January 30, 2025

Announcements

- Talk topic
- HW1

Goals

- assembly
- march
- memory

Review

- Dennard scaling

$\frac{1}{2} \in \mathbb{R}$

power dissipation / area ~ constant

- ⇝ parallel computers
- old-timey architecture
- instructions, assembly

Demos

Demo: intro/Assembly Reading Comprehension

Demo: Source-to-assembly mapping
Code to try:

```
int main()
{
  int y = 0, i;
  for (i = 0; y < 100; ++i)
    y += i*i;
  return y;
}
```

Also try https://godbolt.org for direct source-to-assembly mapping

Outline

Modern Processors?

All of this can be built in about 4000 transistors.
(e.g. MOS 6502 in Apple II, Commodore 64, Atari 2600)

So what exactly are Intel/ARM/AMD/Nvidia doing with the other billions of transistors?

Execution in a Simple Processor

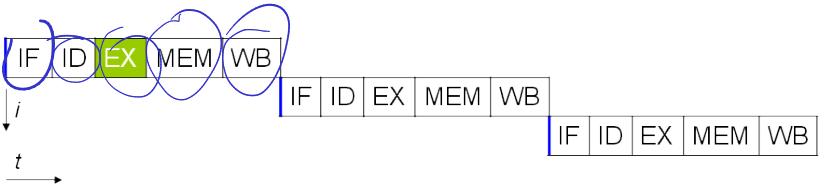

- [IF] Instruction fetch
- [ID] Instruction Decode
- [EX] Execution
- [MEM] Memory Read/Write
- [WB] Result Writeback

[Wikipedia ©]

Solution: Pipelining

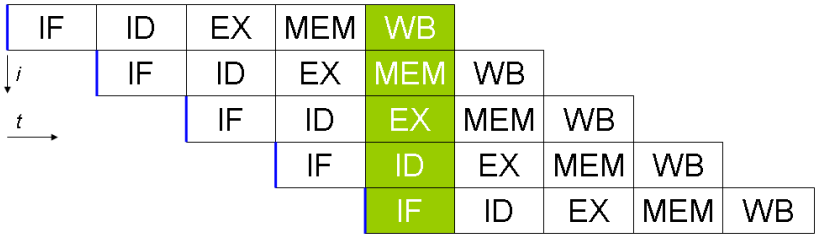

[Wikipedia ©]

MIPS Pipeline: 110,000 transistors

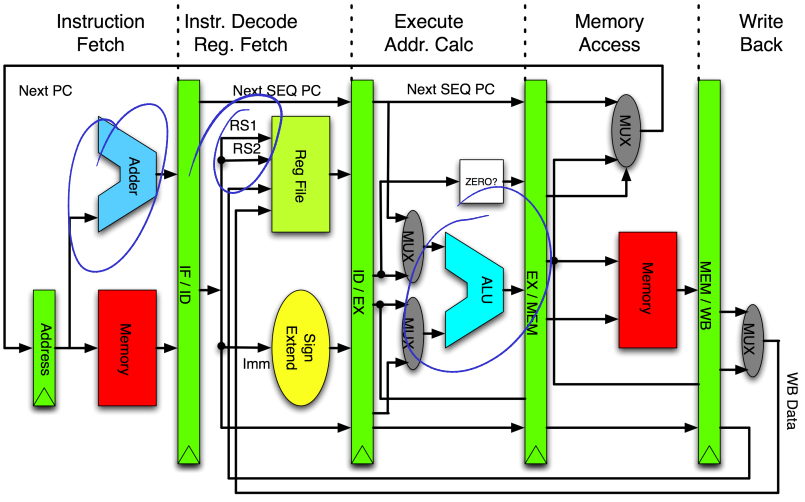

[Wikipedia ©]

Hazards and Bubbles

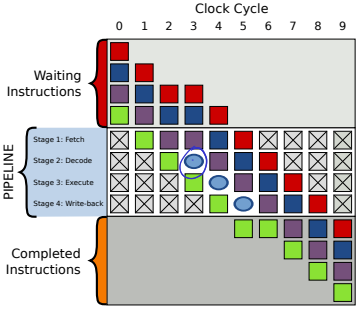

Q: Types of Pipeline Hazards? (aka: what can go wrong?)

