Searching for lower bounds in computational geometry. A survey on methods

T. Recio
Universidad de Cantabria*

1 Models of computation and reduction to a membership problem

Computational geometry provides algorithms to solving some problems dealing with geometric entities. When the algorithms are described in a concrete model of computation one can estimate the "complexity" of the algorithm, as a rough measure of how efficient it will be in practice. Problems usually depend on some parameters (such as the number of points in Convex Hull algorithms or the number of vertices in the problem of testing if a point belongs to a given polyhedra). The estimated size of the algorithm is, thus, regarded as a function of the parameters. One is obviously interested in finding algorithms that are efficient for large values of the parameters; therefore, some effort is devoted in finding suitable algorithms whose complexity is assymptotically low, as a function of the parameter size. But it can turn out that, after extensive search, all known algorithms for one specific problem have a size greater than some given function. One could then wonder if this function is an intrinsic lower bound on the efficiency for solving the given problem (implying, therefore, that there is no sense in continuing the search for improvements if some known algorithm has assymptotically the same complexity as the lower bound).

Thus, finding lower bounds is a task that depends on the problem itself, and not on the methods for solving it. It provides, some times, an accuarate evidence for determining that we have already an optimal algorithm; some other times, it just gives an idea of how difficult, at least, will be the problem. When we have a gap between estimated lower bounds and complexity of known algorithms, it could happen that the lower bounds are not sharp enough, or it could be the case that there is still room for improving algorithms solving the problem. The nicest situation is the one in which one knows exactly the complexity of the problem (i.e. where one can construct algorithms running in the best possible time to solve a problem and a proof that this time is the best one).

Strange as it might seem, in the most common model of computation there is a theoretical method to finding out exactly for every fixed parameter size (for instance, for computing the convex hull of any fixed number of points) and for every computational geometry problem, a best time algorithm. This result uses Model Theoretic methods, but it is unpractical for several reasons: we can hardly ever find the performing algorithm in concrete instances, and moreover, it does not provide precise information about what will be the best algorithms for solving the problem when we consider different and sufficiently large values of the parameters. In some other models of computation, even for this restricted version (fixed parameter value) of the best complexity finding problem, it is unknown if the complexity can be exactly determined: that is why we remark here this theoretical result.

The most popular and complete model of computation in Computational Geometry is, probably, the real Algebraic Computation Tree (ACT) model (see the paper of Ben-Or for a description, "Lower bounds for algebraic computation trees", Proc. 15th Ann. ACM Symp. Theory Comput. (1986), 80-86). It is a non-uniform model, in the sense that for any given instance of a problem, eg. for a number n of points, an ACT depending on n will model the behaviour of the algorithms solving it: for different instances we could have different ACT's. It is assumed that the geometric operations performed by any algorithm have always an algebraic counterpart with the coordinates of the geometric objects given as input. The ACT allows elementary arithmetic operations with the input coordinates and with the result of such operations; plus sign test comparisons and branching, accordingly. Operations with (real) constants are also allowed. The final output of a ACT computation will be a series of leaves, including the expected output in terms of coordinates of the input.

The natural uniformization of this model (i.e. a model able to handle simultaneously all instances of a given problem) has been discussed in Cucker-Montaña-Pardo: "Time Bounded Computations over the Reals", International Journal of Algebra and Computation, vol.2, no.4 (1992), 395-408, analyzing the Blum-Shub-Smale model (cf. "On a theory of computation and complexity over the real numbers". Bull. Amer. Math. Soc. (1) 21, (1989), 1-46).

It is also standard that the search for lower bounds is made just on some weaker version of the same geometric problems: namely, the decisional version (see the book of Preparata-Shamos for a discussion on this point: "Computational Geometry: an introduction", Springer-Verlag, Texts and Monographs in Computer Science, 1985). For instance, instead of looking for the complexity of constructing the Convex Hull of a finite family of n points, we consider the problem of deciding if a collection of n points is the set of extreme points (vertices) of the Convex Hull of the given family. The ACT for such decision problem has the advantage that its leaves are just labeled with a Yes/No label, according to the case where the computation that leads to one of such leaves gives the correct Yes/No answer to the decision problem. Of course, if we have an algorithm (an ACT) for solving the Convex Hull problem, then it will be easy to solve also the decision version of the same problem: namely, we just add some tests to check whether the given points verify some of the outputs and thus to verfy that they are the extreme points. Therefore, lower bounds for Computational Geometry problems can be attacked by finding lower bounds for decision problems (but it could be the case that the converse does not hold). Anyhow, most of what is known about lower bounds in Computational Geometry happens to behave as if always the decision version has the same complexity as the one of the truly proposed problem.

