Regular cylindrical algebraic decomposition

J.H. Davenport, A.F. Locatelli & G.K. Sankaran
Department of Mathematical Sciences,
University of Bath,
Bath BA2 7AY, UK
J.H.Davenport@bath.ac.uk
acyr.locatelli86@gmail.com
G.K.Sankaran@bath.ac.uk

March 5, 2018

Abstract: We show that a strong well-based cylindrical algebraic decomposition $\mathcal{P}$ of a bounded semi-algebraic set is a regular cell decomposition, in any dimension and independently of the method by which $\mathcal{P}$ is constructed. Being well-based is a global condition on $\mathcal{P}$ that holds for the output of many widely used algorithms. We also show the same for S of dimension ≤ 3 and $\mathcal{P}$ a strong cylindrical algebraic decomposition that is locally boundary simply connected: this is a purely local extra condition.

1 Introduction

Cylindrical algebraic decomposition (abbreviated c.a.d.; see Definition 3.1) is a method that decomposes a semi-algebraic subset $S \subseteq \mathbb{R}^n$ into simpler pieces (cells) in a systematic way. It first arose [9] in the context of quantifier elimination, but has since become a useful technique for effective computation of topological invariants, such as homology groups, of semi-algebraic sets. For example, the piano movers' problem (see, inter alia, [20]) asks whether the configuration space of allowable positions of an object in a subset of $\mathbb{R}^3$ is connected. Questions of this nature can have both theoretical and practical importance. [1]

Cell decompositions can be quite pathological, however, and for purposes of computation (again, both theoretical and practical) some further conditions are needed. One would hope, at least, to obtain a representation of S as a CW-complex: better still, a regular cell complex.

[1]Some of the results of this paper formed part of the Bath Ph.D. thesis [16] of the second author, which was funded by the University of Bath. We acknowledge discussions with Matthew England and David Wilson, and EPSRC grant EP/J003247/1 which funded them. GKS thanks Andrew Ranicki and Kenichi Ohshika for education about cobordism.

Question 1.1 *Let $S \subset \mathbb{R}^n$ be a closed and bounded semi-algebraic set.*

(i) Can we find a c.a.d. of S into regular cells?

(ii) Given a c.a.d. of S, can we tell easily whether it is a regular cell decomposition?

A partial answer to Question 1.1(i) was given in [2], where it is shown that the bounded cells of a semi-monotone c.a.d. (see [1]) are regular, and an algorithm is given to construct such a c.a.d. if $\dim S \leq 2$ or $n = 3$. However, being semi-monotone is a strong condition and it is not at present clear whether semi-monotone c.a.d.s exist at all in general. Even if they do, they are likely to be laborious to construct and to have many cells, making them unsuitable for computation.

We can always find a c.a.d. of S with regular cells if we allow a change of coordinates, but this is usually undesirable. From a computational point of view, implementations of c.a.d. algorithms often improve run time by exploiting sparseness, which is destroyed by change of coordinates. Quantifier elimination, the original motivating example for c.a.d. in [9], does not allow arbitrary changes of coordinates, and indeed requires some ordering on the coordinates: we assume, as is usual in c.a.d. theory, a total order $x_1 \prec x_2 \prec \cdots \prec x_n$. Thus it is important to understand which c.a.d.s have good properties such as giving regular cell decompositions.

For these reasons the earlier study of the topological properties of c.a.d.s in [15] remains very relevant to Question 1.1. Lazard describes some much weaker conditions and conjectures (Conjecture 3.13, below) that a c.a.d. satisfying them will have regular cells, and also shows how to construct these c.a.d.s for $n = 3$.

Much earlier, Schwartz and Sharir [20] had proved that a c.a.d. produced by Collins' algorithm [9], the only method known at that time, gives a regular cell complex provided it is well-based (see Definition 3.8).

In Section 3.1 of this paper we prove that any well-based strong c.a.d gives a regular cell complex: see Theorem 3.11 for the precise statement and Definition 3.6 for the meaning of "strong". A well-based c.a.d. produced by Collins' algorithm is always strong, so this is a generalisation of the result of Schwartz and Sharir, but it is entirely independent of the method used to construct the c.a.d. and is thus more widely applicable.

In Section 3.2, we prove a slightly weaker form of Lazard's conjecture for $\dim S = 2$ or $n = 3$: see Theorem 3.27. Our methods also suggest a strategy for $n \geq 4$.

These two results are superficially similar but quite different in detail. In Section 3.1 we consider a c.a.d. $\mathcal{P}$ that is $\mathbf{F}$-invariant (see Definition 3.3) for a large set of polynomials $\mathbf{F} \subset \mathbb{R}[x_1, \ldots, x_n]$, including as a minimum all the polynomials that are used to define S. (Indeed, the term well-based itself already presumes that $\mathcal{P}$ is $\mathbf{F}$-invariant.) In one respect, this restriction is not onerous: algorithms commonly do produce **F**-invariant c.a.d.s by construction. On the one hand, **F**-invariance is a very strong global condition, which may force $\mathcal{P}$ to have many cells even far from S and cannot be checked locally near each cell.

By contrast, in Section 3.2 we are concerned with the topology of c.a.d.s in general, subject only to local conditions. Apart from its theoretical interest, this is potentially important in the context of Brown's NuCAD algorithm [7], which constructs cells that are capable of being cells in c.a.d.s, rather than complete c.a.d.s, and is thus inherently local: global conditions such as **F**-invariance do not arise.

This difference is also reflected in the methods of proof of Theorem 3.11 and Theorem 3.27. For Theorem 3.11, we use largely elementary methods of real algebraic geometry, exploiting the rigidity imposed by the **F**-invariance. The tools used to prove Theorem 3.27 are topological and are anything but elementary as they include the h-cobordism theorem (in effect, the Poincaré Conjecture).

Some of the statements make sense over an arbitrary real closed field, but as our methods are in part topological we work over $\mathbb{R}$ throughout. See [10] for an approach to h-cobordism in the context of real closed fields, which could possibly allow one to remove this restriction.

A subsidiary aim of this paper is to give some consistent terminology for ideas that have appeared in different parts of the literature under various, sometimes incompatible, names. We try to do this in the course of Section 2, which explains the background to the problems. The main results are found in Section 3. Finally, in Section 4, we make some brief observations about another question (Question 4.2) raised by Lazard in the same paper [15].

2 Cells and cell decompositions

Throughout the rest of the paper, we use $\mathbb{B}(p,\varepsilon)$ and $\mathbb{S}(p,\varepsilon)$ to denote the open ball and the sphere, respectively, in $\mathbb{R}^n$ with centre p and radius ε (the dimension will always be clear): we use $\mathbb{B}^n$ and $\mathbb{S}^{n-1}$ for the standard unit ball and sphere in $\mathbb{R}^n$. If $X \subset \mathbb{R}^n$ then $\overline{X}$ denotes the closure of X in the Euclidean topology.

We begin with a well-known example, which motivates Question 1.1.

Example 2.1 *Put* $\Delta = \{(x,y) \in \mathbb{R}^2 \mid 0 < x < 1, -x < y < x\}$ *and consider the semi-algebraic set*

$$W = \{(x,y,z) \in \Delta \times \mathbb{R} \subset \mathbb{R}^3 \mid x^2 z = y^2\},$$

a subset of the Whitney umbrella $\{(x,y,z) \in \mathbb{R}^3 \mid x^2 z - y^2 = 0\}$.

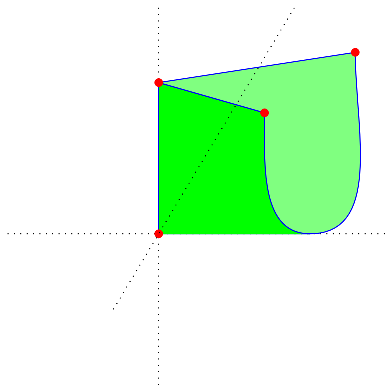

Example 2.1: a CW complex that is not regular

We can partition $\overline{W}$ *into nine disjoint cells: the corners* $(0,0,0)$, $(0,0,1)$, $(1,1,1)$ *and* $(1,-1,1)$*; the edges* $\{(0,0,t)\}$, $\{(t,t,1)\}$ *and* $\{(t,-t,1)\}$ *(for* $0<t<1$*) together with* $\{(1,t,t^2)\}$ *(for* $-1<t<1$*); and* W*.*

This cell decomposition is a c.a.d. of $\overline{W}$ and represents $\overline{W}$ as a CW complex [13, pp. 5 & 519] but not as a regular cell complex (see Definition 2.4).

