

# Fast Square Root Approximation in C++

Tuğrul Kök

January 2026

## Abstract

This document details the design and implementation of a fast square root approximation algorithm. While modern hardware intrinsics (SSE/AVX) are the standard for runtime performance, this implementation demonstrates low-level algorithmic optimization, bit-wise manipulation of IEEE 754 floating-point standards, and numerical refinement via Newton-Raphson iteration.

## Design Philosophy

In modern C++ (circa 2026), the optimal method for square root calculation is typically to utilize hardware intrinsics (e.g., `SQRTSS`) or to leverage `constexpr` to shift the computational cost from runtime to compile-time for constant inputs.

However, the `sqrt_fast` function detailed below is implemented to demonstrate a specific optimization strategy: **algorithmic approximation**. Adapted from the famous *Quake III* fast inverse square root algorithm, this method trades minor floating-point precision for reduced instruction latency.

In High-Frequency Trading (HFT) environments, where execution speed is often measured in nanoseconds, minimizing CPU cycles on the critical path is paramount. This implementation showcases the ability to manipulate memory at the bit level to achieve performance gains when strict IEEE 754 compliance is not required.

## IEEE 754 and Bit Reinterpretation

The core of this optimization lies in the derivation of the magic constant `0x1fbd1df5` and the bitwise shift operation `i = 0x1fbd1df5 + (i >> 1)`. To understand this, we must examine the binary representation of floating-point numbers.

### 1. Floating Point Structure

In the IEEE 754 standard, a 32-bit floating-point number consists of three parts:

- **Sign (1 bit):** Determines if the number is positive or negative (0 for positive).
- **Exponent (8 bits):** Represents the magnitude ( $E$ ).
- **Mantissa (23 bits):** Represents the fractional precision ( $M$ ).

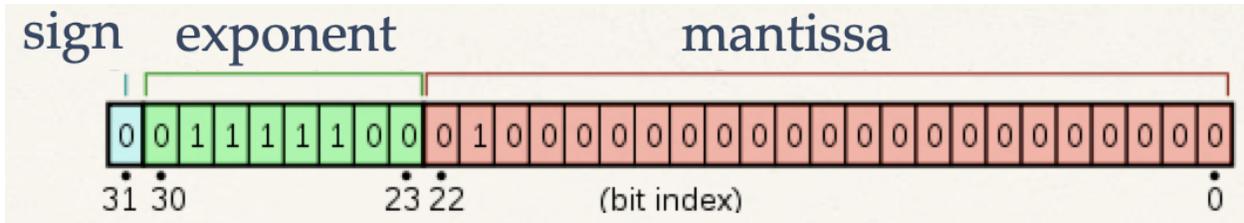


Figure 1: IEEE 754 Float Bit representation

The mathematical value of the number is given by:

$$\text{Value} = \left(1 + \frac{M}{2^{23}}\right) \times 2^{E-127}$$

For example, the number 42 is represented as:

$$42 = 2^5 \times (1 + 0.3125)$$

Binary representation: 0 10000100 0101000000000000000000

## 2. The Logarithmic Transformation

To calculate the square root, we need to divide the exponent by 2. This is difficult in standard floating-point form but trivial in logarithmic form.

Taking the base-2 logarithm on both sides:

$$\log_2(\text{Value}) = \log_2 \left( \left(1 + \frac{M}{2^{23}}\right) \times 2^{E-127} \right)$$

Using logarithm rules ( $\log(a \cdot b) = \log a + \log b$ ):

$$\log_2(\text{Value}) = \log_2 \left( 1 + \frac{M}{2^{23}} \right) + (E - 127)$$

## 3. The Linear Approximation

Computing  $\log_2(1+x)$  is expensive. However, for  $x \in [0, 1)$ , the function is nearly linear. We can approximate it as:

$$\log_2(1+x) \approx x + \sigma$$

Where  $\sigma$  is a free parameter (a tunable constant) used to minimize the error. The optimal value for  $\sigma$  is approximately **0.0450465**.

