

## *ARMv8 - 64-bit Architecture*

First look at the arm64  
architecture and assembly  
language


### *Assembly Language and Architecture*

- Assembly-language programs consist of instructions that manipulate processor hardware elements
  - Goal of manipulations:  
implement high-level-language program behaviors
- Processor designs determine what instructions are available
- Understanding assembly language requires understanding the processor architecture

## Computer Organization - Programmer's View

- Fundamental components:

- CPU
- Memory
- I/O ("everything else")



- Possible connection schemes:

- Single Controller (*above right*)
- Direct Memory Access (DMA) (*below*)



## Memory Organization

- The memory subsystem is a 1-dimensional array of storage locations

- Each location has a unique address (or "array index")
- Each location is made up of  $m$  bits ( $m \geq 1$ )
- Location's contents is any one of  $2^m$  bit patterns

- Memory can hold:

- Programs: bit patterns represent instructions
- Data: bit patterns represent numbers, characters, etc.

- Volatile contents: R/W memory, or RAM


- DRAM is common for main memory

- Nonvolatile (unchangeable): Read-Only, or ROM

- PROM, EEPROM, Flash - erasable/reprogrammable ROM


## How a program occupies memory

- When a program runs it occupies a block of RAM




Your program's memory block

- The block is divided into instructions, data, and stack segments, each in part of the memory



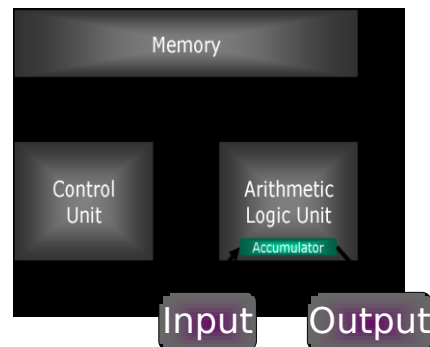
## Dividing up the program memory

- Each function, including main(), has its own unchanging portion of the instruction memory
  - Any global variables occupy memory similar to instruction memory.
- Function-local variables exist in the stack only while the function is executing



## the Central Processing Unit

- a.k.a. CPU or Processor
  - (or microprocessor)
- Reads program instructions
- Accesses program data
- Performs numerical and logical operations
- von Neumann model
  - basis for most modern computers



## John von Neumann's IAS Processor



- Developed at the Institute for Advanced Studies, Princeton, NJ
- Single, shared memory for program instructions and data

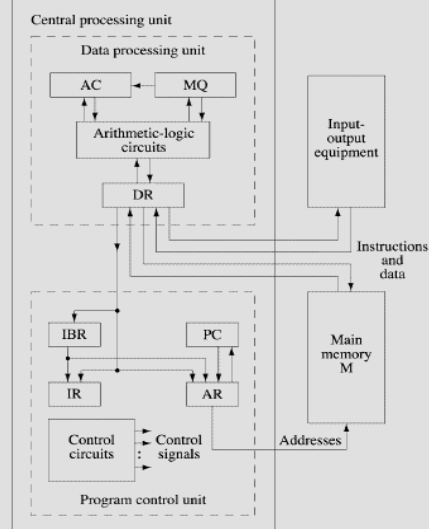
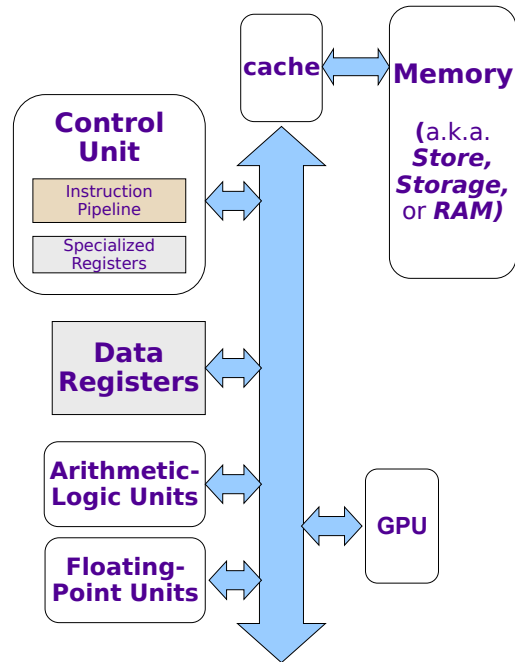


Figure 1

Structure of von Neumann's IAS processor. Reprinted with the permission of the McGraw-Hill Book Company from page 23 of the second edition of *Computer Architecture and Organization* by John P. Hayes, published by the McGraw-Hill Book Company, copyright 1988.