For decision problems, we can link the Computational Geometry problem with a Membership Problem to a specific (i.e. related to the problem) Semialgebraic Set: if one wants to decide if a family of n points in the plane is the set of extreme points of its own convex hull, one can construct the subset S of points $(x_1, y_1, x_2, y_2, \ldots, x_n, y_n)$ in real affine $2n$-space, such that the set of planar points $\{(x_1, y_1), (x_2, y_2), \ldots, (x_n, y_n)\}$ is made of pairwise different poins and such that it is the set of vertices of the convex hull of $\{(x_1, y_1), (x_2, y_2), \ldots, (x_n, y_n)\}$. It is clear, but rather cumbersome to detail, that such subset is in fact a semialgebraic set in $2n$-space. Now the decisional version of Convex Hull is in fact equivalent to the problem of testing if a point $(x_1, y_1, x_2, y_2, \ldots, x_n, y_n)$ belongs to S or not (Membership to the set S problem).

Clearly, in the example above, whatever algorithm is used to solve the Convex Hull decisional version, could be also used to solve the membership problem for S; and conversely. Thus a lower bound for testing membership to S will be also a lower bound for the Convex Hull problem. In fact, almost by definition, all problems in Computational Geometry that can be solved by algorithmic methods, modeled in an ACT, can be related to a semialgebraic set, since the result of operating throught an ACT gives just polynomials and sign tests... and this is the definition of semialgebraic set. In general, one reduces the search of lower bounds in Computational Geometry to the search of lower bounds in the membership problem to the corresponding semialgebraic sets.

As commented above, the problem of estimating the exact complexity of a semialgebraic set under the ACT model is decidable (cf. Montaña-Pardo-Recio, "The complexity of semialgebraic sets" Tech. Report 3/90, Universidad de Cantabria, 1990) : given the semialgebraic set, we can easily construct one ACT that solves the membership problem to this set. Then we consider all possible ACT's with smaller height as a kind of finite parametrized object and we can decide the existence of ACT's of height h solving the membership problem to the given set by means of Quantifier Elimination techniques. Of course this procedure in not efficient, so we can not, in practice, find the complexity of a given problem. Thus some indirect methods are required to obtain hints about the complexity. Most of the methods work following this philosophy: first it is stated that sets defined by means of ACT's of complexity h, have a certain geometric feature that can be numerically quantified by a number p that is bounded by a function of h, $p \leq f(h)$. Thus, if the semialgebraic set has a large p, it implies that h must be larger than some function of p, $f-1(p) \leq h$. On the other hand, in order to obtain the bound $p \leq f(h)$, one must "relate" h with some more mathematical objects, such as the number of polynomials and their degree, that are involved in an ACT of height h: the purely mathematical question of bounding p in terms of these objects is then subject of research. Unfortunately, the final product is usually not a very sharp bound, but it can be still useful in many instances.

There are other ways of defining the complexity of a semialgebraic set, see the book by Benedetti-Risler: "Real algebraic and semialgebraic sets". Hermann 1990. If instead of counting all required arithmetic operations we only consider the number of additions, we can have a variant of ACT, with the so called "additive" complexity. The interest of using weaker complexity measures is that there is usually a big gap, for the stronger measures, between bounds and practical instances: it is enough, for example, to know something related to the additive complexity of one polynomial (number of monomials) in order to estimate the number of its real roots (this is the classical Descartes lemma). One could add, perhaps, some comment on one of the major differences between "number of monomials" and "additive complexity", namely, the fact that the latter behaves well under linear change of variables (as does the degree of one polynomial) while number of monomials does not (see 4.3.6 and 4.6.6 in Benedetti-Risler book). It could seem that this disgression is far from applications to Computational Geometry, but in Ngoc-Minh Le, "On Voronoi diagrams in the L_p metric in higher dimension", STACS 94, Lecture Notes Computer Science 775, Springer-Verlag, 1994, 711-722, upper bounds on the number of Lp-spheres passing through $D+1$ points in general position in D-space and on the sum of Betti numbers of the intersection of bisectors in the L_p-metric, where p is an even positive integer, are established by these means. It is shown that the bounds do not depend on p, but only on D (for the first result) and on D and on the number of bisectors (for the second result). He states both problems as a set of equations whose number of solutions provides the required bound and shows that the additive complexity of such system is bounded just in terms of D (and the number of bisectors for the second result). The fact that the additive complexity bounds the number of solutions is then used to obtain the results.