2.1 Definitions and terminology

Next, we collect some definitions. This is already not quite trivial, because the same or very similar conditions have been introduced by several authors at different times under very different, and sometimes incompatible, names. Before doing any mathematics at all, we propose some terminology which we believe is consistent, flexible and memorable.

Definition 2.2 *A subset* C *of* $\mathbb{R}^n$ *is a* d-cell*, for* $d \in \mathbb{N}_0$*, if there exists a homeomorphism* $\mathbb{B}^d \to C$*, for some* $d \in \mathbb{N}_0$ *called the* dimension $\dim C$ *of* C*. The* boundary *of* $C \subset \mathbb{R}^n$ *(sometimes, for emphasis, the* cell boundary*) is* $\partial_c C = \overline{C} \setminus C$*.*

The cell boundary of a cell C does not, in general, coincide with the topological boundary of C, which is $\overline{C} \setminus \mathrm{Int}(C)$. Also, $\overline{C}$ may have a structure of manifold with boundary in which the manifold boundary $\partial_m \overline{C}$ (see Definition 3.15) might coincide with neither $\partial_c C$ nor the topological boundary.

Definition 2.3 *Let* X *be a subset of* $\mathbb{R}^n$*. A* cell decomposition *of* X *is a partition* $\mathcal{P} = \{C_\alpha\}$ *of* X *into disjoint cells.*

Even in $\mathbb{R}^2$ a cell may have bad boundary: for instance, if we take $C = \{(x,y) \in \mathbb{R}^2 \mid y \cos x < 1\}$ then $\partial_c C$ has infinitely many connected components. We define below some desirable conditions on a cell C and its boundary. Some of these conditions are intrinsic to C; others are related to a cell decomposition.

Recall that if $X \subseteq X'$ and $Y \subseteq Y'$ are inclusions of topological spaces, a homeomorphism $\varphi\colon (X', X) \to (Y', Y)$ is a homeomorphism $\varphi\colon X' \to Y'$ such that $\varphi|_X$ is a homeomorphism $X \to Y$.

Definition 2.4 *We say that two subsets $X \subset \mathbb{R}^n$ and $Y \subset \mathbb{R}^m$ are* equiregular *if there exists a homeomorphism $\varphi\colon (\overline{X}, X) \to (\overline{Y}, Y)$. A d-cell $C \subset \mathbb{R}^n$ is said to be* a regular cell *if C is equiregular with $\mathbb{B}^d$.*

The hypercube $(0,1)^d \subset \mathbb{R}^n$ is a regular cell. On the other hand, the cell $\mathbb{B}^2 \setminus \{(0,y) \in \mathbb{R}^2 \mid y \geq 0\}$ is not regular, even though its closure is $\overline{\mathbb{B}}^2$. Moreover, even if C is regular, a particular homeomorphism $\mathbb{B} \to C$ need not extend to $\overline{\mathbb{B}} \to \overline{C}$ even as a continuous map.

We establish a convention for naming cell decompositions where all cells have a certain property.

Convention 2.5 *If Π is a property of cells we shall say that $\mathcal{P}$ is a Π cell decomposition if $\mathcal{P}$ is a cell decomposition and all cells of $\mathcal{P}$ satisfy Π.*

Thus a regular cell decomposition is a decomposition into regular cells. Not every property of $\mathcal{P}$ can be checked on the cells, however: for example, a finite cell decomposition is simply a decomposition into finitely many cells. There is no ambiguity, because finiteness is not a property of cells.

Definition 2.6 *We say that two cells C and D in $\mathbb{R}^n$ are* adjacent *if either one intersects the closure of the other. We say that C is* subadjacent *to D, written $C \preceq D$, if $C \cap \overline{D} \neq \emptyset$.*

If $D \preceq C$ and $D \cap C = \emptyset$, then by [3, Theorem 5.42], $\dim D \leq \dim C - 1$.

Now we define some extrinsic properties of a cell, in relation to a cell decomposition.

Definition 2.7 *Let C be a cell of a cell decomposition $\mathcal{P}$. We say that C is* closure finite in $\mathcal{P}$ *if $\overline{C}$ (or, equivalently, $\partial_c C$) is the union of finitely many cells of $\mathcal{P}$.*

This condition is found in the literature under different names. It is called *boundary coherent* in [15, Definition 2.7]: elsewhere, sometimes without the finiteness requirement, it is called the *frontier condition*. We use the term *closure finite* as it is more descriptive than either of the terms above, and is the usual term in the topology literature. Indeed, it is the meaning of the "C" of "CW complex": see [13, p. 520].

2.2 Examples and basic properties

We first illustrate some relations among the properties introduced in Section 2.1.

Example 2.8 *Consider the following cell decomposition of* $[0,1]^3$*, in which the end points* $(0,0,1/2)$ *and* $(0,1,1/2)$ *are not cells, and the edges* $(0,0,z)$ *and* $(0,1,z)$ $(z \in (0,1))$ *are not sub-divided by these points.*

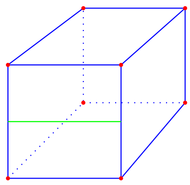

Example 2.8: subadjacency and closure finiteness I

The cube is closure finite; the 1*-cell* $[0,1] \times \{0\} \times \{\frac{1}{2}\}$*, subadjacent to the cube, is not closure finite.*

Example 2.8 also shows that even if C is closure finite in $\mathcal{P}$ its boundary may contain a cell of $\mathcal{P}$ that is not closure finite.

Example 2.9 *The cells* $C_1 = \mathbb{S}^1 \setminus \{(x,y) \in \mathbb{R}^2 \mid x \geq 0 \text{ and } y \leq 0\}$ *and* $C_1 = \{(x,0) \in \mathbb{R}^2 \mid -1 < x < 2\}$ *in* $\mathbb{R}^2$ *are not closure finite in* $\mathcal{P} = \{C_1, C_2\}$*, and they are subadjacent to each other.*

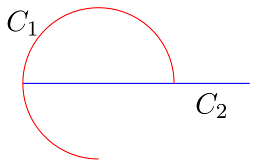

Example 2.9: subadjacency and closure finiteness II

Lemma 2.10 *Let* C *be a cell of a cell decomposition* $\mathcal{P}$*. Then* C *is closure finite if and only if* $D \preceq C$ *implies that* $D \subseteq \overline{C}$*. In particular, if two cells* C *and* D *of* $\mathcal{P}$ *are closure finite and subadjacent to each other, then* $C = D$*.*

Proof: Suppose that C is closure finite in $\mathcal{P}$; that is, $\partial_c C = \bigcup_{i=1}^k D_i$. If $D \preceq C$ then it must intersect some D_i, and thus $D = D_i$. In particular $D \subset \partial_c C$. Conversely, if $\overline{C}$ contains all cells subadjacent to C, then $\overline{C} = \bigcup_{D \preceq C} D$. $\square$

Definition 2.11 *Let* C *be a cell of a cell decomposition* $\mathcal{P}$*. We say that* C *is* well-bordered *in* $\mathcal{P}$ *if there is a finite collection* $\{C_i\} \subset \mathcal{P}$ *of cells of* $\mathcal{P}$ *such that* $\dim C_i = \dim C - 1$ *and* $\partial_c C = \bigcup_i \overline{C_i}$*.*

For instance a 2-sphere minus a point is a cell that is not well-bordered. See [15, Example 2.9] for more examples of how a cell can fail to be well-bordered, and for a closure finite decomposition that is not well-bordered.

Like closure finiteness, the well-bordered property does not permeate to subadjacent cells.

