

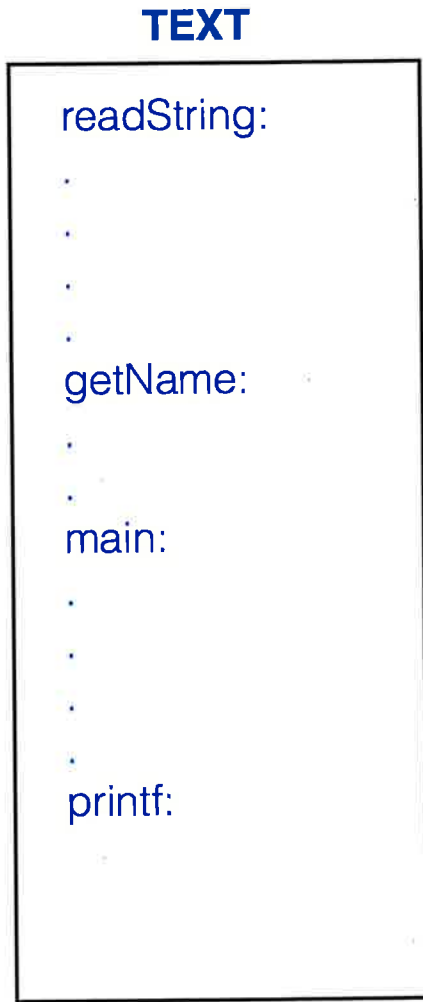



Precept 23
Week 13, Mon/Tue



C6.2.5 ADD (shifted register)

Add (shifted register) adds a register value and an optionally-shifted register value, and writes the result to the destination register.



32-bit variant

Applies when sf == 0.

ADD <Wd>, <Wn>, <Wm>{, <shift> #<amount>}

64-bit variant

Applies when sf == 1.

ADD <Xd>, <Xn>, <Xm>{, <shift> #<amount>}

Decode for all variants of this encoding

```
integer d = UInt(Rd);
integer n = UInt(Rn);
integer m = UInt(Rm);
integer datasize = if sf == '1' then 64 else 32;

if shift == '11' then ReservedValue();
if sf == '0' && imm6<5> == '1' then ReservedValue();

ShiftType shift_type = DecodeShift(shift);
integer shift_amount = UInt(imm6);
```

Assembler symbols

- <wd> Is the 32-bit name of the general-purpose destination register, encoded in the "Rd" field.
- <wn> Is the 32-bit name of the first general-purpose source register, encoded in the "Rn" field.
- <wm> Is the 32-bit name of the second general-purpose source register, encoded in the "Rm" field.
- <xd> Is the 64-bit name of the general-purpose destination register, encoded in the "Rd" field.
- <xn> Is the 64-bit name of the first general-purpose source register, encoded in the "Rn" field.
- <xm> Is the 64-bit name of the second general-purpose source register, encoded in the "Rm" field.
- <shift> Is the optional shift type to be applied to the second source operand, defaulting to LSL and encoded in the "shift" field. It can have the following values:
 - LSL when shift = 00
 - LSR when shift = 01
 - ASR when shift = 10
 - The encoding shift = 11 is reserved.
- <amount> For the 32-bit variant: is the shift amount, in the range 0 to 31, defaulting to 0 and encoded in the "imm6" field.
For the 64-bit variant: is the shift amount, in the range 0 to 63, defaulting to 0 and encoded in the "imm6" field.

Operation

```
bits(datasize) result;  
bits(datasize) operand1 = X[n];  
bits(datasize) operand2 = ShiftReg(m, shift_type, shift_amount);  
  
(result, -) = AddWithCarry(operand1, operand2, '0');  
  
X[d] = result;
```

C6.2.23 B.cond

Branch conditionally to a label at a PC-relative offset, with a hint that this is not a subroutine call or return.

**19-bit signed PC-relative branch offset variant**

B.<cond> <label>

Decode for this encoding

bits(64) offset = SignExtend(imm19:'00', 64);

Assembler symbols

- <cond> Is one of the standard conditions, encoded in the "cond" field in the standard way.
- <label> Is the program label to be conditionally branched to. Its offset from the address of this instruction, in the range +/-1MB, is encoded as "imm19" times 4.

Operation

```
if ConditionHolds(cond) then
    BranchTo(PC[] + offset, BranchType_JMP);
```

SP	See <i>Register names</i> on page C1-145.
Wn	See <i>Register names</i> on page C1-145.
WSP	See <i>Register names</i> on page C1-145.
WZR	See <i>Register names</i> on page C1-145.
Xn	See <i>Register names</i> on page C1-145.
XZR	See <i>Register names</i> on page C1-145.

C1.2.3 Instruction Mnemonics

The A64 assembly language overloads instruction mnemonics and distinguishes between the different forms of an instruction based on the operand types. For example, the following ADD instructions all have different opcodes. However, the programmer must only remember one mnemonic, as the assembler automatically chooses the correct opcode based on the operands. The disassembler follows the same procedure in reverse.

