



Digital Commons@

Loyola Marymount University
LMU Loyola Law School

Mathematics, Statistics and Data Science
Faculty Works

Mathematics, Statistics and Data Science

1996

A Multiple-Precision Division Algorithm

David M. Smith

Loyola Marymount University, dsmith@lmu.edu

Follow this and additional works at: https://digitalcommons.lmu.edu/math_fac



Part of the [Mathematics Commons](#)

Digital Commons @ LMU & LLS Citation

Smith, David M., "A Multiple-Precision Division Algorithm" (1996). *Mathematics, Statistics and Data Science Faculty Works*. 2.

https://digitalcommons.lmu.edu/math_fac/2

This Article is brought to you for free and open access by the Mathematics, Statistics and Data Science at Digital Commons @ Loyola Marymount University and Loyola Law School. It has been accepted for inclusion in Mathematics, Statistics and Data Science Faculty Works by an authorized administrator of Digital Commons@Loyola Marymount University and Loyola Law School. For more information, please contact digitalcommons@lmu.edu.

A MULTIPLE-PRECISION DIVISION ALGORITHM

DAVID M. SMITH

ABSTRACT. The classical algorithm for multiple-precision division normalizes digits during each step and sometimes makes correction steps when the initial guess for the quotient digit turns out to be wrong. A method is presented that runs faster by skipping most of the intermediate normalization and recovers from wrong guesses without separate correction steps.

1. INTRODUCTION

Knuth [3] described classical algorithms for multiplication and division using t digits of precision with base- b arithmetic. These methods require $O(t^2)$ operations. Knuth also surveyed several algorithms with faster asymptotic running time. Bailey's multiple-precision package [1] included routines based on FFT methods for use at very high precision, as well as $O(t^2)$ routines. The FFT methods are faster asymptotically, but because the algorithms are complicated, the classical $O(t^2)$ methods are more efficient when the precision is less than a few thousand digits.

For multiplication, Brent's package [2] improved on the classical methods by restricting b so that integer overflow could not occur for numbers less than $8b^2$. This allowed his routine to do eight steps of the multiplication before normalizing, resulting in significantly faster execution. The FM multiplication routine [7] was based on this idea and reduced the number of normalizations still further. Here "normalizing" refers to the act of restoring the individual digits of an intermediate result to the range $[0, b - 1]$.

It is more difficult to skip normalization during division, because the quotient digit q used in each "multiply and subtract" step is a guess. Early routines ([3, 4, 5, 6, 8]) used integer arithmetic to estimate q . Bailey used double precision to estimate q , which made the probability of an incorrect q much smaller. Brent used Newton iteration and multiplication to avoid the long-division algorithm, but this made division much slower than multiplication.

The early routines normalized the partial result at each step so that cases where q was wrong could be recognized and corrected. This correction took the form of an extra "addback" step that was done so the value of q would be correct before the next multiply and subtract step was taken.

Since a large base was usually used, the correction step was needed only rarely. This means that eliminating the correction step makes the division routine only slightly faster. The time spent normalizing partial results could be greater than the time spent generating them, so minimizing the normalization is important.

Received by the editor June 17, 1994 and, in revised form, February 12, 1995.
1991 *Mathematics Subject Classification*. Primary 65-04, 65D15.

In version 1.0 of the FM package, the division routine normalized only a few leading digits of the partial result so that q could be verified or a correction step taken. Bailey's division routine eliminated the correction step, but normalized each partial result.

This paper presents a division algorithm that eliminates most of the intermediate normalization and does no separate correction steps. This algorithm still requires $O(t^2)$ steps, but is often two or three times faster than previous $O(t^2)$ methods ([1, 2, 7]).

2. EXAMPLES

The method is similar enough to the standard long-division algorithm that the differences can be seen using the example quotient π/e . A floating-point division will be done, although the method works in the same way for multiple-precision integer division.

Commonly used bases are in the range $[10^4, 10^8]$. The examples use $b = 10^4$, $t = 5$, and compute six digits of the quotient to simulate the guard digits carried by division routines.

To begin, the numerator is copied into a work array, W , with zero digits padded at the end to leave room for guard digit calculations. The numbers are stored in a floating-point format where the first word is the exponent (as a power of b) followed by the digits of the fraction part:

n =	1	3	1415	9265	3589	7932			
d =	1	2	7182	8182	8459	0452			

As the division proceeds, the "active" part of the work array shifts to the right, and most routines store the quotient digits in the left part of the array. The first word of W is the exponent of the quotient:

W =	1	0	3	1415	9265	3589	7932	0	0	0
-----	---	---	---	------	------	------	------	---	---	---

In this example, four words of each number are used to generate the floating-point approximations. Since they will be divided to get the quotient digits, both can be scaled by any convenient factor. The numerator value, x_n , will be scaled to be an integer, and the denominator value, x_d , is 271828182.8459. The approximation to each quotient digit is $q = \lfloor x_n/x_d \rfloor$.

