


Computer Architectures

Pipelined instruction execution

Hazards, stages balancing, super-scalar systems

Pavel Píša, Michal Štepanovský, Miroslav Šnorek

Main source of inspiration: Patterson



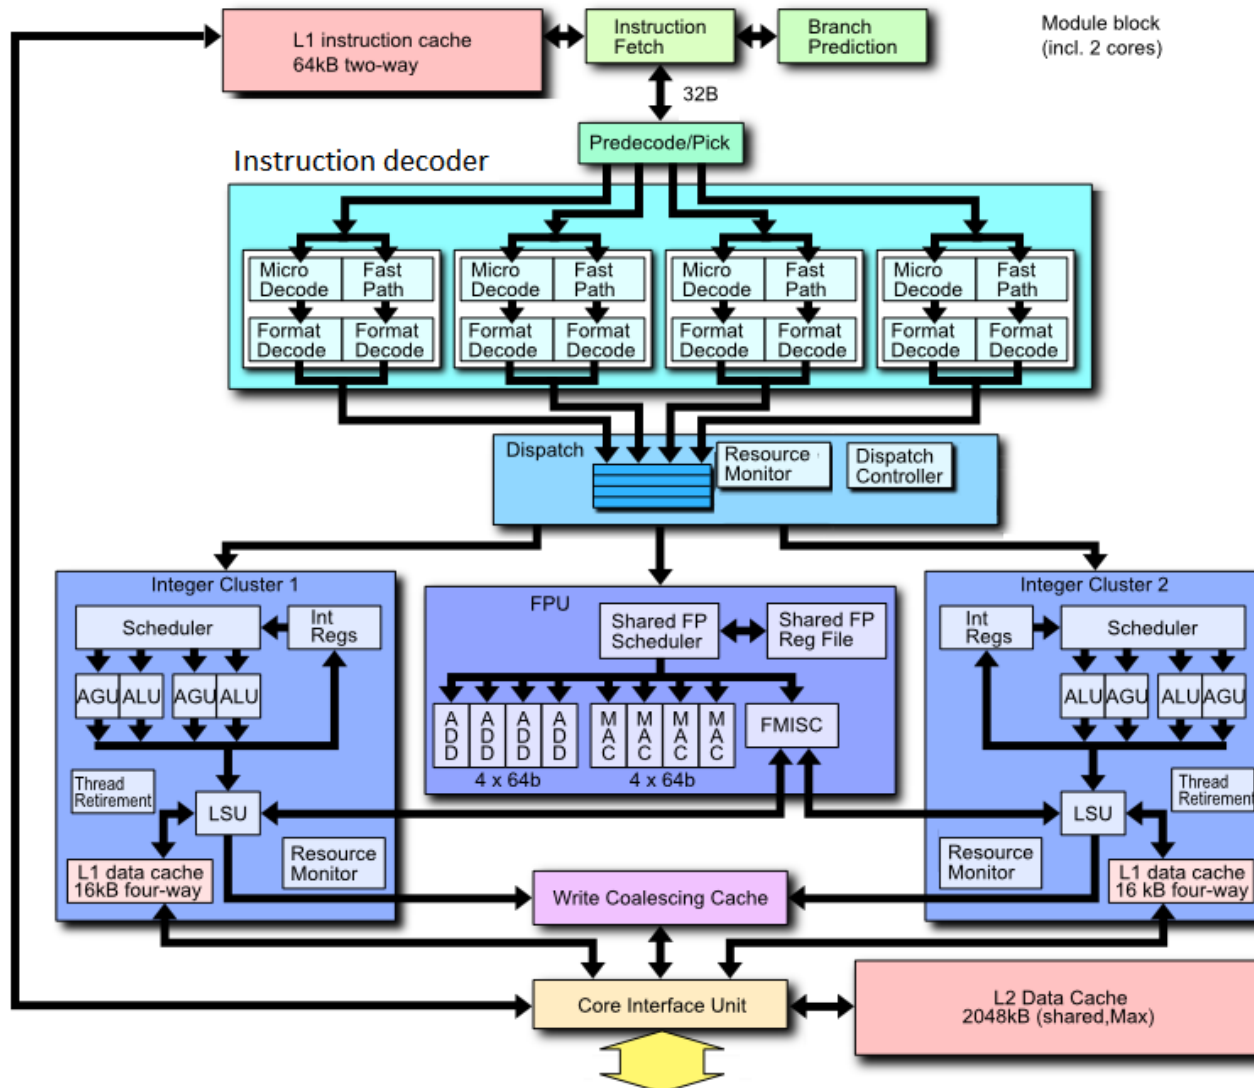
Czech Technical University in Prague, Faculty of Electrical Engineering

English version partially supported by:

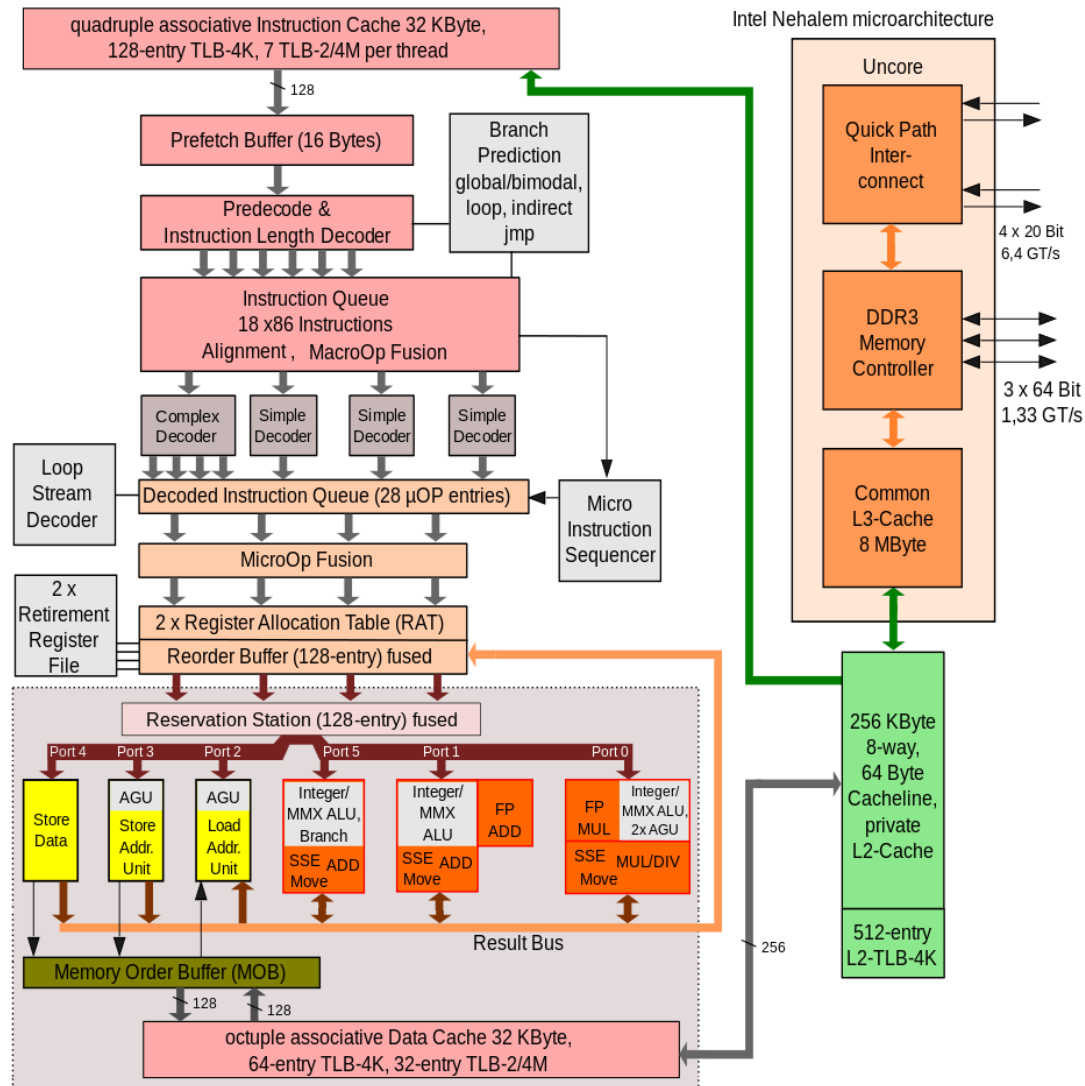
European Social Fund Prague & EU: We invests in your future.



Motivation – AMD Bulldozer 15h (FX, Opteron) - 2011



Motivation – Intel Nehalem (Core i7) - 2008

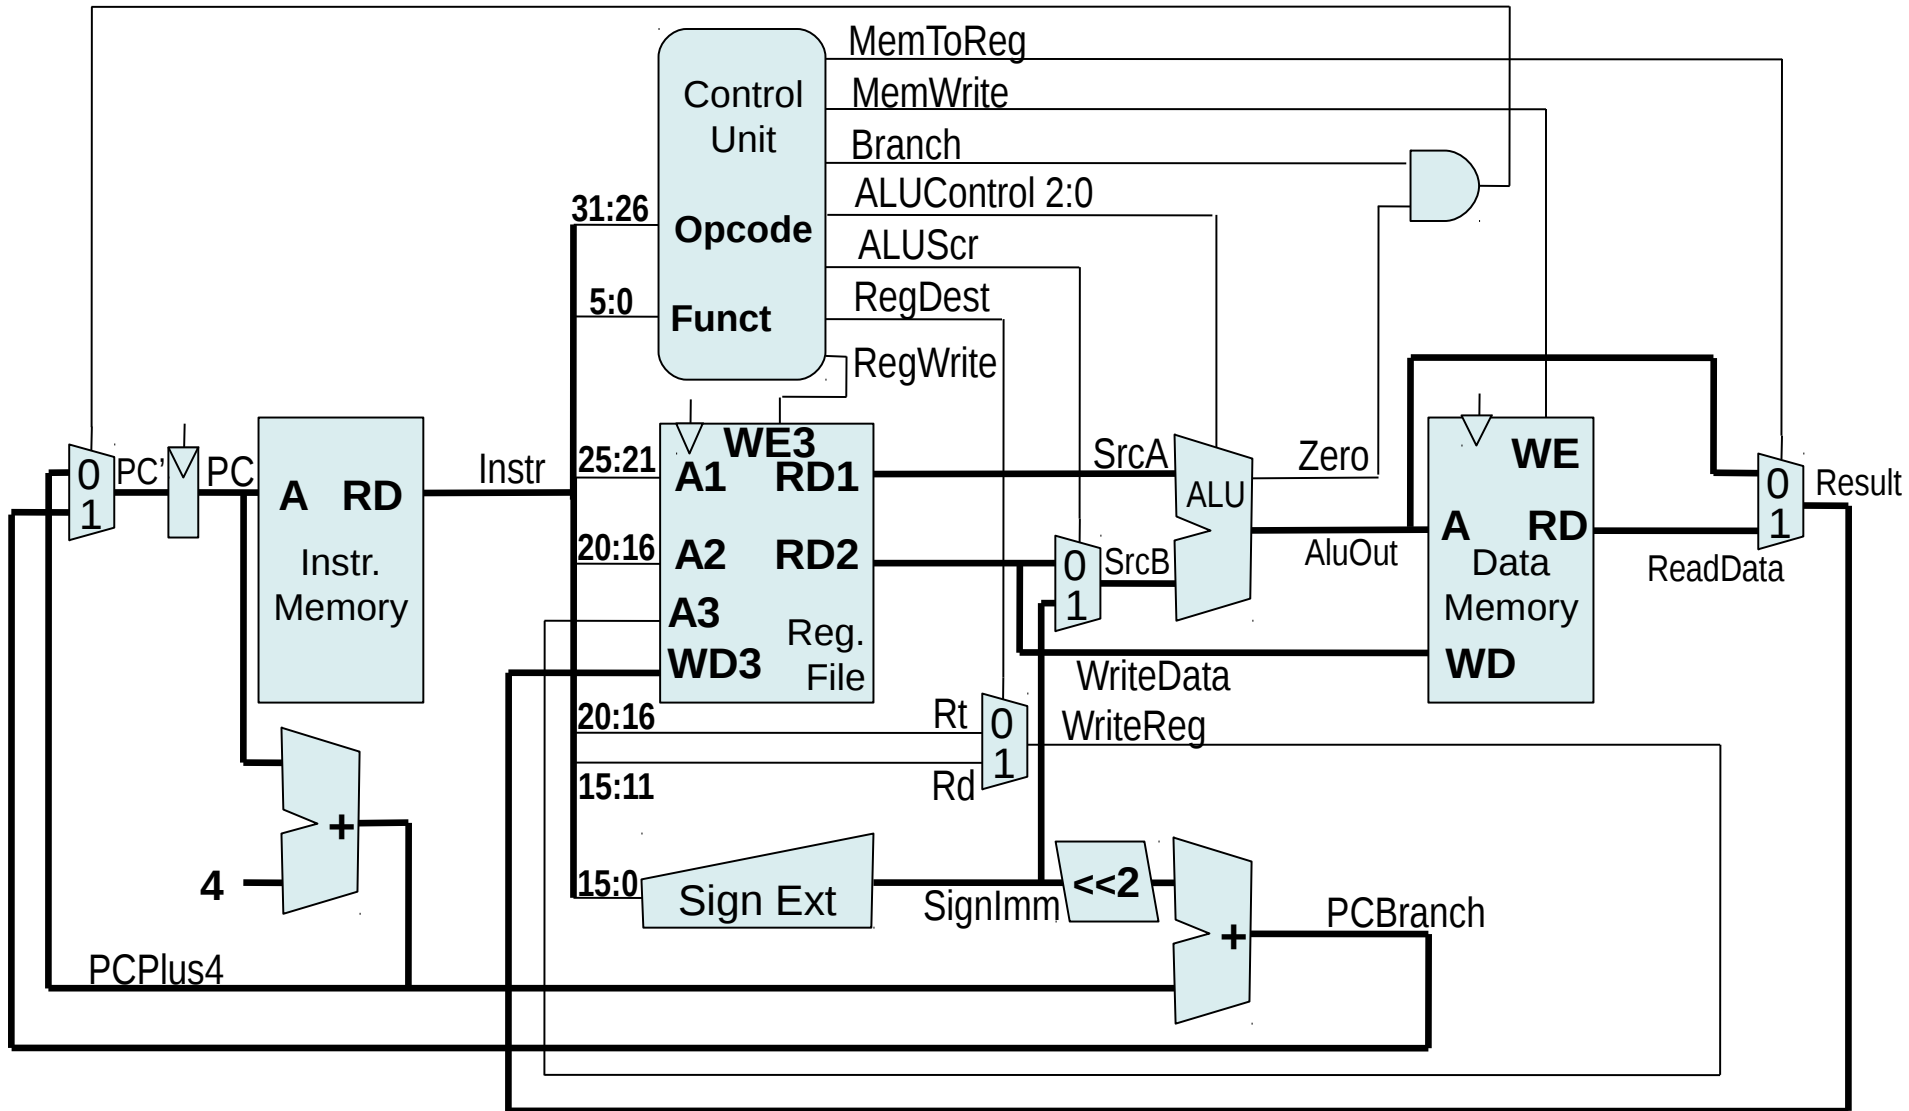


The goal of today lecture

- Convert/extend CPU presented in the lecture 2 to the pipelined CPU design.
- The following instructions are considered for our CPU design:
add, sub, and, or, slt, addi, lw, sw and beq

Typ	31... 0					
R	opcode(6), 31:26	rs(5), 25:21	rt(5), 20:16	rd(5), 15:11	shamt(5)	funct(6), 5:0
I	opcode(6), 31:26	rs(5), 25:21	rt(5), 20:16	immediate (16), 15:0		
J	opcode(6), 31:26	address(26), 25:0				