It is unknown weather additive complexity is decidable; also if the collection of semialgebraic sets with bounded additive complexity has a finite number of topological types (something important if we want to connect topology and complexity, as it will be done in the next sections): see the partial results for dimension 3 of F. Schurmann: "Gradunabhangigen Sktanken fur die Topologie Semi-Algebraischer Mengen" Ph. D. Dissertation, Munster Universitat, (1988). Both claims are true for the total or usual complexity on ACT (Montaña-Pardo-Recio, loc.cit.). One could also mention in this respect the recent paper of Grigoriev-Karpinsiki: "Computability of the additive complexity of algebraic circuits with root extracting" (I just have a draft copy courtesy from the authors). Their notion of generalized additive complexity is computable. They also show that there are polynomials of arbitrarily high additive complexity but of low and fixed generalized additive complexity. Quantifier elimination on the theory of differentially closed fields is involved in order to obtain bounds on the generalized additive complexity: the fact that to explicit these bounds one relies on bounds of Hilbert's Idealbasissatz is a remarkable connection with another research field.

There are other, useful but weaker models of computation (therefore models which in principle could allow solving the same problems but with less complexity), such as variants of the ACT model in which one takes some linear operations for free (see Montaña-Pardo-Recio: "The non-scalar model of complexity in computational semialgebraic geometry", Proc. MEGA 90, Progress in Mathematics 94, Birkhauser (1991), 346-362), or models in which all arithmetic operations are given free and it is only counted the number of sign tests of some fixed degree d polynomials (Algebraic Decision Trees of degree d=d-ADT, see J.M. Steele, A. C-C. Yao: "Lower bounds for Algebraic Decision Trees", J. Algorithms 3, (1982) 1-8). The special case when $d = 1$ (linear decision trees modelling the so called linear search algorithms) is one where most lower bounds for Computational Geometry problems have been first searched and proved.

One could even ask, as does F. Yao (cf. F. Yao: "Computational Geometry", in the Handbook of Theoretical Computer Science: Algorithms and Complexity, Vol. A., edited by J. Van Leeuwen, Elsevier, 1990) for a model of Algebraic Decision Tree with no degree bound. In this respect we must remark that under such model the complexity of any semialgebraic set of dimension p is bounded by some effective functions of p: if no bounds on degrees are set, we can always describe an open (with respect to its Zariski closure) semialgebraic set of dimension p by some $s(p) \times t(p)$ polynomials (where S and t are some functions usually called Brocker's invariants, cf. L. Brocker "Minimale Erzeungung von Positivberich", Geom. Dedicata 16, (1984), 335-350). Then every semialgebraic set can be written as a union of an open set plus a semialgebraic set of smaller dimension, and we can use induction to end the proof. This indicates that such model of computation has some essential "a priori" limitations, since it is quite clear that there exist plenty of semialgebraic sets of fixed dimension but arbitrary high complexity (think of a collection of points in the real line), cf. Montaña-Pardo-Recio: "A note on Rabin's width of a complete proof", Comp. Complex. vol.4, no.1, 1994, pp12-36,. Another consequence of our remark is that, in turn, there can be no bound on the degree of the polynomials described by Brocker for every dimension p: if it would exist, then the also the model of computation d-ADT for all sufficiently large degrees d would have a limited complexity, but we will see below that this is not the case (we can establish arbitrarily high lower bounds in this model). We have a personal communication of P. Vélez-Melón for different proof of this fact, inspired by an observation of A. Prestel.

