Reliable and Efficient Computational Geometry via Controlled Perturbation*

Kurt Mehlhorn, Ralf Osbild, and Michael Sagraloff

Max-Planck-Institut für Informatik, Stuhlsatzenhausweg 85, 66123 Saarbrücken, Germany.

Abstract. Most algorithms of computational geometry are designed for the Real-RAM and non-degenerate inputs. We call such algorithms idealistic. Executing an idealistic algorithm with floating point arithmetic may fail. Controlled perturbation replaces an input x by a random nearby $\widetilde{x}$ in the δ-neighborhood of x and then runs the floating point version of the idealistic algorithm on $\widetilde{x}$. The hope is that this will produce the correct result for $\widetilde{x}$ with constant probability provided that δ is small and the precision L of the floating point system is large enough. We turn this hope into a theorem for a large class of geometric algorithms and describe a general methodology for deriving a relation between δ and L. We exemplify the usefulness of the methodology by examples.

Most algorithms of computational geometry are designed for Real-RAMs and non-degenerate inputs. A Real-RAM computes with real numbers in the sense of mathematics. The notion of degeneracy depends on the problem; examples are collinear or cocircular points or three lines with a common point. We call an algorithm designed under the two simplifying assumptions an *idealistic algorithm*. Implementations have to deal with the precision problem (caused by the Real-RAM assumption) and the degeneracy problem (caused by the non-degeneracy assumption). The *exact computation paradigm* [10, 9, 3, 14, 12, 13] addresses the precision problem. It proposes to implement a Real-RAM tuned to geometric computations. The degeneracy problem is addressed by reformulating the algorithms so that they can handle all inputs. This may require non-trivial changes. The approach is followed in systems like LEDA and CGAL. Halperin et al. [5, 7, 6] proposed *controlled perturbation* to overcome both problems. The idea is to solve the problem at hand not on the input given but on a nearby input. The perturbed input is carefully chosen, hence the name *controlled perturbation*, so that it is non-degenerate and can be handled with approximate arithmetic. They applied the idea to three problems (computing polyhedral arrangements, spherical arrangements, and arrangements of circles) and showed that variants of the respective idealistic algorithms can be made to work. Funke et al. [4] extended their work and showed how to use controlled perturbation in the context of randomized algorithms, in particular randomized incremental constructions, and designed specific schemes for planar Delaunay triangulations and convex hulls and Delaunay triangulations in arbitrary dimensions. We extend their work further. We prove that controlled perturbation and guarded tests are a general conversion strategy for a wide class of geometric algorithms; the papers

* Partially supported by the IST Programme of the EU under Contract No IST-006413, Algorithms for Complex Shapes (ACS).

cited above hint at this possibility but do not prove it. Moreover, we develop a general methodology for analyzing controlled perturbation, in particular, for deriving quantitative relations between the amount of perturbation and the precision of the approximate arithmetic.

1 Controlled Perturbation (Review from [4])

Geometric algorithms branch on geometric predicates. Typically, geometric predicates can be expressed as the sign of an arithmetic formula E. For example, the *orientation predicate* for $d+1$ points in $\mathbb{R}^d$ is given by the sign of a $(d+1)\times(d+1)$ determinant: the determinant has one row for each point and the row for a point contains the coordinates of the point followed by the entry 1 and evaluates to zero iff the $d+1$ points lie in a common hyperplane. This is considered a degeneracy.

When evaluating an arithmetic formula E using floating-point arithmetic, round-off error occurs which might result in the wrong sign being reported. If this stays undetected, the program may enter an illegal state and disasters may happen, see [11] for some instructive examples. In order to guard against round-off errors, we postulate the availability of a predicate $\mathcal{G}_E$ with the following *guard property: If $\mathcal{G}_E$ evaluates to true when evaluated with floating point arithmetic, the floating point evaluation (fp-evaluation) of E yields the correct sign.* In an idealistic algorithm A we now guard every sign test by first testing the corresponding guard. If it fails, we abort. We call the resulting algorithm a *guarded algorithm* and use A_g to denote it.

The controlled perturbation version of idealistic algorithm A is as follows: Let δ be a positive real. On input x, we first choose a δ-perturbation $\widetilde{x}$ of x and then run the guarded algorithm A_g on $\widetilde{x}$. If it succeeds, fine. If not, repeat. What is a δ-perturbation? A δ-perturbation of a point is a random point in the δ-cube (or δ-ball) centered at the point and for a set of points a δ-perturbation is simply a δ-perturbation of each point in the set. For more complex objects, alternative definitions come to mind, e.g., for a a circle one may want to perturb the center or the center and the radius. The goal is now to show experimentally and/or theoretically that A_g has a good chance of working on a δ-perturbation of any input and a small value of δ. More generally, one wants to derive a relation between the precision L of the floating point system (= length of the mantissa), a characteristic of the input set, e.g., the number of points in the set and an upper bound on the maximal coordinate of any point in the input, and δ. Halperin et al. have done so for arrangements of polyhedral surfaces, arrangements of spheres, and arrangements of circles and Funke et al. have done so for Delaunay diagrams and convex hulls in arbitrary dimensions.

We want to stress that a guarded algorithm can be used without any analysis. Suppose we want to use it with a certain δ. We execute it with a certain precision L. If it does not succeed, we double L and repeat. *Our main result states that this simple strategy terminates for a wide class of geometric algorithms. Moreover, it gives a quantitative relation between δ and L and characteristic quantities of the instance.*

Guard predicates must be safe and should be effective, i.e., if a guard does not fire, the approximate sign computation must be correct, and guards should not fire too often unnecessarily. It is usually difficult to analyze the floating point evaluation of $\mathcal{G}_E$

E	$\widetilde{E}$	$\widetilde{E_{\sup}}$	ind_E
$c = const$	c	$\|c\|$	0
$x+y$ or $x-y$	$\widetilde{x} \oplus \widetilde{y}$ or $\widetilde{x} \ominus \widetilde{y}$	$\widetilde{x_{\sup}} \oplus \widetilde{y_{\sup}}$	$1+\max(\text{ind}_x, \text{ind}_y)$
$x \cdot y$	$\widetilde{x} \odot \widetilde{y}$	$\widetilde{x_{\sup}} \odot \widetilde{y_{\sup}}$	$1+\text{ind}_x+\text{ind}_y$
$x^{1/2}$	$\sqrt{\widetilde{x}}$	$\begin{cases} (\widetilde{x_{\sup}} \oslash \widetilde{x}) \odot \sqrt{\widetilde{x}} & \text{if } \widetilde{x} > 0 \\ \sqrt{\widetilde{x_{\sup}}} \odot 2^{p/2} & \text{if } \widetilde{x} = 0 \end{cases}$	$1+\text{ind}_x$