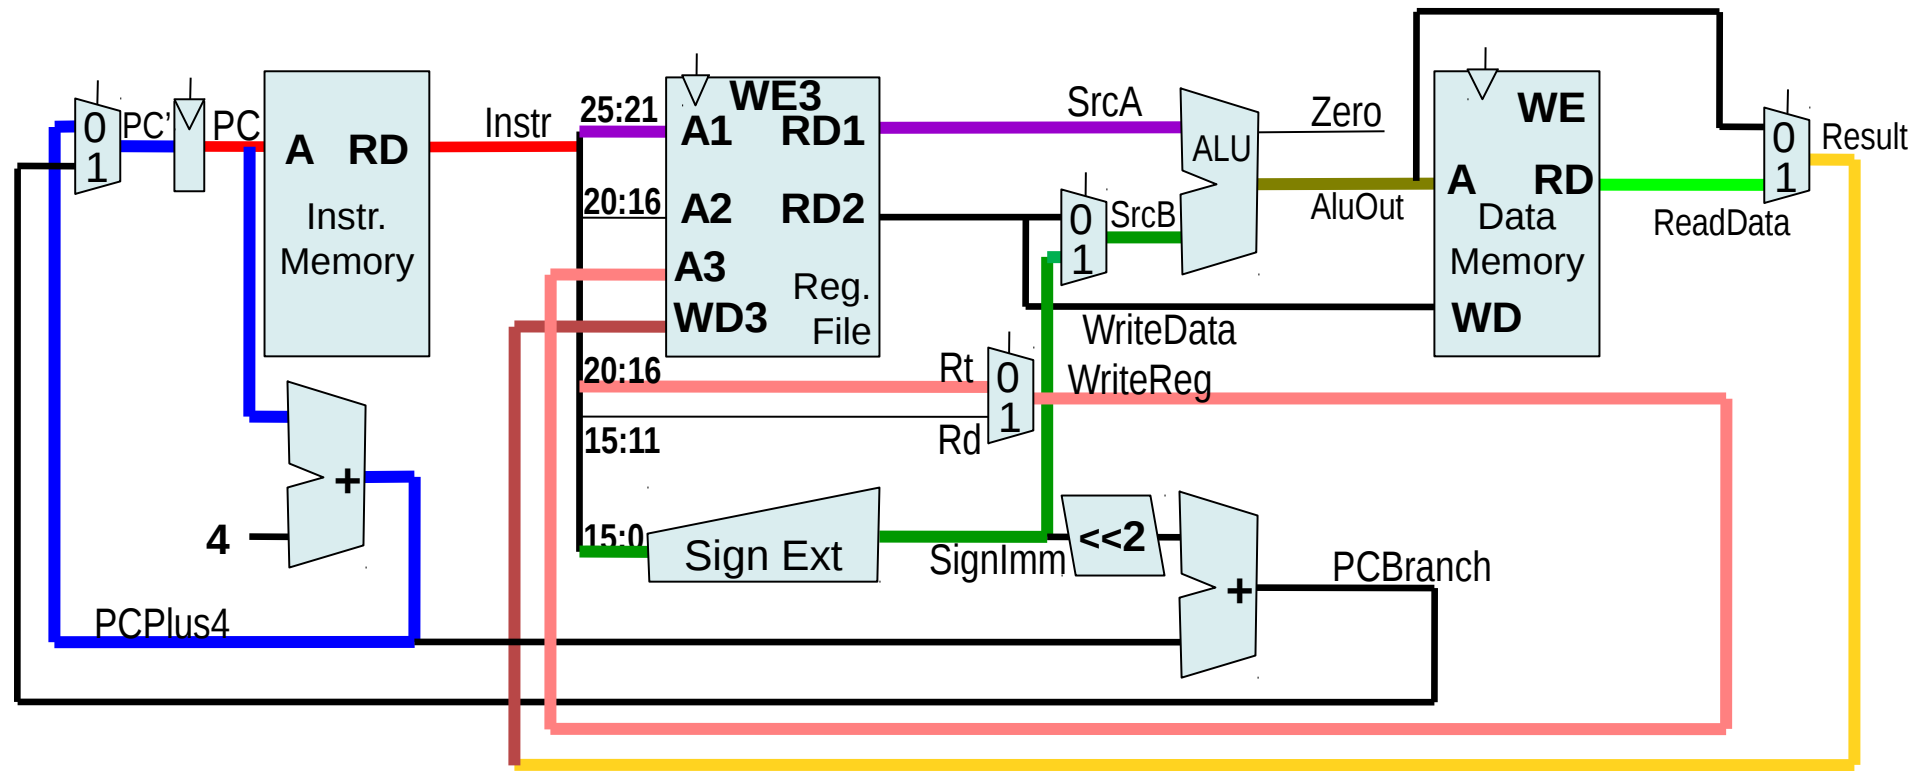
Single cycle CPU together with memories



Single cycle CPU – performance: $IPS = IC / T = IPC_{avg} \cdot f_{CLK}$

- What is the maximal possible frequency of this CPU?
- It is given by latency on the critical path – it is **lw** in our case:

$$T_c = t_{PC} + t_{Mem} + t_{RFread} + t_{ALU} + t_{Mem} + t_{Mux} + t_{RFsetup}$$



Single cycle CPU – throughput: $IPS = IC / T = IPC_{avg} \cdot f_{CLK}$

- $T_c = t_{PC} + t_{Mem} + t_{RFread} + t_{ALU} + t_{Mem} + t_{Mux} + t_{RFsetup}$
- Consider following parameters

$$t_{PC} = 30 \text{ ns}$$

$$t_{Mem} = 300 \text{ ns}$$

$$t_{RFread} = 150 \text{ ns}$$

$$t_{ALU} = 200 \text{ ns}$$

$$t_{Mux} = 20 \text{ ns}$$

$$t_{RFsetup} = 20 \text{ ns}$$

Then $T_c = 1020 \text{ ns} \rightarrow f_{CLK \max} = 980 \text{ kHz}$,

$IPS = 1 \cdot 980e3 = 980 \text{ 000 instructions per second}$

Pipelined instructions execution

Suppose that instruction execution can be divided into 5 stages:



IF – Instruction Fetch, ID – Instruction decode (and Operands Fetch),
EX – Execute, MEM – Memory Access, WB – Write Back

and $\tau = \max \{ \tau_i \}_{i=1}^k$, where τ_i is time required for signal propagation (*propagation delay*) through i -th stage.

IF – setup PC for memory and fetch pointed instruction. Update PC = PC+4

ID – decode the opcode and read registers specified by instruction, check for equality (for possible beq instruction), sign extend offset, compute branch target address for branch case (this is means to extend offset and add PC)


EX – execute function/pass register values through ALU

MEM – read/write main memory for *load/store* instruction case

WB – write result into RF for instructions of register-register class or instruction *load* (result source is ALU or memory)

Instruction-level parallelism - pipelining

IF	I1	I2	I3	I4	I5	I6	I7	I8	I9	I10
ID		I1	I2	I3	I4	I5	I6	I7	I8	I9
EX			I1	I2	I3	I4	I5	I6	I7	I8
MEM				I1	I2	I3	I4	I5	I6	I7
ST					I1	I2	I3	I4	I5	I6



- The time to execute n instructions in the k -stage pipeline:

$$T_k = k \cdot \tau + (n - 1) \tau$$

- Speedup:
$$S_k = \frac{T_1}{T_k} = \frac{nk \tau}{k\tau + (n - 1) \tau} \quad \lim_{n \rightarrow \infty} S_k = k$$

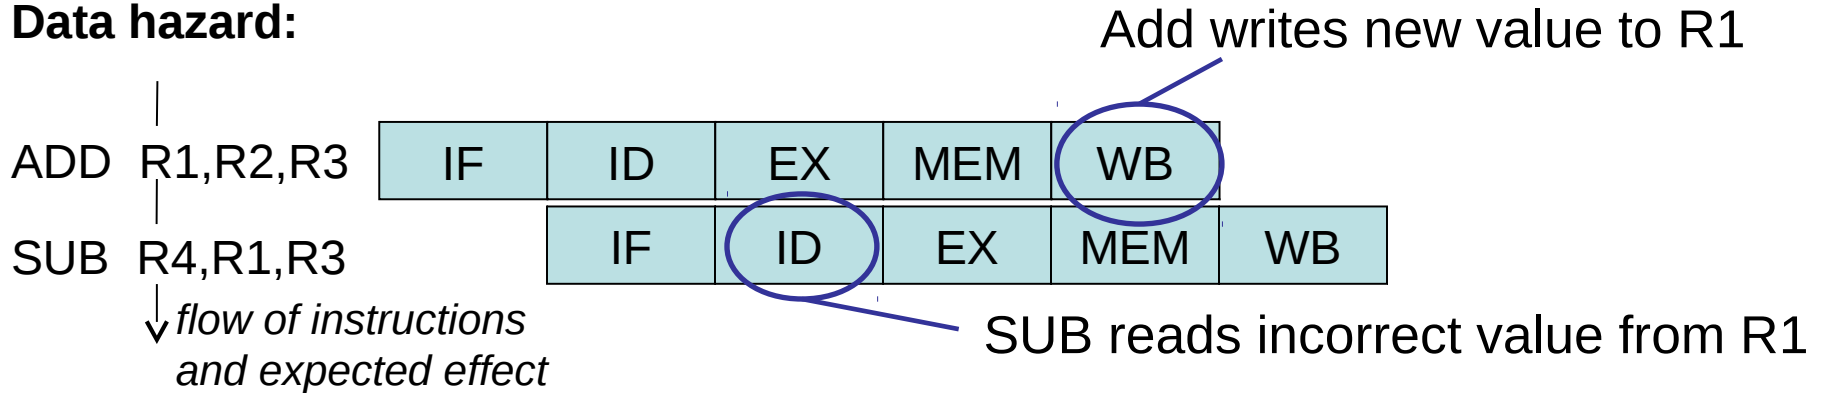
Prerequisite: pipeline is optimally balanced, circuit can arbitrarily divided

Instruction-level parallelism - pipelining

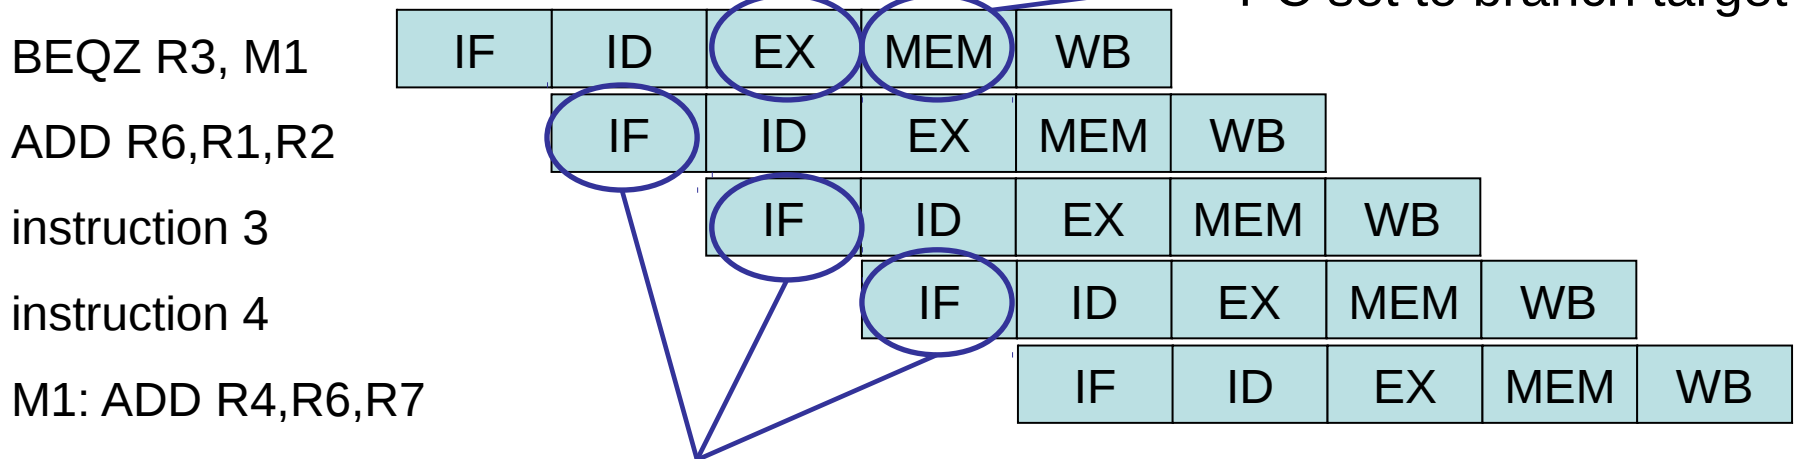
- Does not reduce the execution time of individual instructions, effect is just the opposite...
- Hazards:
 - structural (resolved by duplication),
 - data (result of data dependencies: RAW, WAR, WAW)
 - control (caused by instructions which change PC)...
- Hazard prevention can result in pipeline stall or pipeline flush
- Remark : Deeper pipeline (more stages) results in shorter sequences of gates in each stage which enables to increase the operating frequency of the processor..., but more stages means higher overhead (demand to arrange better instructions into pipeline and result in more significant lag in the case of stall or pipeline flush)

Instruction-level parallelism – Semantics violations

Data hazard:

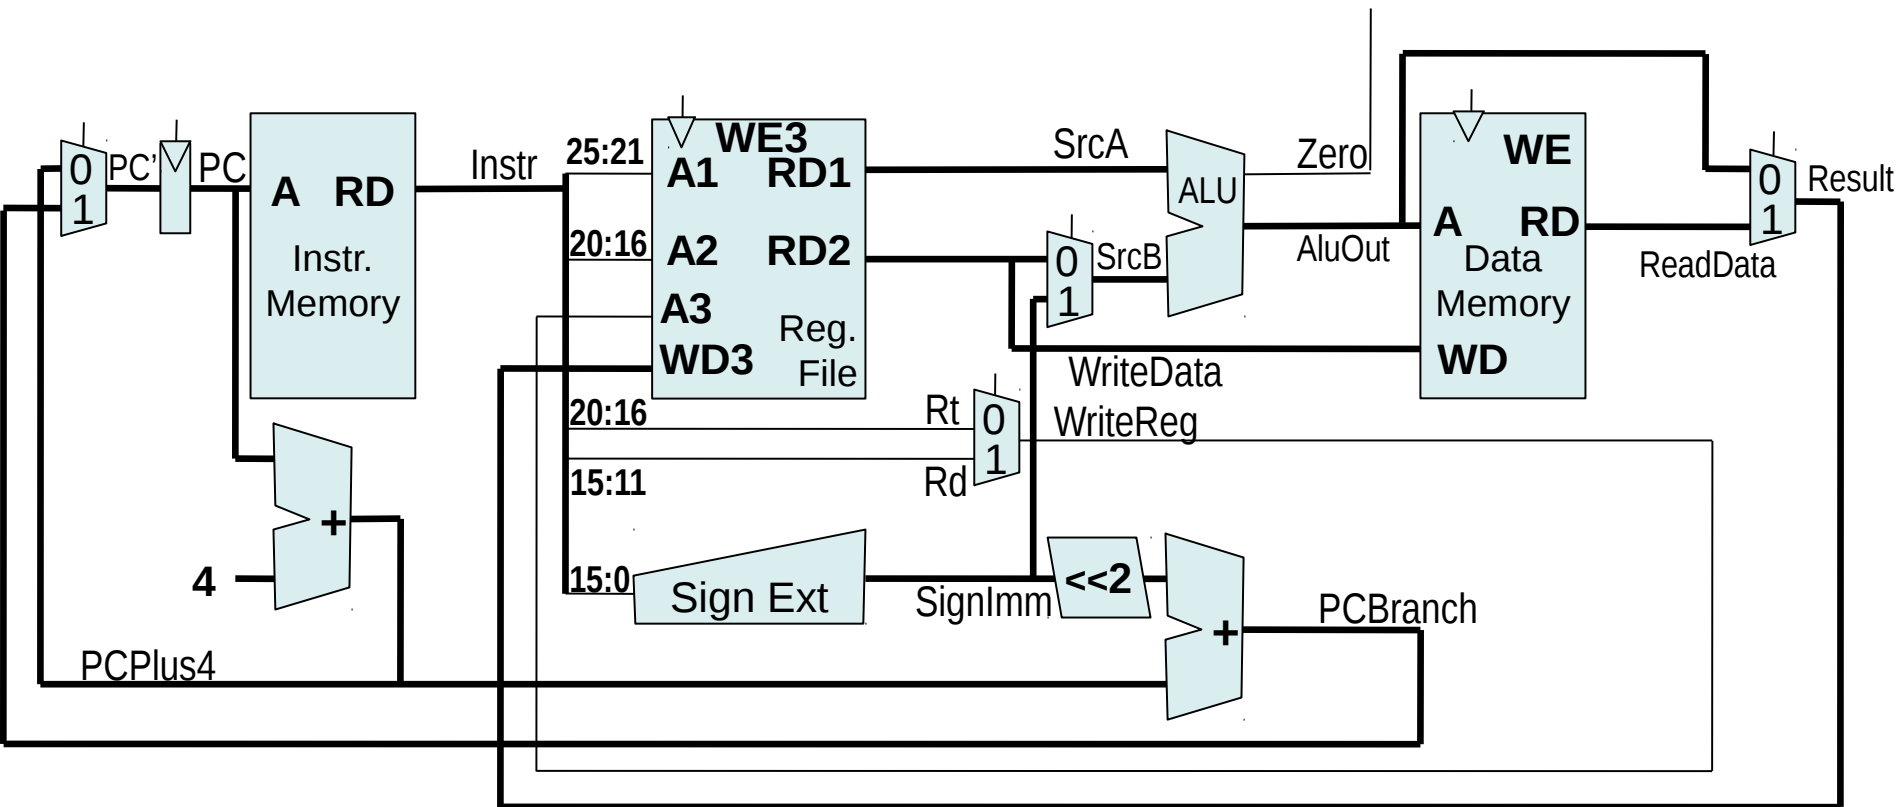


Control hazard:

