Topological Perplexity of Feedback Stabilization

Yu. Baryshnikov*
University of Illinois at Urbana-Champaign
ymb@illinois.edu

August 26, 2021

Abstract

Feedback control is one of the key approaches in the controlled dynamical systems, allowing to stabilize complex constructs on some desirable trajectories (or attractors). It is well known that the disparity between the homotopy type of the configuration space of the system and the attractor precludes existence of a continuous vector field realizing such a stabilization, thus forcing the discontinuity of the control function at some points of the configuration space. How complicated should be this discontinuity set? In this note we present some simple lower bounds on the ranks of homology groups of such sets of discontinuities (the cuts), and discuss some examples.

1 Introduction

Control systems deal with (locally) parameterized dynamical systems. The inherent flexibility of such systems can often be leveraged to realize, approximately and locally, almost arbitrary dynamics, at least in the situation when the system is controllable. There are however some *global* constraints that often force discontinuities of the steering algorithms, especially in the so-called *closed loop* situation, where the control is a function of the system state.

This note deals with some baseline quantification of the topology of those discontinuities, leading to certain quantities (introduced in the Section 3) that bound from below the ranks of homology groups of the set where the control is discontinuous.

1.1 Setup

Let M be the configuration space of the system. We assume that it is a smooth manifold of dimension d, perhaps with corners, equipped with a *control structure*

$$\dot{x} = f(x, u). \tag{1}$$

Formally, (1) can be thought of as the tuple $(M, \mathcal{U}, f)$, where

*Supported in part by the ARO under the SLICE Multidisciplinary University Research Initiatives (MURI) Program.

- $\mathcal{U} \subset E$ is an open subset of the total space of a vector bundle $p : E \to M$ interpreted as admissible controls: $p^{-1}(x) \cap \mathcal{U} =: \mathcal{U}_x$ is the collection of control actions available at the point $x \in M$, and
- $f : \mathcal{U} \to TM$ is a fiberwise mapping taking values in the tangent bundle of M.

The details of this setup are not critically important here, as the focus of this note lies on the restrictions one faces when the control structure is coupled with the feedback function, so that the resulting dynamical system achieves the *feedback stabilization* of the system (1) on an "attractor" $\boldsymbol{A}$.

We will assume that $\boldsymbol{A}$ is a stratified subset of M. One can imagine quite broad variety of setting, including the situation where the attractor is disconnected (say, one is aiming at stabilization on one of a few gaits, depending on the starting point of the trajectory). The components of the attractor can be of different dimensions, and so on.

Most important examples of attractors in practice are, of course, the stationary points, and the periodic trajectories.

1.2 Feedback Stabilization

In terms of the control structure (1), *feedback stabilization* starts with a section

$$k : M \to \mathcal{U}, k(x) \in \mathcal{U}_x$$

of the bundle $\mathcal{U}$ resulting in the dynamical system

$$\dot{x} = g(x) = f(x, k(x)).$$

The map k is called the *feedback control*).

The smooth vector field g is from now on the focus of our attention, so that we can ignore its origin, and concentrate on its properties for the remainder of the paper.

We say that g *stabilizes* on $\boldsymbol{A}$ if

- the configuration space M is preserved by g (i.e. the vector field *points inside* M at the boundary points of M);
- the attractor is an *invariant set* of the vector field g (for example, this is the case when g is tangent to the strata of the attractor); and moreover
- the attractor is *globally asymptotically stable*, in the usual sense: for any open vicinity of the attractor, trajectories starting close enough to $\boldsymbol{A}$ never leave that vicinity, and every trajectory of g approaches $\boldsymbol{A}$ asymptotically (the ω-limit of any trajectory is in $\boldsymbol{A}$).

1.3 Obstacles to Continuous Feedback Stabilization

It is well known that in general such a feedback stabilizing control might fail to exist. Indeed, under some mild regularity conditions, the existence of such a vector field g would imply that $\boldsymbol{A}$ is a deformation retract of the configuration space M, meaning that the attractor and the configuration space have the same homotopy type, a strong topological constraint.

