A CELLULAR DECOMPOSITION ALGORITHM FOR SEMI-ALGEBRAIC SETS

by

Dennis Arnon

Computer Sciences Technical Report #353

June 1979

A CELLULAR DECOMPOSITION ALGORITHM FOR SEMI-ALGEBRAIC SETS[1]

Dennis Arnon[2]
Computer Sciences Department
University of Wisconsin-Madison
Madison, Wisconsin 53706

Abstract

For any $r \geq 1$ and any i, $0 \leq i \leq r$, an i-dimensional cell (in E^r) is a subset of r-dimensional Euclidean space E^r homeomorphic to the i-dimensional open unit ball. A subset of E^r is said to possess a cellular decomposition (c.d.) if it is the disjoint union of finitely many cells (of various dimensions). A semi-algebraic set S (in E^r) is the set of all points of E^r satisfying some given finite boolean combination ϕ of polynomial equations and inequalities in r variables. ϕ is called a defining formula for S. A real algebraic variety, i.e. the set of zeros in E^r of a system of polynomial equations in r variables, is a particular example of a semi-algebraic set. It has been known for at least fifty years that any semi-algebraic set possesses a c.d., but the proofs of this fact have been nonconstructive. Recently it has been noted that G. E. Collins' 1973 quantifier elimination algorithm for the elementary theory of real closed fields contains an algorithm for determining a c.d. of a semi-algebraic set S given by its defining formula, apparently the first such algorithm. Specifically, each cell c of the c.d. C of S is itself a semi-algebraic set, and for every c in C, a defining formula for c and a particular point of c are produced. In the present paper we give a description of this c.d. algorithm, in the form of a constructive proof of the theorem that any semi-algebraic set has a c.d. We then show that the algorithm can be extended to determine the dimension of each cell in a c.d. and the incidences among cells. A computer implementation of the algorithm is in progress.

---

[1]A condensed version of the present paper will appear in the Proceedings of the 1979 European Symposium on Symbolic and Algebraic Manipulation, to be published in 1979 by Springer-Verlag in the series Lecture Notes in Computer Science. Various corrections to the condensed version are incorporated in the present version.

[2]This research has been supported by NSF Grants MCS74-13278 Aol and MCS78-01731.

1. Introduction

For any $i \geq 1$, the i-dimensional open unit ball B^i in i-dimensional Euclidean space E^i is the set of all $\langle x_1, \ldots, x_i \rangle \in E^i$ such that $x_1^2 + \ldots + x_i^2 < 1$. We define B^0 to be a single point. For any $r \geq 1$ and any i, $0 \leq i \leq r$, an i-dimensional cell, or i-cell, in E^r is a subset of E^r homeomorphic to B^i. A subset of E^r is said to possess a cellular decomposition (c.d.) if it is the disjoint union of finitely many cells (of various dimensions). The following figure illustrates a cellular decomposition of the curve $y^2 - x^3 - x^2 = 0$; c_j^i denotes the j^{th} i-cell:

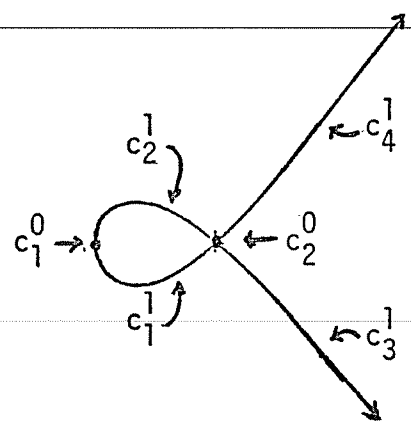

$(c_1^0 = \langle -1,0 \rangle; c_2^0 = \langle 0,0 \rangle)$

Figure 1

An i-cell c^i is incident on a j-cell c^j if $i<j$, $c^i \cap c^j = \emptyset$, and c^i is contained in the closure of c^j. For example, c_1^0 is incident on c_1^1 in Figure 1. Suppose, for a subset X of E^r possessing a c.d., that one knows the dimension of each cell and the incidences among cells

in some c.d. of X. It is intuitively evident that one can then infer certain information of a global, topological character about X; for example, the dimension of X (namely the dimension of the cell of largest dimension in the c.d.) and the number of connected components of X. In algebraic topology one formalizes this intuition by using the dimensions and incidences of the cells of a c.d. of X to define certain groups, called the homology groups of X, from which a good deal more information about X can be obtained. Thus one important reason for interest in cellular decompositions is the possibility, if the dimensions and incidences of cells are known, of using a c.d. as a means of homology group calculation.

For any $m \geq 1$, let I_m denote $Z[x_1, \ldots, x_m]$, the ring of polynomials in m variables over the integers. For any F in any I_m, we call the equation "$F = 0$" and the inequality "$F > 0$" <u>atomic formulas</u>. Any finite combination of atomic formulas built up using the boolean operations "and" ("&") and "or" ("$\vee$") will be called a <u>quantifier-free formula</u> (q.f.f.). The polynomials occurring in the atomic formulas of a q.f.f. ϕ will be called the <u>polynomials of ϕ</u>. If all the polynomials of some q.f.f. ϕ are in I_r, we will also write $\phi(x_1, \ldots, x_r)$ for ϕ. Given a q.f.f. $\phi(x_1, \ldots, x_r)$, the set of all points in Euclidean space E^r satisfying ϕ, which we denote by $S(\phi)$, is the *semi-algebraic set* defined by ϕ. ϕ is a *defining formula* (d.f.) for $S(\phi)$. We remark that the collection of semi-algebraic sets would not be enlarged if we were to allow any of the relations "$<$", "$\leq$", "$\geq$", or "$\neq$" in atomic formulas, or the boolean operation of negation in q.f.f.'s.

An example of a semi-algebraic set and a d.f. for it is the set of points of the line $y = x$ in E^2 having nonnegative coordinates and the q.f.f. $\phi(x,y) = [y - x = 0]$ & $[(y>0) \vee (y - 0)]$. Real algebraic varieties (also called "algebraic sets") are an important particular sort of semi-algebraic set; they are the semi-algebraic sets which have a d.f. which is a conjunction of equations, i.e. a "system of polynomial equations". The topology of real varieties has long been studied; for example, Hilbert's 16^{th} problem is concerned with the mutual disposition of the connected components of real algebraic curves and surfaces.

Fifty years ago van der Waerden [1929, p. 360] published a proof that any semi-algebraic set has a c.d. His argument, however, was nonconstructive, as is that of Hironaka [1974]. In 1973, G. E. Collins [1975], [1976] discovered a new quantifier elimination algorithm for the elementary theory of real closed fields, quite different from previous algorithms. The principal component of Collins' method is an algorithm which, for a given $r \geq 1$, determines a certain decomposition of E^r into connected subsets. Collins called this decomposition a "cylindrical algebraic decomposition" (c.a.d.) and called the connected subsets of E^r which comprised it "cells". For quantifier elimination, however, the relevant property of each "cell" is that certain polynomials in I_r are invariant on it. ($F \in I_r$ is <u>invariant</u> on $U \subset E^r$ if the sign of $F(\alpha)$ is the same for all $\alpha \in U$.). P. Kahn [1978] recently noted that the "cells" of a c.a.d. are in fact cells as defined above, hence a c.a.d. of E^r is a c.d. of E^r. It follows that the c.a.d. algorithm can be used to determine, in a sense we now explain, a c.d. ~~of a given semi-algebraic set.~~

Define a c.d. to be <u>semi-algebraic</u> if each of its cells is a semi-algebraic set. For $r \geq 1$, a point of E^r is <u>algebraic</u> if each of its coordinates is a real algebraic number. For any cell c, a <u>sample point</u> (s.p.) for c is an algebraic point in c. If for every cell of some semi-algebraic c.d. we have a d.f. and an s.p., we will say that we have d.f.'s and s.p.'s for that c.d. (We remark that in the present paper we assume available algorithms for needed arithmetic operations on real algebraic numbers; Rubald [1974] and Loos [1973] discuss such algorithms.)

When applied to a d.f. ϕ of some semi-algebraic set, the output of the c.a.d. algorithm consists of d.f.'s and s.p.'s for a (semi-algebraic) c.d. C of S(ϕ). At first sight there may appear to be little connection between having d.f.'s and s.p.'s for C, and the objective proposed above of obtaining the dimensions and incidences of the cells of C. In Section 3, however, we will present dimension and incidence algorithms which make essential use of d.f.'s and s.p.'s for C. In addition, these d.f.'s and s.p.'s are quite clearly of interest in their own right. For example, if ϕ is a system of polynomial equations, the d.f.'s and s.p.'s for C can be considered a solution of ϕ. By a "cellular decomposition algorithm for semi-algebraic sets", we mean an algorithm which behaves as the c.a.d. algorithm does: given ϕ, d.f.'s and s.p.'s for a (semi-algebraic) c.d. of S(ϕ) are produced.

