multiple stages with pipeline registers inserted in between. The best strategy is to design the CPU such that the delay between registers are nearly equal in all stages, so that the clock frequency is not dominated by the worst stage in the pipeline.

Increasing the number of stages in a pipeline theoretically will not affect the performance in terms of Cycle per Instruction (CPI). It remains slightly higher than 1 due to stalling and flushing as a result of data and control hazards. In an earlier example, we stated that CPI could be around 1.25 – taking an average of 1.25 clock cycles to execute one instruction.

Overall performance increase comes from the increase clock frequency due to reduced cycle time  $T_c$ . Instruction time =  $CPI \times T_c$ .

Why can we not keep increasing the pipeline stages? Two reasons: 1) deeper pipeline results in more complicated logic to detect and mitigate hazards; 2) more hazards will be generated due to data dependency and branches, increasing the CPI. There is a trade off between reducing  $T_c$  resulting in increased CPI due to hazards.

Further, as propagation in the logic decrease due to increased pipeline stages, the setup time of the flip-flops dominates. This results in diminishing return.

As an example, a modern ARM processor as used in most mobile devices uses a 13-stage pipeline architecture.



Let us consider a concrete example.

Assume that your single-cycle RISC-V processor has a propagation delay of 750ps overall. You may be considering what performance gain you might obtain by introducing N stages of pipeline.

The assume here are as stated in the slide above.

The question – how many pipeline stages should one use for best performance?

Given that the register setup time (overhead) is 90ns. If we split the 750ps into N stages with each delay,  $T_c = \frac{750}{N} + 90$  ns.

We are told that minimum stages should be 5, and for this, CPI is 1.25. Every additional stage will add, on average, another 0.1 CPI. Therefore, CPI of N-stage processor = 1.25 + 0.1 (N-5) for N  $\ge 5$ .

The graphs here plots the number of stages vs  $T_c$ . It can seen from the plot that the minimum instruction time is when N=8 and  $T_c \times CPI = 281$ ns. This is only a marginal improvement from the 5-stage architecture, which has an instruction time of 295 ps. Make N larger increases the instruction time and reduces the performance.



So far, we always fetch the **next instruction** even if the current instruction is a branch or jump instruction. This is the cause of control hazard in a pipelined architecture. This is the same as predicting that a branch is NEVER TAKEN.

We can improve the performance of the processor by reducing the number of branch hazard encountered by making better prediction on whether a branch instruction is taken. For example, lots of program loops are written with the conditional branch instruction at the end of the loop branching BACKWARDS. These loop structure suggests that a branch destination addrss that goes backwards is normally taken and should be predicted as such. In contrast, forward branches or jumps are usually NOT taken. This strategy results in a **static branch prediction** depending on the location of the branch destination address. It requires little hardware and does better than NOT having any prediction at all.

However, programs can be more complex than this. Some backward branches are normally NOT taken while forward branches are, depending on how codes are generated and the type of loops that are being executed.

A better strategy is to assume that if the branch was taken last time, it is predicted to be taken next time as well. This requires the processor to establish a table of branch instructions with a 1-bit indicator that starts the most recent result (taken or not taken). This effective is a two-state state machine.

This table is known as a branch target buffer, and it also stores the destination address for fast execution.



Consider the code segment in the slide.

This is a loop that goes around 10 times. With one-bit predictor, the bge correctly predict not taken until the last time around the loop. However, if the code is executed again, the first time would also be mispredicted.

The j instruction is mispredicted both the first and last time.

In general, one-bit predictor mispredict branches the first and last time around a loop.

This can be improved by using a slightly more sophisticated predictor – a twobit predictor giving four states: strongly taken, weakly taken, weakly nottaken and strongly not-taken.

If a branch is taken two or more times, the predictor is in the strongly taken state. If a branch is not taken twice or more, the predictor is in the strongly not-taken state.

For the code segment in the slide, the bge instruction is in the strongly nottaken state. It predicts incorrectly first time the loop is exited (branch to done:), but the predictor stays in the weakly not-taken state. Therefore, next time when the code is executed again, it returns to the strongly not-taken state. In general, the two-bit predictor only gets it wrong once around the loop, not twice.



A simple pipeline processor has a CPI that is 1 or above (due to hazard mitigation). Superscalar processor, on the other hand, can reduce this to lower than 1 by duplicating hardware within a processor.

Show here is a two-way superscalar processor with two ALU working in parallel. To keep both ALU busy, two instructions needs to be issue at each clock cycle. This also means that the Register File needs to double its ports and do does the data memory and other internal hardware.

For example, modern processor such as Intel i9 can handle 6 instructions per cycle and ARM Cortex A78 processor executes up to 3 instructions per cycle.



Here is an example of our 5-stage pipeline processor with two-way superscalar executing six instructions in 3 cycles, giving a CPI of 0.5.