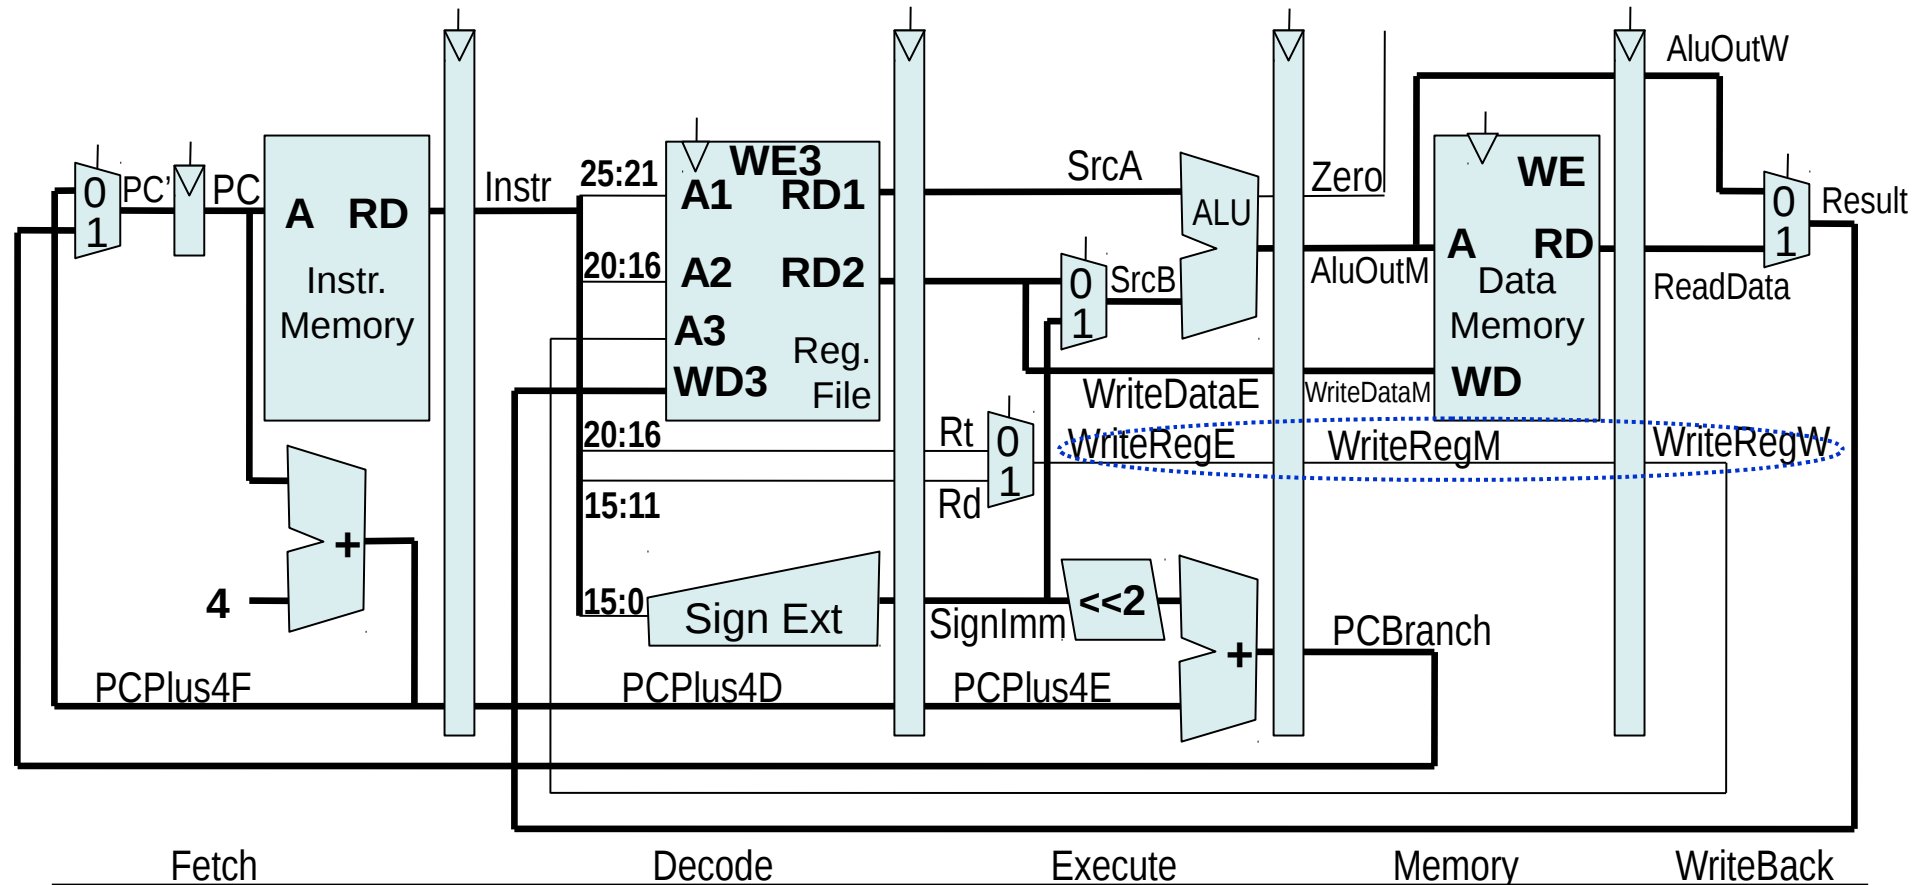


Should be these instructions fetched (and executed then)?

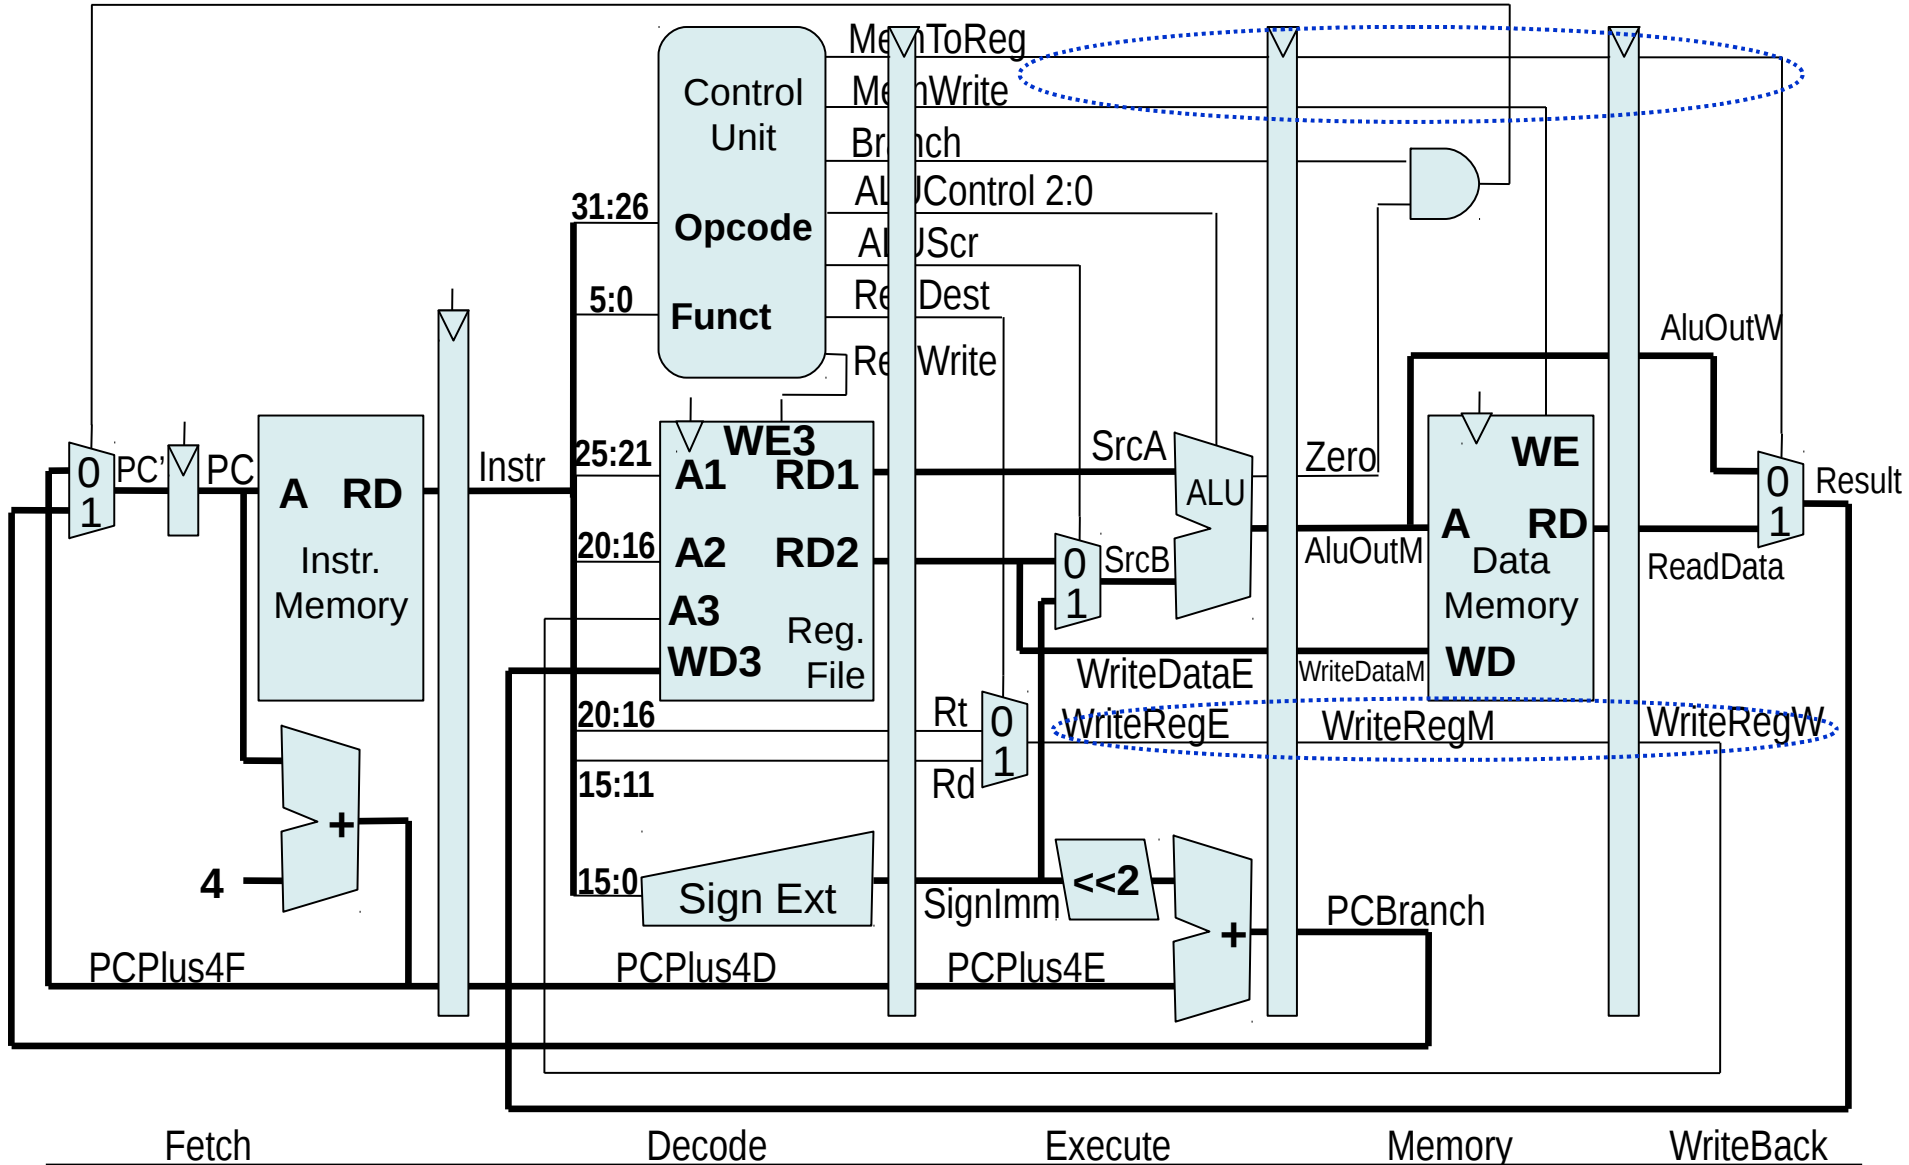
Non-pipelined execution



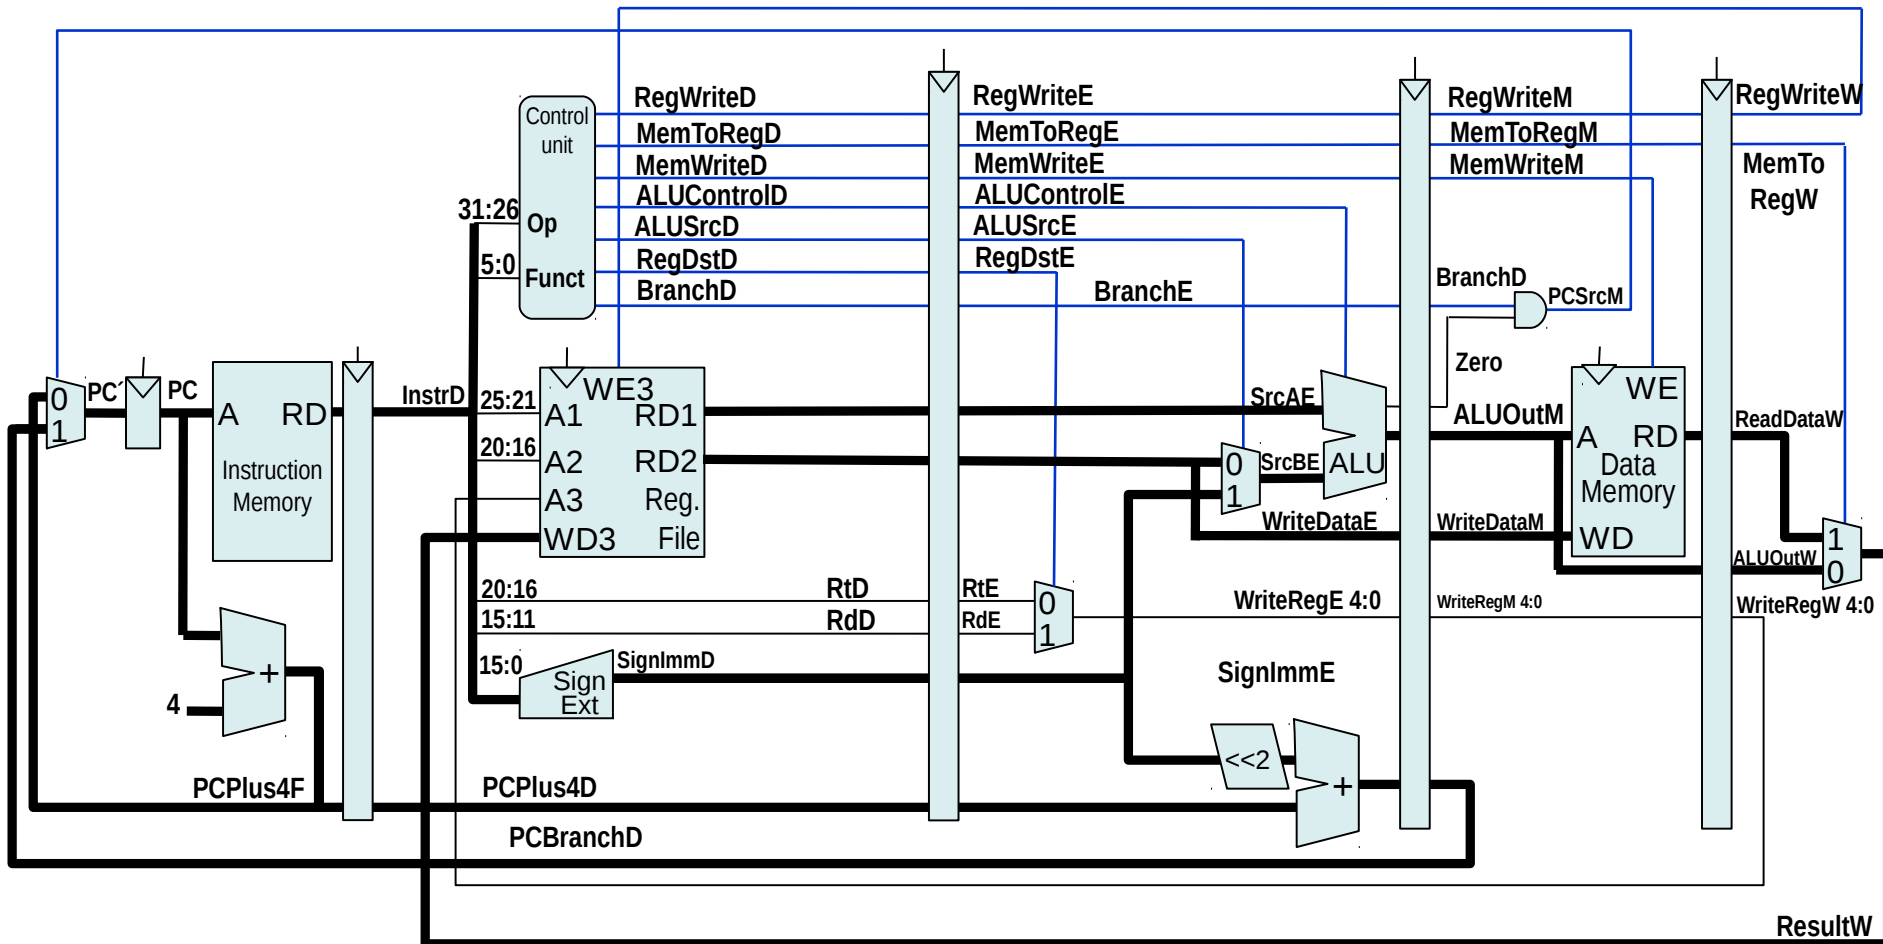
Pipelined execution




Pipelined execution



The same design but drawn scaled down...




