

Speeding Up Floating- Point Division With In- lined Iterative Algorithms

Robert Enenkel, Allan Martin
IBM® Toronto Lab

Outline

- **Hardware floating-point division**
- **The case for software division**
- **Software division algorithms**
- **Special cases/tradeoffs**
- **Performance results**
- **Automatic generation**

Hardware Division

- PPC `fdiv`, `fdivs`

- Advantages

 - f* accurate (correctly rounded)

 - f* handles exceptional cases (Inf, NaN)

 - f* lower latency than SW

- Disadvantages

 - f* occupies FPU completely

 - f* inhibits parallelism

Alternatives to HW division

- **Vector libraries**

 - f* MASS

 - f* higher overhead, greater speedup

- **In-lined software division**

 - f* low overhead, medium speedup

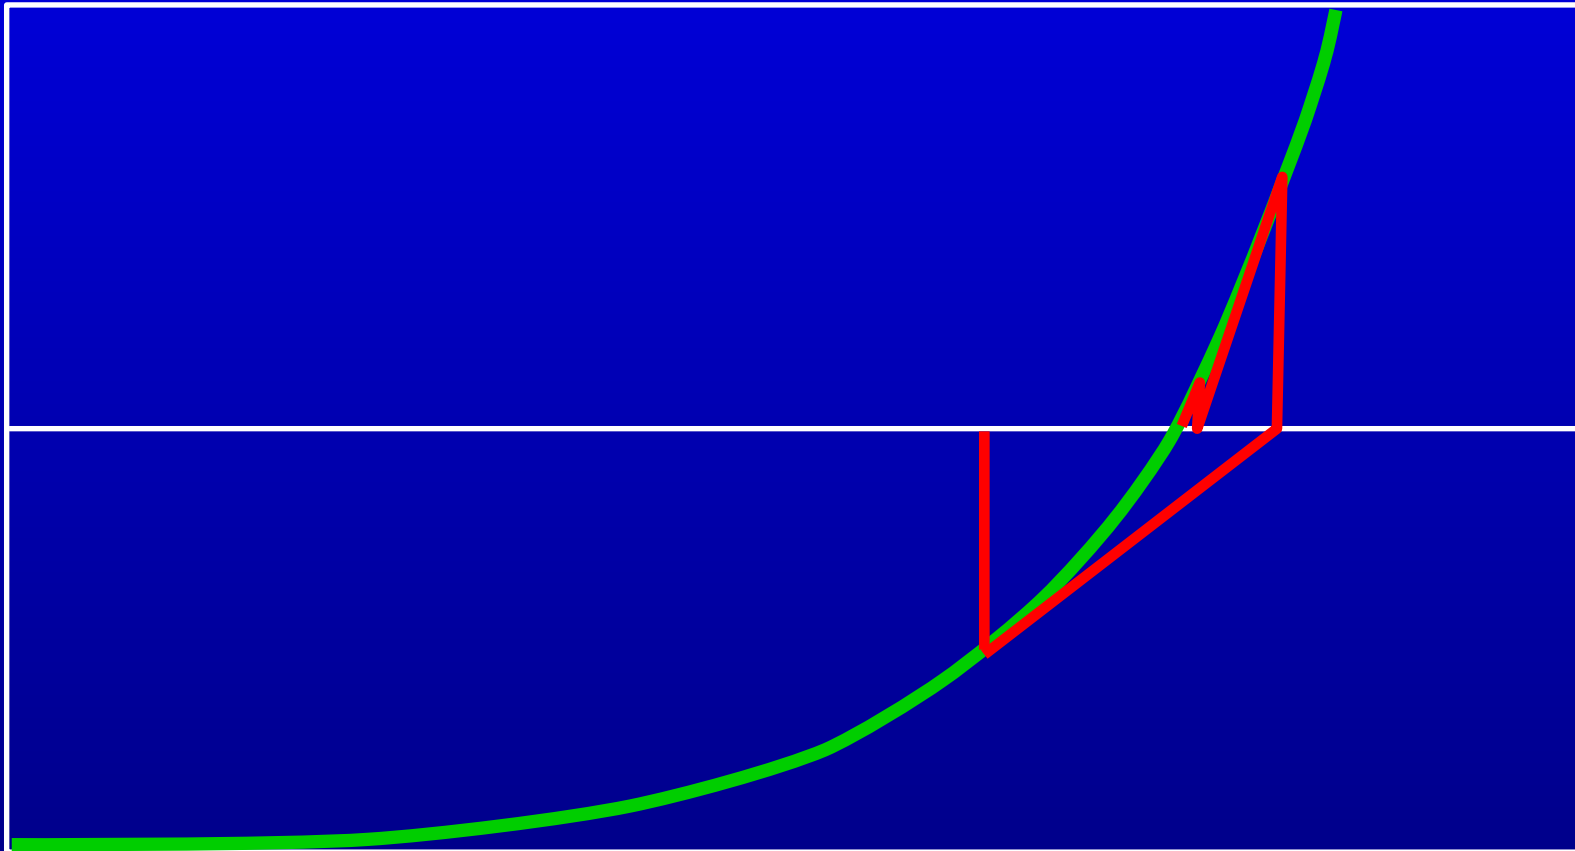
Rationale for Software Division

- Write SW division algorithm in terms of HW arithmetic instructions
 - f Newton's method or Taylor series
- Latency will be higher than HW division
- But...SW instructions can be interleaved, so throughput may be better
- Requires enough independent instructions to interleave
 - f loop of divisions
 - f other work