Curiously, lower bounds obtained with different models are the same in many instances, reinforcing the idea that even under strong assumptions many well known algorithms are optimal: it seems that relaxing the counting of steps in the algorithm does not have an essential effect in the complexity of many problems. Another instance of this phenomenon appears when considering different d's for ADT's: if a problem can be solved with a linear ($d = 1$) tree and the complexity of such problem is estimated for such model, it seems that allowing larger values of d in d-ADT's does not improve efficiency, i.e. the complexity of the problem remains the same. It is an open problem whether this conjecture holds in general (see P. Ramanan:"Obtaining lower bounds using articial components" Information Processing Letters 24 (1987), 243-246); now it is just an experimental evidence, true for the known examples.

2 Lower bounds by estimating connected components and Betti numbers

The now classical way of establishing lower bounds is by counting the number of connected components of the associated semialgebraic set S (see D. Dobkin-R. Lipton: "On the complexity of computations under varying sets of primitives" in Automata Theory and Formal Languages (H. Brandhage, ed.) Lec. Notes in Comput. Sci. vol. 333, Springer-Verlag, (1975), 110-117). The essential idea is that if a set S has many components, then it can not be described with few equalites, inequalites or sign tests, with few variables ...and of low degree ...Putting it in the other way around, if a set is described by few polynomials and with small degree, the number of connected components will be also low. For instance, everybody is aware that conics (planar sets described by one polynomial equality and of degree two) have at most two connected components (the hyperbola case): thus if a planar curve has more than two components we deduce that it must be necessarily described by higher degree polynomials.

Technically speaking we have two steps: one is to bound the topology of semialgebraic sets from the number of involved polynomials, sign tests, number of variables and maximum of polynomial degrees. The result is the so called Milnor-Thom (see the above quoted technical report of Montaña-Pardo-Recio for references and different versions of this result) bound, that estimates the sum of the Betti numbers of a basic semialgebraic set as being less than $d(2d-1)^{n+q-1}$, when the set is defined by at most q polynomials, of degree less than d, with n variables. In particular the number of connected components is bounded by the same quantity. The second technical tool is to interpret such bound in the ACT model of computation. This is Ben-Or result in the mentioned above paper, which implies that the height of a tree computing a semialgebraic set in n-dimensional space with C components, is bounded from below by $\Omega[\log(C) - n]$. The idea is to produce, from a given tree, a description of the semialgebraic set with a number of polynomials and maximum degree which is bounded by the height of the tree. Thus, in the formula of Milnor-Thom, the height will appear in the exponent, and taking logarithms we obtain the reverse bound. We remark that this formula applies only to connected components, since an ACT gives in fact a finite union of basic semialgebraic sets, and it is easy to bound from above the number of connected components of a union of sets, if we have a bound for each set–but this is not longer true for the sum of Betti numbers: glueing together sets with small Betti numbers we can produce large ones (independently of the number of sets).

This procedure has been applied succesfully for almost 10 years now to solve many Computational Geometry problems (cf. the book of Preparata-Shamos or the more recent survey by F. Yao, loc. cit.). It is widely known, but still is worth to make a few comments on some aspects.

First of all we have the surprising fact that the term "-n" in the above bound can be dropped; and thus the bound is actually $\Omega(\log(C))$; of course, involving perhaps different constants, hidden in the assymptotic estimation.

This simplification is achieved in J. L. Montaña, "Cotas inferiores en teoría de la complejidad algebraica", Dissertation, Universidad de Cantabria, 1992; and published in F. Cucker, J.L. Montaña and L.M. Pardo, "Time bounded computations over the reals", International Journal of Algebra and Computation, vol.2, no 4, 1992, pp.395,408. The idea is to bound also the number of involved variables by considering the height of the tree . The same result applies to give lower bounds for the uniform model of computation (time-bounded machines) discussed by the authors in the latter paper, extending Ben-Or's bound.

A second remark comes from the fact that roughly the same lower bound holds (see Ben-Or's paper) for weaker models of computation (therefore models which in principle could allow solving the same problems but with less complexity), such as models in which one takes some linear operations for free, or models in which all arithmetic operations are given free and it is only counted the number of sign tests of some fixed degree d polynomials (Algebraic Decision Trees of degree d=d-ADT). The same thing happens (roughly) for computing the membership problem under restrictions of the input to points in the rational or integer affine space (cf. A.C-C . Yao: "Lower bounds for algebraic computation trees with integer inputs". Proccedings 49th Ann. Symp. on Fundamentals of Computer Science, (1989) 308-313; also M. Hirsch: "Lower bounds for the non-linear complexity of algebraic computation trees with integer inputs", preprint, 1990; and Pardo-Recio: "Arboles Algebraicos: un modelo de computación en Geometría", en Contribuciones Científicas (en honor del profesor E. Villar), Universidad de Cantabria, (1988), 241-249). The interest of these restricted models of computation appears in some problems of Computational Geometry where the inputs are also restricted, for instance, to vertices of a simple polygon (see the paper of A. C-C. Yao mentioned above).