Table 1. Rules for calculating error bounds. $\oplus$, $\ominus$, $\odot$, $\oslash$, and $\sqrt{}$ stand for floating point addition, subtraction, multiplication, division, and square-root, respectively.

directly. For the purpose of the analysis, we therefore postulate the existence of a *bound predicate* $\mathscr{B}_E$ with the property: *If* $\mathscr{B}_E$ *holds,* $\mathscr{G}_E$ *evaluates to true when evaluated with floating point arithmetic.* We next give some concrete examples for guard and bound predicates.

When E is evaluated by a straight-line program, it is easy to come up with suitable predicates $\mathscr{G}_E$ and $\mathscr{B}_E$ using forward error analysis. For example, the rules in Table 1 ([1]) recursively define two quantities $\widetilde{E_{\sup}}$ and ind_E for every arithmetic expression E such that $|E - \widetilde{E}| \leq B_E := \widetilde{E_{\sup}} \cdot \text{ind}_E \cdot 2^{-L}$ where $\widetilde{E}$ denote the value of E computed with floating point arithmetic and L denotes the mantissa length of the floating-point system. (i.e. $L = 52$ for IEEE doubles). We can then use $\mathscr{G}_E \equiv \left(|\widetilde{E}| > B_E\right)$ and $\mathscr{B}_E \equiv (|E| > 2B_E)$, where $\mathscr{B}_E$ is valid since it guarantees that $|\widetilde{E}| = |E| - |E - \widetilde{E}| > 2B_E - B_E = B_E$ by the inverse triangle inequality. For the orientation test of three points in the plane, one obtains $B_{orient} = 24 \cdot M^2 2^{-L}$ and for the incircle test of four points in the plane, one obtains $B_{incircle} = 432 \cdot M^4 2^{-L}$. In both cases, it is assumed that all point coordinates are bounded by M in absolute value.

We assume for this paper that input values are bounded by M in absolute value and that bound predicates are of the form $c_E M^{e_E} 2^{-L}$ where c_E and e_E are constants depending on the predicate expression E. If E is a polynomial, e_E is the degree of the polynomial.

2 The Class of Algorithms

Our result applies to algorithms which can be viewed as decision trees. There is a decision tree T_n for each input size n. We assume that the input consists of a set of n points with coordinates bounded by M in absolute value. Boundedness is essential in some of our arguments and we leave it as a challenge to remove this restriction. The internal nodes of the decision tree are labelled by predicate evaluations $\text{sign} f(x_{i_1}, \ldots, x_{i_k})$ where f stems from a fixed finite set of real-valued functions (for example, orientation of three points or the incircle test of four points) and the x_{i_j} are input points. The tree is ternary and branches according to the sign of f. Observe that predicates can only be applied to input points and not to computed points. This restriction can be relieved somewhat. For example, if the input consists of a set of line segments, each specified by a pair of points, then a predicate applied to an intersection point of two segments is easily reduced to a more complex predicate involving only input points. What predicate functions are allowed? We require that the functions f fulfil the postulates set forth in Section 4.

Many algorithms of computational geometry are within the model, e.g., Delaunay diagram and Voronoi diagram computations, convex hulls, line arrangements, ... It is important to understand the limitations of the model. Gaussian Elimination for $n \times n$ matrices is outside the model since it tests the sign of expressions depending on all n^2 matrix entries. So the number of predicate functions is infinite and their arity is not bounded. Observe however, that Gaussian elimination on $d \times d$ matrices used in an algorithm to compute convex hulls of n points in $\mathbb{R}^d$ is within the model as d does not depend on the input size. Algorithms whose running time depends on actual point coordinates and not just on the number of points are also outside the model. It is the subject of further work to weaken this restriction.

3 The Basic Idea

We concentrate on a single predicate, say $P(x_1, \ldots, x_k) = \text{sign} f(x_1, \ldots, x_k)$, of k points in the plane. The treatment readily generalizes to points in higher dimensions. Forward error analysis gives us an expression B_f which upper bounds the error in the evaluation of f. For this extended abstract, we assume that B_f is a constant as discussed above. We can make B_f arbitrarily small by increasing the precision L of the floating point system.

We want to prove a result of the following form: If each coordinate of any input point is modified by a random number in $[-\delta, +\delta]$ and the program is executed with sufficiently high floating point precision L, the guarded program succeeds with probability at least 1/2. It is clear that such a result is true if f is continuous and the zero set of f is lower-dimensional, because then the set of k-tuples for which $|f| < 2B_f$ is within a small neighborhood of the zero set. In order to obtain a quantitative relation between δ and L, we need to estimate the maximal volume of the set of k-tuples with $|f| < 2B_f$ within an arbitrary axis-oriented 2δ-cube.

We suggest a general approach for deriving such estimates exploiting the fact that functions f underlying geometric predicates have structure. As a first step, we split the arguments of f into $k-1$ points $\mathbf{x}$ and a single point x. We write $f(\mathbf{x}, x)$ even if x is not the last argument of f. We consider the points in $\mathbf{x}$ fixed and the point x variable. Geometric predicates can usually be interpreted as follows: $\mathbf{x}$ defines a partition of the plane into regions and P tells the location of x with respect to this partition. P returns zero if x lies on a region boundary, $+1$ if x lies in the positive regions, and -1 is x lies in the negative regions. We use $C_{\mathbf{x}} = \{x : f(\mathbf{x}, x) = 0\}$ to denote the zero set of f and call it the *curve of degeneracy*. If $\lambda x. f(\mathbf{x}, x)$ is identically zero[1], we call $\mathbf{x}$ *degenerate*. We call it *regular*, otherwise.

