On the Design and Performance of Reliable Geometric Predicates using Error-free Transformations and Exact Sign of Sum Algorithms*

Marc Mörig† Stefan Schirra†

Abstract

We study the relevance of algorithms for exact computation of the sign of a sum of floating-point numbers and error-free transformations of arithmetic expressions on floating-point numbers for the design and implementation of low-dimensional geometric predicates. In a case study, we experimentally compare several implementations of planar orientation test and incircle tests that make use of such utilities.

1 Introduction

Floating-point computation is very fast, but also inherently inaccurate due to rounding errors. Naïvely applied, floating-point computation can cause fatal errors even with simple geometric algorithms [10], since floating-point based implementations of geometric predicates can produce wrong and inconsistent results. Most geometric predicates only involve sign computations for polynomial expressions in input coordinates or can at least be implemented like that. If all sign computations are correct, the predicate yields the correct result. The exact geometric computation paradigm now only asks for exactness in the sign computations, and not for computing exact values of the polynomial expressions. Configurations where a polynomial expression in a geometric predicate G evaluates to zero are called degenerate with respect to G.

Much progress has been made on the efficient implementation of reliable geometric predicates since the importance of the exact geometric computation paradigm has been widely acknowledged more than a decade ago. Efficient floating-point filters have been designed and combined to adaptive evaluation schemes for computing the sign of an arithmetic expression and are used in geometric software libraries like CGAL and LEDA. Floating-point filters compute a sign using fast floating-point arithmetic first and try to verify the computed sign, e.g. using an error bound computation. If the verification fails, they switch to some other method, for example, some exact arithmetic or a better filter, that uses higher precision in the floating-point computation or computes a better error bound. By cascading filters, sign computations can be made adaptive with respect to nearness to degeneracy. Prime examples of efficient cascaded floating-point filters are Shewchuk's predicates for orientation and incircle tests [18].

In the mid nineties Ratschek and Rokne presented an algorithm, called ESSA for exact sign of sum algorithm, to compute the sign of a sum of floating-point values exactly [14]. They advocated the use of their algorithm in computational geometry, see [15] for an overview. Using error-free transformations, any polynomial expression on floating-point numbers can be transformed into a sum of floating-point numbers. Thus an algorithm like ESSA is indeed a useful tool for the reliable implementation of a geometric predicate. Applications of summation algorithms in computational geometry are discussed by Graillat [6].

In Section 2, we look at some error-free transformations, especially those used in Shewchuk's predicates and ESSA. In Section 3 we briefly review ESSA and some more recent summation algorithms, that can be used to compute the exact sign of a sum, and discuss some modifications. Next, in Section 4, we illustrate the use of sign of sum algorithms and error-free transformations for the 2D incircle test. Finally, in Section 6, we report on the results of a case study, where we compare different reliable implementations of 2D incircle and orientation tests in CGAL's 2D Delaunay triangulation algorithm. Since initial floating-point filters for Delaunay predicates are already well studied [4, 5, 9, 18], we are mainly interested in nearly degenerate configurations. An input point set generator designed for this purpose is described in Section 5.

2 Error-Free Transformations

Error-free transformations transform an arithmetic expression involving floating-point numbers into a mathematically equivalent expression that is more suited for a particular purpose, e.g. sign computation. For example, $a + b$ can be transformed into $c^{\mathrm{hi}} + c^{\mathrm{lo}}$, such that $a \oplus b = c^{\mathrm{hi}}$ and $a+b = c^{\mathrm{hi}} + c^{\mathrm{lo}}$. Here we use $\oplus$, $\ominus$ and $\odot$ to denote floating-point addition, subtraction, and multiplication, respectively. Note that c^{lo} is the rounding error involved in computing $a \oplus b$. Efficient algorithms

*Supported by DFG grant SCHI 858/1-1

†Department of Simulation and Graphics, Faculty of Computer Science, Otto von Guericke University Magdeburg, Germany. moerig,stschirr at isg.cs.uni-magdeburg.de

for performing this transformation have been devised for IEEE 754 compliant arithmetic with exact rounding to nearest. The transformations are error-free unless overflow occurs. TWOSUM(a, b), due to Knuth [11], uses six floating-point additions and subtractions to perform this transformation, FASTTWOSUM(a, b), due to Dekker [2], requires $|a| \geq |b|$, but uses only three operations. Analogously, TWOPRODUCT(a, b) computes floating-point values $c^{\rm hi}$ and $c^{\rm lo}$ with $a \odot b = c^{\rm hi}$ and $a \cdot b = c^{\rm hi} + c^{\rm lo}$. TWOPRODUCT is due to Veltkamp and Dekker [2] and uses 17 floating-point operations, where 2×4 operations are required to split a and b respectively. Thus the cost for a sequence of TWOPRODUCT transformations can be reduced a bit if some numbers are involved in more than one product. The transformation is error-free, unless overflow or underflow occurs.