As mentioned above, it has been difficult to estimate how higher Betti numbers of semialgebraic sets contribute to the complexity of the set. A. Bjorner-L. Lovasz-A. Yao ("Linear decision trees: volume estimates and topological bounds" Proc. A.C.M. 24th Symp. on Theory of Comp., (1992), 170-177) and A. Bjorner and L. Lovasz ("Linear decision trees, subspace arrangements and Moebius fuctions". J. of the A.M.S. 7, no3 (1994), 677-705) showed that the Euler characteristic (respectively, the sum of all Betti numbers) provides lower bounds for the restricted class of linear decision trees. It was shown essentially that the number of Yes (respec. No) leaves in a linear decision tree computing a closed polyhedron is bounded from below by either the Euler characteristic (alternate sum of Betti numbers) or the sum itself. These techniques were applied to a variety of combinatorial problems, including a lower bound nlog(2n/k) for the k-equal problem (given n real numbers, decide if some k of them are equal). A. Yao extended the Euler characteristic lower bound for general algebraic decision trees and algebraic computation trees ("Algebraic decision trees and Euler characteristic", in Proc. 33rd. Annual IEEE Sympos. on Foundations of Computer Science, (1992), 269-277), showing that the logarithm of the Euler characteristic is a lower bound for the complexity of the membership to a compact semialgebraic set. Finally, in ("Decision Trees and Betti numbers", Proc. A.C.M. Symp. Theory of Computing STOC 94, 1994, 615-624), A. Yao showed that the the log of each Betti number gives a lower bound for compact semialgebraic sets under the fixed degree algebraic decison tree model or the ACT model. The true possibilities of these new techniques in Computational Geometry problems are still unexplored.

Finally let us comment the work done on topological bounds for a different model of computation, the arithmetic networks model (see von zur Gathen, "Parallel arithmetic computations: a survey" Proc. 13rd Symp. on Mathematical Foundations of Computer Science, (1986), 93-112). It is the parallel counterpart of algebraic decision trees, so we can say that a problem is well parallelizable if if goes from a polynomial complexity in the sequential model to a polylogarithmic complexity in the parallel model. Now the results of Montaña-Pardo ("Lower bounds for arithmetic networks", Journal of A.A.E.C.C. 4, (1993), 1-24) show that some problems, such as the knapsack problem, are not well parallelizable. In fact, they extend Ben-Or result, showing that $\sqrt{\log(b)/n}$, where b is the number of connected components and n is the number of variables, is a lower bound for an arithmetic network accepting a semialgebraic set. Thus they get a $\sqrt{n}$ lower bound for the knapsack problem, and therefore it is not well parallelizable. A similar result $\sqrt{\log(\sum b)/n}$ has been recently obtained, where $\sum b$ is the sum of Betti numbers (Montaña-Morais-Pardo, Proc. 10th European Conference on Computational Geometry, Santander, March 1994, extended version in "Lower bounds for arithmetic networks II: sum of Betti numbers", Journal of A.A.E.C.C., to appear). It is conjectured that the square root relation holds always between bounds in the sequential/parallel case, at least when they regard topological quantities.

3 Lower bounds by introducing artificial components

This technique is essentially a trick to allow the use of other lower bound methods in difficult cases. One considers a semialgebraic set, perhaps with few connected components, and then realizes that it will have many such components if intersected with another semialgebraic set. The point is that the complexity of the given set plus the complexity of the artificially introduced set is bigger than the complexity of the resulting intersection. Therefore, the complexity of the given set is bigger than the difference of the complexities of the artificial component and the intersection. Thus, if the artificial component has low complexity and the intersection is very complex, we obtain good lower bounds for the given set. This technique is a natural generalization of the classical degree of an algebraic variety method (counting the points in common with a suitable generic linear variety) for Algebraic Complexity problems.