Step 1. Use words 2 through 5 of W to get x_n : $x_n = 314159265.0$, $q = 1$.

Multiply and subtract: $W \leftarrow W - q * d$

W is displayed on two lines to show the current quotient and the active part of the array:

W =	1	0							
	1	-5767	1083	-4870	7480	0	0	0	

In step k , W_{k+1} is often nonzero. One word of W is normalized by setting $W_{k+2} = b * W_{k+1} + W_{k+2}$, and then q is stored in W_{k+1} :

W =	1	1							
	1	-5767	1083	-4870	7480	0	0	0	

Step 2. Use words 3 through 6 for x_n : $x_n = 423310825130.0$, $q = 1557$.

W =	1	1	1557						
	1119	-11181291	-12744244	-13163183	-703764	0	0		

Step 3. Use words 4 through 7 for x_n : $x_n = 743444396817.0$, $q = 2734$.

W =	1	1	1557	2734					
	3241	-32379832	-35532771	-23830670	-1235768	0			

Step 4. Use words 5 through 8 for x_n : $x_n = 2661448459330.0$, $q = 9790$.

W = 1 1 1557 2734 9790
 10588 -105844551 -103932450 -84049378 -4425080

Step 5. Use words 6 through 9 for x_n : $x_n = 2505491450622.0$, $q = 9217$.

W = 1 1 1557 2734 9790 9217
 17015 -170128944 -159462872 -82391683

Step 6. Use words 7 through 10 for x_n : $x_n = 510888888317.0$, $q = 1879$.

W = 1 1 1557 2734 9790 9217 1879
 17298 -172957850 -97765661

Now the result can be rounded to five digits and the value 1.1557273497909217 returned.

This example shows that the active part of W can be kept unnormalized. When x_n is computed using double-precision arithmetic, four words of W are essentially normalized during the calculation.

No attempt is made to check that the quotient digit q is correct. The next example shows what would happen if q were one too high in step 3 above. The first two steps are the same as before.

Step 3. Use $q = 2735$.

W = 1 1 1557 2735
 3239 -32387014 -35540953 -23839129 -1236220 0

Step 4. Use words 5 through 8 for x_n : $x_n = -56833369129.0$, $q = -210$.

W = 1 1 1557 2735 -210
 3406 -34032733 -22120909 -540170 -94920

Step 5. Use words 6 through 9 for x_n : $x_n = 2505491450622.0$, $q = 9217$.

W = 1 1 1557 2735 -210 9217
 8833 -88317403 -74873324 -77871683

Step 6. Use words 7 through 10 for x_n : $x_n = 510888888317.0$, $q = 1879$.

W = 1 1 1557 2735 -210 9217 1879
 8839 -88368302 -93245661

This shows that an incorrect q leads to unnormalized quotient digits later, so a final pass is needed to normalize W before returning the result. This gives the same value as in the first example.

Step 4 does the missing correction from step 3 at the same time as the normal multiply and subtract operation from step 4. The negative $q = -210 = -10,000 + 9790$ can be interpreted as subtracting 1 from the previous q while also using $q = 9790$ for the current step. The final normalization pass moves the -1 to the correct word.