Note that the instruction sequence here does not generate any data or control hazards. This is in general NOT the case.



This code snippet has plenty of data hazard.

The 1w instruction cannot provide values for s8 until cycle 5, but the add instruction needs it in cycle 4. Therefore there is no solution except to delay issue of the add instruction to cycle 2, and to insert a stall cycle in cycle 4. There are other dependencies such as the and instruction also requires s8, and the sw instruction requires s11. However, the stall cycle will allow these hazard to be resolved through forwarding after the stall cycle is added.

However, inserting the stall cycle results in 6 instructions executed in 5 cycle, a IPC of 1.2, much lower than the ideal case of IPC = 2.0.



To mitigate this, all modern superscalar processor introduce Out-of-Order execution. Here the processor fetch a number of instructions from program memory into a buffer. This instructions are analyzed for data dependency and issue not in sequence, but in a way that hazard due to data dependency is avoided.

As an example, the same code snippet is executed in the order shown in the slide above.

In cycle 1, add sub and all depend on S8, which is not available. Therefore, or is issued with 1w.

In cycle 2, s8 is still not available, so sw is the only instruction that can be executed and is therefore issued on its own because it relies on s11, which is ready by the time the or has done its job.



Finally, the remaining three instructions are issued in the following two cycles as shown.

This means that the 6 instructions is now executed in four cycles, give us a IPC of 1.5 – better than 1.2 before.

Note that in the slide, the data hazards are labelled as RAW and WAR.

RAW stands for Read-after-Write hazard. The hazard is cause by or instruction writing to s11, and it is read by the following instruction sw. Similar the dependency of add on 1w because of s8 is another RAW hazard.

WAR stands for Write-after-Read hazard. The example is with add on sub because s8 is used by both instructions. sub must not write to s8 before add finishes reading s8.

These hazards can be overcome by a technique known as **Register Renaming**. This is a topic that will be left to the Advanced Computer Architecture elective in the 3<sup>rd</sup> year.



This short one-term module only allows 8 teaching weeks. Not everything that should be in the module are covered due to shortage of time. Here are 5 topics that I wish I have time to include, but not able.



There are also a number of interesting aspects of the RISC-V processor that you may like to read yourself in your spare time.

jal rd,	la	bel	ju	ump a	nd l	ink				PC	= J	TA,				rd = PC +	- 4
Instruction Formats	31	30 29 28	27	26 25	24 2	3 22	21 2	20 19	18	17 16	15	14 13	12	11 10	987	6 5 4 3 2	1
Jump	[20]		ir	mm[10:1	]		[	[11]		imm[	19:12]				rd	opcode	
<ul> <li>JTA =</li> <li>PC is  </li> <li>rd = re</li> </ul>	Jun loac etui	np Tar led wi rn ado	get th t lres	Adc the . ss = l	lres ITA P <b>C</b> ·	ss = + <b>4</b> ,	PC	val . ad	ue dre	+ się ss c	gne of n	ed in ext	nm in:	edia struc	ate o	ffset	,
<ul> <li>JTA =</li> <li>PC is it</li> <li>rd = re</li> <li>Note</li> <li>0. In it</li> </ul>	Jun load etui thai oth	np Tar, led wi rn adc t the f er wo	get tht lres orn rd,	Add the . ss = I nat o offso	Ires ITA PC - of t et is	ss = + <b>4</b> , he i s al	PC i.e imn way	val . ad ned /s a	ue dre iate n e	+ sig ss c e val ven	gne of n lue nu	ed in ext is u mbe	nm ins int er	iedia struc isua	ate o ction I. Bit	ffset t 0 is alwa	ays

Answer some questions from a number of students about JAL and JALR.

jalr rd,	rs1, imm	jump ar	nd link r	egister		PC	C = r	^s1 +	Sigr	Ext(	(imm)	, rd = PC +
Instruction Formats	31 30 29 28	27 26 25 2	24 23 22	21 20	19 18	7 16	15 1	4 13 1	2 11	10 9	8 7	6 5 4 3 2 1
Immediate		imm[11:0]			1	s1		funct3		rd	1	opcode
• JTA = • Note = 32-bit	rs1 + SignI that the im s before a	is also Ext(imn hmedia dding t	n), i.e. te offs o <b>rs1</b>	deriv deriv et is d	ved f only	rom 12 k	so oit a	urce and i	reg t is	giste sigr	er <mark>rs</mark> n-ex	1 tended to
<ul> <li>JTA =</li> <li>Note</li> <li>32-bit</li> <li>Finally</li> </ul>	rs1 + SignI that the in that s before a the stores	Ext(imn nmedia dding t the re	used i n), i.e. te offs o <b>rs1</b> turn ad	deriv deriv et is d ddres	ved f only s	rom 12 k	oit a	urce and i	re <sub>{</sub>	giste sigr	er <mark>rs</mark> n-ex	1 tended to