- Structural
- Data
- Control

[Wikipedia ©]

Demo

v1: 2.6
v2: 1.4

control flow: 0.06
v2 unrolled: 1.3

Demo: intro/Pipeline Performance Mysteries

A Glimpse of a More Modern Processor

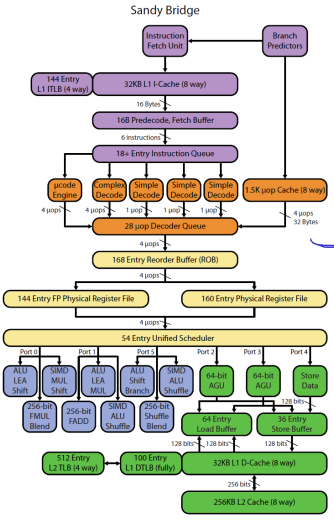

[David Kanter / Realworldtech.com]

A Glimpse of a More Modern Processor: Frontend

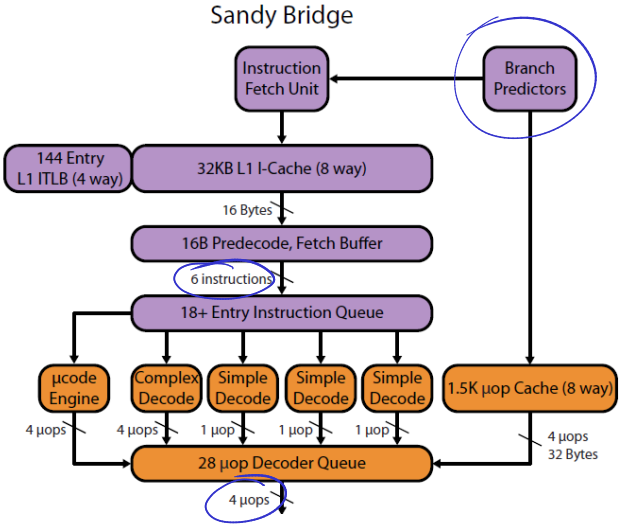

[David Kanter / Realworldtech.com]

A Glimpse of a More Modern Processor: Backend

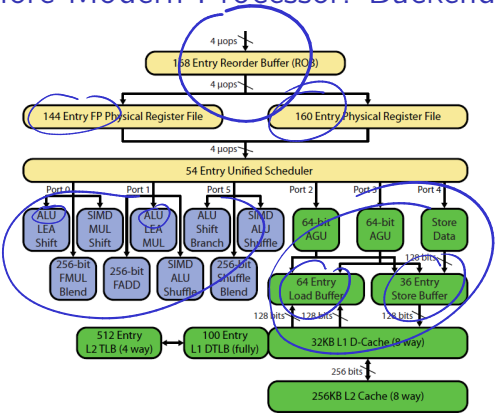

- New concept: Instruction-level parallelism (“ILP”, “superscalar”)
- Where does the IPC number from earlier come from?

[David Kanter / Realworldtech.com]

A Glimpse of a More Modern Processor: Golden Cove

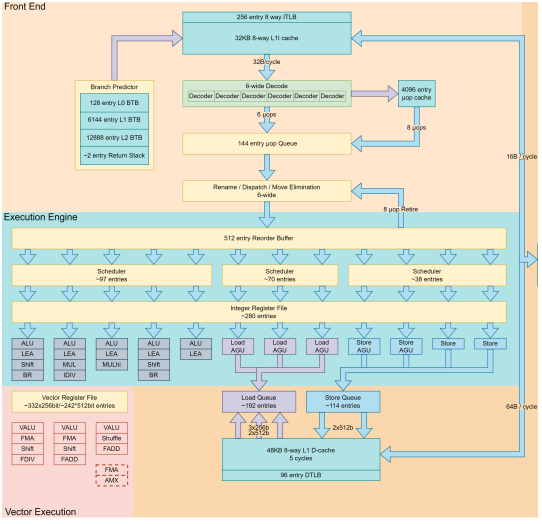

[Wikipedia ©]

Demo

Demo: intro/More Pipeline Mysteries

SMT/“Hyperthreading”

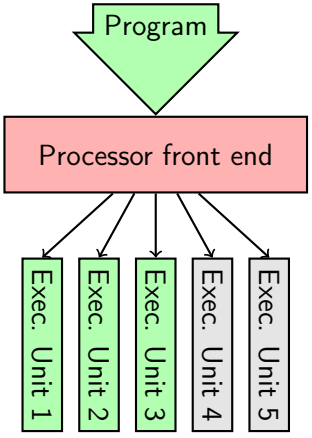

Q: Potential issues?