The transformations above use standard IEEE 754 floating-point operations only and do not require explicit access to mantissa or exponent of a floating-point number. They are key ingredients in Shewchuk's efficient adaptive orientation and incircle predicate implementations [18]. There are also error-free transformations that involve bit-manipulation and/or exponent extraction. For example, Ratschek and Rokne [14] transform $a - b$ for $2^{52}b \geq a \geq b > 0$ into $a' - b'$ where $a' = (a \ominus u)$ and $b' = (b \ominus u)$ and $u = \min\{a, 2^{\lceil \log_2 b \rceil}\}$.

3 Summation Algorithms

Accurate summation of floating-point numbers has always been in the focus of research on numerical computations [7]. In this section we shortly address algorithms used in our predicate implementations to compute the sign of a sum exactly.

ESSA iteratively performs error-free transformations on the largest positive and the smallest negative number in the current sum, thereby decreasing the sum of the absolute values of the summands. The iteration continues until the sum vanishes, or the largest positive number clearly dominates the sum of negative ones, or vice versa. In the original ESSA implementation [12], the error-free transformation described above is used. Both the set of positive summands and the set of negative summands are sorted initially. Apparently, it has not been observed so far, that creating a heap order is sufficient. In our implementation, we do not sort and use FASTTWOSUM for error-free transformation. Our modification, called revised ESSA later on, allows us to prove a better bound on the number of iterations, but there are also examples where the actual number of iterations grows.

SIGNK uses a compensated summation [8]: We add the summands one by one to a current approximation using TWOSUM. Besides an approximation a, this gives us a list of error terms $e_1, e_2, \ldots$ Using standard floating-point addition we sum the error terms and a in this order and compute an error bound as described by Ogita et al. [13]. If the error bound does not allow us to reliably determine the sign, we start over with a and the e_i, after eliminating zeros.

ACCSIGN is based on the ACCSUM algorithm by Rump et al. [16]. ACCSUM computes the value of a sum almost exactly rounded. It computes either the next larger or next smaller floating-point number. The algorithm uses error-free transformations, too, especially transformations based on splitting floating-point numbers, which is done without explicitly accessing mantissa or exponent. The final step of ACCSUM can be omitted in ACCSIGN since we are interested in the sign of the sum only.

4 Case Study Incircle Test

We illustrate the use of error-free transformations and corresponding sign of sum algorithms for the predicates of 2D Delaunay computation using CGAL's implementation based on the Delaunay hierarchy [1, 3]. In this section, we discuss the incircle predicate only, the orientation predicate is implemented analogously. The incircle predicate checks whether $s = (s_x, s_y)$ is contained in the circle through points $p = (p_x, p_y)$, $q = (q_x, q_y)$, and $r = (r_x, r_y)$. This is tantamount to computing the sign of the determinant

$$\begin{vmatrix} p_x & p_y & p_x^2 + p_y^2 & 1 \\ q_x & q_y & q_x^2 + q_y^2 & 1 \\ r_x & r_y & r_x^2 + r_y^2 & 1 \\ s_x & s_y & s_x^2 + s_y^2 & 1 \end{vmatrix}$$

Expanding the determinant results in a polynomial expression $p_x q_y r_x^2 + p_x q_y r_y^2 + \ldots$ consisting of 48 such subterms. By applying TWOPRODUCT six times, a single term is transformed into a sum of 8 floating-point numbers, resulting in a sum of 348 floating-point numbers. Let S_{348} be the corresponding arithmetic expression. Alternatively, we consider

$$\begin{vmatrix} p_x - s_x & p_y - s_y & (p_x - s_x)^2 + (p_y - s_y)^2 \\ q_x - s_x & q_y - s_y & (q_x - s_x)^2 + (q_y - s_y)^2 \\ r_x - s_x & r_y - s_y & (r_x - s_x)^2 + (r_y - s_y)^2 \end{vmatrix}$$

corresponding to a translation that moves the point s to the origin. Using error-free transformations like TWOSUM$(p_x, -s_x) = p_x^{\rm hi} + p_x^{\rm lo}$, we get