## Modern CPU elements

- Control Unit
  - Instruction pipeline
  - Specialized registers
    - » contain execution state
    - » control program flow
- General-Purpose Registers
  - Hold data values for ALU, FPU
- ALU, FPU, GPU
  - Do the computations
- Cache
  - Improves memory accesses
  - Generally not visible to programmer
- Internal buses
  - Connect things

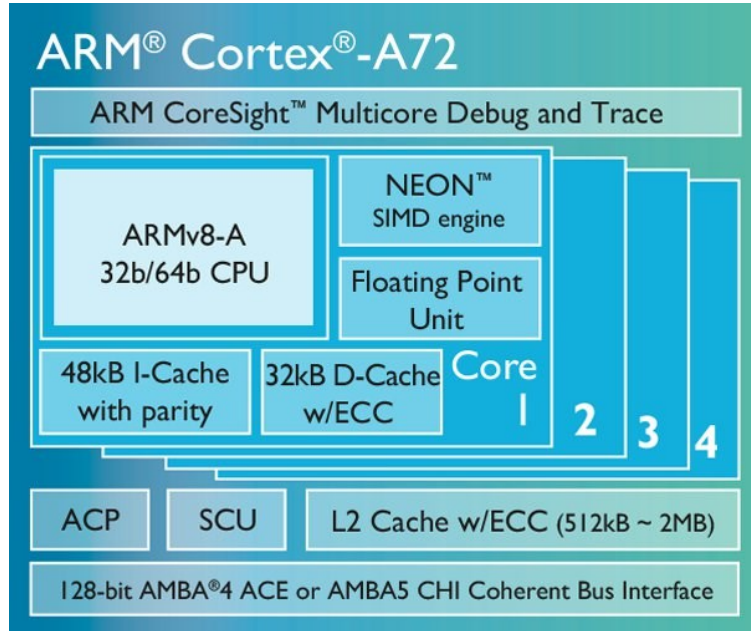


## Multi-core Processors

- Each CPU is a "core"
- One chip holds 2, 4, 8, ... cores
- Cores share memory
- Cores may have local L1 cache, may share L2 and L3 cache
- A CPU with a GPU is a form of *heterogeneous* multicore processor


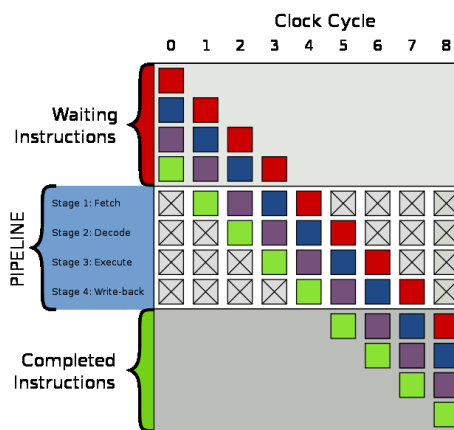


## Raspberry Pi 4B Processor Block Diagram



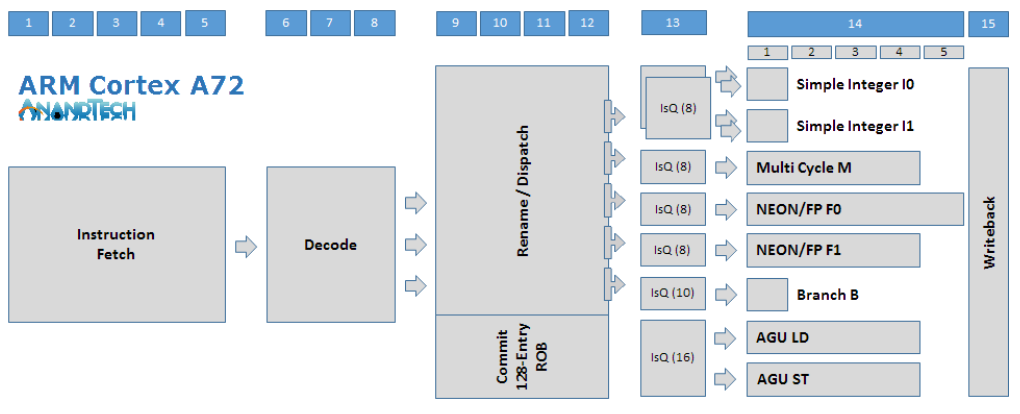
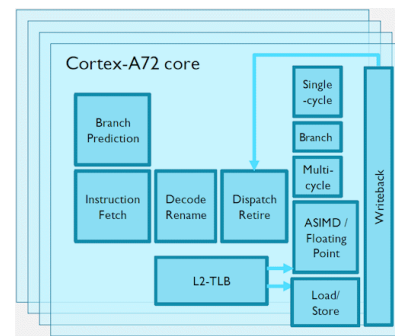
## Pipelined Designs