Example 2.12 *Consider the cell decomposition that consists of the open cube* $(0,1)^3$*, all six of its faces, eleven of its edges (not the* z*-axis) and seven of its corners (not the origin), together with the* 1*-cell* $\{0\} \times \{0\} \times (-1,1)$ *and the* 0*-cell* $(0,0,-1)$*. We observe the following:*

1. *The cell* $(0,1)^3$ *is well-bordered but not closure finite.*

2. *the faces that are adjacent to the cell* $\{0\} \times \{0\} \times (-1,1)$ *are neither well-bordered nor closure finite.*

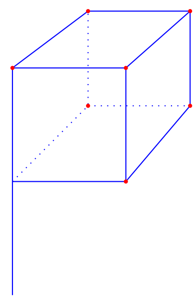

Example 2.12: well-borderedness and closure finiteness

These two conditions are nevertheless related.

Lemma 2.13 *Let* C *be a cell of a cell decomposition* $\mathcal{P}$*. If* C *and all cells subadjacent to* C *are well-bordered, then* C *is closure finite.*

Proof: In view of Lemma 2.10 it suffices to show that $D \subseteq \partial_c C$ if $D \preceq C$ and $D \neq C$. We proceed by induction on $\dim C$: for $\dim C = 0$ there is nothing to prove.

As C is well-bordered, $\partial_c C = \bigcup_i \overline{C_i}$ for some finite collection of cells C_i with $\dim C_i = \dim C - 1$. If $D \cap C_i \neq \emptyset$, for some i, then $D = C_i$; otherwise, $D \cap \overline{C_i} \neq \emptyset$ for some i. Then by induction $D \subset \partial_c C_i \subset \partial_c C$ and the result follows. $\square$

Corollary 2.14 *Any well-bordered cell decomposition is closure finite.*

As we have seen, [15, Example 2.9] shows that the converse is not true. However, for regular cell decompositions the two coincide.

Lemma 2.15 *A regular cell decomposition of a compact set* $S \subset \mathbb{R}^n$ *is closure finite if and only if it is well-bordered.*

Proof: If C is a closure finite regular d-cell then $\partial_c C$ is homeomorphic to $\mathbb{S}^{d-1}$, and decomposes into finitely many cells. Then $\partial_c C$ is the closure of the union of the $(d-1)$-cells in that decomposition.

The other direction is just Corollary 2.14. $\square$

Definition 2.16 *For any topological property* Π, *we say that a set* $C \subset \mathbb{R}^n$ *is* locally boundary Π *if every* $p \in \partial_c C$ *has a base of neighbourhoods* N *in* $\overline{C}$ *such that* Π *holds for each* $N \cap C$.

In many cases one may take the neighbourhoods N to be the intersections $\mathbb{B}(p,\varepsilon) \cap C$ for $0 < \varepsilon \ll 1$: we shall do this without further comment when it is convenient, but one should check that it is permissible to do so. An example of a property Π for which this would not be permissible is disconnectedness: $\mathbb{R}_+ \subset \mathbb{R}$ is locally boundary disconnected (as well as being locally boundary connected!) because we may take for N the sets $[0, \frac{1}{r}) \cup (\frac{2}{r}, \frac{3}{r})$, but the intervals $(0,\varepsilon)$ are all connected. We shall not in fact consider any properties for which the balls are not suitable neighbourhoods.

Lazard [15, Definition 2.7] defines "boundary smooth", which according to Definition 2.16 is the same as "locally boundary connected". We prefer this terminology because it extends to other properties (we shall need "locally boundary simply connected" later, for instance) and because the term "smooth" is already overloaded. In particular, "boundary smooth" has nothing to do with either being C^∞ or the absence of singularities.

2.3 Semi-algebraic cell decompositions

Now we limit ourselves to semi-algebraic cells. We shall make constant use of the conic structure of semi-algebraic sets [5, Theorem 9.3.6]. We also need a slightly stronger relative version (take $X = \emptyset$ to recover the usual version).

Proposition 2.17 *Suppose that* $X \subseteq Y$ *are semialgebraic subsets of* $\mathbb{R}^n$ *and* $p \in Y$. *Then for* $0 < \varepsilon \ll 1$ *there exists a semi-algebraic homeomorphism* $\psi \colon \overline{\mathbb{B}}(p,\varepsilon) \to \overline{\mathbb{B}}(p,\varepsilon)$ *which is the identity on* $\mathbb{S}(p,\varepsilon)$, *such that* $\|\psi(q) - p\| = \|q - p\|$ *for all* $q \in \overline{\mathbb{B}}(p,\varepsilon)$ *and* $\psi\big((Y,X) \cap \overline{\mathbb{B}}(p,\varepsilon)\big)$ *is the cone on* $(Y,X) \cap \overline{\mathbb{S}}(p,\varepsilon)$ *with vertex* p.

Proof: Consider

$$M := \{(y,t) \in \mathbb{R}^n \times \mathbb{R} \mid (y \in Y \text{ and } t = 0) \text{ or } (y \in X \text{ and } 0 \le t \le 1)\}.$$

This is a semi-algebraic set (it is the mapping cylinder of the inclusion $X \hookrightarrow Y$) so we may consider its conic structure near a point $p \in M$ where $t = 0$. Then it is sufficient to take ψ to be the restriction to $t = 0$ of the map $\phi \colon \overline{\mathbb{B}}(p,\varepsilon) \to \overline{\mathbb{B}}(p,\varepsilon)$ in $\mathbb{R}^{n+1}$ guaranteed by [5, Theorem 9.3.6]. $\square$

The following consequence of the local conic structure is also useful.

Proposition 2.18 *Let $C \subset \mathbb{R}^n$ be a semi-algebraic set and $p \in \partial_c C$. Then for $0 < \varepsilon \ll 1$ the intersection $C \cap \mathbb{B}(p, \varepsilon)$ has $C \cap \mathbb{S}(p, \varepsilon)$ as a deformation retract.*

Proof: Applying [5, Theorem 9.3.6] to $C \cup \{p\}$ yields a homeomorphism between $\overline{\mathbb{B}}(p, \varepsilon) \cap C \cup \{p\}$ and the cone on $\mathbb{S}(p, \varepsilon) \cap C$. Away from $\{p\}$, this restricts to a homeomorphism between $(\overline{\mathbb{B}}(p, \varepsilon) \cap C) \setminus \{p\}$ and $(\mathbb{S}(p, \varepsilon) \cap C) \times (0, 1]$, and the latter retracts onto $(\mathbb{S}(p, \varepsilon) \cap C) \times \{\frac{1}{2}\}$. $\square$

Because of Proposition 2.18 we can often replace the ball with a sphere when checking Definition 2.16.

Corollary 2.19 *If Π is a homotopy property for which Definition 2.16 can be checked on balls, and $C \subset \mathbb{R}^n$ is a semi-algebraic cell, then C is locally boundary Π if and only if, for all $p \in \partial_c C$, there exists $\delta > 0$ such that $C \cap \mathbb{S}(p, \varepsilon)$ has property Π for all $0 < \varepsilon < \delta$.*

Clearly the same is also true with $\overline{\mathbb{B}}$ instead of $\mathbb{S}$.

With the definitions we have made, being locally boundary Π is automatically an equiregularity invariant property. In particular, as was pointed out in [20], a regular cell, even if not semi-algebraic, is always locally boundary connected.

3 Cylindrical algebraic decomposition

We think of a cylindrical algebraic decomposition as a finite partition of $\mathbb{R}^n$ into semi-algebraic cells, built inductively, and whose projections onto the first k variables, for $k < n$, are either disjoint or the same. These cells are not just arbitrary semi-algebraic sets homeomorphic to $(0, 1)^d$, for some $d \in \mathbb{N}$, but cells that arise from graphs $\Gamma(g)$ of some semi-algebraic functions g, as below.