3. DETAILS OF THE ALGORITHM

The algorithm will be given using Mathematica [9]. This version does not handle various special cases such as zero denominator, and has not been tuned for efficiency. Both arguments are assumed to be positive, with length at least five. No provision is made for normalizing the digits of W to avoid overflow. If rounding is to be fairly accurate, the base b should be at least 100. For a full implementation, see the Fortran subroutine in FMLIB 1.1.

```
RealW[ j_ ] := N[ ((W[[j-1]]*b + W[[j]])*b + W[[j+1]])*b +
  If[ Length[W]>=j+2 , W[[j+2]] , 0 ] ];
```

```

subtract[ q_ , d_ , ka_ , kb_ ] := Module[ { },
  Do[ W[[j]] = W[[j]] - q*d[[j-ka+2]] , {j,ka,kb} ];
  Return[ W ] ];
normalize[ ka_ , q_ ] := Module[ { },
  W[[ka]] = W[[ka]] + W[[ka-1]]*b;
  W[[ka-1]] = q;
  Return[ W ] ];
FinalNorm[ kb_ ] := Module[ {carry},
  Do[ carry = If[ W[[j]]<0 , Floor[ (-W[[j]]-1)/b ] + 1 ,
      If[ W[[j]]>=b , -Floor[ W[[j]]/b ] , 0 ] ];
      W[[j]] = W[[j]] + carry*b;
      W[[j-1]] = W[[j-1]] - carry
  , {j,kb,3,-1} ];
  Return[ W ] ];
divide[ n_ , d_ ] := Module[ {last,laststep,q,result,round,xd,xn},
  W = Table[ 0 , {Length[n]+4} ];
  W[[1]] = n[[1]] - d[[1]] + 1;
  Do[ W[[j+1]] = n[[j]] , {j,2,Length[n]} ];
  xd = N[ (d[[2]]*b + d[[3]])*b + d[[4]] + d[[5]]/b ];
  laststep = t + 2;
  Do[ xn = RealW[ step+2 ];
      q = Floor[ xn/xd ];
      last = Min[ step+t+1 , Length[W] ];
      W = subtract[ q , d , step+2 , last ];
      W = normalize[ step+2 , q ]
  , {step,1,laststep} ];
  W = FinalNorm[ laststep+1 ];
  laststep = If[ W[[2]]==0 , laststep , laststep-1 ];
  round = N[ W[[laststep+1]]/b ];
  W[[laststep]] = W[[laststep]] + If[ round>=0.5 , 1 , 0 ];
  result = If[ W[[2]]==0 , Table[ W[[j+1]] , {j,1,t+1} ] ,
      Table[ W[[j]] , {j,1,t+1} ] ];
  result[[1]] = If[ W[[2]]==0 , W[[1]]-1 , W[[1]] ];
  Return[ result ] ]

```

The base b , precision t , and working array W are global symbols in this version. The following command performs the division shown in the previous section:

```

b = 10000;
t = 5;
n = { 1, 3, 1415, 9265, 3589, 7932 };
d = { 1, 2, 7182, 8182, 8459, 0452 };
divide[ n , d ]

```

In practice, W may be of type integer or double precision; x_n and x_d are computed in double precision. Let T be the threshold below which integers can be represented exactly in double precision and rounding errors do not affect arithmetic operations with integer arguments and results. The base is chosen so that $b \geq 2$ and so that b^2 can be represented exactly in W . This implies $b^2 < T$.

Using four words of W to compute x_n gives good accuracy for q . Suppose n is the exact value of the active part of W , scaled so that x_n is an integer. Then during step k ,

$$n = W_{k+1} b^3 + W_{k+2} b^2 + W_{k+3} b + W_{k+4} + \frac{W_{k+5}}{b} + \dots$$

The goal is to approximate n/d , where d is the scaled denominator

$$d = d_2 b^2 + d_3 b + d_4 + \frac{d_5}{b} + \frac{d_6}{b^2} + \dots$$

Because d is normalized, $d \geq b^2$.

The quantities x_n and x_d are defined as the first four terms of these sums. We will bound the error in the estimate q .

Suppose that the value q used in step $k - 1$ is no more than one in error. Then the correct quotient digit $q^* = \lfloor n/d \rfloor$ for step k satisfies $-b \leq q^* \leq 2b$. So $q^* \leq n/d \leq q^* + 1$ implies $|n/d| \leq 2b + 1$.

After step k , $|W_i| < kb^2$. This means $|n - x_n| < 2kb$. Since $|d - x_d| < 1/b \leq d/b^3$, we have

$$\left| \frac{x_n}{x_d} \right| < \frac{|n| + 2kb}{d(1 - b^{-3})} < 2b + 3 + \frac{3k}{b}.$$

The error in q is

$$\begin{aligned} \left| \frac{n}{d} - \frac{x_n}{x_d} \right| &= \left| \frac{x_d(n - x_n) - x_n(d - x_d)}{dx_d} \right| \leq \left| \frac{n - x_n}{d} \right| + \left| \frac{x_n}{x_d} \right| \left| \frac{d - x_d}{d} \right| \\ &< \frac{2k}{b} + \frac{2k + 4}{b^2}. \end{aligned}$$

When a large base is used ($b^4 > T$), it often happens that x_n and x_d are not as accurate as assumed above, because of double-precision rounding errors. In this case, we can assume that the two values actually computed, y_n and y_d , agree with x_n and x_d to about double precision accuracy. Then

$$\left| \frac{x_n - y_n}{x_n} \right| < \frac{1}{b^2} \quad \text{and} \quad \left| \frac{x_d - y_d}{x_d} \right| < \frac{1}{b^2}.$$

Here the quotient digit used is $q = \lfloor y_n/y_d \rfloor$. Then

$$\left| \frac{y_n}{y_d} - \frac{x_n}{x_d} \right| \leq \left| \frac{x_n}{y_d} \right| \left| \frac{y_n - x_n}{x_n} \right| + \left| \frac{x_n}{y_d} \right| \left| \frac{y_d - x_d}{x_d} \right| < \frac{4}{b} + \frac{5k + 13}{b^2}.$$

The inequalities above assume $b \geq 2$, and the bound on the error in q could be made slightly smaller for the case where b is large. For the usual case where $b \geq 10^4$, having $|q - q^*| \leq 1$ in step $k - 1$ implies that the quotient digit used in step k is also no more than one in error.