Some examples: (1) in the orientation predicate of three points p, q, and r, the first two points ($\mathbf{x}$ comprises p and q) define an oriented line $\ell(p, q)$ and $orient(p, q, r)$ tells

[1] We use the notation $\lambda x. f(\mathbf{x}, x)$ to emphasize that we view f as a function of x and keep $\mathbf{x}$ fixed.

the location of r (x corresponds to r) with respect to this line. The curve of degeneracy is the line $\ell(p,q)$ if $p \neq q$. The pair (p,q) is degenerate if $p = q$. (2) in the side of circle predicate of four points p, q, r, and s, the first three points define an oriented circle $C(p,q,r)$ and $soc(p,q,r,s)$ tells the location of s with respect to this circle. The curve of degeneracy is $C(p,q,r)$ and the triple (p,q,r) is degenerate if it contains equal points. (3) in the side-of-wedge predicate of four points p, q, r, and s, the first three points define a wedge with boundaries $\ell(p,q)$ and $\ell(p,r)$ and $sow(p,q,r,s)$ tells the location of s with respect to this wedge. The curve of degeneracy is $\ell(p,q) \cup \ell(p,r)$ and the triple (p,q,r) is degenerate if either $q = p$ or $r = p$.

The function $\lambda x. f(\mathbf{x},x)$ is zero on the curve of degeneracy. It will be small near it and larger further away, i.e., $|f(\mathbf{x},x)|$ measures, in some sense, locally the distance of x from the curve of degeneracy.

In our examples this is quite explicit: (1) $orient(p,q,r) = \mathrm{sign} f_o(p,q,r)$ where[2] $|f_o(p,q,r)| = dist(p,q) \cdot dist(r,l(p,q))$, (2) $soc(p,q,r,s) = \mathrm{sign} f_{soc}(p,q,r,s)$ where[3] $|f_{soc}(p,q,r,s)| \geq (1/2) dist(p,q) dist(p,r) dist(q,r) dist(C,s)$ and C denotes the circle or line defined by the first three points, and finally (3) $sow(p,q,r,s) = \mathrm{sign} f_{sow}(p,q,r,s)$ where $|f_{sow}(p,q,r,s)| = |f_o(p,q,s) \cdot f_o(p,r,s)| = dist(p,q) \cdot dist(p,r) \cdot dist(s,\ell(p,q)) \cdot dist(s,\ell(p,r)) \geq dist(p,q) \cdot dist(p,r) \cdot dist(s,\ell(p,q) \cup \ell(p,r))^2$.

Assume we have a function $g(\mathbf{x},d)$ such that $|f(\mathbf{x},x)| \geq g(\mathbf{x}, dist(x,C_\mathbf{x})) \geq 0$ that is non-zero if $dist(x,C_\mathbf{x}) > 0$, i.e., we bound $f(\mathbf{x},x)$ from below by a function in $\mathbf{x}$ and the distance of x from the curve of degeneracy. The requirement $|f(\mathbf{x},x)| \geq 2B_f$ would then translate into the condition $g(\mathbf{x}, dist(C_\mathbf{x},x)) \geq 2B_f$, i.e., if x lies outside a certain tubular neighborhood of the curve of degeneracy $C_\mathbf{x}$, $|f(\mathbf{x},x)|$ is guaranteed to be at least $2B_f$. The width of the tubular region is related to the growth of g and depends on $\mathbf{x}$. What can we say about the growth of g?

Again it is useful to consider our examples. For the orientation predicate, we have $g(p,q,d) = dist(p,q) \cdot d$ and so g grows linearly in d with slope $dist(p,q)$, for the side-of-circle predicate, we have $g(p,q,r,d) \geq dist(p,q) dist(p,r) dist(q,r) \cdot d$ and so g grows at least linearly in d with slope $dist(p,q) dist(p,r) dist(q,r)$, and for the side-of-wedge predicate, we have $g(p,q,r,d) \geq dist(p,q) \cdot dist(p,r) \cdot d^2$ and so g grows at least quadratically in d with factor $dist(p,q) \cdot dist(p,r)$. The slope (factor) is zero for degenerate $\mathbf{x}$ ($p = q$ for the orientation predicate, $|\{p,q,r\}| \leq 2$ for the side-of-circle predicate, and $p \in \{q,r\}$ for the in-wedge-predicate) and grows in the distance of $\mathbf{x}$ from degeneracy. We want to guarantee that the slope (factor) has a certain guaranteed size because this allows us to control the width of the forbidden region for x.

So we proceed as follows. We fix the width of the forbidden region for x at some value γ and then study the function $g(\mathbf{x},\gamma)$. We study the conditions on $\mathbf{x}$ guaranteeing $g(\mathbf{x},\gamma) \geq 2B_f$. Now $g(\mathbf{x},\gamma)$ has one less argument and so continuing in this way k times, we arrive at a trivial case. The non-trivial details are given in the next section.

[2] f_o is the value of a 3×3 determinant. The value of the determinant is twice the signed area of the triangle formed by the three points which in turn is the distance of the first two points times the distance of the third point from the line through the first two points.

[3] We are going to prove this in section 5.

Let us consider our examples: In all three examples the perturbation must guarantee that points have a certain minimum distance. This will guarantee that $dist(p,q)$, $dist(p,q) \cdot dist(p,r)$, and $dist(p,q)dist(p,r)dist(q,r)$ have certain minimum values.

4 The General Scheme

We concentrate on a single predicate f of k point variables.

Requirement 1 *f is continuous and f is not identically zero.*

In order to apply the scheme, one needs to determine a family of non-negative continuous functions f_s, one for every sequence $s = (s_0, s_1, \ldots, s_{\ell-1})$ with $1 \leq \ell \leq k$ and $s_j \in [1 \ldots k-j]$ for $0 \leq j \leq \ell-1$; s describes the order in which we eliminated variables, we first eliminated the s_0-th variable of a k-argument function, then the s_1-th variable of a $k-1$-argument function, and so on. The function f_s depends on $k-\ell$ point variables. For the empty sequence ε, we set $f_\varepsilon = |f|$. One also needs to fix positive constants γ_ℓ. Consider a fixed s with $\ell := |s| < k$; f_s is a function of $k-\ell$ point variables. Let $h \in [1 \ldots k-\ell]$ be arbitrary and let $t = s \circ h$. We use x to denote the h-th variable of f_s and $\mathbf{x}$ to denote the remaining $k-\ell-1$ variables. We write $f_s(\mathbf{x}, x)$ instead of $f_s(\mathbf{x}', x, \mathbf{x}'')$ where $\mathbf{x}'$ comprises the first $h-1$ arguments and $\mathbf{x}''$ comprises the last $k-\ell-h$ arguments. For each $\mathbf{x}$, let $C^t_{\mathbf{x}} = \{x : f_s(\mathbf{x}, x) = 0\}$. We call $C^t_{\mathbf{x}}$ the *curve of degeneracy*. If $C^t_{\mathbf{x}} = \emptyset$, we set $C^t_{\mathbf{x}}$ to an arbitrary singleton set for purely technical reasons. We call $\mathbf{x}$ *degenerate* if $\lambda x. f_s(\mathbf{x}, x)$ is identically zero and *regular* otherwise.