Definition 3.1 *A* cylindrical algebraic decomposition *or* c.a.d. *of $\mathbb{R}^n$ is a finite semialgebraic cell decomposition $\mathcal{P} = \mathcal{P}_n$ of $\mathbb{R}^n$ defined inductively by the following conditions.*

1. *If $n = 1$ then $\mathcal{P} = \mathcal{P}_1$ is any finite cell decomposition of $\mathbb{R}$.*
2. *The projection $\mathrm{pr}_n : \mathbb{R}^n \to \mathbb{R}^{n-1}$ on the last $n-1$ coordinates is* cylindrical*: that is, if C, $C' \in \mathcal{P}_n$ then either $\mathrm{pr}_n(C) \cap \mathrm{pr}_n(C') = \emptyset$ or $\mathrm{pr}_n(C) = \mathrm{pr}_n(C')$.*
3. *$\mathcal{P}_{n-1} = \{\mathrm{pr}_n(C) \mid C \in \mathcal{P}_n\}$ is a c.a.d. of $\mathbb{R}^{n-1}$.*
4. *For each $D \in \mathcal{P}_{n-1}$ there are finitely many continuous semi-algebraic functions $g_1, \ldots, g_k : D \to \mathbb{R}$, satisfying $g_j(p) < g_{j+1}(p)$ for all j and*

all $p \in D$, such that $\mathcal{P} \ni \Gamma(g_j) = \{(p, y) \mid g_j(p) = y\}$ for each j and $\mathcal{P} \ni \Delta_j = \{(p, y) \mid g_j(p) < y < g_{j+1}(p, y)\}$: furthermore both $\Delta_{-\infty} = \{(p, y) \mid y < f_1(p)\}$ and $\Delta_\infty = \{(p, y) \mid f_k(p) < y\}$ also belong to $\mathcal{P}$.

Definition 3.2 *In Definition 3.1, the graphs $\Gamma(g_j)$ are called* sections *and the cells Δ_j are called* sectors *of $\mathcal{P}$.*

A c.a.d. is usually chosen to respect some data, such as some functions on $\mathbb{R}^n$ or subsets of $\mathbb{R}^n$.

Definition 3.3 *Let $\mathbf{F} \subset \mathbb{R}[x_1, \ldots, x_n]$ be a finite set of nonzero polynomials. A c.a.d. $\mathcal{P}$ is said to be* $\mathbf{F}$-invariant *if, for every $f \in \mathbf{F}$, the sign of f is constant on each $C \in \mathcal{P}$.*

This is sometimes called *sign-invariance*: one could instead require other properties of f, such as its order of vanishing, to be constant on each C, but we shall not need any other kind of invariance here.

Definition 3.4 *Let $S \subset \mathbb{R}^n$ be a semi-algebraic set. A c.a.d. $\mathcal{P}$ of $\mathbb{R}^n$ is* adapted to S *if S is a union of cells of $\mathcal{P}$.*

It is sometimes useful to give some more information about $\mathcal{P}$.

Definition 3.5 *A* sampled c.a.d. *is a c.a.d. $\mathcal{P}$ together with a choice of base point $b_C \in C$ for each cell $C \in \mathcal{P}$.*

In general, a c.a.d. is not a CW complex: for instance, a c.a.d. can fail to be closure finite.

Definition 3.6 *We say that a c.a.d. $\mathcal{P}$ of $\mathbb{R}^n$ is a* strong c.a.d. *if $\mathcal{P}$ is well-bordered and locally boundary connected.*

This is equivalent to the definition in [15] in view of Corollary 2.14.

3.1 F-invariant c.a.d.s

The aim of this section is to investigate when an $\mathbf{F}$-invariant c.a.d. adapted to a closed bounded semi-algebraic set S exhibits S as a regular cell complex.

Definition 3.7 *If $\mathcal{P}$ is an $\mathbf{F}$-invariant c.a.d. and $C \in \mathcal{P}$ is a section, we put $\mathbf{F}_C = \{f \in \mathbf{F} \mid f|_C \equiv 0\}$. We say that $(\mathcal{P}, \mathbf{F})$ is* reduced *if $\mathbf{F}_C \neq \emptyset$, for every section $C \in \mathcal{P}$.*

We do not require that C should actually be cut out by $\mathbf{F}_C$, but we will usually impose the next condition, which may be seen as a weaker version.

Definition 3.8 *If $\mathcal{P}$ is an $\mathbf{F}$-invariant c.a.d., we say that a section $C \in \mathcal{P}$, with $\mathrm{pr}_n(C) = D \in \mathcal{P}_{n-1}$, is a* bad cell *for $(\mathcal{P}, \mathbf{F})$ if there is an $f \in \mathbf{F}_C$ such that $f|_{\mathrm{pr}_n^{-1}(D)} \equiv 0$. If there are no bad cells, we say that $(\mathcal{P}, \mathbf{F})$ is* well-based.

Note that the definitions of bad cell and well-based depend on $\mathbf{F}$ as well as $\mathcal{P}$. Note also that we apply the term "bad cell" to C, not D. This makes no difference, since if C is a bad cell then so is any C' with $\mathrm{pr}_n(C') = \mathrm{pr}_n(C)$ (consider the same f), but it does reflect our point of view of starting with a given c.a.d. rather than constructing one inductively.

We aim to show that reduced well-based strong c.a.d.s give regular cell complexes (see Theorem 3.11 for the precise statement). This was shown in [20, Theorem 2] for a c.a.d. $\mathcal{P}$ constructed via Collins' algorithm [9]: by [15, Theorem 4.4], such a $\mathcal{P}$ is strong if it is well-based. Other algorithms are now in use, though, such as c.a.d. via regular chains [8] or via comprehensive Gröbner systems [12], so we want to be able to dispense with the condition on the construction.

We need two lemmas: the first is a variant of [5, Lemma 2.5.6].

Lemma 3.9 *Suppose that $\mathbf{F} \subset \mathbb{R}[x_1, \ldots, x_{n-1}, y]$ is a finite set of non-zero polynomials and denote by*

$$\mathbf{F}' = \left\{ \frac{\partial^r f}{\partial y^r} \mid r \in \mathbb{Z}_{\geq 0},\ f \in \mathbf{F},\ \frac{\partial^r f}{\partial y^r} \not\equiv 0 \right\}$$

its closure under the operator $\frac{\partial}{\partial y}$. Let $C = \Gamma(g)$ be a bounded section in an $\mathbf{F}'$-invariant c.a.d. $\mathcal{P}$ of $\mathbb{R}^n$, for a semi-algebraic continuous bounded function $g\colon D = \mathrm{pr}_n(C) \to \mathbb{R}$. If $\mathbf{F}_C \neq \emptyset$ and $(\mathcal{P}, \mathbf{F})$ is well-based, then g can be extended continuously to $\overline{D}$.

Proof: The proof is similar to step (ii) in the proof of [5, Lemma 2.5.6]. It is enough to show that g extends continuously to $D \cup \{p\}$, for an arbitrary $p \in \overline{D}$.

By the Curve Selection Lemma [5, Theorem 2.5.5], we choose a continuous semi-algebraic path $\eta\colon [0, 1] \to \overline{D} \cap \mathbb{B}(p, 1)$, such that $\eta(0) = p$ and $\eta(t) \in D$ for $t > 0$. Then we define $\widetilde{\eta}(t) = g(\eta(t)) \in \mathbb{R}$, for $t > 0$: since g is bounded by hypothesis, $\widetilde{\eta}\colon (0, 1] \to \mathbb{R}$ is a bounded continuous semialgebraic function and hence extends continuously to $\widetilde{\eta}\colon [0, 1] \to \mathbb{R}$ by [5, Proposition 2.5.3].

Now we extend g to $\overline{g}\colon D \cup \{p\} \to \mathbb{R}$ by putting $\overline{g}(p) = \widetilde{\eta}(0)$, and $\overline{g} = g$ on D. The claim is that $\overline{g}$ is continuous at p. If not, then

$$\exists\, \varepsilon > 0\ \forall \delta > 0\ \exists\, q \in D \text{ such that } \|q - p\| < \delta \text{ and } |g(q) - \overline{g}(p)| \geq \varepsilon,$$

and hence if we define $E = \{q \in D \mid |g(q) - \overline{g}(p)| \geq \varepsilon\}$ then $p \in \overline{E}$. Again applying the Curve Selection Lemma we obtain a path $\theta\colon [0, 1] \to \overline{E}$ with $\theta(0) = p$ and $\theta(t) \in E$ for $t > 0$, so exactly as before we put $\widetilde{\theta}(t) = g(\theta(t))$

and this extends continuously to $\widetilde{\theta}\colon [0,1] \to \mathbb{R}$. By continuity, we have $|\widetilde{\eta}(0) - \widetilde{\theta}(0)| \geq \varepsilon$, and also $(p, \widetilde{\eta}(0)) \in \overline{C}$ and $(p, \widetilde{\theta}(0)) \in \overline{C}$.