- Separate circuits perform each stage of the fetch-execute cycle
- In a *pipelined* design the separate circuits work simultaneously, on sequential instructions



## RPi Instruction Pipeline

- *Right*: simplified
- *Below*: showing clock cycles



## The ARM Architectural Models

- Aarch32 (armhf, armel)
  - Original models, optional floating-point
  - Provides 32-bit registers and instructions
- T32 (Thumb2)
  - 16-bit instructions expand to 32 bits on use
  - Compatible with Aarch32
- Aarch64 (arm64)
  - 64-bit registers, some new instructions
  - Floating-point, SIMD hardware standard
  - Context switch between Aarch64 , Aarch32

## ARMv8 General-Purpose Register Set

X0	W0	r0
X1	W1	r1
X2	W2	r2
X3	W3	r3
X4	W4	r4
X5	W5	r5
X6	W6	r6
X7	W7	r7
X8	W8	r8
X9	W9	r9
X10	W10	r10
X11	W11	r11
X12	W12	r12
X13	W13	r13
X14	W14	r14
X15	W15	r15
X16	W16	r16
X17	W17	r17
X18	W18	r18
X19	W19	r19
X20	W20	r20
X21	W21	r21
X22	W22	r22
X23	W23	r23
X24	W24	r24
X25	W25	r25
X26	W26	r26
X27	W27	r27
X28	W28	r28
X29	W29	r29
X30	W30	r30
Xzr	Wzr	rZR

← 32 bits →  
← 64 bits →

- r0 - r7
  - Arguments, return values
- r8
  - Syscall number / Returned-struct pointer
- r9 - r15
  - Temporary
- r16 - r18
  - Intra-procedure/platform
  - Avoid using these
- r19 - r28
  - Callee preserves these
- r29, r30
  - Frame, link registers:
  - Not general-purpose


Register names:

- X-names
  - 64-bit versions
- W-names
  - Lower 32 bits
- r-names
  - generic references

Annotations:

- r8: Syscall number / Return pointer
- r29: Frame Pointer
- r30: Link Register
- rZR: Zero Register / Stack Pointer

## Aarch64 Special Registers



- Zero Register
  - reads always return 0
  - writes are always discarded
- PC
  - no explicit access
- SP, SPSR, EL registers:
  - Separate set for each Exception level EL0 - EL 3
- SPSR
  - includes NZCV condition-code flags
- EL
  - holds exception return address
- Some instructions treat ZR or SP as "r31"



## Aarch64 VFP/SIMD (NEON) Register Set

- V0 - V7
  - Arguments, return values
- V8 - V15
  - Callee-saved  
(lower 64 bits only)
- V16 - V31
  - Spare temporary registers
    - » Not in Aarch32
- D0 - D31
  - Double-precision registers
    - » D0 - D15 also in Aarch32
- S0 - S31
  - Single-precision registers
    - » S0 - S15 also in Aarch32

V0	D0	S0
V1	D1	S1
V2	D2	S2
V3	D3	S3
V4	D4	S4
V5	D5	S5
V6	D6	S6
V7	D7	S7
V8	D8	S8
V9	D9	S9
V10	D10	S10
V11	D11	S11
V12	D12	S12
V13	D13	S13
V14	D14	S14
V15	D15	S15
V16	D16	S16
V17	D17	S17
V18	D18	S18
V19	D19	S19
V20	D20	S20
V21	D21	S21
V22	D22	S22
V23	D23	S23
V24	D24	S24
V25	D25	S25
V26	D26	S26
V27	D27	S27
V28	D28	S28
V29	D29	S29
V30	D30	S30
V31	D31	S31

Diagram showing bit widths: 128 bits for V0-V15, 64 bits for V16-V31, and 32 bits for S0-S31.

## Exception Levels

- ARMv8 security model
- "Userland" operates at Level 0 (L0)
  - Least privilege
  - Library Functions
- Operating system operates at L1
  - System calls
- Virtualization runs at L2
- Hardware has highest privilege level, L3