We next define the lower bound function. Let $U_d = \{x \in U : dist(x, C^t_{\mathbf{x}}) \geq d\}$ and let d_0 be maximal such that U_d is non-empty. Define $g_t(\mathbf{x}, d) := \min_{x \in U_d} f_s(\mathbf{x}, x)$ for $0 \leq d \leq d_0$ and $g_t(\mathbf{x}, d) := g_t(\mathbf{x}, d_0)$ for $d \geq d_0$.

Lemma 1. *The function $g_t(\mathbf{x}, d)$ is non-decreasing in its second argument, $g_t(\mathbf{x}, d) > 0$ for $d > 0$ if $\mathbf{x}$ is regular, and $g_t(\mathbf{x}, d) = 0$ for all d if $\mathbf{x}$ is degenerate.*

Proof. If $\mathbf{x}$ is degenerate, $\lambda x. f_s(\mathbf{x}, x)$ is identically zero and hence $g_t(\mathbf{x}, d) = 0$ for all d. If $\mathbf{x}$ is regular, $C_{\mathbf{x}}$ is a closed proper subset of U and hence $d_0 > 0$. Also U_d is a closed non-empty subset of U for $0 < d \leq d_0$. Since f_s is continuous, $\inf_{x \in U_d} f_s(\mathbf{x}, x)$ is attained for a point $x \in U_d$ and $f_s(\mathbf{x}, x) > 0$. Thus $g_t(\mathbf{x}, d) > 0$.

Requirement 2 *$f_t(\mathbf{x})$ is a continuous function with $0 \leq f_t(\mathbf{x}) \leq g_t(\mathbf{x}, \gamma_\ell)$ and $f_t(\mathbf{x}) = 0$ iff $\mathbf{x}$ is degenerate.*

In our applications, $g_t(\mathbf{x}, \gamma_\ell)$ is continuous and we may choose $f_t(\mathbf{x}) = g_t(\mathbf{x}, \gamma_\ell)$. Allowing an inequality, gives additional flexibility. However, there are situations where $g_t(\mathbf{x}, \gamma_\ell)$ is not continuous.

Lemma 2. *If $|s| = k$, f_s is a positive constant.*

Proof. Let $(x_1, \ldots, x_k) \in U^k$ be such that $f(x_1, \ldots, x_k) \neq 0$. We may assume without loss of generality that $s = (k, k-1, \ldots, 1)$, i.e., we remove the arguments from the end. We prove $f_{(k,k-1,\ldots,k-i+1)}(x_1, \ldots, x_{k-i}) \neq 0$ by induction on i. For $i = 0$, there is nothing to prove. So assume $i \geq 1$. We have $f_{(k,k-1,\ldots,k-i+2)}(x_1, \ldots, x_{k-i}, x_{k-i+1}) \neq 0$ by induction hypothesis. So, $\mathbf{x} = (x_1, \ldots, x_{k-i})$ is regular and hence $f_{(k,k-1,\ldots,k-i+1)}(x_1, \ldots, x_{k-i}) > 0$.

The Perturbation: Our input are points $q_1, q_2, \ldots$. We move each q_i to a random point p_i in the δ-cube centered at q_i. Assume that we have already chosen p_1 to p_{n-1} and that the following *perturbation property* (PP) holds true: for every sequence s, $\ell := |s|$, and every tuple of distinct indices j_1 to $j_{k-\ell}$ in $[1..n-1]$: $f_s(p_{j_1}, \ldots, p_{j_{k-\ell}}) \geq 2B_f$. For $n = 1$, the conditions are vacuously true if $\ell < k$. For $\ell = k$, f_s is a positive constant and hence by making the mantissa length L large enough, we can satisfy all conditions.

Requirement 3 *Precision L is large enough so that $2B_f \leq f_s$ for all s with $|s| = k$.*

We now choose p_n.

Lemma 3. *If for all ℓ, $0 \leq \ell < k$, any t with $|t| = \ell + 1$, and any tuple of distinct indices j_1 to $j_{k-\ell-1}$ in $[1..n-1]$ and $\mathbf{p} = (p_{j_1}, \ldots, p_{j_{k-\ell-1}})$, p_n does not lie in the γ_ℓ-neighborhood of $C^t_{\mathbf{p}}$, (PP) holds for n. Moreover, if these neighborhoods together cover at most a fraction $1/(2n)$ of the δ-cube centered at q_n, the precondition fails with probability at most $(1/2n)$.*

Proof. Consider the application of any f_s to a $k-\ell$ tuple of perturbed points. If p_n is not among them, (PP) holds by induction hypothesis. If p_n is among them, assume it is the h-th argument where $1 \leq h \leq k-\ell$. Let $t = s \circ h$ and let $\mathbf{p}$ be the remaining arguments. By induction hypothesis we have $f_t(\mathbf{p}) \geq 2B_f$. Also $dist(p_n, C^t_{\mathbf{p}}) \geq \gamma_\ell$ and hence $f_s(\mathbf{p}, p_n) \geq g_t(\mathbf{p}, dist(p_n, C^t_{\mathbf{p}})) \geq g_t(\mathbf{p}, \gamma_\ell) = f_t(\mathbf{p}) \geq 2B_f$.

We also need that the local geometry of the curves of degeneracy is simple.