Newton's Method

- To find x such that $f(x) = 0$,
- Initial guess x_0
- $x_{n+1} = x_n - f(x_n)/f'(x_n)$, $n=0, 1, 2, \dots$
- Provided x_0 is close enough
 - f x_n converges to x
 - f It converges quadratically $|x_{n+1}-x| < c|x_n-x|^2$
 - f Number of bits of accuracy doubles with each iteration

Newton's Method



Newton Iteration for Division

- For $1/b$, let $f(x) = 1/x - b$
- For a/b , use $a*(1/b)$ or $f(x) = a/x - b$
- Algorithm for $1/b$

f $x_0 \sim 1/b$ initial guess

f $e_0 = 1 - b*x_0$

f $x_1 = x_0 + e_0*x_0$

f $e_1 = e_0*e_0$

f $x_2 = x_1 + e_1*x_1$

f etc...

How Many Iterations Needed?

- Power5 reciprocal estimate instructions

 - f* FRES (single precision), FRE (double prec.)

 - f* |relative error| $\leq 2^{-8}$

- Floating-point precision

 - f* single: 24 bits

 - f* double: 53 bits

- Newton iterations

 - f* error: 2^{-16} , 2^{-32} , 2^{-64} , 2^{-128}

 - f* single: 2 iterations for 1 ulp

 - f* double: 3 iterations for 1 ulp

 - f* +1 iteration for correct rounding (0.5 ulps)

Taylor Series for Reciprocal

- $x_0 \sim 1/b$ initial guess
- $e = 1 - b x_0$
- $1/b = x_0 / (b x_0) = x_0 (1 / (1 - e))$
 $= x_0 (1 + e + e^2 + e^3 + e^4 + \dots)$
- Algorithm (6 terms)
 - $f \quad e = 1 - d * x_0$
 - $f \quad t_1 = 0.5 + e * e$
 - $f \quad q_1 = x_0 + x_0 * e$
 - $f \quad t_2 = 0.75 + t_1 * t_1$
 - $f \quad t_3 = q_1 * e$
 - $f \quad q_2 = x_0 + t_2 * t_3$

Speed/Accuracy tradeoff

- IBM compilers have `-qstrict/-qnostrict`
- `-qstrict`: SW result should match HW division exactly
- `-qnostrict`: SW result may be slightly less accurate for speed

Exceptions

- **Even when a/b is representable...**
- **$1/b$ may underflow**
 - f* $a \sim b \sim \text{huge}, a/b \sim 1, 1/b$ denormalized
 - f* Causes loss of accuracy
- **$1/b$ may overflow**
 - f* a, b denormalized, $a/b \sim 1, 1/b = \text{Inf}$
 - f* Causes SW algorithm to produce NaN
- **Handle with tests in algorithm**
 - f* Use HW divide for exceptional cases

Algorithm variations

- **User callable built-in functions**
 - f* **swdiv(a,b): double precision, checking**
 - f* **swdivs(a,b): single precision, checking**
 - f* **swdiv_nochk(a,b): double, non-checking**
 - f* **swdivs_nochk(a,b): single, non-checking**
- **Accuracy of swdiv, swdiv_nochk depends on -qstrict/-qnostrict**
- **_nochk versions faster but have argument restrictions**

Accuracy and Performance

| | Power5 speedup ratio | Power4 speedup ratio | Power5 ulps max error | Power4 ulps max error |
|----------------------|-------------------------|-------------------------|--------------------------|--------------------------|
| swdivs | 1.07 | 1.05 | 0.5 | 0.5 |
| swdivs_nochk | 1.46 | 1.28 | 0.5 | 0.5 |
| swdiv_strict | 1.05 | | 0.5 | |
| swdiv_nostrict | 1.50 | | 1.5 | |
| swdiv_nochk_strict | 1.51 | | 0.5 | |
| swdiv_nochk_nostrict | 1.77 | | 1.5 | |

Automatic Generation of Software Division

- The `swdivs` and `swdiv` algorithms can also be automatically generated by the compiler
- Compiler can detect situations where throughput is more important than latency

Automatic Generation of Software Division

- In straight-line code, we use a heuristic that calculates how much FP can be executed in parallel

f independent instructions are good, especially
if they divide

f dependent instructions are bad (they increase
latency)

Automatic Generation of Software Division

- In modulo scheduled loops software-divide code can be pipelined, interleaving multiple iterations
- Divides are expanded if divide does not appear in a recurrence (cyclic data-dependence)

Summary

- **Software divide algorithms**

 - f* user callable

 - f* compiler generated

- **Loops of divides**

 - f* up to 1.77x speedup

- **UMT2K benchmark**

 - f* 1.19x speedup