

The AArch64 processor (aka arm64), part 1: Introduction

 devblogs.microsoft.com/oldnewthing/20220726-00

July 26, 2022



Raymond Chen

The 64-bit version of the ARM architecture is formally known as AArch64. It is the 64-bit version of classic 32-bit ARM, which has been retroactively renamed AArch32.

Even though the architecture formally goes by the name AArch64, many people (including Windows) call it arm64. Even more confusing, the instruction set is called A64. (The 32-bit ARM instruction sets have also been retroactively renamed: Classic ARM is now called A32, and Thumb-2 is now called T32.)

AArch64 differs from AArch32 so much that I'm going to cover it fresh rather than treating it as an extension of AArch32. That said, I will nevertheless call out notable points of difference from AArch32.

No more Thumb mode

AArch64 is an extension of the classic ARM instruction set, not an extension of Thumb-2. So we're back to fixed-size 32-bit instructions (aligned on 4-byte boundaries). No more gymnastics with low registers and high registers, or using non-intuitive instructions to avoid a 32-bit encoding, or remembering to set the bottom bit on code addresses to avoid accidentally switching into classic mode.

A note for those familiar with the classic ARM instruction set: One thing that did not get carried forward was arbitrary predication. The answers to this StackOverflow question dig into the reasons why predication was removed. Short version: Predication is rarely used, it consumes a lot of opcode space, it doesn't interact well with out-of-order execution, and branch prediction is almost as good.

Data sizes

The architectural terms for data sizes are the same as AArch32.

Term	Size
------	------

byte	8 bits
halfword	16 bits
word	32 bits
doubleword	64 bits

The processor supports both big-endian and little-endian operation. Windows uses it exclusively in little-endian mode. AArch64 lost the AArch32 **SETEND** instruction for switching endianness from user mode. Not that Windows supported it anyway.

Registers

Everything has doubled. The general-purpose registers are now 64 bits wide instead of 32. And the number of such registers has doubled from 16 to 32 okay just 31. The encoding that would correspond to register 31 has been reused for other purposes. So not quite doubled.

Register	Preserved?	Notes
<i>x0</i>	No	Parameter 1, return value
<i>x1</i>	No	Parameter 2
<i>x2</i>	No	Parameter 3
<i>x3</i>	No	Parameter 4
<i>x4</i>	No	Parameter 5
<i>x5</i>	No	Parameter 6
<i>x6</i>	No	Parameter 7
<i>x7</i>	No	Parameter 8
<i>x8</i>	No	
<i>x9</i>	No	
<i>x10</i>	No	
<i>x11</i>	No	
<i>x12</i>	No	
<i>x13</i>	No	

<i>x14</i>	No	
<i>x15</i>	No	
<i>x16 (xip0)</i>	Volatile	Intra-procedure call scratch register
<i>x17 (xip1)</i>	Volatile	Intra-procedure call scratch register
<i>x18 (xpr)</i>	read-only	TEB
<i>x19</i>	Yes	
<i>x20</i>	Yes	
<i>x21</i>	Yes	
<i>x22</i>	Yes	
<i>x23</i>	Yes	
<i>x24</i>	Yes	
<i>x25</i>	Yes	
<i>x26</i>	Yes	
<i>x27</i>	Yes	
<i>x28</i>	Yes	
<i>x29 (fp)</i>	Yes	frame pointer
<i>x30 (lr)</i>	No	link register
register “31” usually represents <i>sp</i> or <i>zr</i> , depending on instruction		

The link register is architectural; the rest are convention.

You can refer to the least significant 32 bits of each 64-bit register by changing the leading *x* to a *w*, so we have *w0* through *w30*. If an instruction targets a *w* register, the result is zero-extended to fill the *x* register.¹

Particularly notable is that the stack pointer *sp* and program counter *pc* are no longer general-purpose registers, like they were in AArch32. The registers still exist, but they are treated as special registers rather than being encoded in the same way as the other general-purpose registers.

In AArch64, the *pc* special register reads as the address of the instruction being executed, rather than being four bytes ahead, as it was in AArch32. The extra +4 in AArch32 was an artifact of the internal pipelining of the original ARM and became a backward compatibility constraint even as the pipeline depth changed.

Windows requires that the stack remain 16-byte aligned, and it enables hardware enforcement of this requirement. The 32-bit subregister of *sp* is called *wsp*, although it is of no practical use. (The 64-bit register is still called *sp*, not *xsp*. Go figure.)

There is a 16-byte red zone below the stack pointer, but it's reserved for code analysis. Intrusive profilers inject assembly language fragments into compiled code to update profiling information, and they need some space to store two registers so they can free up some registers to do their profiling work.