Example C1-1 ADD instructions with different opcodes

```

ADD W0, W1, W2           // add 32-bit register
ADD X0, X1, X2           // add 64-bit register
ADD X0, X1, W2, SXTW     // add 64-bit extended register
ADD X0, X1, #42          // add 64-bit immediate
  
```

C1.2.4 Condition code

The A64 ISA has some instructions that set Condition flags or test Condition codes or both. For information about instructions that set the Condition flags or use the condition mnemonics, see *Condition flags and related instructions* on page C6-525.

Table C1-1 shows the available Condition codes.

Table C1-1 Condition codes

cond	Mnemonic	Meaning (integer)	Meaning (floating-point) ^a	Condition flags
0000	EQ	Equal	Equal	Z = 1
0001	NE	Not equal	Not equal or unordered	Z = 0
0010	CS or HS	Carry set	Greater than, equal, or unordered	C = 1
0011	CC or LO	Carry clear	Less than	C = 0
0100	MI	Minus, negative	Less than	N = 1
0101	PL	Plus, positive or zero	Greater than, equal, or unordered	N = 0
0110	VS	Overflow	Unordered	V = 1
0111	VC	No overflow	Ordered	V = 0
1000	HI	Unsigned higher	Greater than, or unordered	C = 1 && Z = 0
1001	LS	Unsigned lower or same	Less than or equal	!(C = 1 && Z = 0)
1010	GE	Signed greater than or equal	Greater than or equal	N = V
1011	LT	Signed less than	Less than, or unordered	N! = V



Table C1-1 Condition codes (continued)

cond	Mnemonic	Meaning (integer)	Meaning (floating-point) ^a	Condition flags
1100	GT	Signed greater than	Greater than	$Z = 0 \ \&\& \ N = V$
1101	LE	Signed less than or equal	Less than, equal, or unordered	$!(Z = 0 \ \&\& \ N = V)$
1110	AL	Always	Always	Any
1111	NV ^b	Always	Always	Any

a. Unordered means at least one NaN operand.

b. The Condition code *NV* exists only to provide a valid disassembly of the `0b1111` encoding, otherwise its behavior is identical to *AL*.

C1.2.5 Register names

This section describes the AArch64 registers. It contains the following subsections:

- *General-purpose register file and zero register and stack pointer.*
- *SIMD and floating-point register file on page C1-146.*
- *SIMD and floating-point scalar register names on page C1-146.*
- *SIMD vector register names on page C1-147.*
- *SIMD vector element names on page C1-147.*

General-purpose register file and zero register and stack pointer

The 31 general-purpose registers in the general-purpose register file are named R0-R30 and encoded in the instruction register fields with values 0-30. In a general-purpose register field the value 31 represents either the current stack pointer or the zero register, depending on the instruction and the operand position.

When the registers are used in a specific instruction variant, they must be qualified to indicate the operand data size, 32 bits or 64 bits, and the data size of the instruction.

When the data size is 32 bits, the lower 32 bits of the register are used and the upper 32 bits are ignored on a read and cleared to zero on a write.

Table C1-2 shows the qualified names for registers, where *n* is a register number 0-30.

Table C1-2 Naming of general-purpose registers, the zero register, and the stack pointer

Name	Size	Encoding	Description
W _n	32 bits	0-30	General-purpose register 0-30
X _n	64 bits	0-30	General-purpose register 0-30
WZR	32 bits	31	Zero register
XZR	64 bits	31	Zero register
WSP	32 bits	31	Current stack pointer
SP	64 bits	31	Current stack pointer

This list gives more information about the instruction arguments shown in Table C1-2:


- The names X_n and W_n both refer to the same general-purpose register, R_n.
- There is no register named W31 or X31.

C6.2.167 MOV (wide immediate)

Move (wide immediate) moves a 16-bit immediate value to a register.

This instruction is an alias of the **MOVZ** instruction. This means that:

- The encodings in this description are named to match the encodings of **MOVZ**.
- The description of **MOVZ** gives the operational pseudocode for this instruction.



32-bit variant

Applies when sf == 0.

MOV <Wd>, #<imm>

is equivalent to

MOVZ <Wd>, #<imm16>, LSL #<shift>

and is the preferred disassembly when ! (IsZero(imm16) && hw != '00').

64-bit variant

Applies when sf == 1.

MOV <Xd>, #<imm>

is equivalent to

MOVZ <Xd>, #<imm16>, LSL #<shift>

and is the preferred disassembly when ! (IsZero(imm16) && hw != '00').

Assembler symbols

- <Wd> Is the 32-bit name of the general-purpose destination register, encoded in the "Rd" field.
- <Xd> Is the 64-bit name of the general-purpose destination register, encoded in the "Rd" field.
- <imm> For the 32-bit variant: is a 32-bit immediate which can be encoded in "imm16:hw".
For the 64-bit variant: is a 64-bit immediate which can be encoded in "imm16:hw".
- <shift> For the 32-bit variant: is the amount by which to shift the immediate left, either 0 (the default) or 16, encoded in the "hw" field as <shift>/16.
For the 64-bit variant: is the amount by which to shift the immediate left, either 0 (the default), 16, 32 or 48, encoded in the "hw" field as <shift>/16.