For instance, let us consider a regular n-polygon , inscribed in a circle. The polygon is topologically trivial, but intersected with the circle yields n connected components. Thus the complexity of the polygon will be at least log(n) (using Ben-Or bound) minus the complexity of the circle... which is constant. A similar but simpler argument shows that the number of polynomials appearing in any quantifier free formula defining a regular n-polygon must be bigger than n (V. Weispfenning, personal communication).

This method has been succesfully applied by Ramanan (for the algebraic decision tree of fixed degree, loc. cit.) and Montaña-Pardo-Recio (The nonscalar model ...loc.cit.) for the more general algebraic computation tree model, showing for the first time the solution to a series of Computational Geometry problems such as Regular and Star-shaped polygon Inclusion, Largest Empty Circle, Maximum Gap, Even Distribution and Path Testing. In the case of considering the algebraic computation tree model, some technical points must be managed, since usually the artificial component has fixed degree and thus constant complexity under the decision tree model, but not under the stronger, algebraic computation model. A detour to non-scalar complexity is therefore required to lower the complexity of the artificial component to the limits of interest for the considered problem–in the above examples, to less than nlog(n)–since, for instance, a degree two polynomial of n variables can require an order of n2 total operations, but just n non-scalar ones: this is the classical result on diagonalization of quadratic forms.

Variants of this approach considering singular points of the boundary of a semialgebraic planar set S, and showing that log(# singular points) is a lower bound for the membership problem under the algebraic computation model, are due to Grigoriev-Karpinski ("Lower bounds for complexity of testing membership to a polygon for algebraic and randomized computation trees", TR-93-042, august 1993. ICSI, Berkeley). As an application they get the already known log(n) estimation for the Polygon Inclusion problem. Further extensions (considering sharp points of semialgebraic sets) of these methods are being applied to the open problem of estimating the exact complexity of polyhedra (see Grigoriev-Karpinski-Vorobjov "Lower Bounds on Testing Membership to a Polyhedron by Algebraic Decision Trees", Proc. ACM Symp. Theory of Computing, 1994, 635-644 and their improvement submitted to STOC 95).

4 Lower bounds by local methods

It was probably M.O. Rabin the one that first introduced the technique of reduction to a membership problem for finding lower bounds in Computational Geometry problems. In its seminal paper ("Proving simultaneous positivity of linear forms", J. Computer and System Sciences 6 (1972), 639-650), he also proposes a method for establishing lower bounds for some problems involving linearly defined sets (such as finding the maximun or the minimun of a collection of real points). The rough idea is that a set defined by the simultaneous positivity of m linear forms $\{l_1 \geq 0, /dots, l_m \geq 0\}$ that are "sign independent" (i.e. such that for any choice of signs: ? can be any of =, ¿, ¡, there are points verifying the inequalities $l_1?0, \ldots, l_m?0$) can not be described by a collection of sets of inequalities with less than m polynomials in each set. For instance, in n-dimensional space, the set $\{x_1 \geq 0, \ldots, x_m \geq 0\}$

with $m \leq n$, which is obviously sign idependent, can not be written as a finite union of sets of the kind $\{f1 \geq 0, \ldots, ft \geq 0\}$ with $t < m$. The idea of Rabin is, then, that an algebraic decision tree computing the convex set $\{l_1 \geq 0, \ldots, l_m \geq 0\}$ gives a description of the same as a union of sets with h polynomial inequalities in each, where h is the complexity of the tree. Thus if the convex set is sign independent, h must be bigger than m.

Let us see how this works over linear decision trees, following a simplified argument by R. Fleischer (Decision trees: old and new results, Proc. A.C.M. 25th Symp. Theory of Comp. 1993, 468-477). It is quite straightforward that the "sign independence" condition is equivalent to the simpler fact "dimension of the affine variety $\{l1 = 0, \ldots, lm = 0\} = n - m$" i.e. to the independence of the involved linear forms. Thus the convex set defined by m sign independent forms is a kind of mth-orthant, having a "corner point" with m faces such as $\{x_1 \geq 0, \ldots, x_m \geq 0\}$. Next we observe that in a linear decision tree, every path arriving at a YES leave describes a convex set containing a linear variety of dimension at least n-size of path (the equality holds only when all the involved polynomials are independent and the sign tests are always equalities). Therefore $n - m \geq n - h$, where h is the height of the tree, i.e. $h \geq m$.