Now suppose that $f \in \mathbf{F}_C$, so $f \in \mathbf{F}$ and $f|_C \equiv 0$. Consider the polynomial $f_p(y) = f(p, y) \in \mathbb{R}[y]$, and observe that $f_p(\widetilde{\eta}(0)) = f_p(\widetilde{\theta}(0)) = 0$. If f_p is not the zero polynomial, we may consider the set $\{f_p, \frac{df_p}{dy}, \frac{d^2 f_p}{dy^2}, \ldots\}$ of all derivatives of f_p. By $\mathbf{F}'$-invariance, for any given r the sign of $\frac{\partial^r f}{\partial y}$ is the same near $(p, \widetilde{\eta}(0))$ (say at $(\eta(t), \widetilde{\eta}(t))$ for small $t > 0$) as near $(p, \widetilde{\theta}(0)$. Hence $\frac{d^r f_p}{dy^r}$ cannot have opposite signs at $y = \widetilde{\eta}(0)$ and $y = \widetilde{\theta}(0)$, although one might be zero and the other not.

But this contradicts Thom's Lemma [5, Proposition 2.5.4]: at two distinct zeros of a real polynomial in one variable, some derivative must have opposite signs.

Hence, if $\overline{g}$ is discontinuous at p, we must have $f_p \equiv 0$: that is, f is identically zero above p. But then, by cylindricity and $\mathbf{F}'$-invariance, f must be identically zero above the cell in $\mathcal{P}_{n-1}$ containing p, contrary to the assumptions. $\square$

Next we need a lemma that allows us to pass extensions up from subdivisions.

Lemma 3.10 *Let $C = \Gamma(g)$ be a bounded, local boundary connected section in an $\mathbf{F}$-invariant c.a.d. $\mathcal{P}$ of $\mathbb{R}^n$, for a semi-algebraic continuous function $g\colon D = \mathrm{pr}_n(C) \to \mathbb{R}$. Suppose that $\mathcal{P}'$ is a well-based strong $\mathbf{F}$-invariant c.a.d. refining $\mathcal{P}$ (i.e. each cell in $\mathcal{P}'$ is a subset of a cell of $\mathcal{P}$), in which C is partitioned into sections $C_i = \Gamma(g_i)$ for semi-algebraic continuous functions $g_i\colon D_i = \mathrm{pr}_n(C_i) \to \mathbb{R}$. If all the g_i extend continuously to $\overline{D}_i$, then g extends continuously to $\overline{D}$.*

Proof: It is enough to show that if $p \in \partial D \cap \partial D_i \cap \partial D_j$ for some $i \neq j$, then $\overline{g}_i(p) = \overline{g}_j(p)$. Then $\overline{g}\colon \overline{D} \to \mathbb{R}$ is consistently defined by $\overline{g}(p) = \overline{g}_i(p)$ if $p \in \overline{D}_i$.

We first show this with the assumption that $D_j \preceq D_i$. Then, since $\mathcal{P}'$ and therefore $\mathcal{P}'_{n-1}$ are strong and in particular closure finite, we have $D_j \subseteq \overline{D_i}$ by Lemma 2.10. Therefore g_j agrees with $\overline{g}_i$ on D_j (they both agree with g) and hence $\overline{g}_j$ agrees with $\overline{g}_i$ also on $\overline{D_j}$, which is contained in $\overline{D_i}$.

For general i and j, we construct a subadjacency chain from D_i to D_j, by considering a semi-algebraic path between $D_i \cap B(p, \epsilon)$ and $D_j \cap B(p, \epsilon)$; this is possible as $D \cap B(p, \epsilon)$ is semi-algebraic and connected, and thus semi-algebraically path-connected by [5, Prop. 2.5.13]. The path gives us a way of selecting the correct consecutive cells.

Let $\eta\colon [0,1] \to C \cap B(p, \epsilon)$ be a semi-algebraic path with $\eta(0) \in D_i \cap B(p, \epsilon)$ and $\eta(1) \in D_j \cap B(p, \epsilon)$. As $\eta([0,1])$ semi-algebraic, $\eta([0,1])) \cap D_k$ has finitely many connected components, for any D_k; thus $\eta^{-1}(D_k)$ is a finite collection of intervals contained in $[0,1]$.

Considering the preimage of every cell in $\mathcal{P}'_{n-1}$, we get a finite partition $[0,1] = \bigcup_{\ell=1}^{N} I_\ell$ where $\sup I_\ell = \inf I_{\ell+1} = t_\ell$. Denote by $D_{k(\ell)}$ the unique cell that contains $\eta(I_\ell)$. Then $\eta(t_\ell) \in \overline{D}_{k(\ell)} \cap \overline{D}_{k(\ell+1)}$, so $D_{k(\ell)}$ and $D_{k(\ell+1)}$ are adjacent: moreover $\eta(t_\ell)$ belongs to either $D_{k(\ell)}$ or $D_{k(\ell+1)}$, so one is subadjacent to the other.

Now we have a finite chain of not necessarily distinct cells

$$D_i = D_{k(1)} \diamondsuit D_{k(2)} \diamondsuit \ldots \diamondsuit D_{k(N)} = D_j,$$

where each $\diamondsuit$ stands for either $\preceq$ or $\succeq$. Hence $\overline{g}_i(p) = \overline{g}_{k(2)}(p) = \cdots = \overline{g}_j(p)$, so $\overline{g}(p)$ is well-defined. $\square$

Theorem 3.11 *Suppose $\mathcal{P}$ is an* **F***-invariant, reduced, well-based strong c.a.d. of $\mathbb{R}^n$ adapted to a closed and bounded subset S. Then the corresponding decomposition of S is a regular cell complex.*

Proof: A strong c.a.d. $\mathcal{P}$ of $\mathbb{R}^n$ is closure finite, so we just need to show that $C \subset S$ is a regular cell. Moreover, if we can show that the sections of $\mathcal{P}$ are regular, then by [20, Lemma 5], so are the sectors.

We will show that a section $C = \Gamma(g)$ is a regular cell by proving that $(\overline{C}, C)$ is homeomorphic to $(\overline{D}, D)$, where $D = \mathrm{pr}_n(C) \in \mathcal{P}_{n-1}$ is the cell below C; then the result follows by induction. It suffices to show that $g\colon D \to \mathbb{R}$ extends continuously to $\overline{D}$, since then $\overline{g}\colon (\overline{D}, D) \to (\overline{C}, C)$ and $\mathrm{pr}_n\colon (\overline{C}, C) \to (\overline{D}, D)$ are mutually inverse homeomorphisms.

Let $\mathbf{F}'$ be the closure of $\mathbf{F}$ under $\frac{\partial}{\partial x_n}$; that is, the smallest set that contains $\mathbf{F}$ and is closed under partial differentiation with respect to x_n. We may choose a $\mathbf{F}'$-invariant c.a.d. $\mathcal{P}'$ refining $\mathcal{P}$. Then C is partitioned into sections $C_i \in \mathcal{P}'$ and each $C_i = \Gamma(g_i)$ for the continuous semi-algebraic function $g_i = g|_{D_i}$ (with, as usual, $D_i = \mathrm{pr}_n(C_i)$).

Now we apply Lemma 3.9, remembering that since $C_i \subseteq C$ there is an $f \in \mathbf{F}$ that vanishes on C_i. We conclude that each g_i can be extended continuously to $\overline{D}_i$.

Finally, we use Lemma 3.10 to extend g continuously to $\overline{D}$. $\square$

The above, in combination with Lemma 2.15, prompts us to raise the following question.