Substituting this back into our equation:

$$\begin{aligned} \log_2(\text{Value}) &\approx \left( \frac{M}{2^{23}} + \sigma \right) + E - 127 \\ &\approx \frac{M}{2^{23}} + E + \sigma - 127 \end{aligned}$$

We can rewrite this to expose the bit representation. The integer interpretation of the floating point bits ( $I$ ) is roughly  $M + 2^{23} \times E$ .

$$\log_2(\text{Value}) \approx \frac{1}{2^{23}}(M + 2^{23} \times E) + \sigma - 127$$

$$\log_2(\text{Value}) \approx \frac{1}{2^{23}}(I) + \sigma - 127$$

**Conclusion:** The integer representation of a float ( $I$ ) is essentially a scaled and shifted version of its logarithm.

## 4. Deriving the Square Root Operation

We want to find  $\sqrt{y}$ . In logarithms:

$$\log_2(\sqrt{y}) = \frac{1}{2} \log_2(y)$$

Let  $I_y$  be the integer representation of  $y$ , and  $I_{\sqrt{y}}$  be the integer representation of the result. Substituting our approximation into both sides:

$$\frac{1}{2^{23}} I_{\sqrt{y}} + \sigma - 127 = \frac{1}{2} \left( \frac{1}{2^{23}} I_y + \sigma - 127 \right)$$

Multiply by  $2^{23}$  to solve for  $I_{\sqrt{y}}$ :

$$I_{\sqrt{y}} + 2^{23}(\sigma - 127) = \frac{1}{2} I_y + \frac{1}{2} 2^{23}(\sigma - 127)$$

Rearranging the terms:

$$I_{\sqrt{y}} = \frac{1}{2} I_y + \frac{1}{2} 2^{23}(127 - \sigma)$$

## 5. Calculating the Magic Number

The term  $\frac{1}{2} I_y$  corresponds to the bitwise right shift ( $i \gg 1$ ). The second term is a constant which becomes our "Magic Number."

Using  $\sigma = 0.0450465$ :

$$\text{Magic Number} = \frac{1}{2} \cdot 2^{23} \cdot (127 - 0.0450465)$$

$$\text{Magic Number} \approx 532,487,669$$

$$\text{Hexadecimal} \approx 0x1FBD1DF5$$

Thus, the entire square root operation simplifies to a single integer calculation:

$$I_{\sqrt{y}} = 0x1FBD1DF5 + (I_y \gg 1)$$

## Code Listing

Listing 1: Fast Square Root Implementation

```
1 // 3. Fast Square Root Approximation
2 float sqrt_fast(float x) {
3     if (x <= 0.0f) return 0.0f;
4
5     // 1. Bit-level Manipulation (Initial Guess)
6     // Safely reinterpret float bits as int32 for manipulation
7     int32_t i = std::bit_cast<int32_t>(x);
8
9     // Apply the magic constant and bit-shift
10    i = 0x1fbd1df5 + (i >> 1);
11
12    // Reinterpret back to float
13    float y = std::bit_cast<float>(i);
14
15    // 2. Newton-Raphson Refinement
16    // f(y) = y^2 - x = 0 => y' = 0.5 * (y + x/y)
17    y = 0.5f * (y + x / y);
18
19    return y;
20 }
```

## Function Analysis

### Newton-Raphson Refinement

The bit manipulation provides a very good initial guess (approx 3-4% error). To achieve trading-grade accuracy, we apply one iteration of Newton's method.

$$y_{new} = \frac{1}{2} \left( y_{old} + \frac{x}{y_{old}} \right)$$

This step treats the bit-hack result as a starting point  $y_0$  on the curve  $f(y) = y^2 - x$  and follows the tangent line to find a more precise root. This single iteration reduces the error margin significantly (typically to  $< 0.2\%$ ).

## Technical Implementation Notes

### Modern C++ Standards (C++20)

This implementation utilizes `std::bit_cast`, which is the modern, type-safe standard for bit reinterpretation. This replaces legacy methods such as:

- `*(int*)&x`: Violates strict aliasing rules (Undefined Behavior).
- `union`: Technically undefined in C++ for type punning.
- `std::memcpy`: Safe but verbose.

## Conclusion

This implementation demonstrates how classic bit-level optimization and numerical methods can be synthesized to solve performance challenges. While it lacks the absolute precision of hardware intrinsics, it serves as a powerful optimization for latency-critical scenarios where execution speed outweighs the need for decimal-perfect accuracy.