This technique provides estimates for some problems that, having an associated semialgebraic set with very simple topology, can not be bounded by the topological methods described above. Anyhow this procedure is 10 years older than the latter ones; and it is one of the first systematic attempts to find lower bounds. It is a local method, since the obstruction to have a simpler complexity comes, as we will see below, from the behaviour of the semialgebraic or semilinear set around one point (the corner). It is clear that it can not yield but linear lower bounds in most situations, since m is usually less than the dimension of the ambient space, but it shows at least why some basic problems can not be solved in constant time.

Unfortunately there was a substantial error in Rabin's proof and in the subsequent paper by J.W. Jaromczyk ("Lower bounds for problems defined by polynomial inequalities", Symposium on Foundations of Computing Theory, Lect. Notes in Computer Sci. 117, Springer-Verlag, 1981, 165-172) where this technique was extended to problems defined by polynomials of higher degree. A first correct proof for all cases was given in Pardo-Recio ("Rabin's width of a complete proof and the width of a semialgebraic set". Proc. EU-

ROCAL'87, Lec. Notes in Comp. Sci. 378, Springer-Verlag, 1989, 456-463), and also in Fleischer, loc. cit. (1993), but the latter just valid for linear target sets. In the more recent paper by Montaña-Pardo-Recio (in Comp. Complex. loc. cit) a deep analysis of the method is provided, showing that it can be extended to different situations, such as semialgebraic sets having a point Nash-diffeomorphic to a corner. As an application, it is shown that the Directed-Oriented Convex Hull problem for n points in the plane (deciding that they are the oriented convex hull of themselves) has complexity at least n. We must remark that this problem is not solvable by linear decision trees (cf. A. Yao, R. Rivest "On the polyhedral decision problem" SIAM J. of Comp. 9, (1980), 343-347, also in D. Avis "Comments on a lower bound for convex hull determination", Inf. Proc. Letters 11 (1980), 126, and J. Jaromczyk "Linear decision trees are too weak for convex hull problem" Inf. Proc. Letters). Another application shows that any quantifier free formula expressing the existence of real root for a degree d polynomial must involve at least d/2 polynomials. There is also a solution to a question of Yao, about computing the largest k elements of a set of given n real numbers (cf. A. Yao, "On selecting the k largest with median tests" Algorithmica 4/2 (1989), 293-300) that must be solved in general decision trees (and not only linear as in Yao's paper) with at least $n - k + \sum \log(n - i + 1)$, for $1 \leq i \leq (k - 1)$ steps (cf. Fleischer, loc. cit. and the previously commented extensions of Rabin's method).

We comment finally that it is still a conjecture to know if, in the general case, the exact definition of width introduced by Rabin is a lower bound for the complexity –most of the above commented results are based on showing that a weaker concept of width (the generic width) is a lower bound and that, in the particular case of the hypothesis of Rabin's theorem, the width and the generic width coincide.

5 Miscellanea

A collection of specific real geometry methods has been developed by Lickteig and Burgisser for testing membership to algebraic sets (not to semi-algebraic sets). This is relevant for Algebraic Complexity problems, but no applications are known for Computational Geometry, since most problems here have strictly semialgebraic associated sets. See the forthcoming book by Burgisser-Clausen-Shokrollahi "Algebraic Complexity Theory", chapter 11, for a detailed exposition of these methods, including also basic facts and proofs of many results that we have described up to now.

On the other hand there is a kind of "stand alone" method, based on volume estimation, that has been used by Bjorner-Lovasz-Yao (loc. cit.) to show the $n\log(n/k)$ lower bound for the k-equal problem in the linear decision tree model. The interesting idea is the following: given a linearly defined set P, the number of of possible Yes leaves is bounded from below by $Vol(P)/V$ where Vol is the volume and V is the maximum volume of any subset of P defined by a path in a computation tree (thus for a convex subset of P). Then some computations are made, finding estimates for the different volumes. The idea of using metric properties of semialgebraic sets for finding lower bounds seems both very attractive and difficult.

Santander, March 20, 1995