

Glossary

680x0-based Macintosh computer Any computer containing a 680x0 central processing unit that runs Macintosh system software. Compare **PowerPC processor-based Macintosh computer**.

ANSI X3J11.1 A branch of the American National Standards Institute (ANSI) that is working on a numerics standard for the C programming language. This group is also called the *Numerical C Extensions Group* (NCEG) and has produced the Floating-Point C Extensions (FPCE) technical report.

antisymmetric Used to describe a function whose graph is not symmetrical across the y-axis; that is $func(x) \neq func(-x)$ for all x .

atomic operations Operations that pass extra information back to their callers by signaling exceptions but that hide internal exceptions, which might be irrelevant or misleading.

bias A number added to the binary exponent of a floating-point number so that the exponent field will always be positive. The bias is subtracted when the floating-point value is evaluated.

binade The collection of numbers that lie between two successive powers of 2.

binary floating-point number A collection of bits representing a sign, an exponent, and a significand. Its numerical value, if any, is the signed product of the significand and 2 raised to the power of the exponent.

complex expression An expression made up of more than one simple expression, that is, an expression with more than one floating-point operation.

Condition Register A 32-bit PowerPC register used to summarize the states of the fixed-point and floating-point processors and to store results of comparison operations.

decimal format structure A data type for specifying the formatting for decimal (base 10) numbers (of conversions). It specifies the decimal number's style and number of digits. It is defined by the `decform` data type.

decimal structure A data type for storing decimal data. It consists of three fields: sign, exponent, and significand (a C string). It is defined by the `decimal` data type.

default environment The environment settings when a PowerPC Numerics implementation starts up: rounding is to nearest and all exception flags are clear.

denormalized number A nonzero binary floating-point number whose significand has an implicit leading bit of 0 and whose exponent is the minimum exponent for the number's data format. Also called *denorm*. See also **normalized number**.

divide-by-zero exception A floating-point exception that occurs when a finite, nonzero number is divided by zero or some other improper operation on zero has occurred.

double format A 64-bit application data format for storing floating-point values of up to 15- or 16-decimal digit precision.

double-double format A 128-bit application data format made up of two double-format numbers. It has the same range as the double format but much greater precision.

environmental access switch A switch, recommended in the FPCE technical report, that specifies whether a program accesses the rounding direction modes and exception flags.

environmental controls The rounding direction modes and the exception flags.

evaluation format The data format used to evaluate the result of an expression. The evaluation format must be at least as wide as the expression's semantic type. (It may be the same as the semantic type.)

exception An error or other special condition detected by the microprocessor in the course of program execution. The floating-point exceptions are invalid, underflow, overflow, divide-by-zero, and inexact.

exception flag Each exception has a flag that can be set, cleared, and tested. It is set when its respective exception occurs and stays set until explicitly cleared.

exponent The part of a binary floating-point number that indicates the power to which 2 is raised in determining the value of the number. The wider the exponent field in a numeric data format, the greater range the format will handle.

expression evaluation method The method by which an evaluation format is determined for an expression.

floating-point operation An operation that is performed on numbers in floating-point formats. The IEEE standard requires that a numerics environment support addition, subtraction, multiplication, division, square root, remainder, and round-to-integer as the basic floating-point arithmetic operations.

Floating-Point Status and Control Register (FPSCR) A 32-bit PowerPC register used to store the floating-point environment.

flush-to-zero system A system that excludes denormalized numbers. Results smaller than the smallest normalized number are rounded to zero.

FPCE technical report A report authored by the Numerical C Extensions Group (ANSI X3J11.1) that proposes a standard for floating-point operations in the C programming language.

FPSCR See **Floating-Point Status and Control Register**.

fraction A field in a floating-point data format that stores all but the leading bit of the significand of a floating-point number.

gradual underflow A process that occurs on a computer system that includes denormalized numbers.

IEEE standard A term used in this book to mean IEEE Standard 754.

IEEE Standard 754 A standard that defines how computers should perform binary floating-point arithmetic.

IEEE Standard 854 A standard that defines how computers should perform radix-independent floating-point arithmetic.

inexact exception A floating-point exception that occurs when the exact result of a floating-point operation must be rounded.

Infinity A special value produced when a floating-point operation should produce a mathematical infinity or when a floating-point operation attempts to produce a number greater in magnitude than the largest representable number in a given format. Infinities are signed.

integer types System types for integral values. Integer types typically use 16- or 32-bit two's-complement integers. Integer types are not PowerPC Numerics formats but are available to PowerPC Numerics users.

integral value A value, perhaps in a numeric data format, that is exactly equal to a mathematical integer. For example, -2, -1, 0, 1, 2, and so on.

invalid exception A floating-point exception that occurs if an operand is invalid for the operation being performed.

invalid-operation exception See **invalid exception**.