Our main objective in the present paper is to describe Collins' c.a.d. algorithm in its usage as a c.d. algorithm for semi-algebraic sets. Our presentation takes the form of a constructive proof, in Section 2, of the theorem that any semi-algebraic set has a c.d. Thus our proofs of the Main Theorem and its auxiliary theorems in Section 2 often consist of outlines of the various subalgorithms of the c.a.d. algorithm. We remark that proofs by induction correspond to recursive algorithms.

In Section 3 we prove that the c.a.d. algorithm is easily extended so as to determine the dimension of each cell in the c.d. of a semi-algebraic set. We also give an algorithm for determining the incidences among the cells of a c.d., however it seems likely that a significantly better algorithm can be found.

Computing time analysis carried out by Collins [1975] suggests that the c.a.d. algorithm is feasible in practice, at least for semi-algebraic sets whose d.f.'s are "small", e.g. contain only a small number of polynomials in a small number of variables. As will become clear in Section 3, extending the c.a.d. algorithm to determine the dimensions of cells would add little to its computing time. The incidence algorithm we describe would add substantially to the computing time. A computer implementation of the c.a.d. algorithm is in progress.

Algorithms for solving systems of polynomial equations have been implemented in computer algebra systems for nearly twenty years (see e.g. Yun [1973], who has references to most prior work). The algorithms used have been based on the classical elimination theory of algebraic geometry (see e.g. van der Waerden [1950]). For a zero-dimensional variety, i.e. a variety consisting of finitely many points, there is little difference between the output of an elimination theory algorithm and a c.d. algorithm. Both produce the (necessarily algebraic) points of the variety. In the positive-dimensional case, however, the output of an elimination theory algorithm is a parametrization of the variety, rather than a decomposition of it. For example, given the equation $y^2 - x^3 - x^2 = 0$ of the curve of Figure 1, the ALGSYS algorithm of MACSYMA [1977, p. 98-99] produces the parametrization $y = \pm x\sqrt{x + 1}$.

2. Construction of defining formulas and sample points for the cells of a cellular decomposition

In this section we give a constructive proof of the following theorem, which we refer to as the "c.d. theorem"

Cellular Decomposition Theorem for Semi-algebraic Sets
Let S be a semi-algebraic set in E^r, $r\geq 1$, and let $\phi(x_1,..,x_r)$ be a given defining formula for S. Then there is a semi-algebraic c.d. of S for which we can construct d.f.'s and s.p.'s.

The proof will rely on what we call the Main Theorem. For any set $\mathcal{A}$ of polynomials in I_r, and any cell c in E^r, c is said to be $\mathcal{A}$-invariant if every A in $\mathcal{A}$ is invariant on c. A c.d. C is $\mathcal{A}$-invariant if every cell of C is $\mathcal{A}$-invariant.

Main Theorem Let $\mathcal{A} = \{A_1,...,A_n\}, n\geq 1$, be a finite set of nonzero polynomials in I_r, $r \geq 1$. Then there is an $\mathcal{A}$-invariant, semi-algebraic c.d. of E^r for which we can construct d.f.'s and s.p.'s.

Once we have the Main Theorem, the proof of the c.d. Theorem is brief:

Proof of the Cellular Decomposition Theorem Let $\mathcal{A}$ be the set of polynomials of ϕ. By the Main Theorem, we can construct d.f.'s and s.p.'s for the cells of an $\mathcal{A}$-invariant c.d. C of E^r. Clearly the truth value of ϕ is invariant on each cell c of C, i.e. for any cell c of C and for any $\alpha \epsilon c$, $\phi(\alpha)$ is true if and only if $\phi(\beta)$ is true for every $\beta \epsilon c$. Thus by evaluating ϕ at the s.p.'s of the cells of C, we can determine which cells belong to S. Clearly we obtain a c.d. of S in this way, for each cell of which we have a d.f. and an s.p.□

We now turn to the Main Theorem, whose proof will occupy the rest of Section 2 (except that an example of the c.d. algorithm is given following the proof). An outline of the proof by induction is as follows. Having treated the case r=1, we suppose $r \geq 2$. The objective is to construct from $\mathcal{A}$ a set $P(\mathcal{A})$ of polynomials in I_{r-1} such that s.p.'s and d.f.'s for a $P(\mathcal{A})$-invariant c.d. of E^{r-1} can be extended to d.f.'s and s.p.'s for an

$\mathcal{A}$-invariant c.d. of E^r. Then applying the inductive hypothesis to $P(\mathcal{A})$, the theorem will be proved. Theorems 1-5 establish sufficient conditions on a c.d. of E^{r-1} for its d.f.'s and s.p.'s to be extendable to d.f.'s and s.p.'s for an $\mathcal{A}$-invariant c.d. of E^r (conditions 5.1-5.3 of Theorem 5). We then, in the course of Theorems 6-10, define the appropriate $P(\mathcal{A})$ and establish that it satisfies conditions 5.1-5.3. We call the $P(\mathcal{A})$ we define the "augmented projection" of $\mathcal{A}$.

The c.d. algorithm for semi-algebraic sets we obtain from our proofs above of the c.d. theorem and the Main Theorem can be summarized as follows. Let $\phi(x_1,\ldots,x_r)$ be a given q.f.f. For any set $\mathcal{B}$ of polynomials in I_m, $m\geq 2$, let $P(\mathcal{B})$ denote the augmented projection of $\mathcal{B}$. Let $\mathcal{A}$ denote the set of polynomials of ϕ. Let $P^0(\mathcal{A}) = \mathcal{A}$. Then we compute $P^1(\mathcal{A}) = P(\mathcal{A})$, $P^2(\mathcal{A}) = P(P(\mathcal{A}))$, $P^3(\mathcal{A}) = P(P^2(\mathcal{A}))$, ... , $P^{r-1}(\mathcal{A})$. $P^{r-1}(\mathcal{A})$ is a set of polynomials in one variable. We construct d.f.'s and s.p.'s for a $P^{r-1}(\mathcal{A})$-invariant c.d. of E^1. Then we extend these to d.f.'s and s.p.'s for a $P^{r-2}(\mathcal{A})$-invariant c.d. of E^2, and continue extending until we have obtained d.f.'s and s.p.'s for an $\mathcal{A}$-invariant c.d. C of E^r. Then we evaluate ϕ at the sample point of each cell of C. Retaining those cells whose sample points satisfy ϕ, we have obtained d.f.'s and s.p.'s for a c.d. of $S(\phi)$.

Proof of the Main Theorem

We will take occasional liberties with the syntax of q.f.f.'s; in each such case it will be clear how an equivalent, syntatically correct q.f.f. could be obtained. We let R denote the real numbers.

Suppose first that r=1. Let H(x) be the product of all the A_i's in $\mathcal{A}$. If H(x) has multiple roots, then replace it with its greatest squarefree divisor, i.e. a polynomial in I_1 having the same roots with multiplicity one. Now suppose H has $k \geq 0$ real roots $\gamma_1, \ldots, \gamma_k$. Clearly the c.d. of E^1 defined by taking each γ_j as a 0-cell and each of the k+1 open intervals of $E^1 - \{\gamma_j\}$ as a 1-cell is $\mathcal{A}$-invariant and semi-algebraic. We call it the $\mathcal{A}$-induced c.d. of E^1, and denote it with $C_{\mathcal{A}}(E^1)$ or $C(\mathcal{A}, E^1)$. If k=0, we take 0 as the s.p. and $(x<0)$ & $(x = 0)$ & $(x>0)$ as a d.f. for the one cell of the c.d. If $k \geq 1$, we compute disjoint isolating intervals $J_1, \ldots, J_k$ (with rational number endpoints) for the γ_j's. For each j, $1 \leq j \leq k$, we take an exact representation for γ_j as a real algebraic number to be the s.p. for the cell γ_j, and where $J_j = (a_j, b_j)$, we take $(H = 0)$ & $(a_j < x < b_j)$ as a d.f. for γ_j. Suppose for a moment k=2. $C_{\mathcal{A}}(E^1)$ is then as indicated in Figure 2:

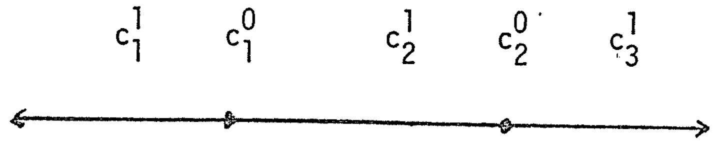

Figure 2

Suppose $J_1=(a_1,b_1)$ isolates $c_1^0 = \gamma_1$ and $J_2=(a_2,b_2)$ isolates $c_2^0 = \gamma_2$, so $a_1<b_1<a_2<b_2$. Then b_1 is in c_2^1; let $\sigma=$ sign$(H(b_1))$. We take b_1 as an s.p. for c_2^1 and $(\sigma H>0)$ & $(a_1<x<b_2)$ as a d.f. for c_2^1. The idea used in this example can be extended to an algorithm for constructing s.p.'s and d.f.'s for the 1-cells of $C_\alpha(E^1)$ for any $k\geq 1$. This completes the proof of the case $r=1$.