Cause of the data hazards



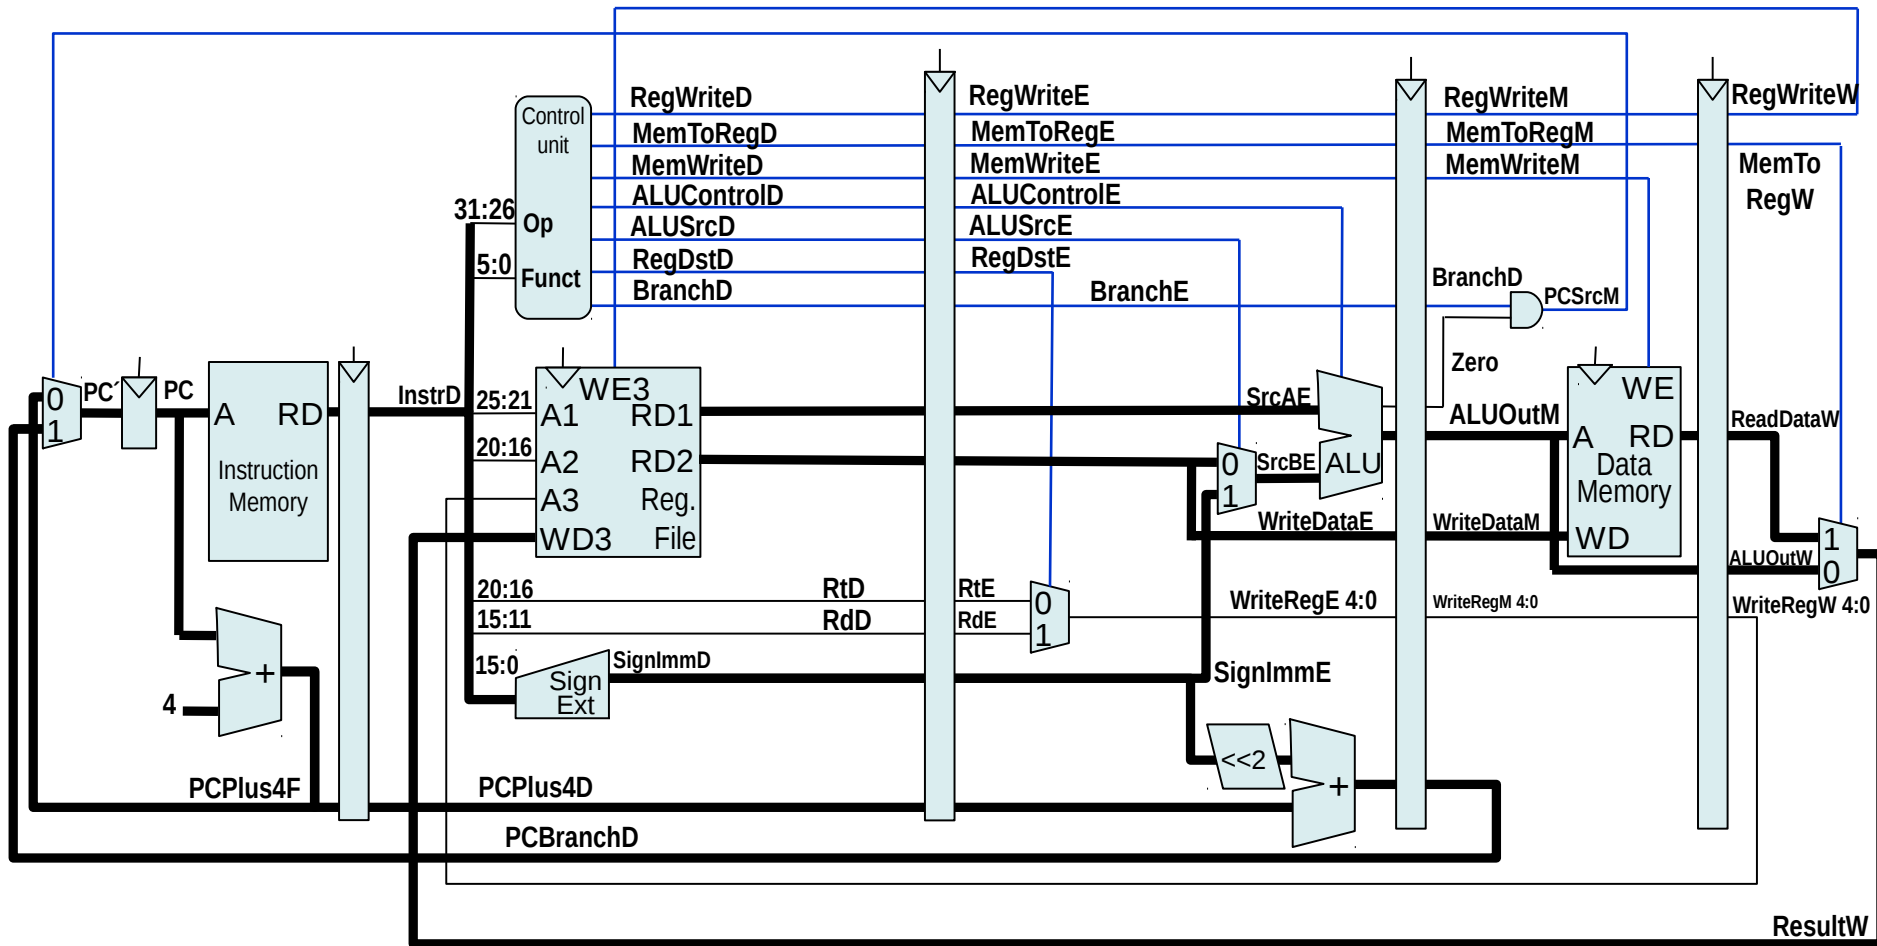
- Register File – access from two pipeline stages (Decode, WriteBack) – actual write occurs at the first half of the clock cycle, the read in the second half \Rightarrow there is no hazard for `sub $s0` input operand
- RAW (Read After Write) hazard – `and` (`or`) requires `$s0` in 3 (4)
- How can such hazard be prevented without pipeline throughput degradation?

Forwarding to avoid data hazards

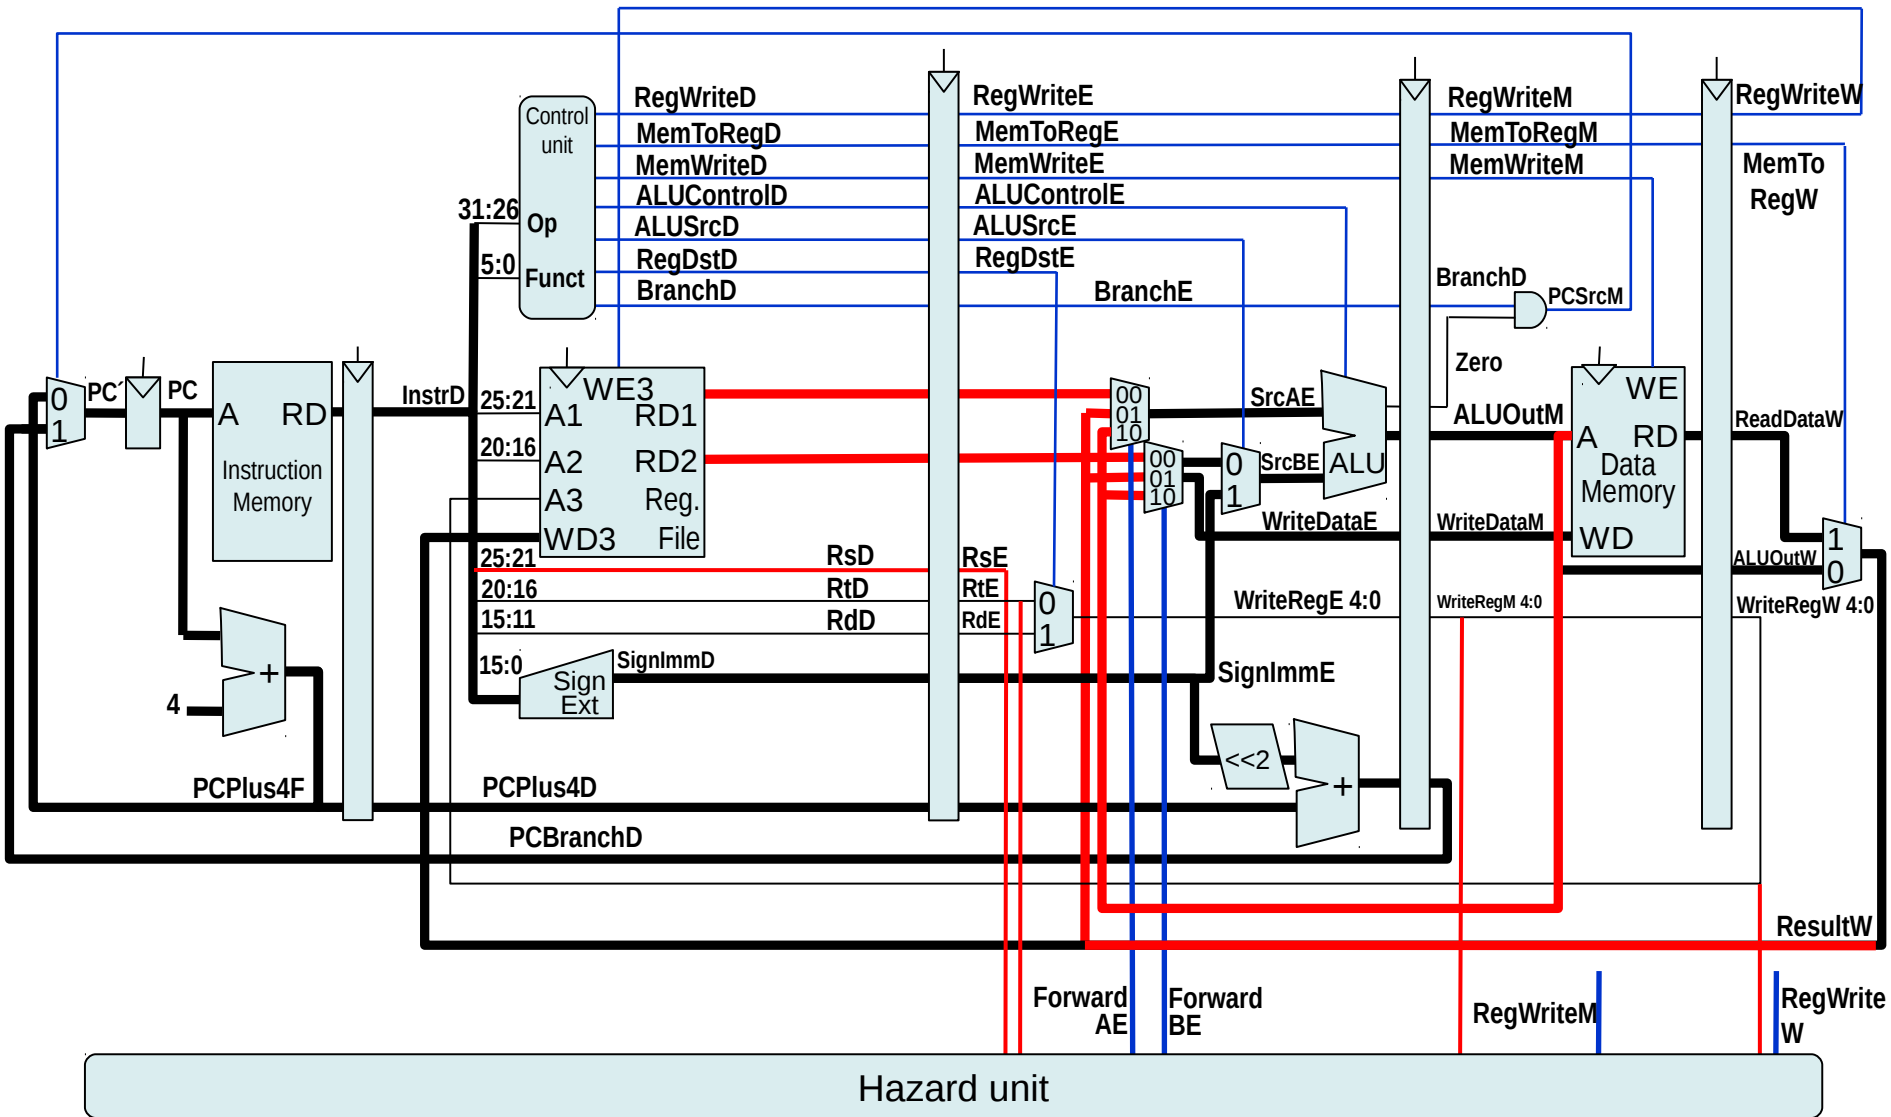


- If a result is available (computed) before subsequent instruction(s) requires the value then data hazard can be avoided by forwarding
- Hazard case is indicated when some of source registers in EX stage is the same as destination register in stage MEM or WB
- The register numbers are fed to the Hazard Unit
- The RegWrite signal from MEM and WB stage has to be monitored as well to check that register number on WriteReg lines takes effect – lw / sw etc.