Requirement 4 *There are constants C and δ_0 such that for $0 \leq \delta \leq \delta_0$ and all $\mathbf{p}$ satisfying (PP), the γ_ℓ-neighborhood of $C^t_{\mathbf{p}}$ covers at most an area $C \cdot \gamma_\ell \cdot \delta$ of any δ-cube.*

The requirement excludes space filling curves and sets $C^t_{\mathbf{p}}$ containing an open set. In some cases, the space estimate can be improved. In particular, if $C^t_{\mathbf{p}}$ consists only of a constant number of points, the estimate can be improved to $C\gamma_\ell^2$. Our final requirement relates δ and the γ_l.

Requirement 5 $2Ck! \cdot \sum_{0 \leq \ell < k} n^{k-\ell} \gamma_\ell \leq \delta$

Theorem 1. *If requirements (1) to (5) hold and the input consists of n points, then with probability $1/2$ the fp-evaluation of f yields the correct sign for any k-tuple of distinct perturbed points.*

Proof. Consider any ℓ with $0 \leq \ell < k$. There are no more than $k!$ sequences t with $|t| = \ell + 1$. Also there are at most $n^{k-\ell-1}$ tuples of $k-\ell-1$ distinct indices in $[1, n]$. Thus the total area covered by the γ_ℓ-neighborhoods of all $C^t_{\mathbf{p}}$ is at most $k! n^{k-\ell-1} \cdot C \cdot \gamma_\ell \cdot \delta$. The sum over all ℓ of this quantity is at most $\delta^2/(2n)$ by requirement (5). Thus the probability that the choice of p_i for any fixed i with $1 \leq i \leq n$ does not support the induction step is at most $1/(2n)$ and hence the probability that some induction step fails is at most $1/2$. Thus with probability at least $1/2$, we have (PP) for n. Since $f_\varepsilon = |f|$, this implies that the fp-evaluation of f yields the correct sign for any k-tuple of distinct perturbed points.

We give an example. $f(p,q,r) = orient(p,q,r) = dist(p,q) \cdot dist(r,\ell(p,q))$. We compute $f_{(3,2,1)}$. Let $t = (3)$. We have $C^t_{\mathbf{x}} = \ell(p,q)$. Then $g_{(3)}(p,q,d) = dist(p,q) \cdot d$ and hence $f_{(3)}(p,q) = dist(p,q) \cdot \gamma_0$. Let $t = (3,2)$. We have $C^t_{\mathbf{x}} = \{p\}$. Then $g_{(3,2)}(p,d) = d \cdot \gamma_0$ and hence $f_{(3,2)}(p) = \gamma_1 \cdot \gamma_0$. Let $t = (3,2,1)$. We have $C^t_{\mathbf{x}} = \emptyset$ and set it to $C^t_{\mathbf{x}} = \{(0,0)\}$. Then $g_{(3,2,1)}(d) = \gamma_1 \cdot \gamma_0$ and hence $f_{(3,2,1)}(p) = \gamma_1 \cdot \gamma_0$. We may use $C = 4$, need $48M^2 2^{-L} \le \gamma_1 \cdot \gamma_0$ to satisfy requirement 3 and $48 \cdot (n^3\gamma_0 + n^2\gamma_1 + n\gamma_2) \le \delta$ to satisfy requirement 5. With $\gamma_2 = 0$, $\gamma_0 = \delta/(96n^3)$, $\gamma_1 = \delta/(96n^2)$, the requirement for L becomes $L \ge 2\log(M/\delta) + 5\log n + O(1)$. This can be improved somewhat by using the fact that requirement 5 can be replaced by $24(n^2\gamma_0\delta + n\gamma_1^2) \le \delta^2/(2n)$ since for $t = (3,2)$, the curve of degeneracy consists of a single point. With $\gamma_0 = \delta/(96n^3)$ and $\gamma_1 = \delta/(\sqrt{96}n)$, the requirement for L becomes $L \ge 2\log(M/\delta) + 4\log n + O(1)$.

We summarize: In order to apply the scheme, one first fixes δ to a value suitable for the application. Then one fixes the γ_ℓ to values obeying requirement 5 and determines suitable functions f_t. This might require some ingenuity and is the subject of the discussion below. Finally, one determines C and makes L large enough to guarantee requirement 3. We next specialize and make the general scheme more concrete.

The function f is frequently *symmetric* in its arguments up to change of sign, i.e., permuting the arguments does not change the absolute value. In the case of symmetric functions, the functions f_t only depend on the length of t and not on the actual structure of t. Writing f_ℓ for f_t with $|t| = \ell$ we obtain a sequence of functions f_0, f_1 to f_k where f_ℓ has $k - \ell$ arguments. For simplicity, we restrict most of the further discussion to symmetric functions.

In our examples, the functions g_ℓ are *separable*, i.e., we have $g_{\ell+1}(\mathbf{x},d) = h_{\ell+1}(\mathbf{x}) \cdot d^{e_\ell} \cdot \prod_{0\le i<\ell} \gamma_i^{e_i}$ for some function $h_{\ell+1}$ and some integers e_i. More frequently, we can locally bound $g_\ell(\mathbf{x},d)$ from below by a separable function, i.e., we have a positive constant d_ℓ and a continuous function $h_{\ell+1}(\mathbf{x})$ with $h_{\ell+1}(\mathbf{x}) = 0$ iff $\mathbf{x}$ is degenerate such that $g_{\ell+1}(\mathbf{x},d) \ge h_{\ell+1}(\mathbf{x}) \cdot d^{e_\ell} \cdot \prod_{0\le i<\ell} \gamma_i^{e_i}$ for $d \le d_\ell$ and $g_{\ell+1}(\mathbf{x},d) \ge h_{\ell+1}(\mathbf{x}) \cdot d_\ell^{e_\ell} \cdot \prod_{0\le i<\ell} \gamma_i^{e_i}$ for $d \ge d_\ell$. With $\gamma_\ell \le d_\ell$ for all ℓ, we obtain $f_{\ell+1}(\mathbf{x}) = h_{\ell+1}(\mathbf{x}) \cdot \prod_{0\le i\le\ell} \gamma_i^{e_i}$ and hence $f_k = c \cdot \gamma_0^{e_0} \cdots \gamma_{k-1}^{e_{k-1}}$ for some constant $c = h_k$. Thus requirement 3 becomes $2B_f \le c \cdot \gamma_0^{e_0} \cdots \gamma_{k-1}^{e_{k-1}}$.