Now suppose $r\geq 2$ and let c be any given i-cell in E^{r-1}, with $0\leq i\leq r-1$. Define $Y(c)$, the <u>cylinder over</u> c, to be $c\times E$, a subset of E^r. When considering $Y(c)$ it will often be convenient to regard c as identical to the subset $c\times\{0\}$ of $Y(c)$. For example, Figure 3 shows the five cylinders in E^2 over the cells of the c.d. of E^1 we had in Figure 2:

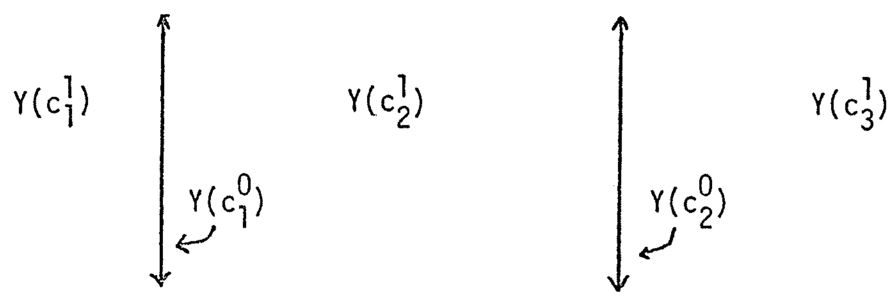

Figure 3

Let F be any given nonzero element of I_r. We write V(F) to denote the real variety defined by F = 0. We say that V(F) is <u>empty</u> on c if V(F) $\cap$ Y(c) = ϕ. For example, if F(x,y)ϵI_2 is the unit circle $y^2 + x^2 - 1$, V(F) is empty on c = $(-\infty,-2)$ in E^1. We say that V(F) is <u>cylindrical</u> on c if V(F)$\cap$Y(c) = Y(c). For example, if F(x,y) = x, then V(F) is <u>cylindrical</u> on c = 0 in E^1. We say that V(F) is <u>delineable</u> on c if V(F)$\cap$Y(c) consists of finitely many disjoint i-cells $L_1,\ldots,L_k$, $k \geq 1$, such that for each j, $1 \leq j \leq k$,

1) L_j is the graph of a continuous function f_j from c to E.
2) there is an integer $e_j \geq 1$ such that for any $\alpha \epsilon c, f_j(\alpha)$ is a root of F (α,x_r) of multiplicity e_j.

The L_j are called the <u>branches</u> of V(F) on c. For example, if F(x,y) = $y^2 - x^3 - x^2$, then V(F) is delineable with two branches on c = (-3/4,-1/4) in E^1:

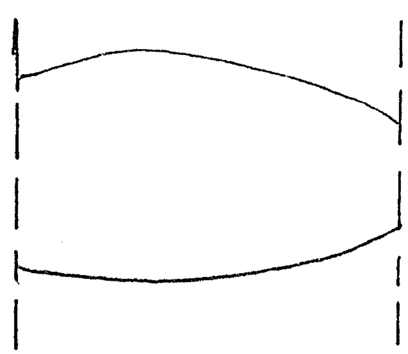

Figure 4

but V(F) is not delineable on c = (-1/2,2) since the "branches are not disjoint":

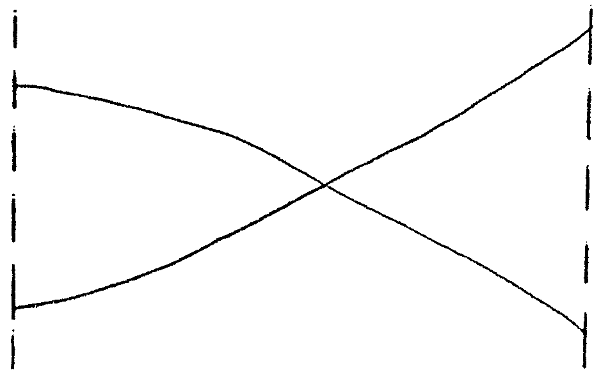

Figure 5

If $F(x,y) = y^3 - x$, then V(F) is not delineable on c = (-1,1):

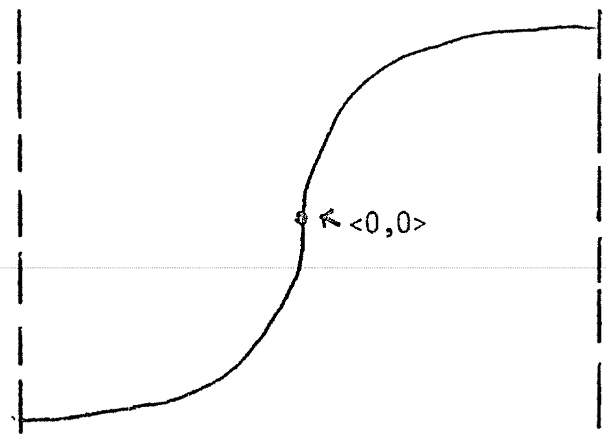

Figure 6

The reason is that $F(1/2,y) = y^3 - 1/2$ has one real root of multiplicity one, but $F(0,y) = y^3$ has one real root of multiplicity three. As Figure 5 illustrates, a given $F \varepsilon I_r$ may be neither empty, cylindrical, nor delineable on any given cell c in E^{r-1}. If V(F) is empty (cylindrical, delineable) on c, we will also say that F is empty (cylindrical, delineable) on c.

If V(F) is delineable on c, there is an obvious F-invariant c.d. of Y(c), consisting of the k branches of V(F) as i-cells and the k+1 connected components of Y(c) - V(F) as (i+1) cells. (It is straightforward to verify that each of the connected components of Y(c) - V(F) is indeed an (i+1)-cell.) We call this the <u>F-induced c.d. of Y(c)</u>, and denote it with $C_F(Y(c))$ or C(F,Y(c)). We also define $C_F(Y(c))$ in case V(F) is empty on c; it then consists of the single (i+1)-cell Y(c). With the following theorem we show that an s.p. for c can be extended to s.p.'s for $C_F(Y(c))$ when V(F) is delineable on c. By an <u>algebraic polynomial</u> we mean a polynomial with real algebraic number coefficients.

<u>Theorem 1</u> Suppose $F \epsilon I_r$ is delineable on a cell c in E^{r-1}. Then given an s.p. α for c, we can construct s.p.'s for $C_F(Y(c))$.

<u>Proof</u> Let $g(x_r)$ be the algebraic polynomial $F(\alpha, x_r)$. Let $b_1, \ldots, b_{2k+1}$, $k \geq 0$, be the s.p.'s for $C_g(E^1)$ constructed as in our argument above for the case r=1 of the Main Theorem. (We assume we can isolate real roots and do other required operations on algebraic polynomials; Rump [1976], for example, presents the algorithms

we need). Then $\langle\alpha,b_1\rangle,\ldots,\langle\alpha,b_{2k+1}\rangle$ are s.p.'s for $C_F(Y(c))$□

For any commutative ring I and any G in I[x], suppose $G\neq 0$ and $G = \Sigma_{i=0}^{m} g_i x^i$, $g_m\neq 0$. Then m is the degree of G, written deg(G), and g_m is the leading coefficient of G, written ldcf(G). If G = 0, then deg(G) = 0 and ldcf(G) = 0. For any k, $0\leq k\leq \deg(G)$, let $\mathrm{der}^k(G)$ denote the k^{th} derivative of G. Hence $\mathrm{der}^0(G) = G$, $\mathrm{der}^1(G)=G'$, etc. We define the derivative set of G, written DER(G), to be

$$\{\mathrm{der}^k(G)\,|\,0\leq d\leq \deg(G)\ \&\ \mathrm{der}^k(G)\neq 0\}.$$

We apply these definitions to I_r by viewing it as $I[x_r]$ with $I = Z[x_1,\ldots x_{r-1}]$. For an $F\varepsilon I_r$ delineable on a cell c, the next theorem establishes the conditions under which a d.f. for c can be extended to d.f.'s for $C_F(Y(c))$.

Theorem 2 Suppose $F\varepsilon I_r$ is delineable on a cell c in E^{r-1}. Suppose also that every element of DER(F) is delineable or empty on c. Then given an s.p. α and a d.f. ϕ_c for c, we can construct d.f.'s for $C_F(Y(c))$.

As an immediate consequence of Theorem 2, we have

Corollary 1 If $F\varepsilon I_r$ is delineable or empty on a cell c in E^{r-1}, then $C_F(Y(c))$ is a semialgebraic c.d.

Proof of Theorem 2 The proof of Theorem 2 is long, and contains Lemmas 1-10.

Let $g(x)$ be any real polynomial, i.e. any polynomial with real coefficients. Suppose g has distinct real roots $\gamma_1<\ldots<\gamma_t$ for some $t\geq 1$. Then

$$c_1 = (-\infty,\gamma_1)$$
$$c_2 = \gamma_1$$
$$\vdots$$
$$c_{2t} = \gamma_t$$
$$c_{2t+1} = (\gamma_t,\infty)$$

are the cells of $C_g(E^1)$. We associate three sets of nonnegative integers with g as follows:

$\mu_i(g)$, for $1\leq i\leq t$, is the multiplicity of γ_i as a root of g.