$$\begin{vmatrix} p_x^{\rm hi} + p_x^{\rm lo} & p_y^{\rm hi} + p_y^{\rm lo} & (p_x^{\rm hi} + p_x^{\rm lo})^2 + (p_y^{\rm hi} + p_y^{\rm lo})^2 \\ q_x^{\rm hi} + q_x^{\rm lo} & q_y^{\rm hi} + q_y^{\rm lo} & (q_x^{\rm hi} + q_x^{\rm lo})^2 + (q_y^{\rm hi} + q_y^{\rm lo})^2 \\ r_x^{\rm hi} + r_x^{\rm lo} & r_y^{\rm hi} + r_y^{\rm lo} & (r_x^{\rm hi} + r_x^{\rm lo})^2 + (r_y^{\rm hi} + r_y^{\rm lo})^2 \end{vmatrix}$$

Note that some of the $\square^{\rm lo}$ values might be zero, because the corresponding subtraction is exact. Since by Sterbenz lemma [19] a floating-point subtraction is exact if the operands differ by a factor of two in size at most, this

	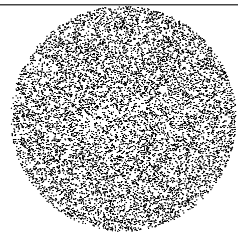 d=0, $r_{\max}$=0, f=0%		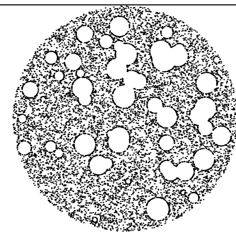 d=63, $r_{\max}$=0.11, f=25%		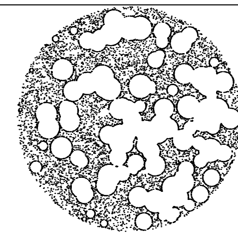 d=125, $r_{\max}$=0.11, f=50%		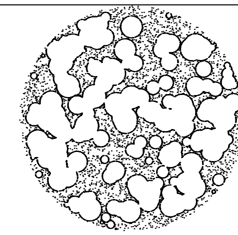 d=188, $r_{\max}$=0.11, f=75%	
	Intel	Sun	Intel	Sun	Intel	Sun	Intel	Sun
org. ESSA S_{348}	7259	53005	7912	62692	8615	72819	9345	83577
rev. ESSA S_{348}	2028	5524	2199	5966	2392	6436	2588	6922
ACCSIGN S_{348}	598	2894	615	2968	636	3076	659	3195
SIGNK S_{348}	504	2749	505	2752	508	2763	510	2775
org. ESSA $S_{96:1152}$	1303	5274	1356	5875	1407	6431	1457	6999
rev. ESSA $S_{96:1152}$	526	1676	539	1726	552	1764	566	1813
ACCSIGN $S_{96:1152}$	260	1342	266	1366	272	1393	279	1433
SIGNK $S_{96:1152}$	208	1044	209	1047	210	1050	213	1057
Shewchuk non-adapt.	546	1943	546	1949	547	1966	552	1985
CGAL	36	122	165	688	266	1116	344	1429
ff org. ESSA S_{348}	31	109	511	5665	885	9884	1177	12965
ff rev. ESSA S_{348}	32	110	162	455	264	719	343	914
ff ACCSIGN S_{348}	31	108	56	222	77	315	93	385
ff SIGNK S_{348}	31	108	49	203	63	279	75	336
ff org. ESSA $S_{96:1152}$	31	108	101	588	153	928	193	1177
ff rev. ESSA $S_{96:1152}$	32	108	55	180	74	237	88	279
ff ACCSIGN $S_{96:1152}$	31	110	43	159	51	200	59	233
ff SIGNK $S_{96:1152}$	31	109	39	148	46	182	53	211
ff Shew. non-adapt.	31	108	49	172	62	220	74	259
Shewchuk adaptive	34	106	52	153	67	198	81	234

Table 1: Times in ms for computing the Delaunay triangulation of 10 000 points, averaged over 25 input sets.