The *xip0* and *xip1* registers are volatile because they are used to assist with branch instructions that try to branch to an address that is out of range. We'll see later that these registers are also used by function prologues and epilogues.

There is a new *xzr* pseudo-register (and its 32-bit alias *wzr*) which reads as zero, and writes are ignored. As I noted in the above table, if an instruction encodes a register number of 31, then a special behavior kicks in, typically by treating mythical register 31 as an alias for *sp* or *zr*. Generally speaking, when being used as a base address register, imaginary register 31 represents *sp*, but when used for arithmetic or as a destination register, it represents *zr*.²

In instruction descriptions, I will use these shorthands:

Shorthand	Meaning
<i>Xn</i>	Any <i>x#</i> register
<i>Xn/zr</i>	Any <i>x#</i> register or <i>xzr</i>
<i>Xn/sp</i>	Any <i>x#</i> register or <i>sp</i>
<i>Wn</i>	Any <i>w#</i> register
<i>Wn/zr</i>	Any <i>w#</i> register or <i>wzr</i>
<i>Wn/sp</i>	Any <i>w#</i> register or <i>wsp</i>
<i>Rn</i>	Any <i>x#</i> or <i>w#</i> register
<i>Rn/zr</i>	Any <i>x#</i> register, <i>w#</i> register, <i>xzr</i> or <i>wzr</i>

The floating point registers have been reorganized. They have doubled in size (to 128 bits) as well as in number, and the single-precision registers are no longer paired up.

Register	Preserved?	Notes
<i>v0</i>	No	Parameter 1, return value
<i>v1</i>	No	Parameter 2
<i>v2</i>	No	Parameter 3
<i>v3</i>	No	Parameter 4
<i>v4</i>	No	Parameter 5
<i>v5</i>	No	Parameter 6
<i>v6</i>	No	Parameter 7
<i>v7</i>	No	Parameter 8
<i>v8</i> through <i>v15</i>	Low 64 bits only	Upper 64 bits are not preserved
<i>v16</i> through <i>v31</i>	No	

Each floating point register can be viewed in multiple ways. The partial registers are stored in the least significant bits of the full register.

Name	Meaning	Notes
<i>v#</i>	SIMD vector	
<i>q#</i>	128-bit value	quad precision
<i>d#</i>	64-bit value	double precision
<i>s#</i>	32-bit value	single precision
<i>h#</i>	16-bit value	half precision
<i>b#</i>	8-bit value	

The flags register is formally known as the Application Program Status Register (APSR). The flags available to user mode are the same as in AArch32:

Mnemonic	Meaning	Notes
N	Negative	Set if the result is negative
Z	Zero	Set if the result is zero
C	Carry	Multiple purposes
V	Overflow	Signed overflow
Q	Saturation	Accumulated overflow
GE[n]	Greater than or equal to	4 flags (SIMD)

The overflow flag records whether the most recent operation resulted in signed overflow. The saturation flag is used by multimedia instructions to accumulate whether any overflow occurred since it was last cleared. The GE flags record the result of SIMD operations. By convention, flags are not preserved across calls.

There are a number of AArch64 features that you are extremely unlikely to see in Windows code, such as tagged pointers, tagged memory, and pointer authentication, so I won't cover them here. I also won't cover floating point instructions or SIMD instructions.

Next time, we'll look at some of the weird transformations that can be performed inside an instruction.

Additional references:

- Code in ARM Assembly: Registers explained. An analogous series looking at AArch64 from the Apple point of view rather than Windows.
- Writing ARM64 Code for Apple Platforms: The Apple ABI specification for AArch64.

¹ The Windows debugger isn't quite sure which name to use for these registers. The disassembler calls the registers *xip0*, *xip1*, and *xpr*, but the expression evaluator doesn't understand those names; you have to call them `@x16`, `@x17`, and `@x18`. On the other hand, the expression evaluator does understand `@fp` and `@lr` and refuses to acknowledge the existence of the names `@x29` and `@x30`. Furthermore, the expression evaluator doesn't understand any of the *w* aliases.

² AArch64's register 31 is similar to PowerPC's register 0, which changes meaning depending on the instruction. In PowerPC assembly, it was on you to keep track of which encodings treat register 0 as a value register, and which treat it as a zero register. At least AArch64 expresses the two cases differently: If an encoding uses pseudo-register 31 to mean *sp*, then you really must write *sp*. If you write *xzr*, you get an error.

PowerPC on the other hand would happily let you specify *r0* even if the instruction treats it as zero. Which was one of the jokes from the short-lived parody twitter account that mocked PowerPC.

PowerPC Instructions 
@ppcinstructions · [Follow](#)

mscdf - Means Something Completely Different For r0
10:31 AM · Jan 21, 2015 

 88  Reply  Share this Tweet

[Read 2 replies](#)

Raymond Chen

Follow