Question 3.12 *Suppose that $\mathcal{P}$ is a c.a.d. of $\mathbb{R}^n$: is it true that $\mathcal{P}$ is closure finite if and only if it is well-bordered?*

Some of the results in this section can be strengthened slightly, by weakening global conditions so that they only apply where they are needed. For example, in Theorem 3.11 it would be enough for $\mathcal{P}$ to give an $\mathcal{F}$-invariant c.a.d. of some open cylinder containing S rather than of the whole of $\mathbb{R}^n$. Similarly, in Lemma 3.9 it is enough for $(\mathcal{P}, \mathbf{F}_C)$ to be well-based (even just near C), and the condition of $\mathbf{F}'$-invariance can also be relaxed analogously.

A related question is whether (or how far) the results of this section apply to c.a.d.s invariant for a formula, as in [17] and [6].

3.2 Topology of strong c.a.d.s

The following conjecture is made in [15, p. 94].

Conjecture 3.13 *Suppose $\mathcal{P}$ is a strong c.a.d. of $\mathbb{R}^n$ adapted to a closed and bounded semi-algebraic set S. Then the corresponding decomposition of S is a regular cell complex.*

We prove some cases of this using techniques from topology. By the nature of the proofs, they are valid only for $\mathbb{R}$, not for arbitrary real closed fields. The idea is that a regular cell is automatically a manifold with boundary: conversely, it is sometimes possible to give conditions on a manifold with boundary that are sufficient to ensure that it is a regular cell, and in low dimension we are able to verify that these conditions always hold.

Definition 3.14 *Let X be a compact d-manifold. We say that X is a* homology d-sphere *if X has the same homology groups as $\mathbb{S}^d$. We say that X is a* homotopy d-sphere *if X has the same homotopy type as $\mathbb{S}^d$.*

It follows from the Hurewicz theorem [13, Thm 4.32] and Whitehead's theorem [13, Thm 4.5] that any simply connected homology sphere Y is a homotopy sphere. One must heed the warning given in [13] after the proof of Whitehead's theorem: we need a weak homotopy equivalence, that is, a map that induces isomorphisms between the homotopy groups. However, if $B \subset Y$ is an open ball then a map that identifies $B \cong \mathbb{B}^d$ with the complement of a point $q \in \mathbb{S}^d$ and sends all of $Y \setminus B$ to q is a weak homotopy equivalence between Y and $\mathbb{S}^d$.

Definition 3.15 *A topological space X is a n-*dimensional manifold with boundary *if, for each $x \in X$, there exists a neighbourhood V of x that is homeomorphic to an open set in either $\mathbb{R}^n$ or $(\mathbb{R}_{\geq 0})^n$. The* manifold boundary *of X, denoted $\partial_m X$, is the set of points of X with no neighbourhood homeomorphic to an open set in $\mathbb{R}^n$.*

If $X = \overline{C}$ for some $C \subset \mathbb{R}^m$ we say that $\overline{C}$ is compatibly a manifold with boundary *if $\partial_m \overline{C} = \partial_c C$.*

For more on manifolds with boundary see [13, p 252]. Our main use of them is based on the following easy result.

Lemma 3.16 *Let X be a d-dimensional compact, contractible manifold with boundary. Then the boundary $\partial_m X$ of X is a homology $(d-1)$-sphere.*

Proof: The homology sequence of $(X, \partial_m X)$ is

$$\cdots \to H_i(\partial_m X) \to H_i(X) \to H_i(X, \partial_m X) \to H_{i-1}(\partial_m X) \to \cdots$$

and Lefschetz duality [13, Theorem 3.43] gives $H_k(X, \partial_m X) \cong H^{d-k}(X)$. Since X is contractible we have $H_0(X) \cong H^0(X) \cong \mathbb{Z}$ and $H_i(X) = H^i(X) = 0$ for $i \neq 0$, and hence $H_0(\partial_m X) \cong H_n(\partial_m X) \cong \mathbb{Z}$ and $H_i(\partial_m X) = 0$ for $i \neq 0$, n. □

In order to apply this to c.a.d.s we need a result (Theorem 3.18) on the contractibility of $\overline{C}$, for a c.a.d. cell C. First we recall the following theorem of Smale [22].

Theorem 3.17 *Suppose that $f\colon X \to Y$ is a proper surjective continuous map between connected, locally compact separable metric spaces, and X is locally contractible. If all the fibres of f are contractible and locally contractible, then f is a weak homotopy equivalence.*

We use Theorem 3.17 to deduce the following, which may be of independent interest.

Theorem 3.18 *Suppose $\mathcal{P}$ is a c.a.d. of $\mathbb{R}^n$ and the induced c.a.d. $\mathcal{P}_{n-1}$ of $\mathbb{R}^{n-1}$ is strong. Then the closure $\overline{C}$ of any bounded cell C of $\mathcal{P}$ is contractible.*

Proof: As often in Section 3.1, let $D = \mathrm{pr}_n(C) \in \mathcal{P}_{n-1}$. We shall show that $\overline{C}$ and $\overline{D}$ have the same homotopy type: then $\overline{C}$ is contractible by induction on n. The result is true for $n = 1$, so we assume $n > 1$.

As C and D are bounded, $\overline{C}$ and $\overline{D}$ are compact semi-algebraic sets and thus, by [5, Thm 9.4.1], $\overline{C}$ and $\overline{D}$ each admits a CW-complex structure. By Whitehead's theorem, it suffices to show that $\overline{C}$ and $\overline{D}$ are weakly homotopy equivalent, and for that it is enough to verify that $\mathrm{pr}_n|_{\overline{C}}$ satisfies the conditions of Theorem 3.17.

The spaces $\overline{C}$ and $\overline{D}$ are connected as they are the closures of connected spaces. Moreover, as they are also Hausdorff compact metric spaces, they are locally compact and separable. It follows immediately from the local conic structure theorem [5, Thm 9.3.6] that any semi-algebraic set, in particular $\overline{C}$ and every fibre of $\mathrm{pr}_n|_{\overline{C}}$, is locally contractible. The map pr_n is a continuous map between a compact space and a Hausdorff space, so it is closed and proper.

Finally, because $\mathcal{P}_{n-1}$ is a strong c.a.d., by [15, Proposition 5.2] the fibres of $\mathrm{pr}_n|_{\overline{C}}$ are closed segments and thus contractible. □

Note that in Theorem 3.18 we do not need the c.a.d. $\mathcal{P}$ to be strong. We do need $\mathcal{P}_{n-1}$ to be strong in order to apply [15, Proposition 5.2].

Question 3.19 *What conditions on a c.a.d. $\mathcal{P}$ are necessary to ensure that the closures of its bounded cells are contractible? Could it be true for an arbitrary $\mathcal{P}$?*

In fact a compact contractible manifold with boundary X is a d-cell ($d \geq 3$) if and only if $\partial_m X$ is simply connected. For this one uses the topological h-cobordism theorem, which is a consequence of the (generalised) Poincaré conjecture: see [19] for some comments on this, and also [14, Conjecture 3.5′].

More precisely: X is contractible so $\partial_m X$ is a homology $(d-1)$-sphere by Lemma 3.16, so $\partial_m X$ is a homotopy $(d-1)$-sphere because it is simply-connected. Then $\partial_m X$ is cobordant with the sphere $\mathbb{S}_0$ bounding a small ball $\mathbb{B}_0 \subset X \setminus \partial_m X$. Because this cobordism is an h-cobordism it is homeomorphic to $\mathbb{S}_0 \times [0,1]$, and we get that X is homeomorphic to $\mathbb{B}_0 \cup_{\mathbb{S}_0} \mathbb{S}_0 \times [0,1]$ which is $\overline{\mathbb{B}}^d$.

Consequently, to show that a c.a.d. $\mathcal{P}$ of a compact d-dimensional semi-algebraic set is a regular cell complex it is enough to show that every cell closure $\overline{C}$ is compatibly a compact contractible manifold with boundary and, for $d > 3$, that $\partial_c C$ is simply connected.

In general, the closure of a c.a.d. cell is not compatibly a manifold with boundary. The cell W in Example 2.1 is an example of this: $\overline{W} \cong \overline{\mathbb{B}}^2$ but $\partial_m \overline{W}$ is strictly contained in $\partial_c W$. Another example is the non-regular cell $\mathbb{B}^2 \setminus \{(0,y) \mid y \geq 0\}$ mentioned after Definition 2.4.

In very low dimension the position is simple.