$\sigma_i(g)$, for $1\leq i\leq t+1$, is the sign of g on c_{2i-1}.

$\lambda_i(g)$, for $1\leq i\leq t$, is the number of distinct real roots of $g'(x)$ less than or equal to γ_i, and $\lambda_{t+1}(g)$ is the total number of distinct real roots of $g'(x)$.

We now show that given the μ's, σ's, and λ's for g and certain of its derivatives, we can construct d.f.'s for $C_g(E^1)$ of a special form.

Lemma 1 Let $g(x)$ be a real polynomial with $\deg(g)\geq 1$. If $g(x)$ has no real roots let $m = -1$; otherwise let $m\geq 0$ be maximal such that $\mathrm{der}^j(g)$ has at least one real root for $0\leq j\leq m$. Suppose that for $0\leq j\leq m$, we are given the μ's, σ's, and λ's associated with $der^j(g)$. Then we can construct d.f.'s for $C_g(E^1)$ such that every polynomial occurring in any d.f. is either $\pm$ h for some element h of DER(g) or the zero polynomial.

Proof We proceed by induction on deg(g). If deg(g) = 1, then $g(x) = ax + b$ for real numbers a and b, $a \neq 0$, g has one real root $\gamma_1 = -b/a$, and $C_g(E^1)$ consists of the three cells $c_1 = (-\infty, \gamma_1)$, $c_2 = \gamma_1$, and $c_3 = (\gamma_1, \infty)$. Let σ_1 and σ_2 be the σ's for g. Then where ϕ_i denotes a d.f. for c_i, we may set $\phi_1 \leftarrow (\sigma_1 g > 0)$, $\phi_2 \leftarrow (g = 0)$, and $\phi_3 \leftarrow (\sigma_2 g > 0)$.

Suppose now that $\deg(g) > 1$. Let $\gamma_1 < \gamma_2 < \ldots < \gamma_k, k \geq 0$, be the distinct real roots of g. If $k = 0$, then C_g consists of the single cell $c_1 = E^1$. (For the remainder of the proof of Lemma 1, we will write C_g for $C_g(E^1)$ and $C_{g'}$ for $C_{g'}(E^1)$). Hence we can set $\phi_1 \leftarrow (0 = 0)$. If $k \geq 1$, then C_g consists of the $2k+1$ cells $c_1 = (-\infty, \gamma_1)$, $c_2 = \gamma_1, c_3 = (\gamma_1, \gamma_2), \ldots, c_{2k} = \gamma_k, c_{2k+1} = (\gamma_k, \infty)$. By the inductive hypothesis, we can construct d.f.'s for $C_{g'}$ such that the only polynomials occurring, up to sign, are elements of $DER(g') \subset DER(g)$ or the zero polynomial. Let ϕ_i denote the d.f. we wish to construct for the cell c_i of C_g for $1 \leq i \leq 2k+1$. Let μ_i $(1 \leq i \leq k)$, σ_i $(1 \leq i \leq k+1)$, and λ_i $(1 \leq i \leq k+1)$ denote the μ's, σ's and λ's associated with g. We recall Rolle's theorem: between any two real roots of g is a real root of g'. We note also that, by definition of λ_{k+1}, $C_{g'}$ has $2\lambda_{k+1}+1$ cells. We let c_j' denote the j^{th} cell of $C_{g'}$, and ϕ_j' the d.f. we assume we have for it.

Consider first the cells $c_{2i} = \gamma_i, 1\le i\le k$. If $\mu_i>1$, then γ_i is also a root of g', in fact the $\lambda_i{}^{th}$ real root of g'. Hence $\phi_{2\lambda_i}'$ is a d.f. for γ_i as a cell of $C_{g'}$, so we may simply set $\phi_{2i}\leftarrow\phi_{2\lambda_i}'$. If $\mu_i=1$, then γ_i is not a root of g'. The $\lambda_i{}^{th}$ real root δ (take $\delta = -\infty$ if $\lambda_i=0$) of g' is less than γ_i, and the $(\lambda_i+1)^{st}$ real root ξ (take $\xi = \infty$ if there is no such real root of g') of g' is greater than γ_i, hence γ_i is in $(\delta,\xi) = c_{2\lambda_i+1}'$. Furthermore, by Rolle's theorem, γ_i is the unique root of g in $c_{2\lambda_i+1}'$. Hence we may set $\phi_{2i} \leftarrow (g = 0)\ \&\ \phi_{2\lambda_i+1}'$.

Consider now the cells $c_{2i-1} = (\gamma_{i-1},\gamma_i)$, $2\le i\le k$. Suppose for some c_{2i-1} that $\mu_{i-1}>1$ and $\mu_i=1$. Intuitively what we want to do is form the disjunction of the d.f.'s for the cells of $C_{g'}$ which overlap c_{2i-1}. At the lower end of c_{2i-1} this works out neatly, because since $\mu_{i-1}>1$, $c_{2i-2} = c_{2\lambda_{i-1}}' = \gamma_{i-1}$, so for c_{2i-1} we can start off with

$$\phi_{2\lambda_{i-1}+1}' \vee \phi_{2\lambda_{i-1}+2}' \vee \dots$$

But at the upper end of c_{2i-1} we have γ_i in the interior of $c'_{2\lambda_i+1}$. By Rolle's theorem, however, there are exactly three cells of C_g meeting $c'_{2\lambda_i+1}$, namely c_{2i-1}, $c_{2i} = \gamma_i$, and $c_{2i+1} = (\gamma_i, \gamma_{i+1})$. By definition of the σ's, $\mu_i = 1$ implies $\sigma_i = -\sigma_{i+1}$, i.e. the sign of g is different on each of c_{2i-1}, c_{2i}, and c_{2i+1}. Hence the q.f.f. $(\sigma_i g>0)$ & $\phi'_{2\lambda_i+1}$ is a d.f. for $c_{2i-1} \cap c'_{2\lambda_i+1}$. Thus our d.f. for c_{2i-1} is

$$\phi_{2i-1} = [\phi'_{2\lambda_{i-1}+1} \vee \phi'_{2\lambda_{i-1}+2} \vee \ldots \vee \phi'_{2\lambda_i}] \vee [(\sigma_i g>0) \,\&\, \phi'_{2\lambda_i+1}].$$

From the above discussion it should be clear how to handle the other three combinations of values of μ_{i-1} and μ_i which can occur: $\mu_{i-1}>1$, $\mu_i>1$ is the easiest, $\mu_{i-1}=1, \mu_i>1$ is similar to the above, and $\mu_{i-1}=1$, $\mu_i=1$ requires special attention to both ends of c_{2i-1}.

There remain only the cells $c_1 = (-\infty, \gamma_1)$ and $c_{2k+1} = (\gamma_k, \infty)$ to consider. For ϕ_1, one begins by forming the disjunction $\phi'_1 \vee \phi'_2 \vee \ldots$ and, by exactly the considerations we had above, terminates it according to the value of μ_1. For ϕ_{2k+1}, it is the beginning of the disjunction that requires care; one then takes the disjunction of all remaining d.f.'s of C_g, up to and including $\phi'_{2\lambda_{k+1}+1}$. This completes the proof of Lemma 1 □

Returning to the $F\epsilon I_r$ and the cell c we are given in the hypotheses of Theorem 2, let $g_\beta(x_r) = F(\beta, x_r)$ for any $\beta\epsilon c$. Then since F is delineable on c, there is a $k\geq 1$ such that for every $\beta\epsilon c$, g_β has k distinct real roots $\gamma_{\beta,1} <\ldots< \gamma_{\beta,k}$. For any $\beta\epsilon c$, let

$$c_1^\beta = (-\infty, \gamma_{\beta,1})$$

$$c_2^\beta = \gamma_{\beta,1}$$

$$\vdots$$

$$c_{2k+1}^\beta = (\gamma_{\beta,k}, \infty)$$

be the cells of $C(g_\beta, E^1)$. For each i, $1\leq i\leq 2k+1$, let $\phi_i^\beta(x_r)$ denote the d.f. for c_i^β we could construct by Lemma 1 if we had the necessary μ's, σ's, and λ's associated with g_β and its derivatives. Every polynomial occurring in ϕ_i^β is either $\pm h$ for some $h\epsilon DER(g_\beta)$ or the zero polynomial. Suppose now that for some $\beta\epsilon c$ we are actually given $\phi_i^\beta(x_r)$ for $1\leq i\leq 2k+1$. For each i, $1\leq i\leq 2k+1$, define $\phi_{\beta,i}(x_1,\ldots,x_r)$ to be the q.f.f. we obtain by replacing every occurrence of $der^j(g_\beta)$ in ϕ_i^β with $der^j(F)$, for each j, $0\leq j\leq deg(g_\beta)$. Where $F^{(j)} = F^{(j)}(x_1,\ldots,x_r)$ denotes $der^j(F)$ for any j, it is easily verified that $F^{(j)}(\xi, x_r) = der^j(F(\xi, x_r))$ for any j and and for any $\xi\epsilon E^{r-1}$. It follows that for each i,

$1 \le i \le 2k+1$, $\phi_{\beta,i}(\beta, x_r) = \phi_i^{\beta}(x_r)$. We now show with Lemmas 2-7 that, given the hypotheses of Theorem 2, $\phi_{\beta,i}$ does not depend on the particular choice of $\beta \epsilon c$ (for each i, $1 \le i \le 2k+1$). We can then state, in Corollary 3, the key property of the $\phi_{\beta,i}$'s that enables us to use them in d.f.'s for $C_F(Y(c))$.