Operation

The description of **MOVZ** gives the operational pseudocode for this instruction.

C6.2.9 ADR

Form PC-relative address adds an immediate value to the PC value to form a PC-relative address, and writes the result to the destination register.



Literal variant

ADR <Xd>, <label>

Decode for this encoding

```
integer d = UInt(Rd);
bits(64) imm;

imm = SignExtend(immhi:immlo, 64);
```

Assembler symbols

- <Xd> Is the 64-bit name of the general-purpose destination register, encoded in the "Rd" field.
- <label> Is the program label whose address is to be calculated. Its offset from the address of this instruction, in the range +/-1MB, is encoded in "immhi:immlo".

Operation

```
bits(64) base = PC[];
X[d] = base + imm;
```


C6.2.244 STRB (immediate)

Store Register Byte (immediate) stores the least significant byte of a 32-bit register to memory. The address that is used for the store is calculated from a base register and an immediate offset. For information about memory accesses, see *Load/Store addressing modes* on page C1-149.

Post-index



Post-index variant

STRB <Wt>, [<Xn|SP>], #<imm>

Decode for this encoding

```
boolean wback = TRUE;
boolean postindex = TRUE;
bits(64) offset = SignExtend(imm9, 64);
```

Pre-index




Pre-index variant

STRB <Wt>, [<Xn|SP>], #<imm>!

Decode for this encoding

```
boolean wback = TRUE;
boolean postindex = FALSE;
bits(64) offset = SignExtend(imm9, 64);
```

Unsigned offset



Unsigned offset variant

STRB <Wt>, [<Xn|SP>], #<pimm>}

Decode for this encoding

```
boolean wback = FALSE;
boolean postindex = FALSE;
bits(64) offset = LSL(ZeroExtend(imm12, 64), 0);
```

Notes for all encodings

For information about the CONSTRAINED UNPREDICTABLE behavior of this instruction, see [Appendix K1 Architectural Constraints on UNPREDICTABLE behaviors](#), and particularly *STRB (immediate)* on page K1-6425.

Assembler symbols

<Rt>	Is the 32-bit name of the general-purpose register to be transferred, encoded in the "Rt" field.
<Rn SP>	Is the 64-bit name of the general-purpose base register or stack pointer, encoded in the "Rn" field.
<imm>	Is the signed immediate byte offset, in the range -256 to 255, encoded in the "imm9" field.
<pinm>	Is the optional positive immediate byte offset, in the range 0 to 4095, defaulting to 0 and encoded in the "imm12" field.

Shared decode for all encodings

```
integer n = UInt(Rn);
integer t = UInt(Rt);
```

Operation for all encodings

```
bits(64) address;
bits(8) data;
boolean rt_unknown = FALSE;

if wback && n == t && n != 31 then
  c = ConstrainUnpredictable();
  assert c IN {Constraint_NONE, Constraint_UNKNOWN, Constraint_UNDEF, Constraint_NOP};
  case c of
    when Constraint_NONE   rt_unknown = FALSE; // value stored is original value
    when Constraint_UNKNOWN rt_unknown = TRUE;  // value stored is UNKNOWN
    when Constraint_UNDEF  UnallocatedEncoding();
    when Constraint_NOP    EndOfInstruction();

if n == 31 then
  CheckSPAlignment();
  address = SP[];
else
  address = X[n];

if !postindex then
  address = address + offset;

if rt_unknown then
  data = bits(8) UNKNOWN;
else
  data = X[t];
Mem[address, 1, AccType_NORMAL] = data;

if wback then
  if postindex then
    address = address + offset;
  if n == 31 then
    SP[] = address;
  else
    X[n] = address;
```


Precept Activity Instructions:

- Translate **adr x1, label2** to machine language
(Assume the offset [address of label2] – [address of adr instruction]
is ...00000000 00000000 00110000 11110000 (bin):
- immlo:
- immhi:
- Rd:
- What is the machine instruction in hexadecimal representation?

Precept Activity Solution:

- **Translate adr x1, label2 to machine language**
 (Assume the offset [address of label2] – [address of adr instruction] is
 ...00000000 00000000 00110000 11110000 (bin):

- **33222222222211111111110000000000**
- **10987654321098765432109876543210**
- 0 (op)
- 00 (immlo)
- 10000 (op)
- 0000000110000111100 (immhi)
- 00001 (Rd)

- So the instruction in binary is this:

- 00010000000000011000011110000001

- convert to hexadecimal...

- After clustering into groups of 4 bits:

- 0001 0000 0000 0001 1000 0111 1000 0001

- After converting to hexadecimal:

- 10018781 (hex)