CPU after previous design steps



Data hazards solved by forwarding




Data hazard avoided by pipeline stall



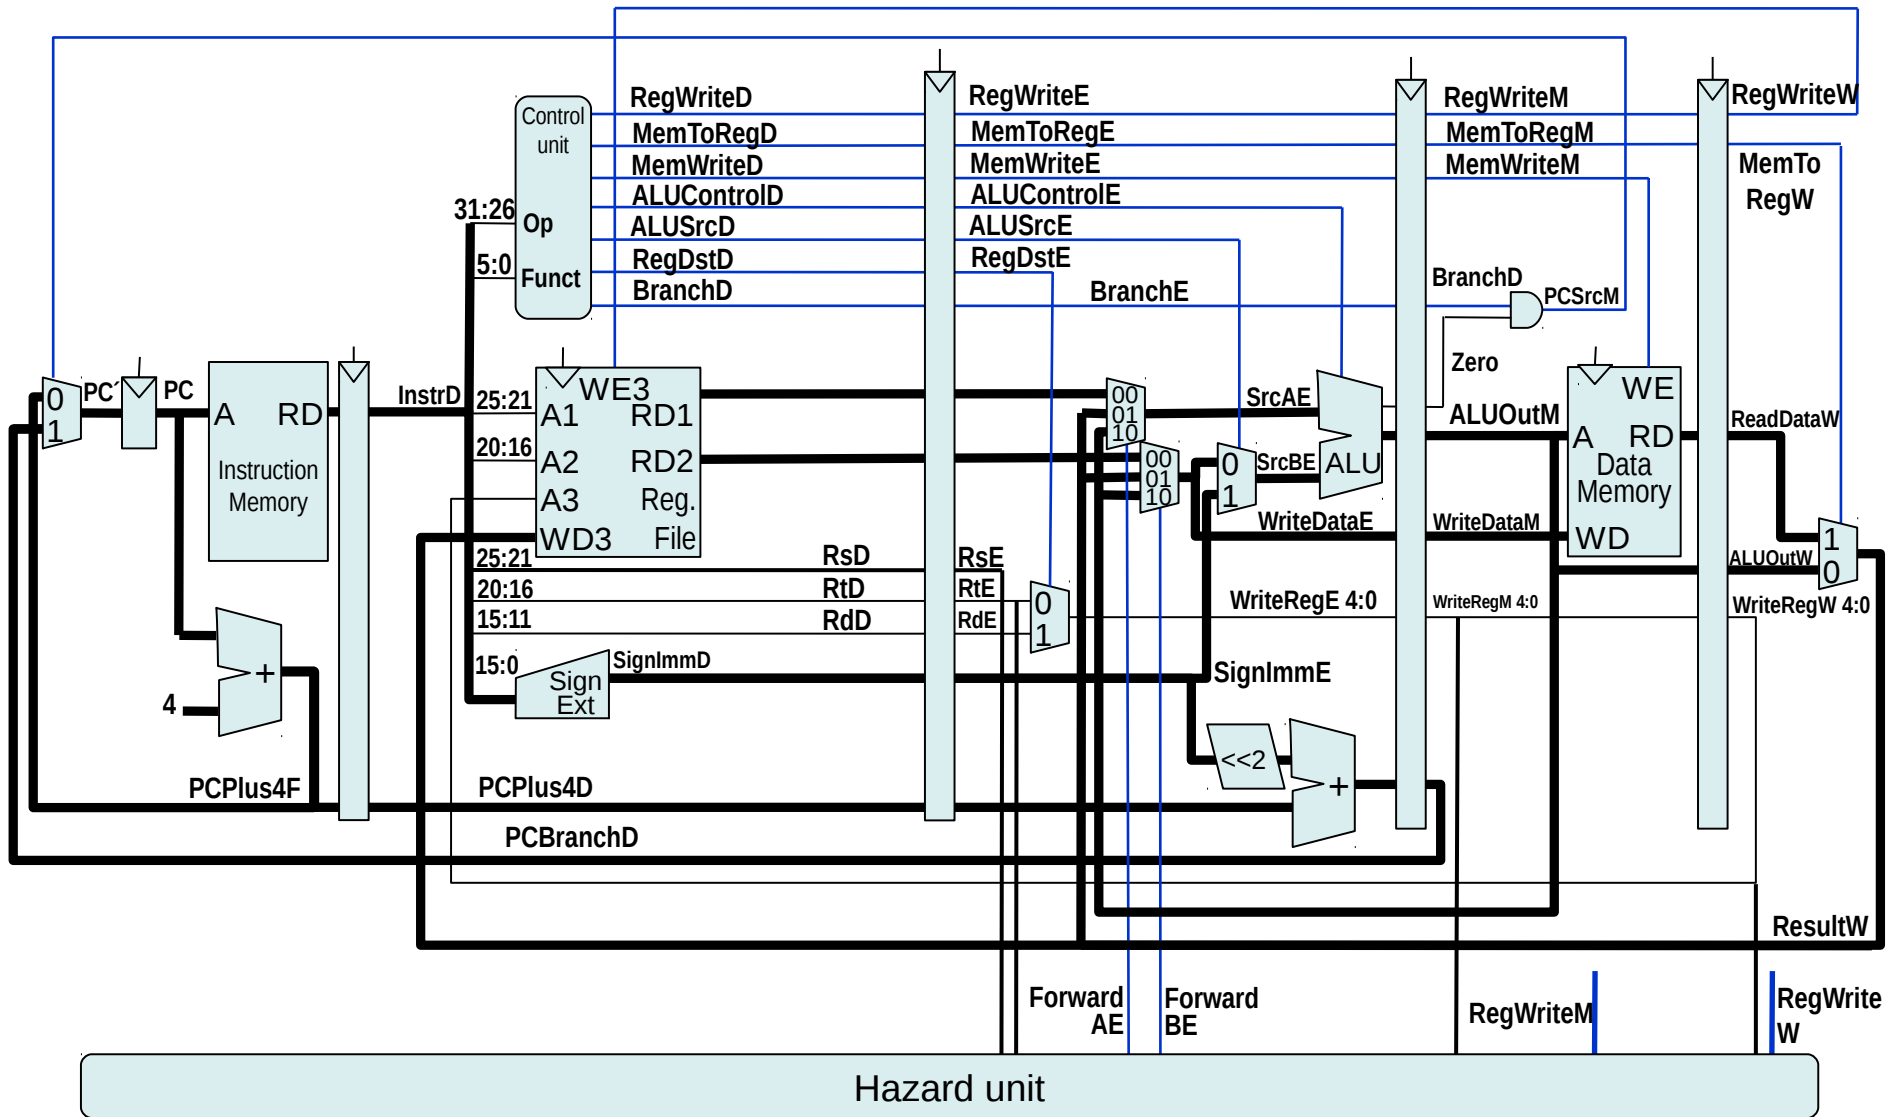
- If subsequent instructions require result before it is available in CPU then the pipeline has to be stalled (stall state inserted)
- The stall is mean to solve hazard but affect system throughput
- Pipeline stages preceding that one which is affected by the hazard are stalled until all results required by subsequent instructions are available – results are forwarded to the sink which required their value

Data hazard avoided by pipeline stall

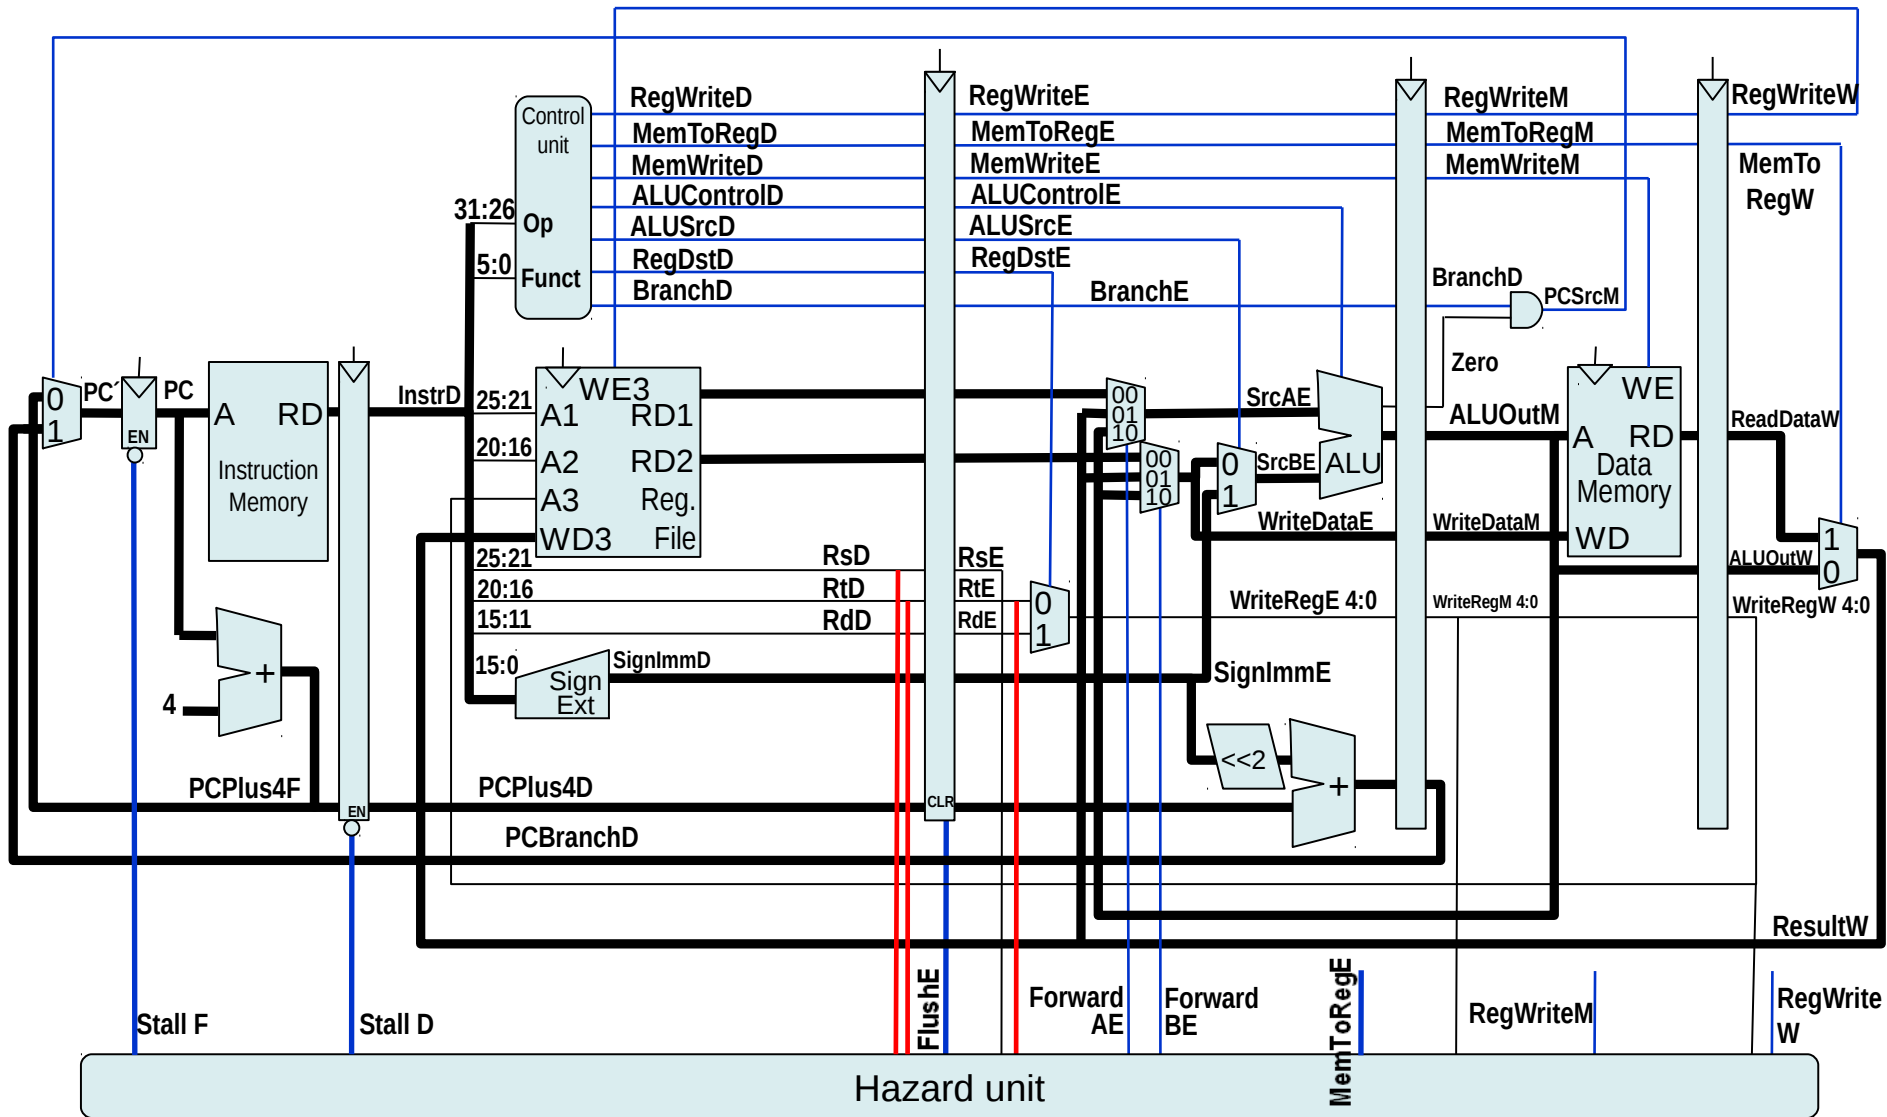


- The stall is realized by the holding content of the inter-stage registers (gating their clocks or blocking their latch enable signals)
- Results from colliding stages have to be „discarded“ – certain control signals in CPU (RF or memory write enable, branch gating) are reset (held low)
- Both is achieved by introduction of control signals to hold and/or reset inter-stages registers

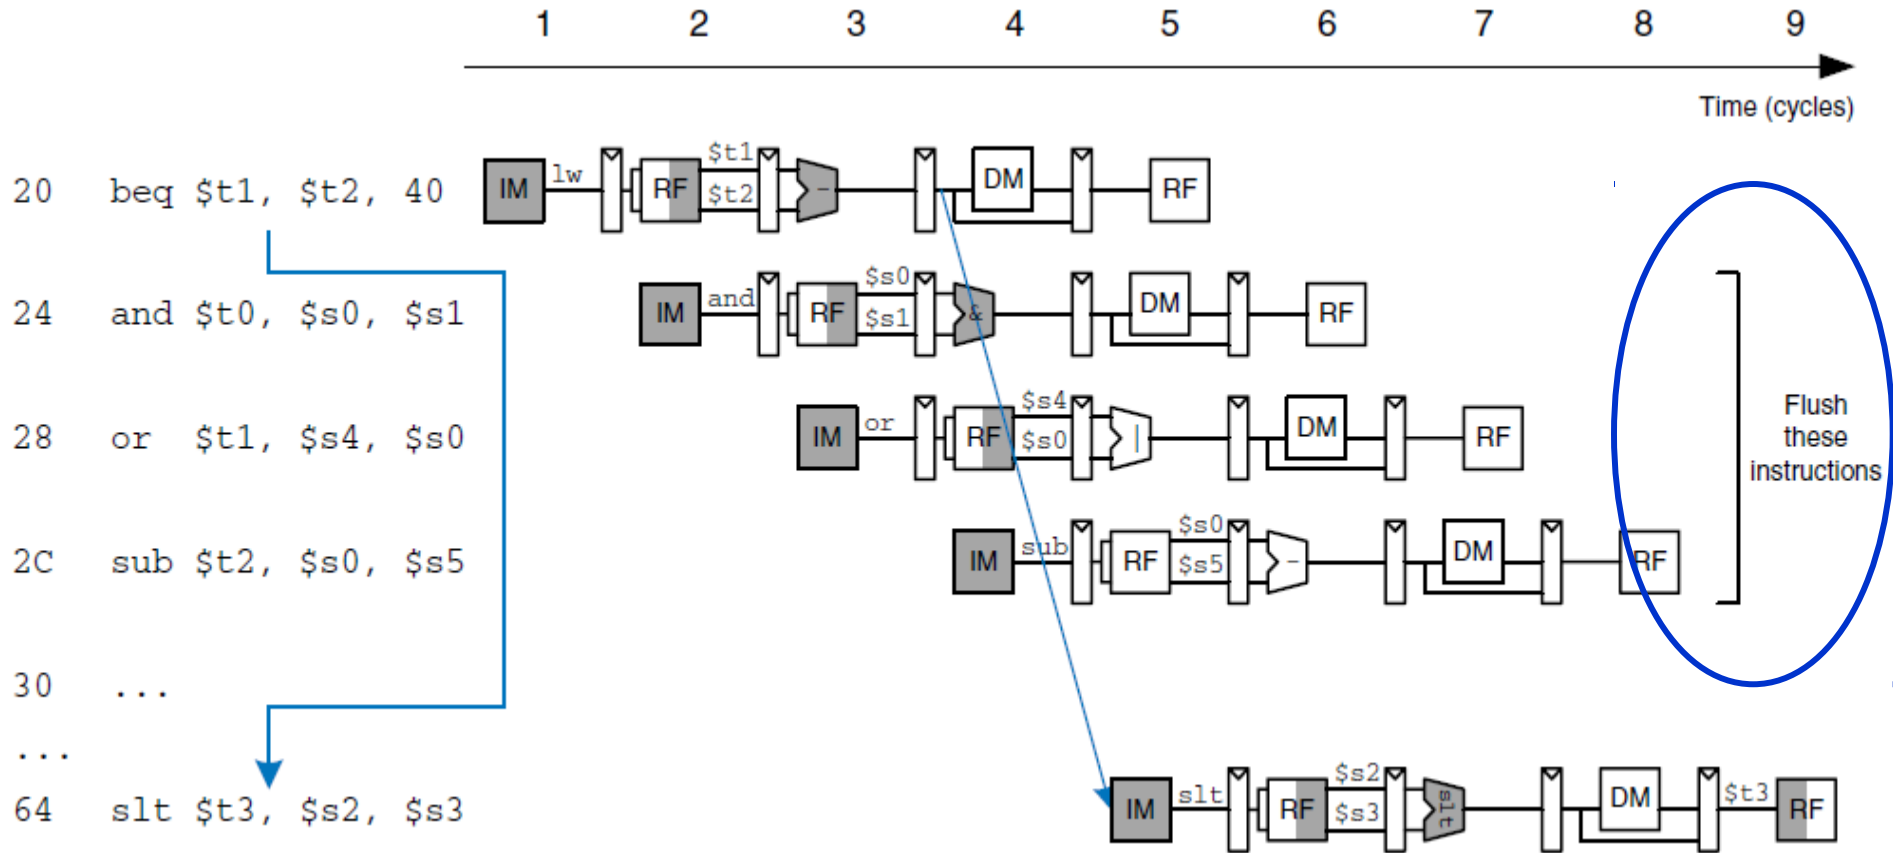
Processor design build till now



Processor with data hazards avoided by stall




Control hazards (branch and jump)