The fact that the topology might serve as an obstacle to existence of continuous stabilizing feedback has been noted by many authors (see [RT09,KR90,Son12] among others). In most of the literature the authors offer various schemes to overcome those obstacles, essentially by exempting a part of the configuration space from being potential starting point of a trajectory.

A classical manifestation of this idea is the fact that the complement to the cut locus of a point in a Riemannian manifold M is contractible [Whi35]. Recall that the cut locus of a point x_* is the set of points in M where the geodesic distance to x_* is not smooth. As an example, for a sphere the cut locus is just the point opposite to x_*, if the metric on the sphere is round.

Generically (at least for real analytic manifolds) the cut locus is a stratified set of codimension 1 [Buc77].

The cut locus of a point in a sphere is always contractible. For other manifolds, the topology of the cut locus is different, but the complement to it can be always retracted to x_* by the flow along the antigradient of the distance to x_*.

1.4 Obstacles in Control Theory

In [RT09], the authors form a non-holonomic version of the cut-locus construction. Away from the excluded set, the authors construct a globally asymptotically stable vector field. Their vector field is smooth only on the complement to the excluded set which is, however, flow-invariant: once out of the excluded set, the trajectory never reenters it and is therefore defined unambiguously.

In another setting motivated by the motion planning problem in robotics the authors of [KR90] relied on the gradient descent derives from an appropriate generic Morse function that explodes near the boundary of the configuration space. This provides the stabilization, providing that one excludes the union of the stable manifolds of critical points of positive indices. The vector field they define is smooth, and the integral trajectories of the vector field exist for any starting point.

Remark 1.1. *The constructions cited above can be contrasted in terms of the time required to reach some open vicinity of the attractor: while the time it takes the flow along the gradient of the distance to x_* to take any point of the complement to the cut locus to a fixed vicinity of x_* is uniformly bounded (if the configuration space is compact), one can experience* stalling *if one uses some artificial Morse function with critical points of positive index: a trajectory could spend arbitrarily long time (of order $\log(1/\varepsilon)$ if the starting point was at the distance ε away from the stable manifold of the non-minimal critical point) before leaving a fixed vicinity of the excluded set.*

The approach of this note is to embrace the necessity of sets where the feedback control is discontinuous, and to start explorations of *how complex* that set should be. We ignore the question of the *size* of the excluded set - be it its volume, or some other integral invariants. We are rather interested in the topology of the set: its properties up to homotopy equivalence. As we will see, the natural topological invariants of the excluded set are, under some natural regularity conditions, are rather rigid, for a given pair of configuration space and attractor.

We work here with singular homologies.

2 Cuts and Their Topology

Motivated by the discussion above we introduce the notion of *cuts*.

Let us assume that the feedback control k and, through it, the vector field g are discontinuous at some subset of the configuration space (and is smooth elsewhere).

2.1 Cuts

Definition 2.1. *We will refer to the closed subset $\boldsymbol{C}$ of M as the* cut *if*

- *the feedback control function k can be defined on the* basin

$$M^{\odot} := M - (\partial M \cup \boldsymbol{C})$$

- *so that the basin is preserved by the flow defined by the vector field $g = f(x, k(x))$, and*
- *the vector field g stabilizes on the attractor $\boldsymbol{A} \subset M^{\odot}$.*

In general, the cuts can be pretty wild (compare [Zam98] for the cut locus case), so we will impose some degree of tameness by fiat. Assuming that the union of the cut and the boundary of M can be Whitney stratified is enough to make sure no pathological situations arise. This assumption is readily satisfied in most applied settings.

As we mentioned above, the topology of cuts is normally not addressed in any detail, besides noting that the excluded sets are (in most constructions) of measure zero, thus rendering them, presumably, unnoticeable in the applied situations.

Our main result is a simple observation on the topology of the cuts, forced by the mismatch between the topologies (homotopy types) of the configuration space and the attractor.

To exclude possible anomalies we introduce a condition of *perfect coupling* between the attractor and the cut (which covers most practical applications one can think of).

2.2 Perfect Coupling

Let us introduce a few extra notations.