Lemma 2 For any $G \epsilon I_r$, suppose G is delineable with $m \ge 1$ branches on c. For any $\beta \epsilon c$, let $h_\beta(x_r) = G(\beta, x_r)$, and let $\delta_{\beta,1} < \ldots < \delta_{\beta,m}$ be the distinct real roots of h_β. Then for any j, $1 \le j \le m$, and for any $\beta, \bar{\beta} \epsilon c$, $\mu_j(h_\beta) = \mu_j(h_{\bar{\beta}})$.

Proof Immediate from the definition of delineability □

Lemma 3 Suppose $G \epsilon I_r$ satisfies the hypotheses of Lemma 2. Then for any j, $1 \le j \le m+1$, and for any $\beta, \bar{\beta} \epsilon c$, $\sigma_j(h_\beta) = \sigma_j(h_{\bar{\beta}})$.

Proof For any j, $1 \le j \le m+1$, c_{2j-1}^{β} and $c_{2j-1}^{\bar{\beta}}$ are both contained in the same cell of $C_G(Y(c))$. Then since $C_G(Y(c))$ is G-invariant, $\sigma_j(h_\beta) = \sigma_j(h_{\bar{\beta}})$ □

Lemma 4 For any $G \epsilon I_r$, suppose G and G' are delineable on c. Let L be any branch of V(G) on c, and let $V_c(G') = V(G') \cap Y(c)$. If $L \cap V_c(G') \ne \emptyset$, then $L \cap V_c(G') = L$.

Proof Let $f: c \to E$ be the continuous function whose graph is L. By hypothesis there is some $\beta \epsilon c$ such that $f(\beta)$ is a root of $G(\beta,x_r)$ of multiplicity $e \geq 2$. But then by definition of delineability, $f(\xi)$ is a root of $G(\xi,x_r)$ of multiplicity e for every $\xi \varepsilon c$, i.e. $f(\xi)$ is a root of $G'(\xi,x_r)$ for every $\xi \varepsilon c$, i.e. $L \cap V_c(G') = L$ □

Lemma 5 For any $G \varepsilon I_r$, suppose G and G' are delineable on c. Let L be any branch of $V(G)$ on c, and suppose $L \cap V_c(G') \neq \emptyset$. Then there is some branch L' of $V(G')$ on c such that $L = L'$.

Proof The assertion is trivial if c is a 0-cell, so assume c is an s-cell with $s \geq 1$. Suppose $m \geq 2$ branches $L_1', \ldots, L_m'$ of $V(G')$ on c meet L. Let $f, f_1', \ldots, f_m'$ be the continuous functions from c to E whose respective graphs are L, $L_1', \ldots, L_m'$. Choose $\beta, \xi \varepsilon c$ such that $f(\beta) = f_1'(\beta)$ and $f(\xi) \neq f_1'(\xi)$. Since c is homeomorphic to an open connected subset of E^s, c is pathwise connected. Let $t:[0,1] \rightarrow c$ be a path from β to ξ, i.e. $t(0)=\beta$, $t(1)=\xi$. Let $V = \{x \varepsilon [0,1] \mid f(t(x)) = f_1'(t(x))\}$. V is nonempty and bounded above, hence has a least upper bound v. Suppose $v \varepsilon V$. Then $v<1$, and there is some $f_j' \neq f_1'$ such that for all $\delta>0$, there is a $y>v$ with $y - v<\delta$, and $f(t(y)) = f_j'(t(y))$. Let $\rho = |f(t(v) - f_j'(t(v))|$. If $\rho>0$, then choosing $\varepsilon=\rho/2$, by continuity of f and f_j' at v, there is a $\delta>0$ such that $|x - v|<\delta$ implies $|f(t(x) - f(t(v))|<\varepsilon$ and $|f_j'(t(x)) - f_j'(t(v))|<\varepsilon$. But this contradicts the existence of a $y>v$ with $y - v<\delta$ and $f(t(y)) = f_j'(t(y))$.

Hence we must have $\rho=0$, i.e. $f_j'(t(v)) = f(t(v)) = f_1'(t(v))$. But this contradicts the disjointness of L_1' and L_j'. If $v \notin V$, by a similar argument we again arrive at a contradiction. Hence we must have $m=1$, i.e. only one branch of $V(G')$ on c meets L. But then by Lemma 4, there is some branch L' of $V(G')$ on c such that $L = L'$ □

Corollary 2 For any $G \epsilon I_r$. suppose G and G' are delineable on c. Then every branch of G' on c is either a branch of G on c or disjoint from all branches of G on c.

Lemma 6 For any $G \epsilon I_r$, suppose G is delineable with $m \geq 1$ branches on c and G' is delineable or empty on c. For any $\beta \epsilon c$, let $h_\beta(x_r) = G(\beta, x_r)$, and let $\delta_{\beta,1} < \ldots < \delta_{\beta,m}$ be the distinct real roots of h_β. Then for any j, $1 \leq j \leq m+1$, and for any $\beta, \bar{\beta} \epsilon c, \lambda_j(h_\beta) = \lambda_j(h_{\bar{\beta}})$.

Proof If G' is empty on c, then for every j, $1 \leq j \leq m+1$, and for every $\beta \epsilon c, \lambda_j(h_\beta) = 0$. So assume G' is delineable on c. By Corollary 2, every branch of G' is either a branch of G on c or disjoint from all branches of G on c. Thus for any j, $1 \leq j \leq m$, and any $\beta, \bar{\beta} \epsilon c$, the number of roots of h_β' less than or equal to $\delta_{\beta,j}$ is the same as the number of roots of $h_{\bar{\beta}}'$ less than or equal to $\delta_{\bar{\beta},j}$, i.e. $\lambda_j(h_\beta) = \lambda_j(h_{\bar{\beta}})$. By the delineability of G' on c, $\lambda_{m+1}(h_\beta) = \lambda_{m+1}(h_{\bar{\beta}})$ □

Lemma 7 Assume the hypotheses of Theorem 2, i.e. $F \epsilon I_r$ is delineable on a cell c in E^{r-1} and every element of

DER(F) is delineable or empty on c. For any $\beta\epsilon c$, $\gamma_{\beta,1}<\ldots<\gamma_{\beta,k}$ are the distinct real roots of $g_\beta(x_r) = F(\beta,x_r)$. Then for each i, $1\le i\le 2k+1$, and any $\beta,\bar{\beta}\epsilon c$, $\phi_{\beta,i} = \phi_{\bar{\beta},i}$.

Proof For any $j\ge 0$, let $F^{(j)}$ denote $der^j(F)$. Since every element of DER(F) is delineable or empty on c, $F^{(j)}(\beta,x_r)$ has the same number of real roots for every $\beta\epsilon c$, for any $F^{(j)}\epsilon DER(F)$. Also, for any delineable $F^{(j)}$ in DER(F), the μ's, σ's, and λ's associated with $F^{(j)}(\beta,x_r) = der^j(g_\beta(x_r))$ for any $\beta\epsilon c$ are independent of the particular choice of β, by Lemmas 2-6. Hence when we examine the recursive algorithm implicit in the proof of Lemma 1, we see that for any $\beta,\bar{\beta}\epsilon c$ and any i, $1\le i\le 2k+1$, when we apply this algorithm to construct ϕ_i^β from g_β, and $\phi_i^{\bar{\beta}}$ from $g_{\bar{\beta}}$, we end up with only the following difference between ϕ_i^β and $\phi_i^{\bar{\beta}}$: every occurrence of $der^j(g_\beta)$, for any j, in ϕ_i^β is replaced by an occurrence of $der^j(g_{\bar{\beta}})$ in $\phi_i^{\bar{\beta}}$. But then by definition of $\phi_{\beta,i}$ and $\phi_{\bar{\beta},i}$, $\phi_{\beta,i} = \phi_{\bar{\beta},i}$ for each i, $1\le i\le 2k+1$ □

Henceforth, for any i, $1\le i\le 2k+1$, and any $\beta\epsilon c$, we shall simply write ϕ_i instead of $\phi_{\beta,i}$.

Corollary 3 Assume the hypotheses of Lemma 7. Then for $1\le i\le 2k+1$, and for any $\beta,\xi\epsilon c$, $\phi_{\beta,i}(\xi,x_r) = \phi_i(\xi,x_r) = \phi_i^\xi(x_r)$.

We now show, with Lemmas 8-10, that we can construct the μ's, σ's, and λ's associated with $F(\alpha,x_r)$.