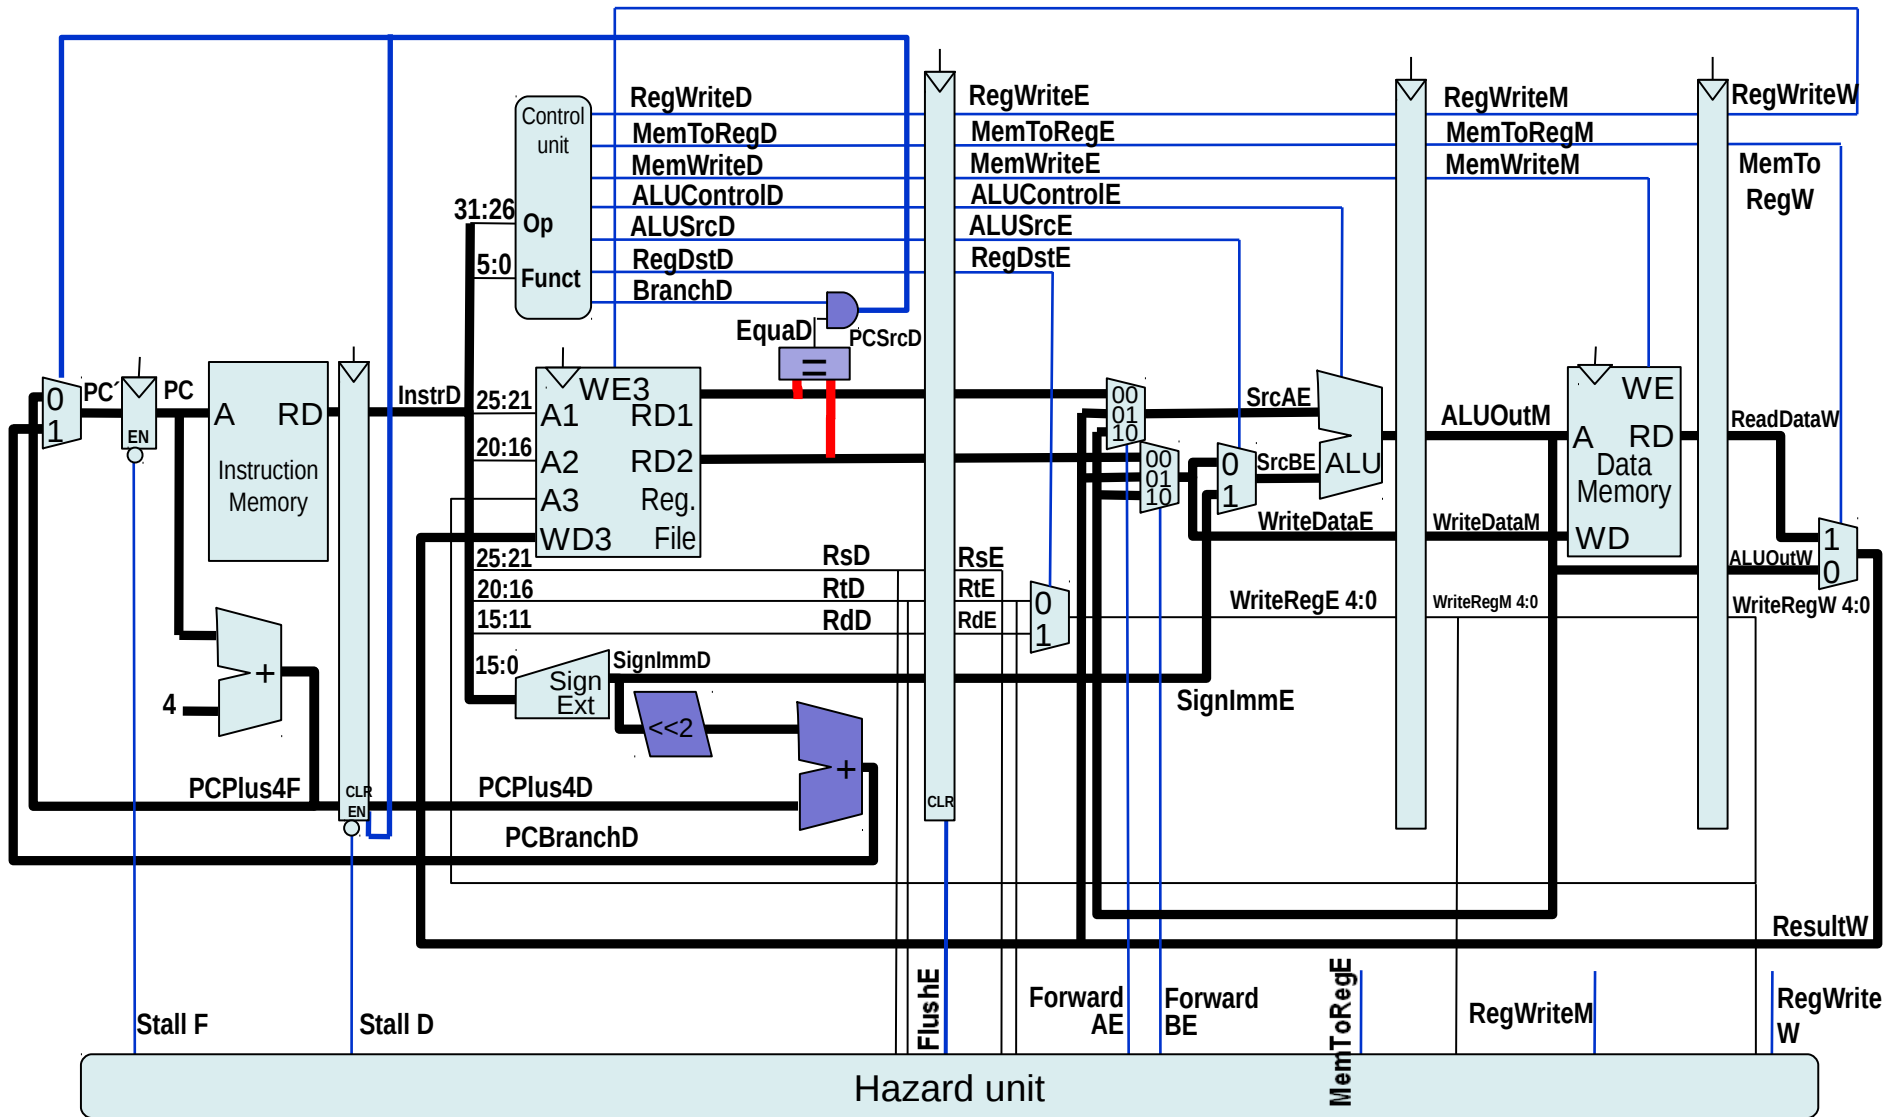
- Result is not known before 4th cycle. Why?

Control hazards – better to know result earlier...

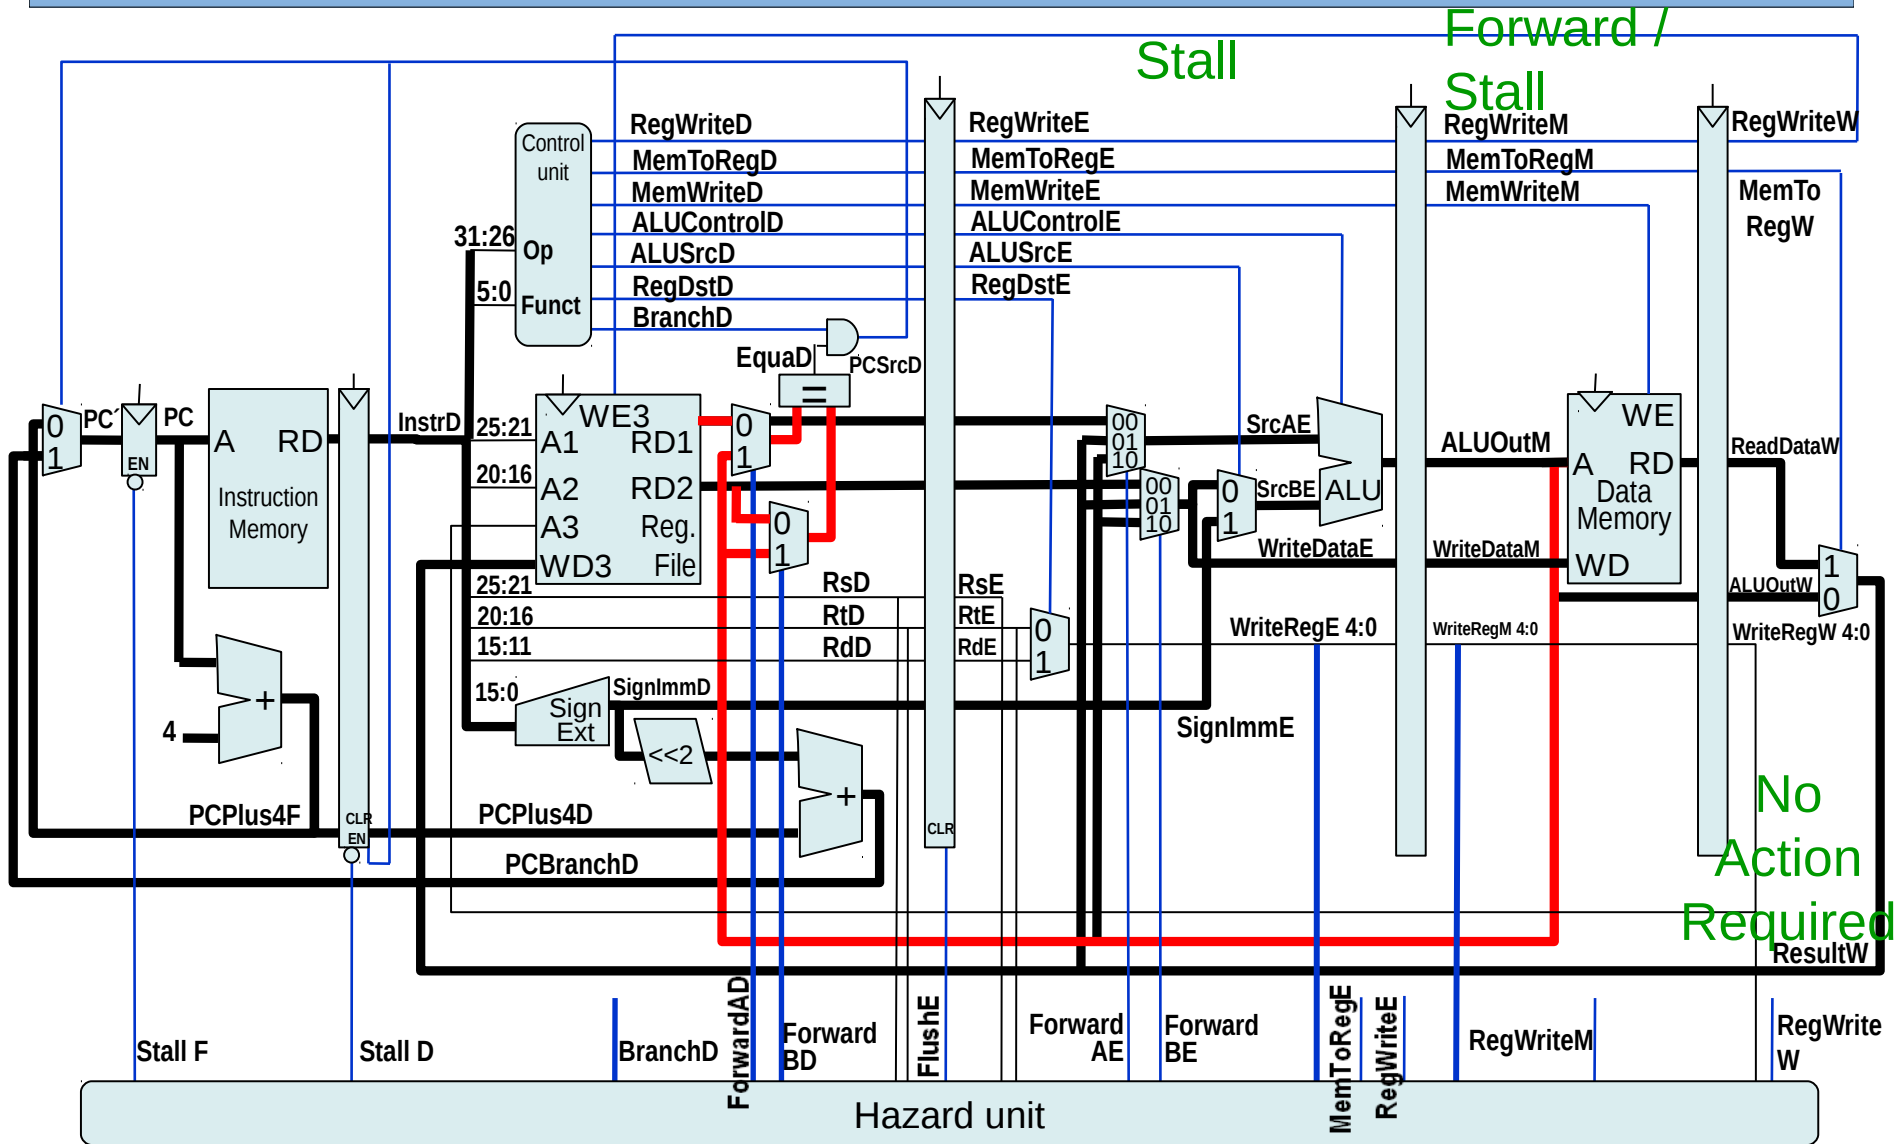


- If the result of comparison can be evaluated in the 2nd cycle **misprediction penalty** can be reduced
- But the processing of the comparison at earlier stage can induce new RAW hazards...!!!