What is a good choice for the γ_ℓ? The condition $\gamma_\ell \le d_\ell$ makes it difficult to give a general answer. We therefore assume $d_\ell = \infty$ for all ℓ. We want to minimize $\sum_\ell n^{k-\ell}\gamma_\ell$ subject to the constraint $2B_f \le c \cdot \gamma_0^{e_0} \cdots \gamma_{k-1}^{e_{k-1}}$. There is an extremal point where the inequality is an equality. The Kuhn-Tucker conditions tell us that at an extremal point the partial derivatives of the objective function and the constraint with respect to the γ_ℓ must line up, i.e., there is a λ such that $n^{k-\ell} = \lambda 2B_f e_\ell/\gamma_\ell$ for all ℓ. Thus $\gamma_\ell = \lambda 2B_f e_\ell / n^{k-\ell}$ and hence $\lambda = \left((2B_f/c)n^S/\prod_\ell e_\ell^{e_\ell}\right)^{1/E}/(2B_f)$ where $E = e_0 + \ldots + e_{k-1}$ and $S = \sum_\ell (k-\ell)e_\ell$. Then $\delta \ge 2Ck!\sum_\ell n^{k-\ell}\gamma_\ell = 2Ck!E\left((2B_f/c)n^S/\prod_\ell e_\ell^{e_\ell}\right)^{1/E}$. This becomes $\delta \ge 2Ck!E\left(2(c_f/c)M^d 2^{-L} n^S/\prod_\ell e_\ell^{e_\ell}\right)^{1/E}$ for f a polynomial of degree d in k point variables and hence $B_f = c_f M^d 2^{-L}$. Thus we need

$$L \ge E\log(1/\delta) + S\log n + d\log M + O(1) .$$

The major terms in this lower bound can be explained intuitively. We have to compute with numbers as large as M^d and this requires $d\log M$ bits before the binary point. We want to perturb by as little as δ and hence we need at least $\log(1/\delta)$ after the binary point. This is multiplied by the sum E of exponents. The number of potential predicate evaluations grows like n^k and hence there should be a term $O(k \log n)$ to account for them. We have no intuitive explanation for the term $S \log n$.

A key step in applying our methodology is to find the appropriate functions f_t. We give some guidelines on how to find them.

Consider f and a regular $\mathbf{x}$. For any $x \in U$, let x_0 be the point on $C_\mathbf{x}$ closest to x. Define $g_\mathbf{x}(d) = f(\mathbf{x}, x_0 + d(x - x_0)/||x - x_0||)$ where $d \in \mathbb{R}_{\geq 0}$ and consider the Taylor or Puiseux expansion of $g_\mathbf{x}$ at 0. If the Taylor expansion exist, there are $D > 0$, $e \geq 1$, and $c > 0$ depending on $\mathbf{x}$ and x_0 such that $|g_\mathbf{x}(d)| \geq c \cdot d^e$ provided that $d \leq D$. We may choose e as the index of the first non-zero coefficient in the Taylor expansion and c as one half of this Taylor coefficient. If e and D can be chosen indecently of $\mathbf{x}$ and x and $c = c(\mathbf{x})$ depends only on $\mathbf{x}$ but not on x, we have a locally valid bound of the desired form: $|f(\mathbf{x}, x)| \geq c(\mathbf{x}) \cdot dist(x, C_\mathbf{x})^e$ for $dist(x, C_\mathbf{x}) \leq D$.

A frequently occurring case is that $C_\mathbf{x}$ has no singularities whenever $\mathbf{x}$ is regular. Let ∇f be the vector of partial derivatives of f with respect to the coordinates of x and let $|\nabla f|(\mathbf{x}, x)$ be the length of the gradient vector at $(\mathbf{x}, x)$. If $C_\mathbf{x}$ has no singularities, $|\nabla f|(\mathbf{x}, x_0) > 0$ for all $x_0 \in C_\mathbf{x}$ and hence the minimum length of the gradient over all points on $C_\mathbf{x}$ is positive. Let $h(\mathbf{x}) = \min_{x_0 \in C_\mathbf{x}} |\nabla f|(\mathbf{x}, x_0)$. Then $h(\mathbf{x}) > 0$ if $\mathbf{x}$ is regular and $h(\mathbf{x}) = 0$ if $\mathbf{x}$ is degenerate. Also, $g_\mathbf{x}(d) \approx |\nabla f|(\mathbf{x}, x_0) \cdot d \geq (1/2)h(\mathbf{x}) \cdot d$ and we have a separable representation which is linear in the distance and is valid for small d.

Another frequently occurring case is that f is a polynomial in the point coordinates. If $\mathbf{x}$ is regular, the curve $C_\mathbf{x}$ has a finite number of singularities. Let S be the set of singularities. For all points x such that x_0 is at least ε away from any singularity, we can proceed as above and obtain a linear estimate in d. Near singularities, we proceed as follows. Let s be a singularity and assume w. l. o. g that s is the origin. Let m be all terms of minimal degree in f. Then $f(\mathbf{x}, x) \approx m(\mathbf{x}, x)$ for x near s. The terms in m have common degree e in the point coordinates of x; the coefficients are polynomials in the point coordinates of the points in $\mathbf{x}$. Over the reals, m factors into a product of linear factors and irreducible quadratic factors. Each linear factor ℓ_i defines a line through the origin whose coefficients are functions in $\mathbf{x}$ and hence $|\ell_i(\mathbf{x}, x)| = c_i(\mathbf{x}) dist(x, \ell_i)$. An irreducible quadratic factor q_j contains only a single real point, namely the origin, and the function value of q_j grows quadratically in the distance of $\mathbf{x}$ from s, i.e., $|q_j(\mathbf{x}, x)| \geq c_j(\mathbf{x}) dist(x, s)^2$ for some $c_j(\mathbf{x})$. Thus

$$
\begin{aligned}
m(\mathbf{x}, x) &\geq \prod_i c_i(\mathbf{x}) \cdot dist(x, \ell_i) \cdot \prod_j c_j(\mathbf{x}) \cdot dist(\mathbf{x}, s)^2 \\
&\geq \prod_i c_i(\mathbf{x}) \cdot \prod_j c_j(\mathbf{x}) \cdot dist(x, \cdots \cup \ell_i \cup \cdots)^e \approx \prod_i c_i(\mathbf{x}) \cdot \prod_j c_j(\mathbf{x}) \cdot dist(x, C_\mathbf{x})^e ,
\end{aligned}
$$

and there is hope for a separable bound which grows like d^e.