This enables us to construct ϕ_1, $\phi_2, \ldots, \phi_{2k+1}$, and thereby to construct the d.f.'s for $C_F(Y(c))$. Lemmas 8-10 are concerned with an arbitrary algebraic polynomial $g(x)$ with $k \geq 1$ real roots $\gamma_1 < \ldots < \gamma_k$. We let c_1, $c_2, \ldots, c_{2k+1}$ be the cells of $C_g(E^1)$, defined as was done prior to Lemma 1.

Lemma 8 For $1 \leq i \leq k$, we can compute $\mu_i(g)$.

Proof We compute a squarefree factorization

$$g_1(x)^{e_1} g_2(x)^{e_2} \ldots g_m(x)^{e_m}$$

of $g(x)$. That is, the g_s's are nonconstant, pairwise relatively prime polynomials without multiple roots. Then, e.g. by isolating the real roots of each g_s, we determine, for $1 \leq i \leq k$, which g_s has γ_i as a root, and set $\mu_i \leftarrow e_s$ □

Lemma 9 For $1 \leq i \leq k+1$, we can compute $\sigma_i(g)$.

Proof Clearly for sufficiently large values of x, $\text{sign}(g(x)) = \text{sign}(\text{ldcf}(g))$. Hence $\sigma_{k+1} = \text{sign}(\text{ldcf}(g))$. Suppose for any j, $1 \leq j \leq k$, that we have determined σ_{j+1}. Let $h_j(x)$ be such that

$$g(x) = (x - \gamma_j)^{\mu_j} h_j(x).$$

$h_j(x)$ is invariant (and nonzero) on $c_{2j-1} \cup c_{2j+1}$. Hence we may set

$$\sigma_j \leftarrow (-1)^{\mu_j} \sigma_{j+1} \; \square$$

Lemma 10 For $1 \leq i \leq k+1$, we can compute $\lambda_i(g)$.

Proof Let $t(x) = g'(x) / \gcd(g(x), g'(x))$. Then the real roots of t are precisely the real roots of g' which are not also roots of g. By obtaining isolating intervals for the real roots of t disjoint from any isolating interval for a γ_i, we can determine nonnegative integers n_i, $1 \leq i \leq k+1$, as follows. n_1 is the number of real roots of t less than γ_1; for $2 \leq i \leq k$, n_i is the number of real roots of t between γ_{i-1} and γ_i; and n_{k+1} is the number of real roots of t greater than γ_k. Then the γ_i's are determined with the following algorithm:

$\lambda_1 \leftarrow n_1$; if $\mu_1 > 1$ then $\lambda_1 \leftarrow \lambda_1 + 1$;

for $j=2,\ldots,k$ do

$\{\lambda_j \leftarrow \lambda_{j-1} + n_j$; if $\mu_j > 1$ then $\lambda_j \leftarrow \lambda_j + 1\}$;

$\lambda_{k+1} \leftarrow \lambda_k + n_{k+1}$. □

Corollary 4 Assume the hypotheses of Theorem 2, i.e. F is delineable with $k \geq 1$ branches on c, every element of $DER(F)$ is delineable or empty on c, and for any $\beta \varepsilon c$, we set $g_\beta(x_r) = F(\beta, x_r)$. Furthermore, we are given a sample point α for c. Then we can construct $\phi_1, \phi_2, \ldots, \phi_{2k+1}$.

Proof By Lemmas 8-10, we can construct the μ's, σ's, and λ's corresponding to $F^{(j)}(\alpha, x_r) = der^j(g_\alpha)$ for any delineable $F^{(j)}$ in $DER(F)$. Then by Lemma 1, we can construct $\phi_1^\alpha, \ldots, \phi_{2k+1}^\alpha$. Then by applying the definition of

$\phi_{\alpha,i}$ for $1 \le i \le 2k+1$, we can construct $\phi_1, \ldots, \phi_{2k+1}$ □

Continue proof of Theorem 2

We now conclude the proof of Theorem 2. For $1 \le i \le 2k+1$, define $\hat{c}_i = \{<\beta,b> | \beta \varepsilon c \ \& \ b \varepsilon \ c_i^{\beta}\}$. Clearly $\hat{c}_1, \hat{c}_2, \ldots, \hat{c}_{2k+1}$ are the cells of $C_F(Y(c))$. Then for each i, $1 \le i \le 2k+1$,

$$\hat{c}_i = \{<\beta,b> | \beta \varepsilon c \ \& \ b \varepsilon S(\phi_i^{\beta}(x_r))\}$$

$= \{<\beta,b> | \beta \varepsilon c \ \& \ b \varepsilon S(\phi_i(\beta,x_r))\}$, by Corollary 3,

$= S(\phi_c \ \& \ \phi_i)$, where ϕ_c is the given d.f. for c.

Hence $\phi_c \ \& \ \phi_i$ is a d.f. for c_i □

We return now to consideration of the given set $\mathcal{A} = \{A_1, \ldots, A_n\}$ of polynomials in I_r. For any cell $c \subset E^{r-1}$, we let $\pi^c A_i$ denote the product of those A_i's in $\mathcal{A}$ which are delineable on c; if no A_i is delineable on c we set $\pi^c A_i = 1$.

Theorem 3 Suppose that for every $A_i \varepsilon \mathcal{A}$, $V(A_i)$ is either cylindrical, empty, or delineable on a cell c in E^{r-1}. Let $H = \pi^c A_i$, and suppose V(H) is delineable or empty on c. Then $C_H(Y(c))$ is $\mathcal{A}$-invariant.

Proof Clearly any A_i for which $V(A_i)$ is empty or cylindrical on c is invariant on $C_H(Y(c))$. V(H) is empty on c if and only if H=1, i.e. if and only if every $V(A_i)$ is empty or cylindrical on c, in which case the single cell $Y(c)$ of $C_H(Y(c))$ is clearly $\mathcal{a}$-invariant. Suppose $V(H)$ is delineable on c. For every delineable A_i, any branch L_i of $V(A_i)$ on c is contained in one or more branches of $V(H)$ on c. Then by an argument similar to that given in Lemma 5, L_i is a branch of $V(H)$ on c. Hence each branch of $V(H)$ on c either is a branch of $V(A_i)$ or does not meet any branch of $V(A_i)$ on c. Hence $C_H(Y(c))$ is A_i-invariant, since $C(A_i,Y(c))$ is A_i-invariant □

The next theorem establishes that, when the hypotheses of Theorem 3 are met, to obtain d.f.'s for $C_H(Y(c))$ it is enough to have d.f.'s for $C(A_i,Y(c))$ for each delineable A_i in $\mathcal{a}$. For we can then "merge" the d.f.'s from the various c.d.'s $C(A_i,Y(c))$ into d.f.'s for $C_H(Y(c))$. Obtaining s.p.'s for $C_H(Y(c))$ is straightforward.

Theorem 4 Suppose that for every $A_i \varepsilon \mathcal{a}$, $V(A_i)$ is either cylindrical, empty, or delineable on a cell c in E^{r-1}. Let $H = \pi^c A_i$, and suppose $V(H)$ is delineable or empty on c. If $V(H)$ is delineable on c, suppose we are given d.f.'s for $C(A_i,Y(c))$ for each A_i delineable on c. If $V(H)$ is empty on c, suppose we are given a d.f. $\phi_c(x_1,\ldots,x_{r-1})$ for c. Then if we are also given an s.p. α for c, we can construct s.p.'s and d.f.'s for

$C_H(Y(c))$.

Proof If $V(H)$ is empty on c, then $C_H(Y(c))$ consists of the single cell $Y(c)$. We may take $\langle\alpha,0\rangle$ as an s.p. for $Y(c)$ and ϕ_c as a d.f. for $Y(c)$. Suppose H is delineable on c. By Theorem 1, we can obtain s.p.'s for $C_H(Y(c))$. As we saw in the proof of Theorem 3, the branches of any delineable A_i on c are a subset of the branches of H on c. For any delineable A_i, by comparing the isolated real roots of $A_i(\alpha,x_r)$ with the isolated real roots of $H(\alpha,x_r)$, we can determine which branches of H are also branches of A_i. Thus we can use the d.f.'s from $C(A_i,Y(c))$ for the branches of H which are also branches of A_i. We now indicate with an example how one constructs d.f.'s for the cells of $C_H(Y(c))$ which are not branches of H. Suppose $\mathcal{A} = \{A_1,A_2\}$, r=2, c is a 1-cell, A_i is delineable on c with two branches, and A_2 is delineable on c with one branch. Let c_1^2 denote the cell of $C(A_1,Y(c))$ between the branches of A_1, and c_2^2 the cell of $C(A_2,Y(c))$ below the branch of A_2. So we have

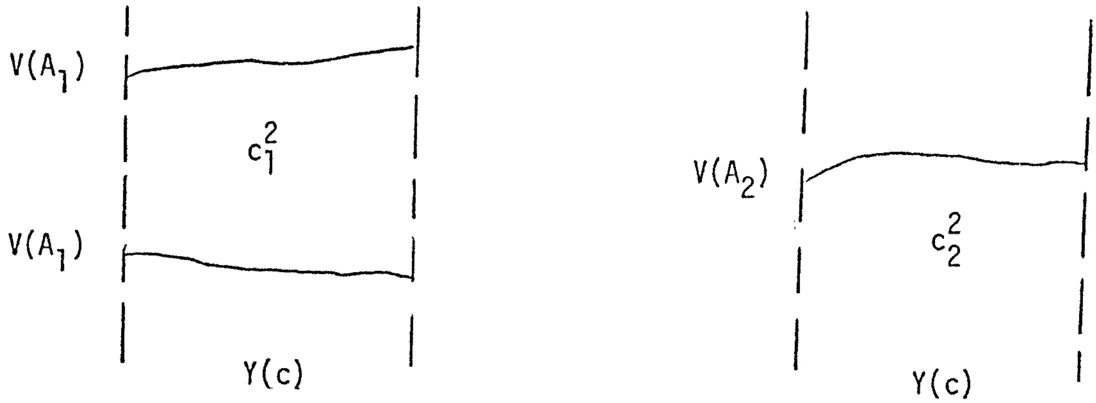

Figure 7

Suppose now that in V(H), the branch of A_2 lies between the branches of A_1, and we wish to obtain a d.f. for the shaded cell in $C_H(Y(c))$.