This indicates that when b is large, using the unnormalized values in W gives an estimate q that is never more than one in error, and the probability of q being wrong is reasonably small.

The algorithm is still correct when the base is small, even though quotient digits may be used that are more than one in error. The final normalization pass takes

longer because more unnormalized digits are generated initially. Another possible strategy for dealing with a small base is to use more than four words to generate x_n and x_d .

The other practical matter that must be considered for a routine using this algorithm involves avoiding overflow or loss of precision. This can occur when the unnormalized intermediate results become too large. Suppose W is implemented as an integer array, $b = 10^4$, and the overflow threshold is $2^{31} - 1 = 2,147,483,647$. Since $21b^2$ will not overflow and each q is less than $2b$, 10 steps can be taken before normalizing the active part of W . As implemented in [7], the routine uses the digits being multiplied to get a sharper upper bound on the elements of W . The average q is about $b/2$, so there are approximately 40 steps between normalization.

In a more recent version of [7], W is a double precision array. The threshold for representing integers exactly in double precision is $T = 2^{53} - 1 \approx 9.01 \times 10^{15}$ on most 32-bit machines. In this case the commonly used base is 10^7 , so about 180 steps can be done before normalizing. Each time the digits of W are normalized to avoid overflow, the value of k is reset to 1 in the above bounds on the error in q . This means the accuracy of the quotient digits does not degrade much, even if high precision is used.

The main reason for guaranteeing that q is within one of the correct value is to avoid cases where the next q cannot correct an error in the previous step. Suppose in the double-precision case that $b = 10^7$ and q is 100 in error. In the next step, q will be about $100b$ in magnitude. Since the elements of d to be multiplied by q may be as large as $b - 1$, these products might exceed T . This could result in some products not being exact. To normalize W at this point would not solve the problem — a more complicated correction step would be needed. When smaller bases are used for the arithmetic, correspondingly larger errors in q can be tolerated by the algorithm.

4. CONCLUSION

This technique for eliminating most of the intermediate digit normalization operations can be used to perform multiple-precision division faster than previous methods. A routine implementing this method was compared to the division algorithm used in FMLIB 1.0, which normalizes several digits of the intermediate result at each step. Running on a 68040 Macintosh, the new version was 2.2 times faster at 40 significant digits, and 1.8 times faster at 500 digits. Similar speed improvements were found in testing on other types of computer.

The theoretical complexity for division is the same as that for multiplication, but in most multiple-precision packages division is much slower, even at high precision. In tests of this algorithm, the ratio of running times for division and multiplication ranged from 1.54 at 40 decimal digits to 1.19 at 1,000 digits. Since the intermediate results are not normalized often, the main difference in running time is due to the double precision calculation of x_n and q at each step. This means that at high precision, division is almost as fast as multiplication.

REFERENCES

1. D.H. Bailey, *Algorithm 719: Multiprecision translation and execution of FORTRAN programs*, ACM Trans. Math. Software **19** (1993), 288–319.
2. R.P. Brent, *A Fortran multiple-precision arithmetic package*, ACM Trans. Math. Software **4** (1978), 57–70.

3. D.E. Knuth, *The art of computer programming, Vol. 2: Seminumerical algorithms*, 2nd ed., Addison-Wesley, Reading, MA, 1981. MR **83i**:68003
4. E.V. Krishnamurthy and S.K. Nandi, *On the normalization requirement of divisor in divide-and-correct methods*, Comm. ACM **10** (1967), 809–813.
5. C.J. Mifsud, *A multiple-precision division algorithm*, Comm. ACM **13** (1970), 666–668.
6. D.A. Pope and M.L. Stein, *Multiple precision arithmetic*, Comm. ACM **3** (1960), 652–654. MR **22**:7277
7. D.M. Smith, *Algorithm 693: A Fortran package for floating-point multiple-precision arithmetic*, ACM Trans. Math. Software **17** (1991), 273–283.
8. M.L. Stein, *Divide-and-correct methods for multiple precision division*, Comm. ACM **7** (1964), 472–474.
9. S. Wolfram, *Mathematica: A system for doing mathematics by computer*, 2nd ed., Addison-Wesley, Redwood City, CA, 1991.

DEPARTMENT OF MATHEMATICS, LOYOLA MARYMOUNT UNIVERSITY, LOS ANGELES, CALIFORNIA 90045

E-mail address: `dsmith@lmumail.lmu.edu`