We will call the union of the boundary of the configuration space and the cut the *extended cut*,

$$\boldsymbol{E} := \partial M \cup \boldsymbol{C}.$$

Definition 2.2. *We say that the vector field $g = f(x, k(x))$ couples the extended cut and the attractor perfectly if*

- *there exist open neighborhoods*

$$\boldsymbol{E}^{\circ} \supset \boldsymbol{E}, \quad \boldsymbol{A}^{\circ} \supset \boldsymbol{A}$$

  of the extended cut and the attractor respectively, such that

- *inclusions $\boldsymbol{E} \hookrightarrow \boldsymbol{E}^{\circ}$ and $\boldsymbol{A} \hookrightarrow \boldsymbol{A}^{\circ}$ are homotopy equivalences, and*

- *the boundaries $\partial \boldsymbol{E}^\bullet$ and $\partial \boldsymbol{A}^\bullet$ are hypersurfaces in $M^\odot$, and the trajectories of g foliate $M - (\boldsymbol{E}^\circ \cup \boldsymbol{A}^\circ)$ establishing a diffeomorphism between $\partial \boldsymbol{E}^\bullet$ and $\partial \boldsymbol{A}^\bullet$ (we denote the closures of $\boldsymbol{E}^\circ$ and $\boldsymbol{A}^\circ$ as $\boldsymbol{E}^\bullet$ and $\boldsymbol{A}^\bullet$, respectively).*

These conditions are not as burdensome as one can suspect. The existence of collars (open neighborhoods retracting onto a closed set and having smooth boundaries) is well-known in a wide range of situations, such as Whitney stratified subsets. The only essential condition is that shifts along the trajectories of g retract: either the complement to a vicinity of the attractor to the extended cut, or the complement of a vicinity of the extended cut to the attractor. Either of those deformation rectractions is conceptually what one expects from the feedback stabilization, and can be verified in most applied settings.

2.3 Long Exact Sequence

Under the assumption of the perfect coupling, the following result is almost immediate:

Proposition 2.3. *Assume that the vector field g perfectly couples the attractor and the extended cut.*

Then there is a long exact sequence of homology groups

$$\ldots \to H_k(\boldsymbol{C}, \partial \boldsymbol{C}) \quad \to \quad H_k(M, \partial M) \to \ldots \to H_k(\boldsymbol{A}^\bullet, \partial \boldsymbol{A}^\bullet) \to H_{k-1}(\boldsymbol{C}, \partial \boldsymbol{C}) \to \ldots$$

(here $\partial \boldsymbol{C} = \boldsymbol{C} \cap \partial M$).

In particular, if the attractor is a connected manifold, the Thom's isomorphism implies that one can replace $H_k(\boldsymbol{A}^\bullet, \partial \boldsymbol{A}^\bullet)$ in (2) with $H^r_{k-m+a}(\boldsymbol{A})$, the reduced homologies of $\boldsymbol{A}$ shifted by the codimension of the attractor in M.

The importance of this exact sequence is in that it sandwiches the homologies of the cut (modulo its boundary) between the homologies of the known objects, namely of the configuration space and the attractor.

Proof. Consider the long exact sequence of the triple $\partial M \hookrightarrow \boldsymbol{E}^\bullet \hookrightarrow M$. The corresponding long exact sequence of homology groups reads

$$\ldots \to H_k(\boldsymbol{E}^\bullet, \partial M) \to H_k(M, \partial M) \to\to H_k(M, \boldsymbol{E}^\bullet) \to H_{k-1}(\boldsymbol{E}^\bullet, \partial M) \to \ldots$$

The terms $H_k(\boldsymbol{E}^\bullet, \partial M)$ are isomorphic to $H_k(\boldsymbol{E}, \partial M)$ (by the perfect coupling), and then to $H_k(\boldsymbol{C}, \partial \boldsymbol{C})$, by excision.

Similarly, the terms $H_k(M, \boldsymbol{E}^\bullet)$ arexs isomorphic to $H_k(M, M - \boldsymbol{A}^\circ)$, by perfect coupling, and then to $H_k(\boldsymbol{A}^\bullet, \partial \boldsymbol{A}^\bullet)$, by excision. □

3 Betti Numbers and Perplexity

The Proposition 2.3 implies several inequalities on the ranks of the homology groups of the cut.