is not unlikely! We transform $(p_x^{\mathrm{hi}}+p_x^{\mathrm{lo}})^2+(p_y^{\mathrm{hi}}+p_y^{\mathrm{lo}})^2 = p_x^{\mathrm{hi}}p_x^{\mathrm{hi}} + 2p_x^{\mathrm{hi}}p_x^{\mathrm{lo}} + p_x^{\mathrm{lo}}p_x^{\mathrm{lo}} + p_y^{\mathrm{hi}}p_y^{\mathrm{hi}} + 2p_y^{\mathrm{hi}}p_y^{\mathrm{lo}} + p_y^{\mathrm{lo}}p_y^{\mathrm{lo}}$ and analogous terms into a sum using TWOPRODUCT, where a product incurring $\Box^{\mathrm{lo}}$ is computed only if $\Box^{\mathrm{lo}} \neq 0$. The resulting sum has between 4 and 12 summands. Then we use cofactor expansion on the last column. We transform $(q_x^{\mathrm{hi}} + q_x^{\mathrm{lo}})(r_y^{\mathrm{hi}} + r_y^{\mathrm{lo}}) - (q_y^{\mathrm{hi}} + q_y^{\mathrm{lo}})(r_x^{\mathrm{hi}} + r_x^{\mathrm{lo}})$ and resembling expressions into a sum with between 4 and 16 summands. Each remaining product of two sums is then transformed into a single sum by applying TWOPRODUCT to each pair of summands, resulting in a sum with between 32 to 384 summands. Overall, we get an expression with between 96 to 1152 summands, called $S_{96:1152}$. We combine both expressions S_{348} and $S_{96:1152}$ with the sign of sum algorithms listed in the previous section, optionally with an initial float-filter from CGAL.

We compare the resulting incircle predicate implementations with Shewchuk's adaptive incircle code [17, 18]. Shewchuk uses error-free transformations as well to transform the determinant into a sum. However, he uses a clever staged evaluation strategy that computes approximations with increasing accuracy, using so-called floating-point expansions. Comparison with an error bound is used to verify the sign of the current approximation. At each stage, some computations from previous stages are reused, leading to an adaptive predicate. Shewchuk also provides non-staged and hence non-adaptive predicate implementations. Furthermore, we compare with CGAL's *exact predicates inexact constructions kernel*. CGAL first uses a semi-static filter, and then a dynamic filter [4]. If both fail, CGAL's exact number type *MP_Float* is used. It is this well-engineered cascaded filter that we use in the initial phase of the filtered versions (ff . . .) of our predicate implementations.

5 Test Data Generator

We are most interested in the performance of implementations of the incircle predicate for difficult instances, since state-of-the-art float-filters can be used to solve easy cases efficiently. In order to force a Delaunay triangulation algorithm to perform more difficult tests the generated test data contains points almost on a circle with no other points in its interior: First, we create a set $\mathcal{D}$ of d disks with a random radius $0 < r \leq r_{\max}$ and place a certain percentage f of the points (almost) on the boundary of their union, $\mathrm{bd}(\cup\mathcal{D})$, cf. Table 1. Next, the remaining points are generated uniformly in the complement of the disks. All points are generated inside the unit circle. In order to get nearly degenerate point sets we use exact arithmetic to compute a point on a circular arc of $\mathrm{bd}(\cup\mathcal{D})$ and then round it to a nearby floating-point point closest to the circular arc.

6 Results and Conclusions

We run experiments on both a PC with an Intel Core 2 Duo T5500 processor with 1.66 Ghz, using `g++ 4.1` and CGAL 3.2.1, and on a Sun Blade Station 1000 with 0.9 Ghz, using `g++ 3.3.3` and CGAL 3.2. Average running times for 25 input sets with 10 000 points each are shown in Table 1. The timings for the most competitive algorithms are visualized in Figure 1. Interestingly, the rankings for the two platforms differ. The experiments

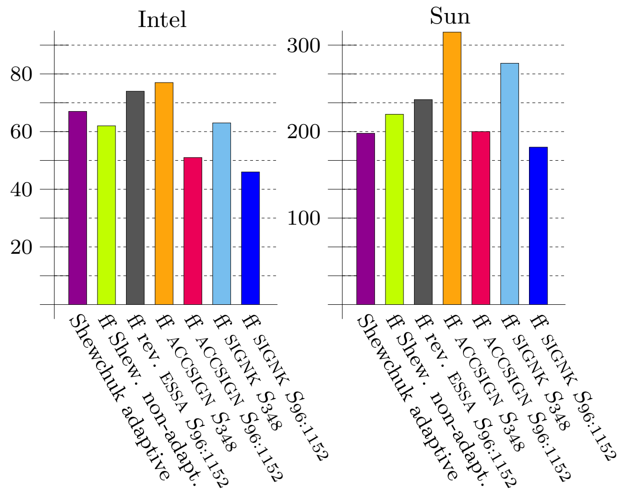

Figure 1: Times for 10 000 points, with 50% points on the boundary.