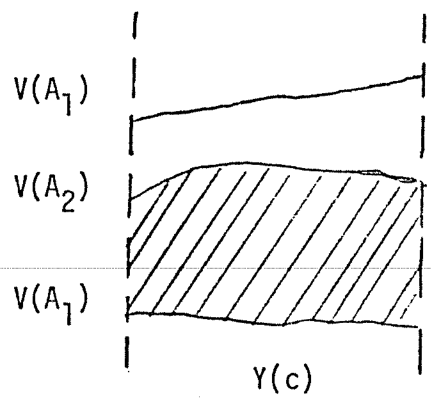

Figure 8

Then where ϕ_1 is a d.f. for c_1^2 in $C(A_1,Y(c))$ and ϕ_2 is a d.f. for c_2^2 in $C(A_2,Y(c))$, ϕ_1 & ϕ_2 is a d.f. for the shaded cell in $C_H(Y(c))$. The idea used in this particular example can be extended to an algorithm for constructing d.f.'s for $C_H(Y(c))$ for any $\mathcal{A}$ □

We now summarize the point to which Theorems 1-4 have brought us.

<u>Theorem 5</u> Given a finite set $\mathcal{A} = \{A_1,\ldots,A_n\}$ of polynomials in I_r, $r \geq 2$, suppose we have d.f.'s and s.p.'s for a (semi-algebraic) c.d. C of E^{r-1} such that the following conditions are met on every cell c of C:

5.1) Each A_i in $\mathcal{A}$ is delineable, cylindrical, or empty on c,

5.2) $\pi^c A_i$ is delineable or empty on c.

5.3) For each A_i in $\mathcal{A}$, if A_i is delineable on c then every element of $DER(A_i)$ is delineable or empty on c.

Then we can construct d.f.'s and s.p.'s for a semi-algebraic $\mathcal{A}$-invariant c.d. of E^r.

<u>Proof</u> It is sufficient to show that, for every cell c if C, we can construct d.f.'s and s.p.'s for a semi-algebraic $\mathcal{A}$-invariant c.d. of Y(c). For any c in C, let $H = \pi^c A_i$. By Corollary 1, $C_H(Y(c))$ is a semi-algebraic c.d. By Theorem 2, for each A_i delineable on c, we can construct d.f.'s for $C(A_i, Y(c))$. Then by Theorem 4, we can construct d.f.'s and s.p.'s for $C_H(Y(c))$. By Theorem 3 $C_H(Y(c))$ is $\mathcal{A}$-invariant □

Detailed proofs of Theorem 6 and Lemmas 3 and 4 are given in [Collins 1975, p. 6-9] and so are omitted here.

Theorem 6 For any $F \in I_r$, if ldcf(F) vanishes nowhere on a cell c in E^{r-1}, and the number of distinct (real and complex) roots of $F(\alpha, x_r)$ is the same for every $\alpha \in c$, then V(F) is delineable or empty on c.

(We note that Collins uses a stronger definition of delineability than ours; if F is delineable by his definition then it is delineable by our definition, but the converse need not be true).

For any $m \geq 1$, any nonzero $F, G \in I_m$, and any j with $0 \leq j \leq \min(\deg(F), \deg(G))$, let $S_j(F,G)$ denote the j^{th} subresultant of F and G. Define the j^{th} principal subresultant coefficient of F and G, denoted $psc_j(F,G)$, to be the coefficient of x_r^j in $S_j(F,G)$. From the fundamental theorem of polynomial remainder sequences [Brown and Traub 1971] we obtain

Lemma 11 For any $m \geq 1$, let F and G be nonzero elements of I_m. Then deg(gcd(F,G)) = k if and only if k is the least j such that $psc_j(F,G) \neq 0$.

From Leibnitz's rule we obtain

Lemma 12 Let g(x) be a univariate polynomial with real coefficients, such that deg(g) = $d \geq 1$, and let k= deg(gcd(g,g')). Then d-k is the number of distinct (real and complex) roots of g.

For any $F,G\varepsilon I_m, m\geq 1$, define the psc set of F and G, denoted PSC(F,G), to be

$\{psc_j(F,G) \mid 0\leq j<\min(\deg(F),\deg(G))\ \&\ psc_j(F,G)\neq 0\}$,

From Lemmas 11 & 12 and Theorem 6 we have immediately

Theorem 7 Let F be a nonzero element of I_r, and let c be a PSC(F,F')-invariant cell in E^{r-1} such that ldcf(F) vanishes nowhere on c. Then V(F) is delineable or empty on c.

For any commutative ring I and any nonzero $G\varepsilon I[x]$, suppose $G = \Sigma_{i=0}^{m} g_i x^i, g_m\neq 0$. For $0\leq k\leq m$, define the k^{th} reductum of G, denoted $red^k(G)$, to be $\Sigma_{i=0}^{m-k} g_i x^i$. The reducta set of G, denoted RED(G), is defined to be

$\{red^k(G)\ 0\leq k\leq \deg(G)\ \&\ red^k(G)\neq 0\}$.

Suppose for some nonzero $F\varepsilon I_r$ that for some $k\geq 1$, $ldcf(red^j(F))$ vanishes everywhere on a cell c in E^{r-1} for $0\leq j<k$, but $ldcf(red^k(F))$ vanishes nowhere on c. Then setting $K = red^k(F)$, F = K on Y(c), hence if c is PSC(K,K')-invariant, V(F) is delineable or empty on c by Theorem 7. Thus given a nonzero $F\varepsilon I_r$, we define a set of polynomials in I_{r-1}, called the projection of F and denoted PROJ(F), as follows. PROJ(F) is the set consisting of ldcf(K) and the elements of PSC(K,K'), for

every K in RED(F).

<u>Theorem 8</u> Let c be a PROJ(F)-invariant cell in E^{r-1}, F a nonzero element of I_r. Then V(F) is delineable, empty, or cylindrical on c.

<u>Proof</u> Suppose there is some k, $0 \le k \le \deg(F)$, such that $\mathrm{ldcf}(\mathrm{red}^k(F))$ does not vanish everywhere on c. Then since $\mathrm{ldcf}(\mathrm{red}^k(F)) \varepsilon \mathrm{PROJ}(F)$, it is invariant on c, hence vanishes nowhere on c. Let $H = \mathrm{red}^k(F)$ for the smallest such k. By definition of PROJ(F), c is PSC(H,H')-invariant. Hence by Theorem 7, V(H) is delineable or empty on c, and so since $F = H$ on $Y(c)$, V(F) is delineable or empty on c. If there is no k as specified above, then F vanishes at every point of Y(c), hence V(F) is cylindrical on c □

Clearly if we took $P(\mathcal{A})$ to be the union, over all A_i in $\mathcal{A}$, of the sets $\mathrm{PROJ}(A_i)$, we would satisfy condition 5.1 of Theorem 5. We now define a still larger set of polynomials with which condition 5.2 will also be satisfied. Let $\mathcal{B}$ be the union, over all A_i in $\mathcal{A}$, of the sets $\mathrm{RED}(A_i)$. We define $\mathrm{PROJ}(\mathcal{A})$, the <u>projection of $\mathcal{A}$,</u> to be the union, over all A_i in $\mathcal{A}$, of the sets $\mathrm{PROJ}(A_i)$, together with the union, over all B_1, B_2 in $\mathcal{B}$, of the sets $\mathrm{PSC}(B_1, B_2)$.

<u>Theorem 9</u> Let c be a $\mathrm{PROJ}(\mathcal{A})$-invariant cell in E^{r-1}. Then

9.1) Each A_i in $\mathcal{A}$ is delineable, cylindrical, or empty on c.

9.2) $\pi^c A_i$ is delineable or empty on c.