Lemma 3.20 *If C is a locally boundary connected* 0- *or* 1-*cell of* $\mathbb{R}^n$, *then $\overline{C}$ is compatibly a manifold with boundary.*

This is immediate from the definition of manifold with boundary. Using a more involved argument, we can show that, under certain conditions, closures of 2-cells are compatibly manifolds with boundary.

Definition 3.21 *A locally compact space X is a* homology n-manifold *if, for all $p \in X$*

$$H_i(X, X \setminus \{p\}) = \begin{cases} \mathbb{Z} & \textit{if } i = n, \\ 0 & \textit{otherwise.} \end{cases}$$

A homology manifold is not a manifold in general, but homology n-manifolds are n-manifolds for $n \leq 2$. In some ways, they behave better than manifolds, as the following fact (stated in [21], and easily checked by the Künneth formula) suggests.

Lemma 3.22 *Suppose X and Y are topological spaces. If $X \times Y$ is a topological manifold, then X and Y are homology manifolds.*

If C is a semi-algebraic cell and $p \in \partial C$ then $\mathbb{B}(p,\varepsilon) \cap C$ is a open subset of a manifold and thus a manifold. As in the proof of Proposition 2.18, the local conic structure gives a homeomorphism

$$\mathbb{B}(p,\varepsilon) \cap C \to (\mathbb{S}(p,\varepsilon) \cap C) \times (0,1),$$

so $\mathbb{S}(p,\varepsilon) \cap C$ is a homology manifold for $0 < \varepsilon \ll 1$.

We can determine when locally boundary connected 2-cells of a cell decomposition are manifolds with boundary.

Proposition 3.23 *Let C be a locally boundary connected semi-algebraic 2-cell in $\mathbb{R}^n$. If, for all $p \in \partial_c C$ and $0 < \varepsilon \ll 1$, the set $\mathbb{S}(p,\varepsilon) \cap C$ is homeomorphic to $(0,1)$, then $\overline{C}$ is compatibly a manifold with boundary.*

Proof: A semi-algebraic homeomorphism $\gamma\colon (0,1) \to \mathbb{S}(p,\varepsilon) \cap C$ extends continuously to a map $\overline{\gamma}\colon [0,1] \to \overline{C}$ by [5, Proposition 2.5.3]. If $\overline{\gamma}(0) = \overline{\gamma}(1)$ (for $\varepsilon \ll 1$) then $\mathbb{S}(p,\varepsilon) \cap C$ is not locally boundary connected, which by Corollary 2.19 would contradict the assumption on C.

Therefore $\overline{\gamma}(0) \neq \overline{\gamma}(1)$, giving a homeomorphism $\overline{\gamma}\colon ([0,1],(0,1)) \to \mathbb{S}(p,\varepsilon) \cap (\overline{C}, C)$.

Hence by the local conic structure of semi-algebraic sets, $\overline{\mathbb{B}}(p,\varepsilon) \cap (\overline{C}, C)$ is homeomorphic to the cone K on $([0,1],(0,1))$, by a map sending p to the vertex. This is a manifold with boundary, and $p \in \partial_m K$, so $\overline{C}$ is a manifold with boundary and $\partial_c C \subseteq \partial_m \overline{C}$. But $C \cap \partial_m \overline{C} = \emptyset$ because C is a manifold, so $\overline{C}$ is compatibly a manifold with boundary. □

We can apply this to strong c.a.d.s.

Lemma 3.24 *Let C be a 2-cell of a strong c.a.d. of $\mathbb{R}^n$. Then $\overline{C}$ is compatibly a manifold with boundary.*

Proof: By the previous discussion, $\mathbb{S}(p,\varepsilon) \cap C$ is a connected 1-manifold so it is homeomorphic to either $\mathbb{S}^1$ or $(0,1)$. As a strong c.a.d. is well-bordered, there exists a 1-cell $C' \subset \partial C$, such that $p \in \overline{C'}$. If $\mathbb{S}(p,\varepsilon) \cap C$ is homeomorphic to $\mathbb{S}^1$, then $\mathbb{B}(p,\varepsilon) \cap \overline{C}$ is the cone with vertex p on $\mathbb{S}^1$, but that has an isolated boundary point at p. □

We have proved the following result.

Corollary 3.25 *Let $S \subset \mathbb{R}^n$ be a 2-dimensional compact semi-algebraic set. If $\mathcal{P}$ is a strong c.a.d. adapted to S, then $\mathcal{P}$ represents S as a regular cell complex.*

To prove Conjecture 3.13 for $n = 3$ by this method, we would need to show that 3-cells of a strong c.a.d. of $\mathbb{R}^3$ are manifolds with boundary. We have not been able to do this but we can do so under the additional assumption that the cells of the c.a.d. are locally boundary simply connected: see Definition 2.16.

Theorem 3.26 *Let $\mathcal{P}$ be a strong c.a.d. of $\mathbb{R}^3$. If C is a locally boundary simply connected 3-cell of $\mathcal{P}$, then $\overline{C}$ is compatibly a manifold with boundary.*

Proof: For $p \in \partial C$ and $0 < \varepsilon \ll 1$, we know that $\mathbb{S}(p, \varepsilon) \cap C$ is a homology 2-manifold and therefore a 2-manifold.

We claim that $\mathbb{S}(p, \varepsilon) \cap C$ is a 2-cell that satisfies the conditions of Proposition 3.23. Then, as before, $\mathbb{S}(p, \varepsilon) \cap C$ is a regular cell and the interior of the cone on its closure is homeomorphic to $[0, 1)^3$, which is compatibly a manifold with boundary.

As in the proof of Proposition 3.23, $\mathbb{S}(p, \varepsilon) \cap C$ is locally boundary connected because of Corollary 2.19.

If $q \in \partial_c(\mathbb{S}(p, \varepsilon) \cap C)$, then $(\mathbb{S}(p, \varepsilon) \cap C) \cap \mathbb{S}(q, \varepsilon')$ is a connected manifold. If it is homeomorphic to $\mathbb{S}^1$ then $(\mathbb{S}(p, \varepsilon) \cap C)$ is not locally boundary 1-connected so C is also not locally boundary 1-connected, again by Corollary 2.19.

Lastly, we prove that $\mathbb{S}(p, \varepsilon) \cap C$ is a cell: as the dimension is 2, it is enough to show that it is contractible. It is a metrisable manifold, so by [18, Corollary 1] it has the homotopy type of a CW complex. Thus, by Whitehead's theorem, it suffices to show that $\mathbb{S}(p, \varepsilon) \cap C$ has trivial homotopy groups: π_0 and π_1 are trivial by assumption. By Hurewicz's theorem it is enough to show that all the homology groups vanish. Certainly $H_i(\mathbb{S}(p, \varepsilon) \cap C) = 0$ for all $i \geq 3$, as $\mathbb{S}(p, \varepsilon) \cap C$ is a 2-manifold.

But $\mathbb{S}(p, \varepsilon) \cap C$ is a non-compact connected 2-manifold, otherwise p is an isolated point of $\partial_c C$ exactly as in Lemma 3.24, so $H_2(\mathbb{S}(p, \varepsilon) \cap C) = 0$ by [11, Cor VIII.3.4]. ☐

Theorem 3.27 *Suppose that $S \subset \mathbb{R}^3$ is semi-algebraic and $\mathcal{P}$ is a strong c.a.d. adapted to S, such that every 3-cell of $\mathcal{P}$ is locally boundary simply connected. Then $\mathcal{P}$ yields a regular cell complex of S.*

Proof: This follows from Theorem 3.26 and the discussion after Question 3.19. ☐

In the light of the above arguments, we consider it likely that Conjecture 3.13 in full generality requires that the cells of the c.a.d. should be assumed to be locally boundary contractible, not just locally boundary connected.

Extending our approach, even with that stronger hypothesis, to higher dimension, we encounter at least two difficulties. First, if $[0, 1] \times M$ is a 3-manifold then M is a 2-manifold, and we used this in the proof of Theorem 3.26, but the corresponding assertion in higher dimension is false. Second, we also used the fact that a contractible 2-manifold is a cell, which also fails in higher dimension: it would be enough to show that in addition the boundary is a sphere, the main obstacle being to determine whether the boundary is simply connected.