3.1 Betti Numbers

Assume now that the homologies have coefficients in a field, and denote by

$$\beta_k(\cdot) := \mathrm{rk}\, H_k(\cdot), \quad \mu_k := \beta_k(\cdot) - \beta_{k-1}(\cdot) + \beta_{k-2}(\cdot) - \ldots \pm \beta_0(\cdot),$$

the ranks of the homology groups (i.e. *Betti numbers*) and their alternating sums (obviously, $(-1)^n \mu_n$ stabilizes at the Euler characteristic of the space).

In these notations we have

Corollary 3.1. *In the conditions of Proposition 2.3,*

$$\begin{array}{rcl} \mu_k(\boldsymbol{C}, \partial \boldsymbol{C}) + \mu_k(\boldsymbol{A}^\bullet, \partial \boldsymbol{A}^\bullet) & \geq & \mu_k(M, \partial M) \\ \mu_k(M, \partial M) + \mu_{k-1}(\boldsymbol{C}, \partial \boldsymbol{C}) & \geq & \mu_k(\boldsymbol{A}^\bullet, \partial \boldsymbol{A}^\bullet) \\ \mu_{k+1}(\boldsymbol{A}^\bullet, \partial \boldsymbol{A}^\bullet) + \mu_k(M, \partial M) & \geq & \mu_k(\boldsymbol{C}, \partial \boldsymbol{C}) \end{array} \tag{2}$$

These inequalities become equalities for $k \geq \dim M$.

Also, somewhat weaker but easier to apply inequalities hold:

$$\begin{array}{rcl} \beta_k(\boldsymbol{C}, \partial \boldsymbol{C}) & \geq & \beta_k(M, \partial M) - \beta_k(\boldsymbol{A}^\bullet, \partial \boldsymbol{A}^\bullet) \\ \beta_k(\boldsymbol{C}, \partial \boldsymbol{C}) & \geq & \beta_{k+1}(\boldsymbol{A}^\bullet, \partial \boldsymbol{A}^\bullet) - \beta_{k+1}(M, \partial M) \end{array} \tag{3}$$

Proof. If one has a long exact sequence of vector spaces,

$$\ldots C_{k+1} \to A_k \to B_k \to\to C_k \to A_{k-1} \to \ldots \to A_0 \to B_0 \to C_0 \to 0$$

one can severe its left end at any place, substituting cokernel D of the mapping where the cut was made, so that the new long sequence is exact as well:

$$0 \to D \to \ldots \to A_0 \to B_0 \to C_0 \to 0.$$

The alternating sum of the ranks of the finite long exact sequence vanishes. Using that and the fact that $\mathrm{rk}\, D \geq 0$, we arrive to inequalities in (2) after specializing to various terms in (2).

To obtain (3) one just adds the first and second inequalities of (2) with indices k and $k-1$. □

3.2 Topological Perplexity

These inequalities provide *lower bounds* on how nontrivial the topology of the cut needs to be.

Consider the standard example, of stabilization on a single equilibrium point. In this case,

$$\beta_k(\boldsymbol{A}^\bullet, \partial \boldsymbol{A}^\bullet) = 1 \text{ if } k = \dim M \text{ and } 0 \text{ otherwise.}$$

Correspondingly, the

$$\beta_k(\boldsymbol{C}, \partial \boldsymbol{C}) \geq \beta_k(M, \partial M)$$

in all dimensions up to $\dim M - 1$: the cut needs to suppress all homologies of the configuration space with exception of the top one. In particular, this forces the dimension of the cut has to be at least the highest dimension of nontrivial homology class of the configuration space other than the fundamental class.

Using the Corollary 3.1 as motivation, we introduce the following definition:

Definition 3.2. *The vector*

$$\boldsymbol{p}_k := |\beta_k(M, \partial M) - \beta_k(\boldsymbol{A}^\bullet, \partial \boldsymbol{A}^\bullet)|$$

of mismatches *between the ranks of the* (*relative*) *homologies of* $(M, \partial M)$ *and* $(\boldsymbol{A}^\bullet, \partial \boldsymbol{A}^\bullet)$ *is called the* topological perplexity *of the feedback stabilization problem* (1).