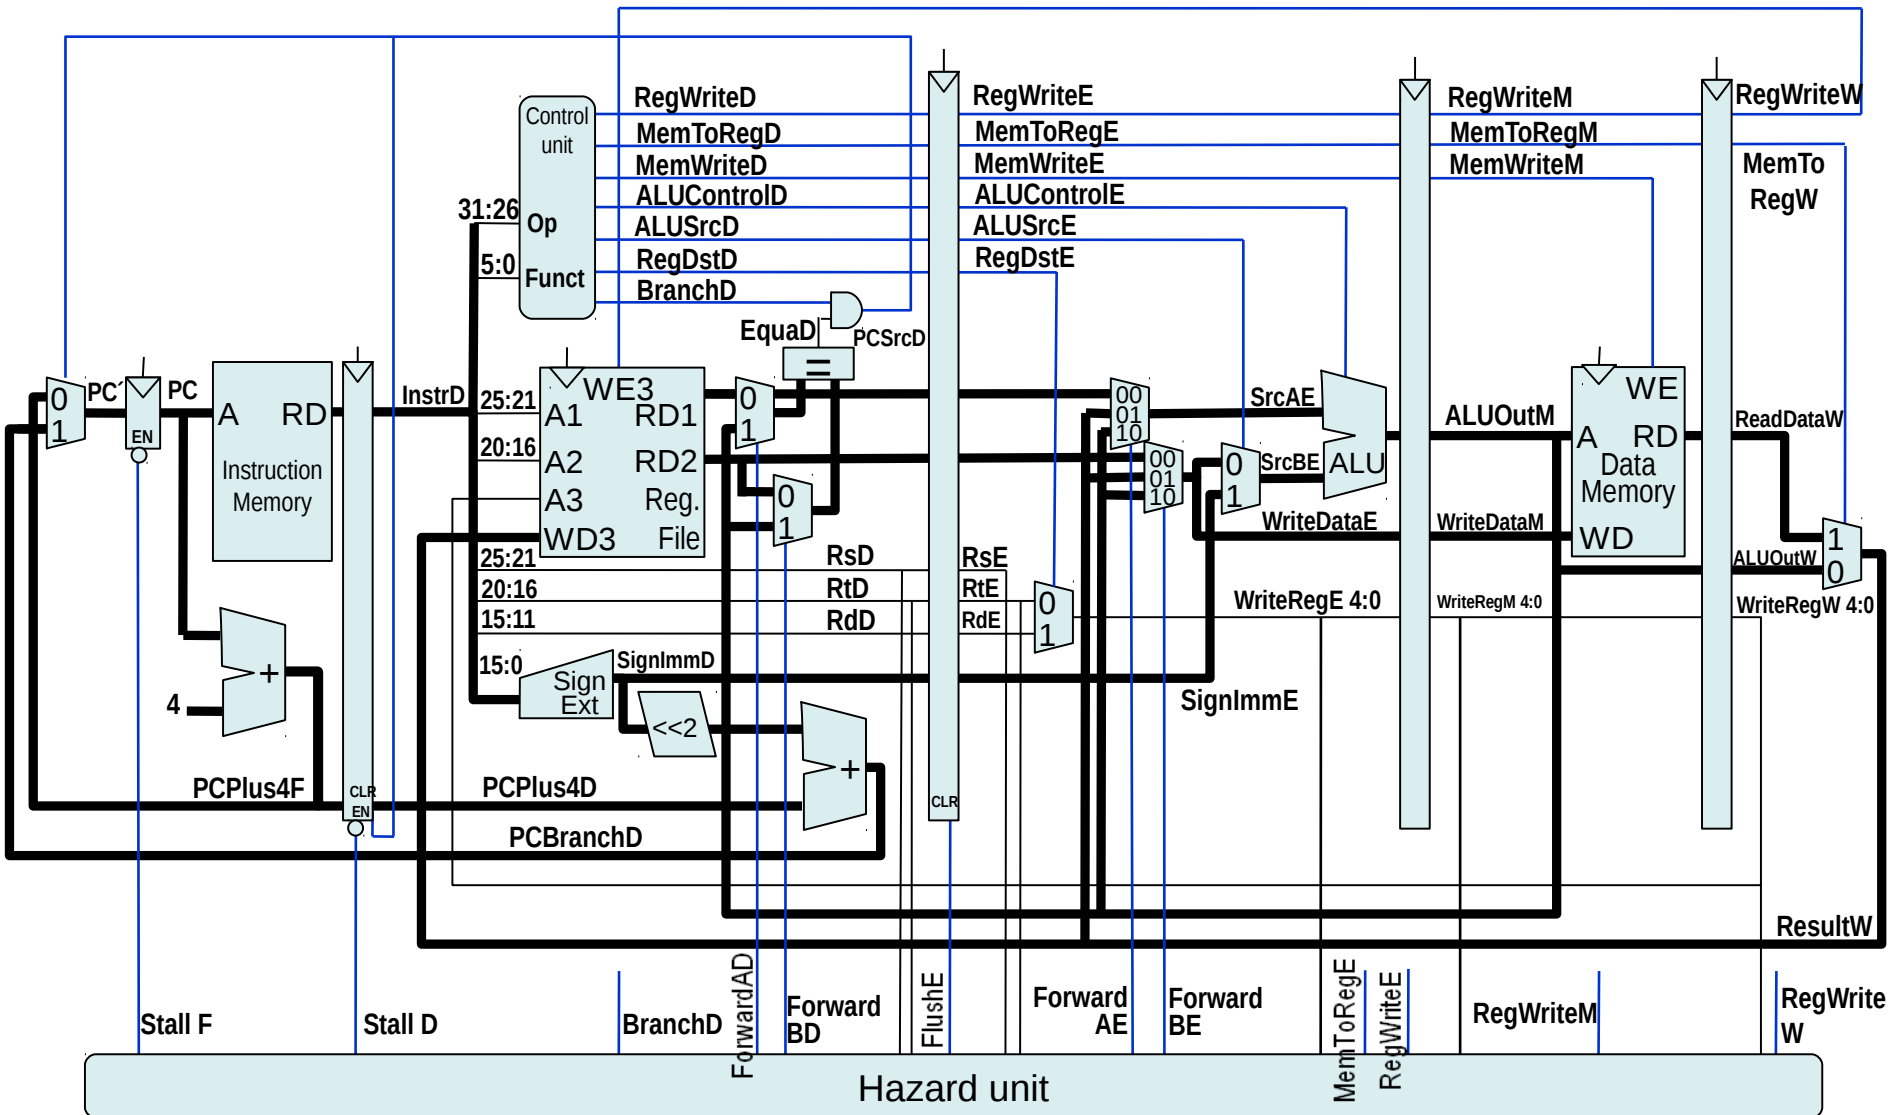
Resolve control hazards by early evaluate and flush



Resolve RAW hazards by forwarding or stalling



We are finished – pipelined processor is designed



Pipelined CPU – performance: $IPS = IC / T = IPC_{avg} \cdot f_{CLK}$

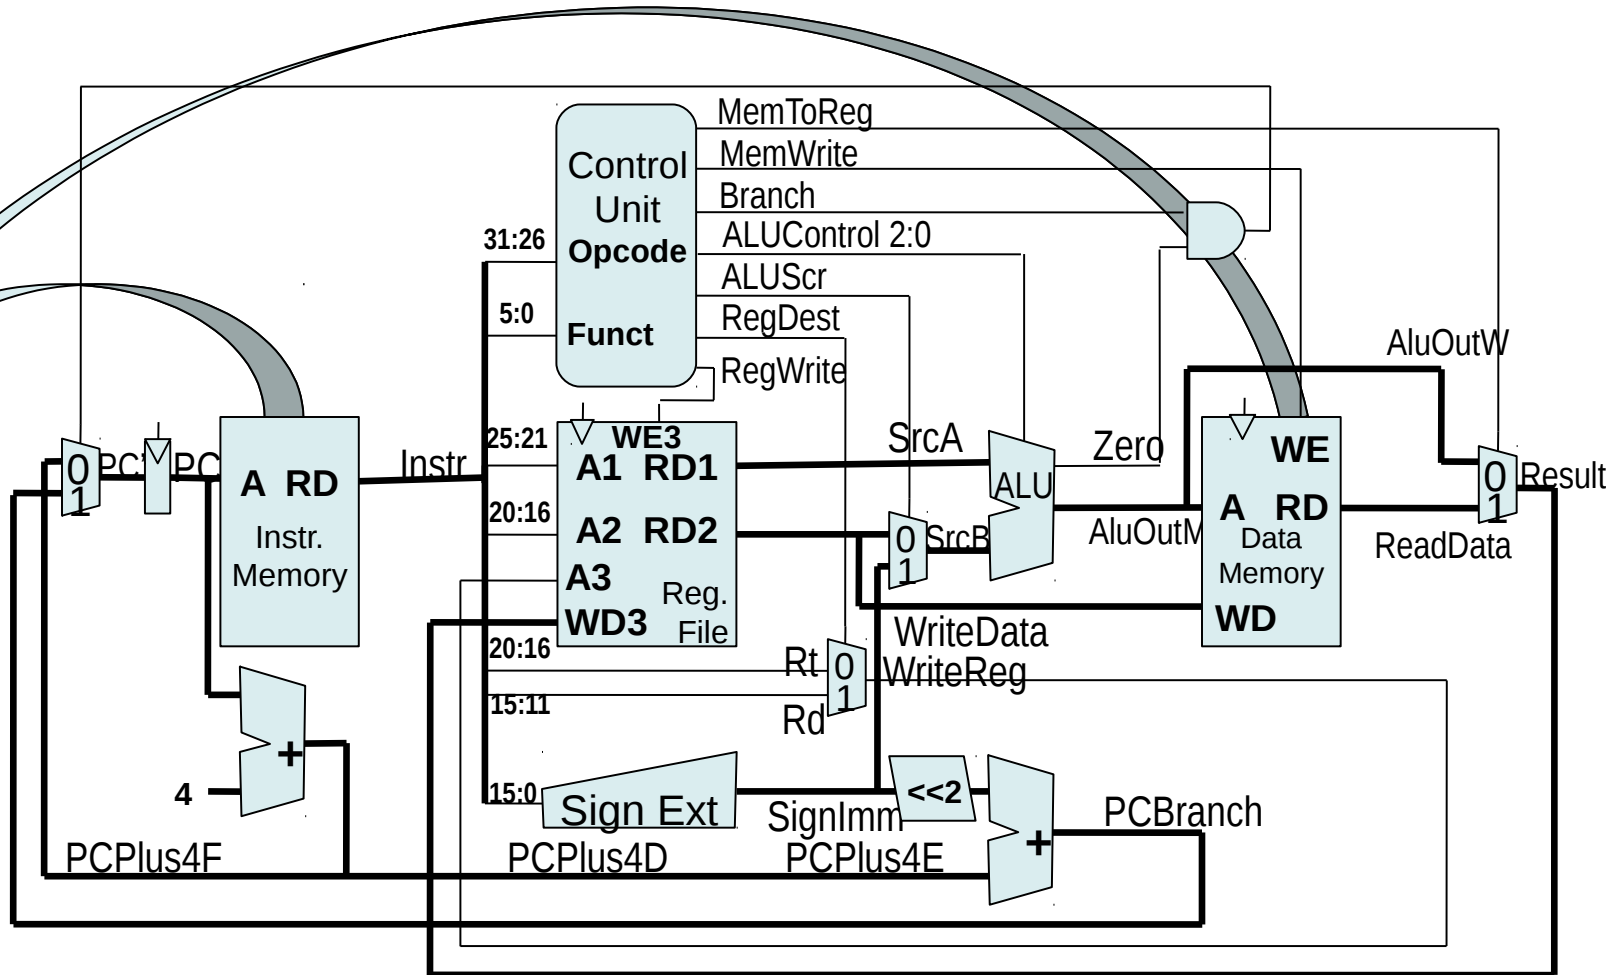
- What is maximal acceptable frequency for the CPU?
- Which stage is the slowest one?
- The cycle time is determined by the slowest stage
- For our case:
 $T_c = 300 \text{ ns} \rightarrow 3\,333 \text{ kHz}$

If the pipeline fill overhead is neglected (i.e. no pipeline stalls and flushes are considered) then ideal $IPC = 1$.
 $IPS = 1 \cdot 3\,333\text{e}3 = 3\,333\,000$ instructions per second

- Introduction of the 5-stage pipeline increases performance (throughput) $3\,333\,000 / 980\,000 = 3.4$ times! (considering $IPC=1$)

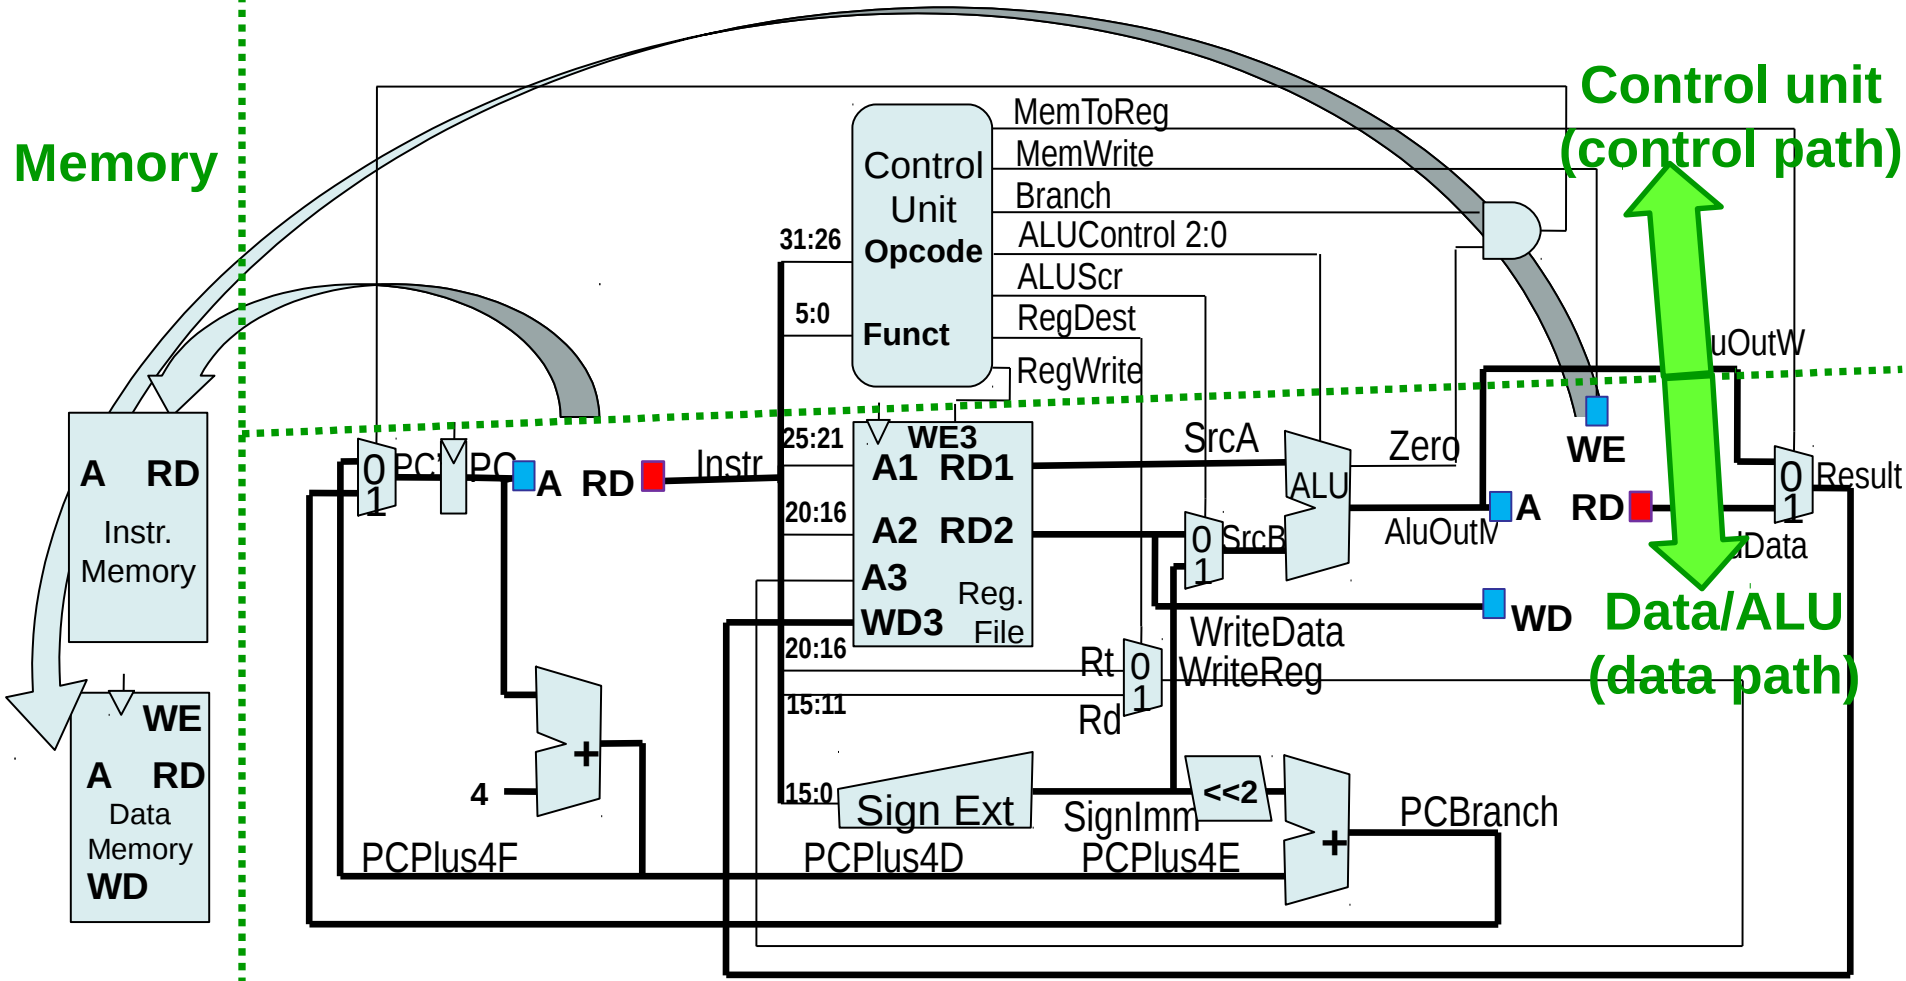
What is result of the design?