show that an initial filter step is absolutely necessary, because otherwise the predicates based on the sign of sum algorithms, which are apparently non-adaptive, are too slow for non-degenerate point sets. Original ESSA is obviously not competitive whereas our revised version comes close to the best ones, especially on the Sun platform. Filtered SIGNK for $S_{96:1152}$ performs best on both platforms. On Sun, Shewchuk's adaptive predicates are next, while they are at rank 4 on Intel. We observed, that the predicates based on $S_{96:1152}$ perform relatively better on larger input sets. Because of the locality of the Delaunay triangulation algorithm and Sterbenz' lemma, the fraction of sign evaluations of type $S_{96:1152}$ where the actual number of summands is small increases with the size of the input point sets, as illustrated in the figure below.

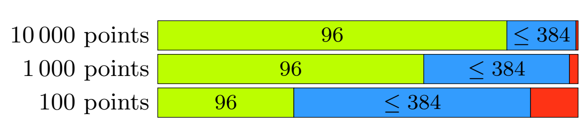

The figure shows the fraction of sign evaluations of type $S_{96:1152}$ with 96, with at most 348, and with more than 348 summands for the generated input sets. Overall, filtered SIGNK for $S_{96:1152}$ is always a good choice, only for small points sets of 100 points Shewchuk's adaptive predicates perform somewhat better on Sun.

References

[1] CGAL, Computational Geometry Algorithms Library. `http://www.cgal.org`.

[2] T. J. Dekker. A floating-point technique for extending the available precision. *Num. Math.*, 18(2):224–242, 1971.

[3] O. Devillers. The Delaunay hierarchy. *Int. J. of Found. of Comput. Sci.*, 13(2):163–180, 2002.

[4] O. Devillers and S. Pion. Efficient exact geometric predicates for Delaunay triangulations. In *ALENEX 2003*, pages 37–44.

[5] S. Fortune. Numerical stability of algorithms for 2-d Delaunay triangulations. *Int. J. of Comput. Geom. and Appl.*, 5:193–213, 1995.

[6] S. Graillat. Applications of fast and accurate summation in computational geometry. Research Report 03, Laboratoire LP2A, Université de Perpignan, 2005.

[7] N. J. Higham. *Accuracy and Stability of Numerical Algorithms*. SIAM, 2. edition, 2002.

[8] W. Kahan. Further remarks on reducing truncation errors. *Comm. of the ACM*, 8(1):40, 1965.

[9] M. Karasick, D. Lieber, and L. R. Nackman. Efficient delaunay triangulation using rational arithmetic. *ACM Trans. Graph.*, 10(1):71–91, 1991.

[10] L. Kettner, K. Mehlhorn, S. Pion, S. Schirra, and C.-K. Yap. Classroom examples of robustness problems in geometric computations. In *ESA 2004*, LNCS 2321, pages 702–713.

[11] D. E. Knuth. *Seminumerical Algorithms*, volume 2 of *The Art Of Computer Programming*. Addison-Wesley, 3. edition, 1997.

[12] G. Mackenbrock, H. Ratschek, and J. G. Rokne. Experimental reliable code for 2d convex hull construction. `http://pages.cpsc.ucalgary.ca/~rokne/convex/`.

[13] T. Ogita, S. M. Rump, and S. Oishi. Accurate sum and dot product. *SIAM J. on Sci. Comput.*, 26(6):1955–1988, 2005.

[14] H. Ratschek and J. G. Rokne. Exact computation of the sign of a finite sum. *Appl. Math. Computation*, 99(2-3):99–127, 1999.

[15] H. Ratschek and J. G. Rokne. *Geometric computations with interval and new robust methods: applications in computer graphics, GIS and computational geometry*. Horwood Publishing, Ltd., 2003.

[16] S. M. Rump, T. Ogita, and S. Oishi. Accurate floating-point summation. Technical Report 05.1, Faculty of Information and Communication Science, Hamburg University of Technology, 2005.

[17] J. R. Shewchuk. Companion web page to [18]. `http://www.cs.cmu.edu/~quake/robust.html`.

[18] J. R. Shewchuk. Adaptive precision floating-point arithmetic and fast robust geometric predicates. *Discrete Comput. Geom.*, 18(3):305–363, 1997.

[19] P. H. Sterbenz. *Floating-Point Computation*. Prentice-Hall, 1974.