The topological perplexity forces the cut, - barring any anomalies in the system, - to have nontrivial topology, and in many situations fixes its structure (implying, e.g. the contractibility of the cut locus of a point in a sphere). It is a powerful numerical characteristic of how complex the algorithms of the stabilization need to be. We won't address the details here, but in an appropriate setting this complexity can be quantified using the ideas of lower bounds of the decision problems of membership in a set forced by the nontrivial topology of the set, compare [GV17] and work quoted there.

4 Examples

Cut Locus

Let M a connected manifold without boundary, and $\boldsymbol{A} = x_* \in M$ a point. The vector of topological perplexities is

$$\boldsymbol{p}_k = (0, \beta_{d-1}(M), \ldots, \beta_0(M) = 1)\,,$$

and automatically defines the lower bounds on the cut. In the simplest situation, when the configuration space is a sphere, we see that the lower bound implies that the cut is a contractible *nonempty* set.

If the configuration space is a d-dimensional torus, the topology of the cut becomes quite nontrivial, with the perplexity vector given by

$$\boldsymbol{p}_k = \binom{d}{k} \text{ for } d > k \geq 0.$$

RHex and Its Legs

In [BS14], we considered a stylized model of *RHex* [SBK01], a versatile multi-legged robotic platform. In that model one starts with the d-torus (describing leg orientations, - it is convenient to think of each leg coordinate as a point on the unit circle in the plane) and remove from it a *forbidden region* corresponding to the positions of the legs physically incompatible with the robotic motion (think of conditions like "not all legs on the right side are up").

In that model, the forbidden region is a collar neighborhood of a *coordinate toric arrangement*, defined by family I of subsets of the collection $\{1, \ldots, d\}$ of the robotic legs. Semantically, a subset I of the legs is in the family if no matter what are their positions, as soon as all other legs are pointing up, the resulting configuration is forbidden. So, the skeleton of the forbidden region is the union of the coordinate tori corresponding to the subset in the family, and the forbidden region is its open vicinity retractable to that union. Clearly, the family I is closed under taking subsets.

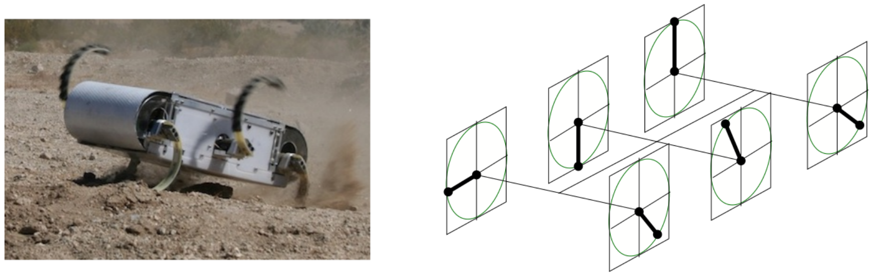

Figure 1: RHex: in real life on the left picture; an abstraction on the right. The configuration space of the model is the product of circles (corresponding to the orientations of the legs) from which certain forbidden configurations are excluded.

It is easy to describe the homologies of the configuration space obtained using this construction: each subset of size k in I kills one generator of $H_k(\boldsymbol{T}^d)$: the basis in the homology groups of the space $(M_I, \partial M_I)$ can be identified with the complementary family of all subsets not in I. In other words,

$$\beta_k(M, \partial M) = \binom{d}{k} - I_k,$$

where I_k is the number of subsets of size k in I.

In practice, one often wants to stabilize on a periodic gait, the most natural ones being a periodic trajectory in the torus $\boldsymbol{T}^d$ outside of the forbidden region, homologous to the diagonal.

In [BS14] we described a cut that, perhaps surprisingly, solves the feedback stabilization problem for *all* possible families I. (The impossible families are those containing subsets of size $(d-1)$, for which one of the legs cannot make full rotation at all - clearly one cannot realize a gait homologous to the diagonal avoiding this forbidden set.)

The cut we constructed consists of all configurations where among the legs with the highest vertical coordinate, there is a pair $k, l \neq k+1$ such that the horizontal position of the leg l is *right* of the horizontal position of the leg k (*admissible* configurations are those compatible with the counterclockwise rotation of the cyclically ordered legs $1-2-\cdots-d-1$). It is a $(d-1)$-dimensional simplicial complex; one can think of it as a physical implementation of a sorting algorithm: after one rotation, the legs, if their phases are distinct, are cyclically sorted.

