

# Double-precision floating-point format

---

**Double-precision floating-point format** is a computer number format, usually occupying 64 bits in computer memory; it represents a wide dynamic range of numeric values by using a floating radix point.

Floating point is used to represent fractional values, or when a wider range is needed than is provided by fixed point (of the same bit width), even if at the cost of precision. Double precision may be chosen when the range or precision of single precision would be insufficient.

In the IEEE 754-2008 standard, the 64-bit base-2 format is officially referred to as **binary64**; it was called **double** in IEEE 754-1985. IEEE 754 specifies additional floating-point formats, including 32-bit base-2 *single precision* and, more recently, base-10 representations.

One of the first programming languages to provide single- and double-precision floating-point data types was Fortran. Before the widespread adoption of IEEE 754-1985, the representation and properties of floating-point data types depended on the computer manufacturer and computer model, and upon decisions made by programming-language implementers. E.g., GW-BASIC's double-precision data type was the 64-bit MBF floating-point format.

## Contents

---

### IEEE 754 double-precision binary floating-point format: binary64

Exponent encoding

Endianness

Double-precision examples

Execution speed with double-precision arithmetic

### Implementations

C and C++

Fortran

Common Lisp

Java

JavaScript

### See also

### Notes and references

## IEEE 754 double-precision binary floating-point format: binary64

---

Double-precision binary floating-point is a commonly used format on PCs, due to its wider range over single-precision floating point, in spite of its performance and bandwidth cost. It is commonly known simply as *double*. The IEEE 754 standard specifies a **binary64** as having:

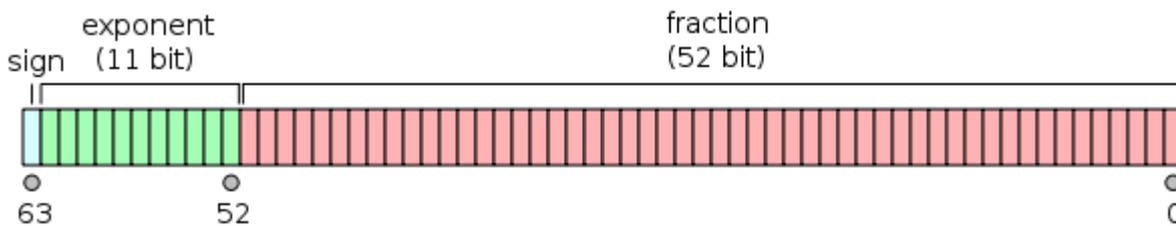
- Sign bit: 1 bit
- Exponent: 11 bits
- Significand precision: 53 bits (52 explicitly stored)

The sign bit determines the sign of the number (including when this number is zero, which is signed).

The exponent field is an 11-bit unsigned integer from 0 to 2047, in biased form: an exponent value of 1023 represents the actual zero. Exponents range from  $-1022$  to  $+1023$  because exponents of  $-1023$  (all 0s) and  $+1024$  (all 1s) are reserved for special numbers.

The 53-bit significand precision gives from 15 to 17 significant decimal digits precision ( $2^{-53} \approx 1.11 \times 10^{-16}$ ). If a decimal string with at most 15 significant digits is converted to IEEE 754 double-precision representation, and then converted back to a decimal string with the same number of digits, the final result should match the original string. If an IEEE 754 double-precision number is converted to a decimal string with at least 17 significant digits, and then converted back to double-precision representation, the final result must match the original number.<sup>[1]</sup>

The format is written with the significand having an implicit integer bit of value 1 (except for special data, see the exponent encoding below). With the 52 bits of the fraction (F) significand appearing in the memory format, the total precision is therefore 53 bits (approximately 16 decimal digits,  $53 \log_{10}(2) \approx 15.955$ ). The bits are laid out as follows:



The real value assumed by a given 64-bit double-precision datum with a given biased exponent  $e$  and a 52-bit fraction is

$$(-1)^{\text{sign}} (1.b_{51}b_{50}\dots b_0)_2 \times 2^{e-1023}$$

or

$$(-1)^{\text{sign}} \left( 1 + \sum_{i=1}^{52} b_{52-i} 2^{-i} \right) \times 2^{e-1023}$$

Between  $2^{52}=4,503,599,627,370,496$  and  $2^{53}=9,007,199,254,740,992$  the representable numbers are exactly the integers. For the next range, from  $2^{53}$  to  $2^{54}$ , everything is multiplied by 2, so the representable numbers are the even ones, etc. Conversely, for the previous range from  $2^{51}$  to  $2^{52}$ , the spacing is 0.5, etc.

The spacing as a fraction of the numbers in the range from  $2^n$  to  $2^{n+1}$  is  $2^{n-52}$ . The maximum relative rounding error when rounding a number to the nearest representable one (the machine epsilon) is therefore  $2^{-53}$ .