5 Applications

We apply the methodology to the side-of-circle test of four points in the plane. It tells the side of a query point with respect to an oriented circle defined by three points. We have three points $p_i = (x_i, y_i)$, $1 \le i \le 3$, and a query point $p = (x, y)$. Let us assume first, that the three points are not collinear. Let R be the radius of the circle C defined by the first three points. We may assume w. l. o. g. that the circle is centered at the origin. Then $x_i^2 + y_i^2 = R^2$ for all i. Let Δ be the signed area of the triangle (p_1, p_2, p_3). The side-of-circle test is given by the sign of the determinant

$$f_0(p_1,p_2,p_3,p) = \begin{vmatrix} 1 & x_1 & y_1 & x_1^2+y_1^2 \\ 1 & x_2 & y_2 & x_2^2+y_2^2 \\ 1 & x_3 & y_3 & x_3^2+y_3^2 \\ 1 & x & y & x^2+y^2 \end{vmatrix} = -\begin{vmatrix} x_1 & y_1 & R^2 \\ x_2 & y_2 & R^2 \\ x_3 & y_3 & R^2 \end{vmatrix} + (x^2+y^2)\cdot \begin{vmatrix} 1 & x_1 & y_1 \\ 1 & x_2 & y_2 \\ 1 & x_3 & y_3 \end{vmatrix}$$
$$= -R^2 2\Delta + (x^2+y^2)2\Delta = 2\Delta(x^2+y^2-R^2) \;.$$

This predicate was already analyzed in [4] using non-trivial geometric reasoning. The purpose of this section is to show that the same result, in fact, a slightly better result, can be obtained by generic reasoning. The curve of degeneracy is the cycle C and the normal vector at $p_0 = (x_0, y_0) \in C$ is $(4x_0, 4y_0)\Delta$ and has norm $4\sqrt{x_0^2+y_0^2}|\Delta| = 4R|\Delta|$. This is independent of p_0. So the first order approximation of f_0's absolute value is $4R|\Delta| \cdot dist(p,C)$. In fact, one half of this is even a global lower bound. Namely,

$$|f_0(p_1,p_2,p_3,p)| = 2|\Delta|\cdot|\sqrt{x^2+y^2}-R|\cdot(\sqrt{x^2+y^2}+R) \ge 2R|\Delta|\cdot dist(p,C) \;.$$

Let $a = dist(p_1,p_2)$, $b = dist(p_1,p_3)$, $c = dist(p_2,p_3)$, and let α be the angle at p_3 in the triangle (p_1,p_2,p_3). Then $2R = a/\sin\alpha$ and $|\Delta| = (1/2)bc\sin\alpha$ and hence $2R|\Delta| = 1/2\cdot abc$. Thus

$$|f_0(p_1,p_2,p_3,p)| \ge \frac{1}{2} dist(p_1,p_2) dist(p_1,p_3) dist(p_2,p_3) dist(C,p)$$

and by continuity of the determinant the latter inequality is also true if the points p_1, p_2, and p_3 are collinear. In this case, C is the line passing through the first three points. The formula also tells us that the triple (p_1,p_2,p_3) is regular iff the points are pairwise distinct.

We next consider $f_1(p_1,p_2,p)/\gamma_0 = (1/2)\cdot dist(p_1,p_2)dist(p_1,p)dist(p_2,p)$. We consider p_1 and p_2 as fixed and treat $p = p_3 = (x,y)$ as a variable. If $p_1 = p_2$, the function is identically zero. If $p_1 \ne p_2$ we have $f_1(p_1,p_2,p) = 0$ iff $p = p_1$ or $p = p_2$. Thus the curve of degeneracy consists of the two isolated points p_1 and p_2 and its tubular neighborhood is two circles. We want to bound f_1 from below. We may assume that p is closer to p_1 than to p_2. Then $dist(p,p_2) \ge dist(p_1,p_2)/2$ and hence $f_1(p_1,p_2,p)/\gamma_0 \ge dist(p_1,p_2)^2/4\cdot dist(p,\{p_1,p_2\})$. Therefore $f_2(p_1,p)/(\gamma_0\gamma_1) = dist(p_1,p)^2/4$ and further $f_3(p_1)/(\gamma_0\gamma_1\gamma_2^2) = 1/4$. In fact, there is no real reason to go down to f_3. We can also argue about f_0 directly. If any two points have a certain minimum distance m, $f_0(p_1,p_2,p_3,p) \ge m^3/2\cdot dist(p,C)$. In [4], $\Delta^{3/2}$ was considered instead of ΔR and then the curve of degeneracy is the line spanned by p_1 and p_2. The tubular neighborhood is then a strip and the forbidden region is larger.

The computation of Voronoi diagrams of line segments is computationally difficult. The available exact algorithms [2] are slow, the fast algorithm of M. Held [8] is not guaranteed to work for all inputs. The key test in the algorithms in the side-of-circle test: A circle C is specified by three sites (points or lines) and the position of a fourth site (point or line) with respect to C is to be determined, see Figure 5. We discuss the situation where the fourth site is a line segment given by points p and q. Let C have center c and radius R and let p be outside C. The query point is $q=(x,y)$. We want to know whether the line $\ell(p,q)$ intersects, touches, or misses C. There are different ways of realizing this test. For simplicity let us put p at $(0,0)$ and c at $(c_0,0)$.

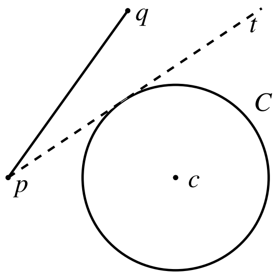

In [2] the test is realized by comparing R and $dist(\ell(p,q),c)$, the distance of c from the line ℓ. The line has equation $y\tilde{x}-x\tilde{y}=0$ (here x and y are the coefficients and $\tilde{x}$ and $\tilde{y}$ are the variables. The signed distance of c from this line is $yc_0/\sqrt{x^2+y^2}$ and hence the test is realized by the formula $yc_0 - \pm R\sqrt{x^2+y^2}$. For each choice of sign, the curve of degeneracy is one of the tangents t from p at C; the equations for the tangents are $y=\pm Rx/\sqrt{c_0^2-R^2}$. We leave it to the reader to verify that the norm of the normal vector has the same value for all points on the curve of degeneracy.