Proof

9.1 is immediate since PROJ($\mathcal{A}$) contains PROJ(A_i) for every A_i in $\mathcal{A}$. Collins [1975, p. 10-11] gives the proof that $\pi^c A_i$ is delineable on c if it is not empty on c □

We now enlarge PROJ($\mathcal{A}$) in such a way that condition 5.3 of Theorem 5 is satisfied. Let $\mathcal{B}$ again be the union, over all A_i in $\mathcal{A}$, of the sets RED(A_i). Let $\bar{\mathcal{B}}$ be the union, over all B in $\mathcal{B}$, of the sets DER(B). Define APROJ($\mathcal{A}$), the augmented projection of $\mathcal{A}$, to be the union, over all $\bar{B}$ in $\bar{\mathcal{B}}$, of the sets PROJ($\bar{B}$), together with the union, over all B_1, B_2 in $\mathcal{B}$, of the sets PSC(B_1, B_2).

Theorem 10 Let c be an APROJ($\mathcal{A}$)-invariant cell in E^{r-1}. Then

10.1) Each A_i in $\mathcal{A}$ is delineable, cylindrical, or empty on c.

10.2) $\pi^c A_i$ is delineable or empty on c.

10.3) For each A_i in $\mathcal{A}$, if A_i is delineable on c then every element of DER(A_i) is delineable or empty on c.

Proof

10.1 and 10.2 are immediate since PROJ($\mathcal{A}$) $\subset$ APROJ($\mathcal{A}$).

Suppose A_i is delineable on c. Then there exists a k, $0 \le k \le \deg(A_i)$, such that $\mathrm{ldcf}(\mathrm{red}^k(A_i))$ vanishes nowhere on c; let $H = \mathrm{red}^k(A_i)$ for the smallest such k. Then $A_i = H$ on $Y(c)$, hence $\mathrm{der}^j(A_i) = \mathrm{der}^j(H)$ on $Y(c)$ for all j, $0 \le j \le \deg(H)$. Since ldcf(H) vanishes nowhere on c, ldcf(K) vanishes nowhere on c for every $K \in \mathrm{DER}(H)$. By definition of APROJ($\mathcal{A}$), c is PSC(K,K')-invariant for every $K \in \mathrm{DER}(H)$. Hence by Theorem 7, every element of DER(H) is delineable or empty on c. Hence every element of $\mathrm{DER}(A_i)$ is delineable or empty on c □

This completes the proof of the Main Theorem □

As a simple example of the c.d. algorithm, consider again the curve $F(x,y) = y^2 - x^3 - x^2$ in E^2. We have $\mathcal{A} = \{F\}$, and omitting inessential elements, APROJ($\mathcal{A}$) = $\{x^3 + x^2\} = \{x^2 (x + 1)\}$. So we obtain the c.d. of E^1 we saw above in Figure 2 with $c_1^0 = -1$ and $c_2^0 = 0$. V(F) is empty on c_1^1 and delineable on each of the other four cells. Thus we obtain the following $\mathcal{A}$-invariant semi-algebraic c.d. of E^2:

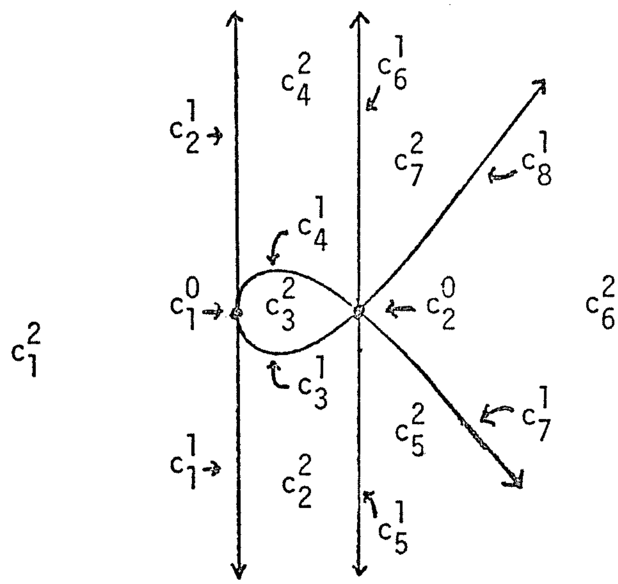

Figure 9

Evaluating F at the sample point for each cell and retaining those cells on which F = 0, we get the c.d. we saw in Figure 1 for V(F).

3. Dimension and Incidence Determination for the cells of a cellular decomposition

Recalling the proof of the c.d. Theorem, we see that the following theorem implies that we can determine the dimension of every cell in the c.d. of a semi-algebraic set.

Theorem 11 Let $\mathcal{A} = \{A_1, \ldots, A_n\}$, $n \geq 1$, be a finite set of nonzero polynomials in I_r, $r \geq 1$. Let C be the $\mathcal{A}$-invariant c.d. of E^r determined by applying the c.d. algorithm described by the proof of the Main Theorem. Then we can determine the dimension of every cell of C.

Proof By induction on r. If r=1, by the definition of $C_{\mathcal{A}}(E^1)$ we know the dimension of each cell in it. Suppose that r>1, and that we know the dimension of every cell in the APROJ($\mathcal{A}$)-invariant c.d. C_{r-1} of E^{r-1} determined by the algorithm. Let c be any cell in C_{r-1}, and let $H = \pi^c A_i$. If H = 1, then Y(c) will be a cell of C, and dimension(Y(c)) = dimension(c) + 1. If $H \neq 1$, then the cells of $C_H(Y(c))$ will be cells of C. By the definition of $C_H(Y(c))$, if the dimension of c is known, the dimension of every cell of $C_H(Y(c))$ is known □

We now show that if we have d.f.'s for the cells of a c.d. C of E^r, and know the dimension of each cell in C, we can determine the incidences among cells. We emphasize that our algorithm is probably far from optimal. Given an i-cell c^i and a j-cell c^j in C with i<j, we show that the statement "c^i is incident on c^j" is decidable. Following a suggestion of Collins, we do so by expressing this statement as a sentence in the elementary theory of real closed fields. Collins' quantifier elimination algorithm could then be applied to decide the truth of the sentence.

Let ϕ^i and ϕ^j be the d.f.'s we assume we have for

c^i and c^j. By definition of incidence, c^i is incident on c^j if and only if every point of c^i is a limit point of c^j. A point $x = \langle x_1, \ldots, x_r \rangle$ of c^i is a limit point of c^j if and only if every open ball of radius $\varepsilon > 0$ with center x contains a point $y = \langle y_1, \ldots, y_r \rangle$ of c^j with $x \neq y$. We let $d(x,y)$ denote the Euclidean distance from x to y. Then the statement "every point of c^i is a limit point of c^j" can be expressed in the elementary theory of real closed fields as follows ("$\supset$" denotes logical implication):

$$(\forall x \forall \varepsilon)\{[\phi^i(x) \ \& \ (\varepsilon>0)] \supset (\exists y)[\phi^j(y) \ \& \ (0<d(x,y)<\varepsilon)]\}.$$

References

Brown, W. S. and Traub, J. F., On Euclid's Algorithm and the Theory of Subresultants, J. Assoc. Comp. Mach., 18, 4 (1971), p. 505-514.

Collins, G. E., Quantifier Elimination for Real Closed Fields by Cylindrical Algebraic Decomposition, 2nd GI Conf. on Automata Theory and Formal Lang., Lect. Notes in Comp. Sci. 33, Springer-Verlag, Berlin, 1975, p. 134-183.

Collins, G. E., Quantifier Elimination for Real Closed Fields by Cylindrical Algebraic Decomposition - A Synopsis, SIGSAM Bulletin of the Assoc. Comput. Mach., 10, 1 (1976), p. 10-12.

Hironaka, H., Triangulations of Algebraic Sets, Proc. Symposia in Pure Math., 29, American Math. Society, Providence, 1975, p. 165-185.

Kahn, P., private communication to G. Collins, 1978.

Loos, R. G. K., A Constructive Approach to Algebraic Numbers, manuscript, 1973.

Mathlab Group, MACSYMA Reference Manual, Version 9, Laboratory for Computer Science, Mass. Inst. of Tech., 1977.

Rubald, C. M., Algorithms for Polynomials over a Real Algebraic Number Field, Tech. Rept. #206, Computer Sciences Dept., Univ. of Wisconsin-Madison, 1974.

Rump, S., Ein Algorithmus zur Isolierung der reellen Nullstellen eines Polynoms mit algebraischen Koeffizienten, Rechenzeitanalyse and Implementierung, Diplomarbeit, Fach. Mathematik, Univ. Kaiserlautern, 1976.

Waerden,, B. L. van der, Topologische Bergründung des Kalküls der abzählenden Geometrie, Math. Ann. 102 (1929), p. 337-362.

Waerden, B. L., van der, Modern Algebra vol. II, tr. from the second revised German edn., Frederick Ungar, New York, 1950.

Yun, D. Y. Y., On Algorithms for Solving Systems of Polynomial Equations, SIGSAM Bulletin of the Assoc. Comput. Mach., #27 (Sept. 1973), p. 19-25.