For that cut, the inequalities in Corollary 3.1 become exact (if one wonders how the same cut can match different topologies of $(M, \partial M)$, recall that the constraint is on $(\boldsymbol{C}, \partial \boldsymbol{C})$, and the boundary $\partial \boldsymbol{C}$ depends on I).

Configuration Spaces of Hard Disks

The familiar construct of *configuration spaces for* n (*distinguishable*) *particles* in V is defined as the Cartesian products of the configuration spaces for each particle from which the diagonals, - configurations where the particles overlap, - are removed:

$$M(n) = \{\mathbf{x} = (x_1, \ldots, x_n) \in V^n, x_k \in V : x_k \neq x_l, \text{ for } k \neq l\}.$$

Most of the literature is focused on the case of $V = \mathbb{R}^s$.

This setup might seem as not immediately fitting our model (open manifold, no boundary); however one can tweak the space by

- fattening the diagonals: replacing the condition that the points are distinct with the condition that some *disks of fixed radii* around those points do not overlap, and
- by compactifying $V^n = \mathbb{R}^{ns}$ by restricting the particles to large enough box. We note that both changes make the model more realistic from the engineering viewpoint, withour changing the homotopy type of the configuration space, yet making it homeomorphic to a manifold with boundary if the bounding box is large enough (and does not have some small nooks): compare [BBK14].

Altogether, this represent a quite natural applied setting: one wants to stabilize a flock of the disk-like agents, say force them to a particular configuration.

The Betti numbers β_k for the configuration spaces in the Euclidean spaces were studied in great details (see, e.g. [Bj4]): thus, their Poincare polynomials for the configuration space of particles in the plane are know to be

$$\sum \beta_k t^k = \prod_{l=0}^{n-1}(1+lt)$$

The homologies we need are somewhat different, of course: we need not the homologies of the space M, but rather that modulo its boundary.

Let us add to V^n the point at infinity compactifying it (the added part of the space will lie outside of the box to which the realistic positions of the agents are confined). Let Δ be the excluded space, i.e. the (closure) of the union of the diagonals we exclude (so that $\partial M = M \cap \Delta$), and the "configurations near the infinity". Then by Alexander duality,

$$H_k(M) = H_{d-k-1}(\Delta)$$

(here $d = ns$, the dimension of the configuration space), and

$$H_{d-k-1}(\Delta) \cong H_{d-k}(\boldsymbol{S}^d, \Delta) \cong H_{d-k}(M, \partial M).$$

This allows one to express the needed Betti numbers of $(M, \partial M)$ in terms of the Betti numbers of the configuration space of n particles in V, and thus to find the perplexity for the problem of non-colliding flock stabilization.

Say, to stabilize a flock of n particles in the plane at a point, one needs to to have a cut of dimension at least $(2n-1)$ as

$$\operatorname{rk} H_1(M) = \binom{n}{2}$$

(the classes are generated by turns of the particle k around particle l, $k < l$).

A cut on the complement to which such a stabilizing feedback control does exist is easy to construct, at lest in the fully actuated situation, where the control set U is the full tangent space (and f is the identity). One should, namely, just to exclude all configurations where the particles

k and l (with $k > l$) are aligned vertically, with the particle l strictly above the particle k. The complement to such configurations is contractible (by linear homotopy to the configuration where k-th particle is placed at the point $(0, k)$ of the plane).

5 Conclusion

Is Topological Perplexity Useful? One might wonder, how significant is the notion of topological perplexity from the viewpoint of the theory and practice of controlled systems. Is knowing the lower bounds on the ranks of homology of excluded regions, cuts, where the feedback control loses continuity useful?

We believe it is. At the very minimum, the perplexity bounds from below the dimension of the cuts, giving some measure-theoretic intuition on the size of those excluded sets.

More importantly, those bounds hint at the complexity of the algorithms implementing the stabilization. If the continuous feedback is unavailable, and switched ("hybrid", in the modern control-theoretic nomenclature) behavior is unavoidable, any implementation would at the very least require the detection whether the system is located at the cut. As we mentioned above, the complexity of the algorithms detecting the membership in a semialgebraic set (the typical setup for the questions of computational complexity over the reals) is lower bounded by the total Betti number of that set, in many natural computational models.