The 11 bit width of the exponent allows the representation of numbers between  $10^{-308}$  and  $10^{308}$ , with full 15–17 decimal digits precision. By compromising precision, the subnormal representation allows even smaller values up to about  $5 \times 10^{-324}$ .

## Exponent encoding







## C and C++

C and C++ offer a wide variety of arithmetic types. Double precision is not required by the standards (except by the optional annex F of C99, covering IEEE 754 arithmetic), but on most systems, the `double` type corresponds to double precision. However, on 32-bit x86 with extended precision by default, some compilers may not conform to the C standard and/or the arithmetic may suffer from double rounding.<sup>[5]</sup>

## Fortran

Fortran provides several integer and real types, and the 64-bit type `real64`, accessible via Fortran's intrinsic module `iso_fortran_env`, corresponds to double precision.

## Common Lisp

Common Lisp provides the types SHORT-FLOAT, SINGLE-FLOAT, DOUBLE-FLOAT and LONG-FLOAT. Most implementations provide SINGLE-FLOATs and DOUBLE-FLOATs with the other types appropriate synonyms. Common Lisp provides exceptions for catching floating-point underflows and overflows, and the inexact floating-point exception, as per IEEE 754. No infinities and NaNs are described in the ANSI standard, however, several implementations do provide these as extensions.

## Java

On Java before version 1.2, every implementation had to be IEEE 754 compliant. Version 1.2 allowed implementations to bring extra precision in intermediate computations for platforms like x87. Thus a modifier `strictfp` was introduced to enforce strict IEEE 754 computations.

## JavaScript

As specified by the ECMAScript standard, all arithmetic in JavaScript shall be done using double-precision floating-point arithmetic.<sup>[6]</sup>

## See also

---

- IEEE 754, IEEE standard for floating-point arithmetic

## Notes and references

---

1. William Kahan (1 October 1997). "Lecture Notes on the Status of IEEE Standard 754 for Binary Floating-Point Arithmetic" (<http://www.cs.berkeley.edu/~wkahan/ieee754status/IEEE754.PDF>) (PDF). Archived (<https://web.archive.org/web/20120208075518/http://www.cs.berkeley.edu/~wkahan/ieee754status/IEEE754.PDF>) (PDF) from the original on 8 February 2012.
2. Savard, John J. G. (2018) [2005], "Floating-Point Formats" (<http://www.quadibloc.com/comp/cp0201.htm>), *quadibloc*, archived (<https://web.archive.org/web/20180703001709/http://www.quadibloc.com/comp/cp0201.htm>) from the original on 2018-07-03, retrieved 2018-07-16
3. "pack – convert a list into a binary representation" (<http://www.perl.com/doc/manual/html/pod/perlfunc/pack.html>).

4. "Nvidia's New Titan V Pushes 110 Teraflops From A Single Chip" (<https://www.tomshardware.com/news/nvidia-titan-v-110-teraflops,36085.html>). *Tom's Hardware*. 2017-12-08. Retrieved 2018-11-05.
5. "Bug 323 – optimized code gives strange floating point results" ([https://gcc.gnu.org/bugzilla/show\\_bug.cgi?id=323](https://gcc.gnu.org/bugzilla/show_bug.cgi?id=323)). *gcc.gnu.org*. Archived ([https://web.archive.org/web/20180430012629/http://gcc.gnu.org/bugzilla/show\\_bug.cgi?id=323](https://web.archive.org/web/20180430012629/http://gcc.gnu.org/bugzilla/show_bug.cgi?id=323)) from the original on 30 April 2018. Retrieved 30 April 2018.
6. *ECMA-262 ECMAScript Language Specification* (<http://www.ecma-international.org/publications/files/ECMA-ST-ARCH/ECMA-262%205th%20edition%20December%202009.pdf>) (PDF) (5th ed.). Ecma International. p. 29, §8.5 *The Number Type*. Archived (<https://web.archive.org/web/20120313145717/http://www.ecma-international.org/publications/files/ECMA-ST-ARCH/ECMA-262%205th%20edition%20December%202009.pdf>) (PDF) from the original on 2012-03-13.

---

Retrieved from "[https://en.wikipedia.org/w/index.php?title=Double-precision\\_floating-point\\_format&oldid=975978799](https://en.wikipedia.org/w/index.php?title=Double-precision_floating-point_format&oldid=975978799)"

---

**This page was last edited on 31 August 2020, at 14:00 (UTC).**

Text is available under the Creative Commons Attribution-ShareAlike License; additional terms may apply. By using this site, you agree to the Terms of Use and Privacy Policy. Wikipedia® is a registered trademark of the Wikimedia Foundation, Inc., a non-profit organization.