Return back to non-pipelined CPU version

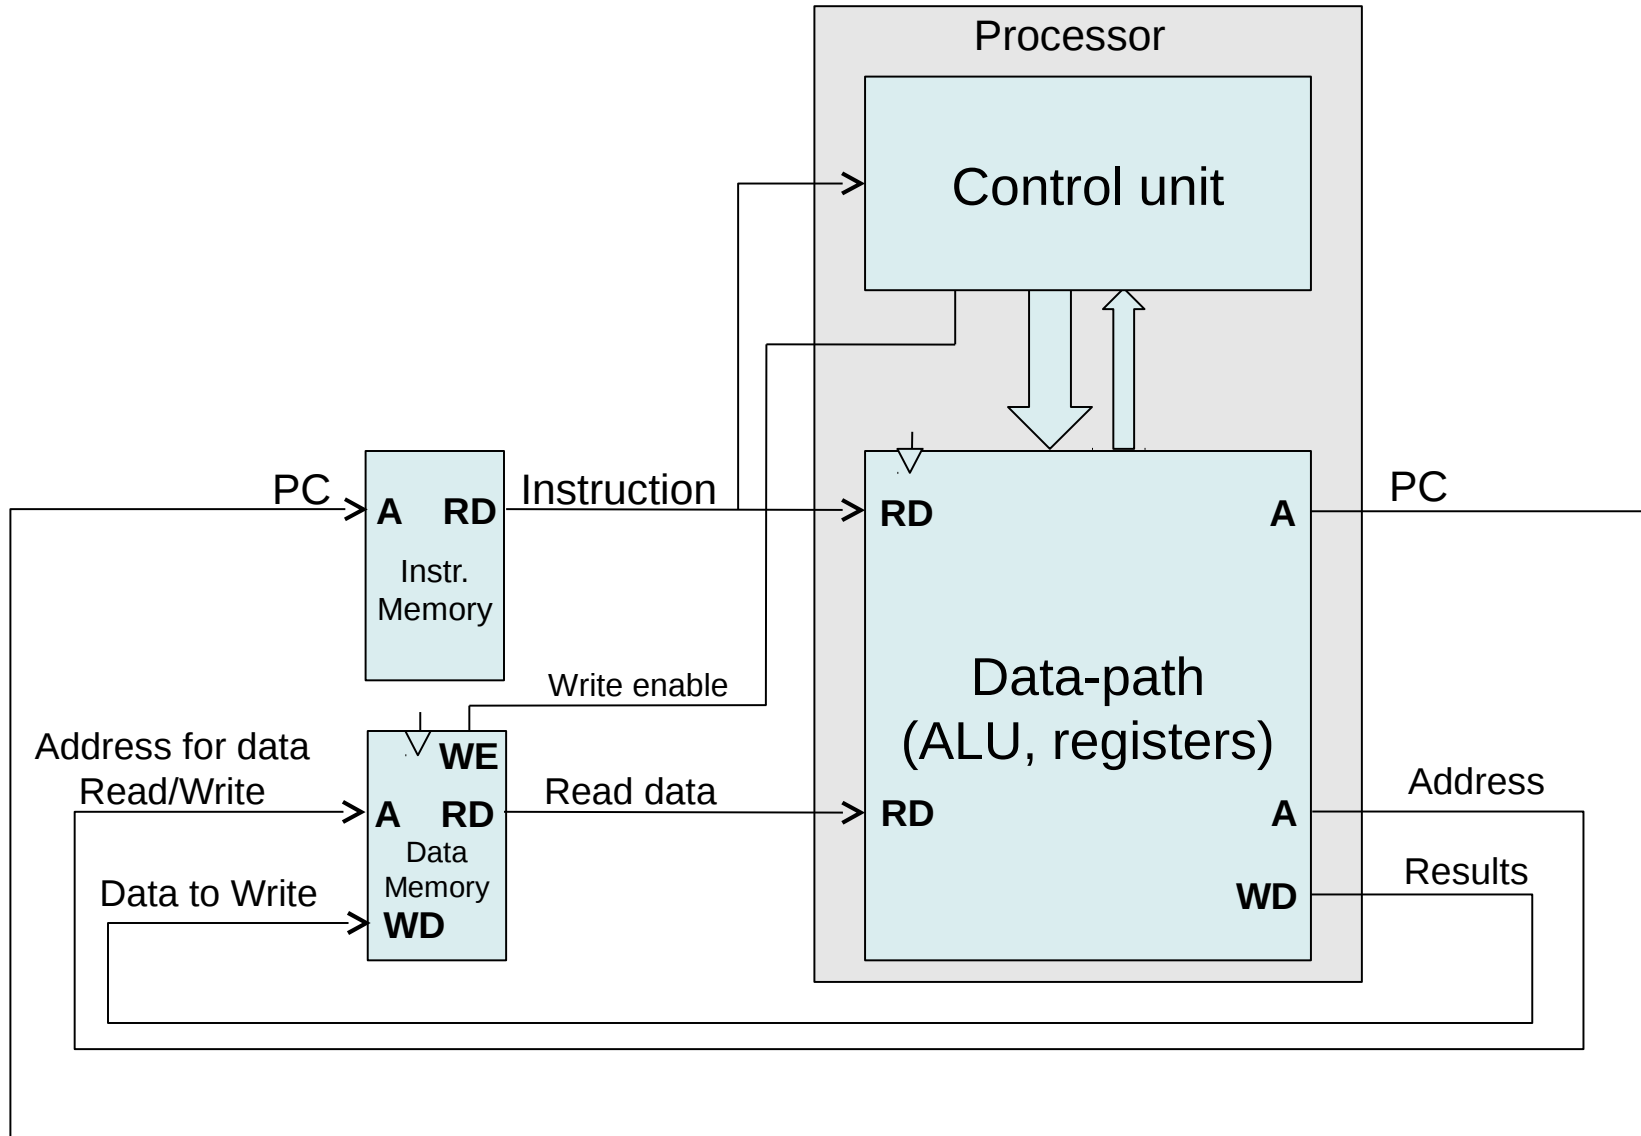


What is result of the design?

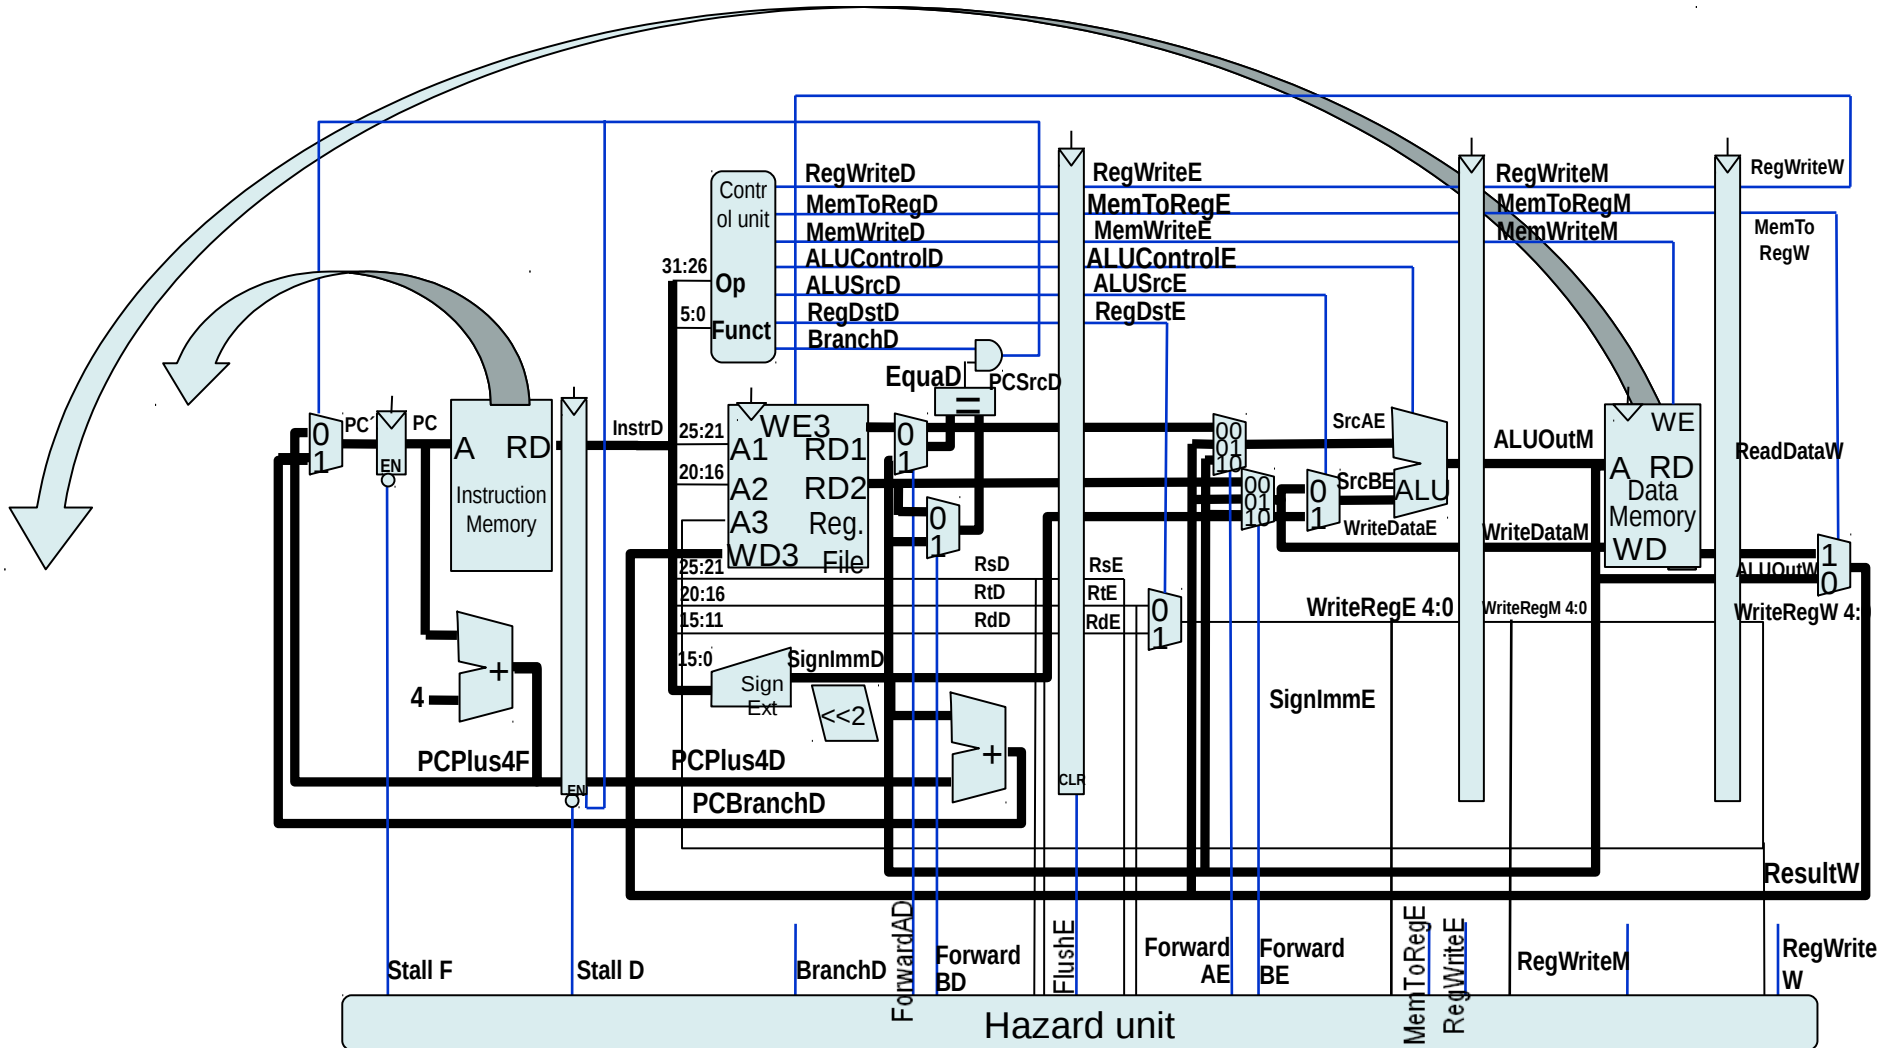
Return back to non-pipelined CPU version



What is result of the design?




CPU design result – pipelined version




Pipelined CPU – timing

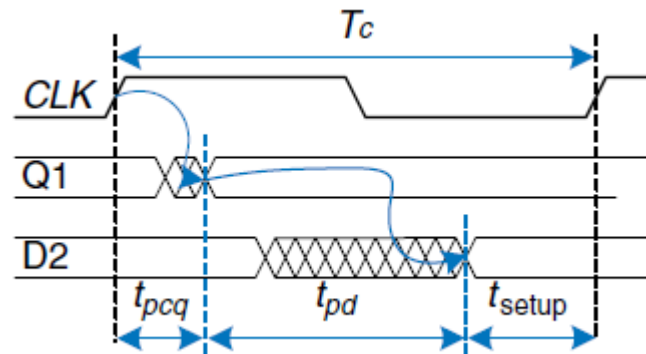
- The timing/AC characteristics of synchronous sequential circuit :



- Signal integrity constrain for the setup time before the clock:





$$T_c \geq t_{pcq} + t_{pd} + t_{setup}$$



t_{pd} – combinational logic propagation delay

Pipelined processor – timing

- Constraint for the setup time (consider the clock distribution jitter):




$$T_c \geq t_{pcq} + t_{pd} + t_{setup} + t_{skew}$$

Clock distribution jitter is limiting factor,
if it reaches or exceeds value of t_{pd}
(too deep pipeline / too many stages...)

Pipeline stages balancing

Linear pipelining:




(applies to tree based adder, multiplier, (unrolled) iterative divider..)


- **Balancing:** the goal is to divide the processing into N stages in such way, that stage propagation delays are roughly the same...
- The number of stages reflects preference of performance (throughput) versus latency.

Superpipeline and beyond

- Not well balanced 5-stage pipeline:



- Deeper pipeline is result of decomposing stages into more stages



- It allows CPU to work at higher frequencies but introduces many problems as well..
- Complex forwarding, more pipeline stalls, hazards need to be solved by complex logic

Typical pipeline depths in today's CPUs

P5 (Pentium) : **5**
P6 (Pentium 3): **10**
P6 (Pentium Pro): **14**
NetBurst (Willamette, 180 nm) - Celeron, Pentium 4: **20**
NetBurst (Northwood, 130 nm) - Celeron, Pentium 4, Pentium 4 HT: **20**
NetBurst (Prescott, 90 nm) - Celeron D, Pentium 4, Pentium 4 HT, Pentium 4 ExEd: **31**
NetBurst (Cedar Mill, 65 nm): **31**
NetBurst (Presler 65 nm) - Pentium D: **31**
Core : **14**
Bonnell: **16**

K7 Architecture - Athlon : **10-15**
K8 - Athlon 64, Sempron, Opteron, Turion 64: **12-17**

ARM 8-9: **5**
ARM 11: **8**
Cortex A7: **8-10**
Cortex A8: **13**
Cortex A15: **15-25**

- The Optimum Pipeline Depth for a Microprocessor:
<http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.93.4333&rep=rep1&type=pdf>

Branch stall discussion and delay slots

- The instruction memory read and fetch is expensive and result of condition evaluation in branch instructions (even worse target in indirect branch instructions) has to be evaluated before next fetch and execute. The **stall** state is waste of cycles. Options to use that cycle(s) are:
 - Start fetch and execution of instruction(s) following branch and flush/discard results if it is resolved that it should not be executed
 - Extend above by adding condition results/branch predictor (**taken/not-taken**) and branch target cache (BTB)
 - Execute one or more instructions after branch unconditionally in (so called) **delay slot**
- Delay slots unconditional execution is common for many DSP (digital signal processor) and some RISC architectures (MIPS, SPARC)