Machine State Register A 32-bit PowerPC supervisor-level register that records the state of the processor, including if floating-point instructions and floating-point exceptions are enabled.

mantissa See **significand**.

MathLib See **PowerPC Numerics library**.

minimum evaluation format The narrowest format in which a floating-point operation can be performed. Each implementation of PowerPC Numerics defines its own minimum evaluation format.

multiply-add instruction A type of instruction unique to the PowerPC architecture. Multiply-add instructions perform a multiply plus an addition or subtraction operation with at most a single roundoff error.

NaN (Not-a-Number) A special bit pattern produced when a floating-point operation cannot produce a meaningful result (for example, $0/0$ produces a NaN). NaNs propagate through arithmetic operations.

NCEG (Numerical C Extensions Group) See **ANSI X3J11.1**.

nextafter functions Functions that return the next value after the input value that is representable in one of the floating-point data formats. For example, `nextafterd(0, +∞)` returns the value that comes immediately after 0 in the direction of $+∞$ in double format.

normalized number A binary floating-point number in which all significand bits are significant: that is, the leading bit of the significand is 1. Compare **denormalized number**.

Numerical C Extensions Group (NCEG) See **ANSI X3J11.1**.

overflow exception A floating-point exception that occurs when the magnitude of a floating-point result is greater than the largest finite number that the destination data format can represent.

PowerPC Numerics The floating-point environment on PowerPC processor-based Macintosh computers. This environment provides floating-point data formats and operations plus some advanced numerical functions, such as logarithmic and trigonometric functions.

PowerPC Numerics library A C library that implements floating-point transcendental functions and contains type definitions and macros used for floating-point operations. It is contained in the file `MathLib`.

PowerPC processor Any member of the family of PowerPC microprocessors. The MPC601 processor is the first PowerPC central processing unit.

PowerPC processor-based Macintosh computer Any computer containing a PowerPC central processing unit that runs Macintosh system software. Compare **680x0-based Macintosh computer**.

precision The number of digits required to accurately represent a number. For example, the value 3.2 requires two decimal digits of precision, and the value 3.002 requires four decimal digits. In numeric data formats, the precision is equal to the number of bits (both implicit and explicit) in the significand.

quiet NaN A NaN that propagates through arithmetic operations without signaling an exception.

rounding An action performed when a result of an arithmetic operation cannot be represented exactly in a numeric data format. With rounding, the computer changes the result to a close value that can be represented exactly.

rounding direction modes Modes that specify the direction a computer will round when the result of an arithmetic operation cannot be represented exactly in a numeric data format. Under PowerPC Numerics, the computer resolves rounding decisions in one of the four directions chosen by the user: to nearest (the default), upward, downward, and toward zero.

roundoff error The difference between the exact result of an IEEE arithmetic operation and the result as it is represented in the numeric data format if the result has been rounded.

SANE See **Standard Apple Numerics Environment**.

semantic type The widest type of the operands of an expression.

signaling NaN A NaN that signals an invalid exception when the NaN is an operand of an arithmetic operation. If no halt occurs, a quiet NaN is produced for the result. No PowerPC Numerics operation creates signaling NaNs.

sign bit The bit of a single, double, or double-double number that indicates the number's sign: 0 indicates a positive number; 1, a negative number.

significand The part of a binary floating-point number that indicates where the number falls between two successive powers of 2. The wider the significand field in a numeric format, the more precision the format has.

simple expression An expression containing one floating-point operation.

single format A 32-bit application data format for storing floating-point values that have a precision of up to seven or eight decimal digits. It is used by engineering applications, among others.

Standard Apple Numerics Environment (SANE) The floating-point environment on 680x0-based Macintosh computers. This environment provides floating-point data formats and operations as well as some advanced numerical functions such as logarithmic and trigonometric functions.

sticky Used to describe a condition in which a bit stays set until it is explicitly cleared. Floating-point exception flags in the FPSCR are sticky, so if one instruction sets an exception flag and another instruction is performed before the flag is tested, it is impossible to tell which instruction caused the exception.

subnormal number A denormalized number.

symmetric Used to describe a function whose graph looks the same on both sides of the y-axis; that is, $func(x) = func(-x)$ for all x .

tiny Used to describe a number whose magnitude is smaller than the smallest positive normalized number in the format of the number.

transcendental functions Functions that can be used as building blocks in numerical functions. All of the functions contained in the PowerPC Numerics library are transcendental functions.

trigonometric functions Functions that perform trigonometric operations, such as cosine, sine, and tangent.

truncate To chop off the fractional part of a real number so that only the integer part remains. For example, if the real number 1.9999999999 is truncated, the truncated value is 1.

underflow exception An exception that occurs when the result of an operation is both tiny and inexact.

usual arithmetic conversions Automatic conversions performed in the C programming language. The ANSI C specification defines these conversions.

widest-need evaluation An evaluation method in which the widest format of all of the operands in a complex expression is used as the format in which the expression is evaluated.