Alternatively, we may locate q with respect to the tangents from p at C. We further discuss this method. Let us concentrate on one of the tangents. We refer to it as t. Then the location of $q=(x,y)$ is given by the sign of $E=\sqrt{c_0^2-R^2}\cdot y - R\cdot x$. Observe that c_0 is the distance between c and p. Hence the general form for arbitrary p and c is given by $E=\sqrt{dist(p,c)^2-R^2}\cdot(y-y_p)+R\cdot(x-x_p)$.

The circle C is defined by three sites. We treat the case of three points sites, the other cases are somewhat more involved. Let our three points be $p_i=(x_i,y_i)$, $1\le i\le 3$. The center c has coordinates (it is the intersection of two bisectors)

$$x_c = \frac{\begin{vmatrix} (x_2^2-x_1^2+y_2^2-y_1^2)/2 & y_2-y_1 \\ (x_3^2-x_1^2+y_3^2-y_1^2)/2 & y_3-y_1 \end{vmatrix}}{2\Delta} \qquad y_c = \frac{\begin{vmatrix} x_2-x_1 & (x_2^2-x_1^2+y_2^2-y_1^2)/2 \\ x_3-x_1 & (x_3^2-x_1^2+y_3^2-y_1^2)/2 \end{vmatrix}}{2\Delta}$$

where Δ is the area of the triangle (p_1,p_2,p_3). Write $x_c=A/(2\Delta)$ and $y_c=B/(2\Delta)$. The radius of the circle is given by $R=\sqrt{(x_1-x_c)^2+(y_1-y_c)^2}=\sqrt{D}/(2|\Delta|)$, where $D=(2x_1\Delta-A)^2+(2y_1\Delta-B)^2$. Next observe $dist(p,c)^2=(x_p-x_c)^2+(y_p-x_c)^2=((2x_p\Delta-A)^2+(2y_p\Delta-B)^2)/(4\Delta^2)$. Plugging into our expression E and multiplying by 2Δ yields the simplified expression

$$E = \sqrt{(2x_p\Delta-A)^2+(2y_p\Delta-B)^2-(2x_1\Delta-A)^2-(2y_1\Delta-B)^2}\cdot(y-y_p) + \sqrt{(2x_1\Delta-A)^2+(2y_1\Delta-B)^2}\cdot(x-x_p)\ .$$

Table 1 yields $B_E=c_EM^42^{-L}$ for some constant c_E. Next observe that $|E|=|H(y-y_p)+G(x-x_p)|=\sqrt{H^2+G^2}\cdot dist(t,q)$ because E is a linear function in x and y and hence the first order approximation is exact. Also the norm of the normal vector is

$\sqrt{H^2+G^2}$. Finally, observe $H^2+G^2=(2x_p\Delta-A)^2+(2y_p\Delta-B)^2=4\Delta^2 dist(p,c)^2$ and hence $\sqrt{H^2+G^2}=2|\Delta|dist(p,c)$. So the requirement $|E|\geq 2B_E$ boils down to

$$|dist(q,t)|\geq \frac{2B_E}{2|\Delta|dist(p,c)}\geq \frac{B_E}{|\Delta|R}=\frac{c^E M^4 2^{-L}}{|\Delta|R}$$

and the quantity ΔR is familiar to us from the side-of-circle test for points. In order to guarantee a lower bound for it, it suffices to guarantee a minimum distance for the defining points of C.

References

1. C. Burnikel, S. Funke, and M. Seel. Exact arithmetic using cascaded computation. In *SocCG*, pages 175–183, 1998.
2. C. Burnikel, K. Mehlhorn, and S. Schirra. How to compute the Voronoi diagram of line segments: Theoretical and experimental results. In *ESA*, volume 855 of *LNCS*, pages 227–239, 1994.
3. S. Fortune and C. van Wyk. Efficient exact integer arithmetic for computational geometry. In *7th ACM Conference on Computational Geometry*, pages 163–172, 1993.
4. S. Funke, Ch. Klein, K. Mehlhorn, and S. Schmitt. Controlled perturbation for Delaunay triangulations. SODA, pages 1047–1056, 2005.
5. Halperin and Shelton. A perturbation scheme for spherical arrangements with application to molecular modeling. *CGTA*, 10, 1998.
6. D. Halperin and E. Leiserowitz. Controlled perturbation for arrangements of circles. *IJCGA*, 14(4):277–310, 2004.
7. D. Halperin and S. Raab. Controlled perturbation for arrangements of polyhedral surfaces with application to swept volumes. available from Halperin's home page; a preliminary version appeared in SoCG 1999, pages 163–172.
8. M. Held. VRONI: An engineering approach to the reliable and efficient computation of Voronoi diagrams of points and line segments. *Comput. Geom.*, 18(2):95–123, 2001.
9. M. Jünger, G. Reinelt, and D. Zepf. Computing correct Delaunay triangulations. *Computing*, 47:43–49, 1991.
10. M. Karasick, D. Lieber, and L.R. Nackman. Efficient Delaunay triangulation using rational arithmetic. *ACM Transactions on Graphics*, 10(1):71–91, January 1991.
11. L. Kettner, K. Mehlhorn, S. Pion, S. Schirra, and C. Yap. Classroom examples of robustness problems in geometric computations. In *ESA*, volume 3221 of *LNCS*, pages 702–713, 2004.
12. K. Mehlhorn and S. Näher. The implementation of geometric algorithms. In *Proceedings of the 13th IFIP World Computer Congress*, volume 1, pages 223–231. Elsevier, 1994.
13. K. Mehlhorn and S. Näher. *The LEDA Platform for Combinatorial and Geometric Computing*. Cambridge University Press, 1999. 1018 pages.
14. C.K. Yap. Towards exact geometric computation. In *Proceedings of the 5th Canadian Conference on Computational Geometry (CCCG'93)*, pages 405–419, 1993.