We would encounter these difficulties, for instance, if we tried to prove a version of Proposition 3.23 (e.g. with C being locally boundary contractible and $\mathbb{S}(p,\varepsilon)\cap C$ a $(d-1)$-cell), but we could also try to exploit the fact that C is a c.a.d. cell, rather than just semi-algebraic.

4 Subadjacency and order complex

In [15], Lazard poses the following question (as rephrased, but not altered, by us).

Question 4.1 *Let $\mathcal{P}$ be a strong sampled c.a.d. adapted to a compact semi-algebraic set S and let $\mathcal{E}$ be the set of cells of $\mathcal{P}$ contained in S. For each subadjacency chain $\mathbf{E} = (E_0 \prec E_1 \prec \cdots \prec E_k)$ with $E_j \in \mathcal{E}$, we let $\sigma_{\mathbf{E}}$ be the convex hull of the sample points $b_{E_0}, \ldots, b_{E_k}$.*

(i) Is $\{\sigma_{\mathbf{E}}\}$ a simplicial complex?

(ii) Is $\bigcup_{\mathbf{E}} \sigma_{\mathbf{E}}$ homeomorphic to S?

Some general position condition on the sample points is needed, otherwise $\sigma_{\mathbf{E}}$ could even fail to be a k-simplex. Even so, the answer to (i) as posed above is no, because there may be intersections, as the example below illustrates.

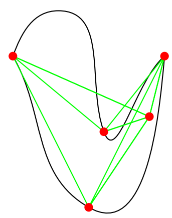

Question 4.1(i): not a simplicial complex

However, this is easily corrected: instead of working inside $\mathbb{R}^n$ one should replace $\{\sigma_{\mathbf{E}}\}$ by the order complex $\Delta(\mathcal{E})$ of the poset $(\mathcal{E}, \preceq)$ (it follows immediately from Lemma 2.10 that $\preceq$ is transitive) and $\bigcup_{\mathbf{E}} \sigma_{\mathbf{E}}$ by the geometric realisation $\|\Delta(\mathcal{E})\|$. See [4] for details of order complexes and barycentric subdivision.

Question 4.2 *Let $\mathcal{P}$ be a strong c.a.d. adapted to a compact semi-algebraic set S and let $\mathcal{E}$ be the set of cells of $\mathcal{P}$ contained in S. Is $\|\Delta(\mathcal{E})\|$ homeomorphic to S, i.e. is $\Delta(\mathcal{E})$ a triangulation of S?*

A regular cell complex can be thought of as a CW-complex that is one barycentric subdivision from being a triangulation.

Lemma 4.3 *Let Σ be a regular cell complex and let Σ^* be the set of all closed cells ordered by inclusion. Then $\|\Sigma\| \cong \|\Delta(\Sigma^*)\|$.*

For details, see [4, 12.4 (ii)]. It is easy to see that in general a regular cell complex is not a triangulation (e.g. [3, Example 5.4]).

Theorem 4.4 *Let $\mathcal{P}$ be a strong c.a.d. of $\mathbb{R}^n$. The partially ordered set $(\mathcal{P}, \preceq)$ of cells with respect to sub-adjacency is isomorphic to the partially ordered set $(\mathcal{P}^*, \subseteq)$ of closed cells with respect to inclusion.*

Proof: The bijection $\mathcal{P} \to \mathcal{P}^*$ is given by closure, $C \mapsto \overline{C}$. The inverse is given by relative interior, taking a closed cell Z to its interior in its Zariski closure, or by taking Z to the cell (unique, by Lemma 2.10) of dimension equal to $\dim Z$ that meets Z. We need to show that $C \preceq D$ if and only if $\overline{C} \subseteq \overline{D}$, which follows immediately from Lemma 2.10. $\square$

Thus Conjecture 3.13 would imply an affirmative to Question 4.2. Note also that in Theorem 4.4 the cylindricity is not used, so $\mathcal{P}$ does not need to be a c.a.d., only a semi-algebraic cell decomposition.

References

[1] S. Basu, A. Gabrielov & N. Vorobjov, Semi-monotone sets. *J. Eur. Math. Soc.* **15** (2013), no. 2, 635–657.

[2] S. Basu, A. Gabrielov & N. Vorobjov, Triangulations of monotone families I: two-dimensional families. *Proc. Lond. Math. Soc.* **111** (2015), no. 5, 1013–1051.

[3] S. Basu, R. Pollack & M.-F. Roy, *Algorithms in real algebraic geometry.* Algorithms and Computation in Mathematics **10**. Springer-Verlag, Berlin, 2003.

[4] A. Björner, *Topological methods.* Handbook of combinatorics, Vols. **1–2**, 1819–1872, Elsevier, Amsterdam 1995.

[5] J. Bochnak, M. Coste & M.-F. Roy, *Real algebraic geometry.* Ergebnisse der Mathematik und ihrer Grenzgebiete **36**. Springer-Verlag, Berlin, 1998.

[6] R.J. Bradford, J.H. Davenport, M. England, S. McCallum & D.J. Wilson, Truth table invariant cylindrical algebraic decomposition, *J. Symbolic Computation* **76** (2016), 1–35.

[7] C.W. Brown, Open non-uniform cylindrical algebraic decompositions, *Proceedings ISSAC 2015* (D. Robertz, Ed.) (2015), 85–92.

[8] C. Chen & M. Moreno Maza, Cylindrical algebraic decomposition in the RegularChains library, *Proceedings Mathematical Software — ICMS 2014*, (2014), 425–433.

[9] G.E. Collins, Quantifier elimination for real closed fields by cylindrical algebraic decomposition. *Automata theory and formal languages (Second GI Conf., Kaiserslautern, 1975)*, 134–183. Lecture Notes in Comput. Sci. **33**, Springer, Berlin 1975.

[10] M. Demdah Kartoue, *h-cobordism and s-cobordism theorems: transfer over semialgebraic and Nash categories, uniform bound and effectiveness.* Ann. Inst. Fourier (Grenoble) **61** (2011), no. 4, 1573–1597 (2012).

[11] A. Dold, *Lectures on algebraic topology.* Grundlehren der Mathematischen Wissenschaften **200**. Springer-Verlag, Berlin-New York 1980.

[12] R. Fukasaku, H. Iwane & Y. Sato, Real quantifier elimination by computation of comprehensive Gröbner systems, *Proceedings ISSAC 2015* (D. Robertz, Ed.) (2015), 173–180.

[13] A. Hatcher, *Algebraic topology.* Cambridge University Press, Cambridge 2002.

[14] J. Hempel, *3-Manifolds.* Ann. of Math. Studies **86**, Princeton University Press, Princeton, NJ 1976.

[15] D. Lazard, CAD and topology of semi-algebraic sets. *Math. Comput. Sci.* **4** (2010), no. 1, 93–112.

[16] A. Locatelli, *On the regularity of cylindrical algebraic decompositions*, PhD Thesis, Bath 2015.

[17] S. McCallum, On projection in CAD-based quantifier elimination with equational constraints, *Proceedings ISSAC '99* (S. Dooley, Ed.) (1999), 145–149.

[18] J. Milnor, On spaces having the homotopy type of a CW-complex. *Trans. Amer. Math. Soc.* **90** (1959) 272–280.

[19] R. Schultz, *Between lower and higher dimensions* in: Proceedings of a symposium in recognition of the retirement of Terry Lawson from Tulane University, 2007. `http://math.ucr.edu/~res/miscpapers/tlawson.pdf`

[20] J.T. Schwartz, M. Sharir, On the "piano movers" problem II. General techniques for computing topological properties of real algebraic manifolds. *Adv. in Appl. Math.* **4** (1983), no. 3, 298–351.

[21] E.G. Sklyarenko, Homology Manifold, *Encyclopedia of Mathematics*, Vol. **6**. English translation edited by M. Hazewinkel. Kluwer, Dordrecht 1990. https://www.encyclopediaofmath.org/index.php/Homology_manifold

[22] S. Smale, A Vietoris mapping theorem for homotopy. *Proc. Amer. Math. Soc.* **8** (1957), 604–610.