It is instructive to consider the hard disks configuration spaces ('swarms of Roombas') example from this perspective. We note that the *total Betti number* of the configuration space of n points in a Euclidean space is $n!$ and so grows faster than exponential: the amount of intelligence per agent in large flocks is necessarily grows to infinity (if slowly).

Topological Complexity It behooves us to clarify relation of the *Topological Perplexity* and *Topological Complexity*, an invariant of a path connected topological space introduced by M. Farber in the context of motion planning [Far08]. Recall that Farber's TC is the Schwarz genus of the fibration sending the path space of a configuration space to its end points: the smallest size of the covering of $M \times M$ by open sets such that there is a continuous section of the path fibration: a *motion planning* algorithm depending on the start- and end-points continuously. When the input of the motion-planning problem, - the (start, finish) pair in $M \times M$, - leaves an open set over which the continuous motion planning mapping exists, the output can change drastically.

While this discontinuity is certainly resembling the discontinuities of the feedback-controlled vector fields g at the cuts, it is rather different in the setup (the attractor is fixed!), and in the outcome: not only existence of the discontinuities is asserted, but their topology is quantified.

Topological perplexity is also providing somewhat richer granularity of the answers: while the TC is always bounded by roghly twice the Lusternik - Schnirelman category of the configuration space (which in turns cannot exceed $\dim(M) + 1$), the topological perplexity can be arbitrarily large.

Acknowledgment This project was started and developed under the influence of Dan Koditschek, telling the stories about robots, and interpreting them as mathematical queries. I am most indebted

to him.

References

[BBK14] Yuliy Baryshnikov, Peter Bubenik, and Matthew Kahle. Min-Type Morse Theory for Configuration Spaces of Hard Spheres. *International Mathematics Research Notices*, 2014(9):2577–2592, 2014.

[Bj4] Anders Björner. Subspace arrangements. In *First European Congress of Mathematics, Vol. I* (*Paris, 1992*), volume 119 of *Progr. Math.*, pages 321–370. Birkhäuser, Basel, 1994.

[BS14] Yuliy Baryshnikov and Boris Shapiro. How to Run a Centipede: a Topological Perspective. In Gianna Stefani, Ugo Boscain, Jean-Paul Gauthier, Andrey Sarychev, and Mario Sigalotti, editors, *Geometric Control Theory and Sub-Riemannian Geometry*, volume 5, pages 37–51. Springer, 2014.

[Buc77] Michael A. Buchner. Simplicial structure of the real analytic cut locus. *Proc. Amer. Math. Soc.*, 64(1):118–121, 1977.

[Far08] Michael Farber. *Invitation to topological robotics*. Zurich lectures in advanced mathematics. European Mathematical Society, Zürich, 2008. OCLC: ocn261176490.

[GV17] Andrei Gabrielov and Nicolai Vorobjov. On topological lower bounds for algebraic computation trees. *Found. Comput. Math.*, 17(1):61–72, 2017.

[KR90] Daniel E. Koditschek and Elon Rimon. Robot navigation functions on manifolds with boundary. *Advances in applied mathematics*, 11(4):412–442, 1990.

[RT09] L Rifford and E Trelat. On the stabilization problem for nonholonomic distributions. *Journal of European Mathematical Society*, 11:223–255, 2009.

[SBK01] Uluc Saranli, Martin Buehler, and Daniel E Koditschek. RHex: A simple and highly mobile hexapod robot. *The International Journal of Robotics Research*, 20(7):616–631, 2001. Publisher: SAGE Publications.

[Son12] Eduardo D. Sontag. Stability and feedback stabilization. In *Mathematics of complexity and dynamical systems. Vols. 1–3*, pages 1639–1652. Springer, New York, 2012.

[Whi35] J. H. C. Whitehead. On the covering of a complete space by the geodesics through a point. *Ann. of Math.* (2), 36(3):679–704, 1935.

[Zam98] Tudor Zamfirescu. Extreme points of the distance function on convex surfaces. *Trans. Amer. Math. Soc.*, 350(4):1395–1406, 1998.