Algorithmic aspects of homotopy theory and embeddability

by

Stephan Zhechev

August, 2019

A thesis presented to the
Graduate School
of the
Institute of Science and Technology Austria, Klosterneuburg, Austria
in partial fulfillment of the requirements
for the degree of
Doctor of Philosophy

IST AUSTRIA

Institute of Science and Technology

The thesis of Stephan Zhechev, titled *Algorithmic aspects of homotopy theory and embeddability*, is approved by:

Supervisor: Uli Wagner, IST Austria, Klosterneuburg, Austria

Signature: ____________________________

Committee Member: Herbert Edelsbrunner, IST Austria, Klosterneuburg, Austria

Signature: ____________________________

Committee Member: Martin Čadek, Masaryk University, Brno, Czech Republic

Signature: ____________________________

Defense Chair: Tamás Hausel, IST Austria, Klosterneuburg, Austria

Signature: ____________________________

signed page is on file

IST Austria Thesis, ISSN: 2663-337X

I hereby declare that this thesis is my own work and that it does not contain other people's work without this being so stated; this thesis does not contain my previous work without this being stated, and the bibliography contains all the literature that I used in writing the dissertation.

I declare that this is a true copy of my thesis, including any final revisions, as approved by my thesis committee, and that this thesis has not been submitted for a higher degree to any other university or institution.

I certify that any republication of materials presented in this thesis has been approved by the relevant publishers and co-authors.

Signature: ______________________

Stephan Zhechev
August, 2019
signed page is on file

Abstract

The first part of the thesis considers the computational aspects of the *homotopy groups* $\pi_d(X)$ of a topological space X. It is well known that there is no algorithm to decide whether the *fundamental group* $\pi_1(X)$ of a given finite simplicial complex X is trivial. On the other hand, there are several algorithms that, given a finite simplicial complex X that is *simply connected* (i.e., with $\pi_1(X)$ trivial), compute the higher homotopy group $\pi_d(X)$ for any given $d \geq 2$.

However, these algorithms come with a caveat: They compute the isomorphism type of $\pi_d(X)$, $d \geq 2$ as an *abstract* finitely generated abelian group given by generators and relations, but they work with very implicit representations of the elements of $\pi_d(X)$. We present an algorithm that, given a simply connected space X, computes $\pi_d(X)$ and represents its elements as simplicial maps from suitable triangulations of the d-sphere S^d to X. For fixed d, the algorithm runs in time exponential in $\mathrm{size}(X)$, the number of simplices of X. Moreover, we prove that this is optimal: For every fixed $d \geq 2$, we construct a family of simply connected spaces X such that for any simplicial map representing a generator of $\pi_d(X)$, the size of the triangulation of S^d on which the map is defined, is exponential in $\mathrm{size}(X)$.

In the second part of the thesis, we prove that the following question is algorithmically undecidable for $d < \lfloor 3(k+1)/2 \rfloor$, $k \geq 5$ and $(k, d) \neq (5, 7)$, which covers essentially everything outside the meta-stable range: Given a finite simplicial complex K of dimension k, decide whether there exists a piecewise-linear (i.e., linear on an arbitrarily fine subdivision of K) embedding $f\colon K \hookrightarrow \mathbb{R}^d$ of K into a d-dimensional Euclidean space.

Acknowledgments

I would like to express my gratitude to my advisor Uli Wagner. I am thankful for his support, patience and all the mathematical insight he has taught me.

I would like to thank Peter Franek for his support throughout my PhD and fruitful collaboration.

I would like to thank Marek Filakovský for patiently sharing his knowledge and fruitful collaboration.

I would like to thank Kristóf Huszár, Sergey Avvakumov, Pascal Wild and Fedya Manin for the countless helpful conversations.

I would like to thank Chethan Kamath for his personal support and warm friendship.

I would like to express my gratitude to my research group for creating a very pleasant atmosphere and my family and friends.

The present thesis is dedicated to prof. Vasil Valdemarov Tsanov, whose devotion and enlightening influence have shone light through a very cloudy sky.

About the Author

After completing a BSc and MSc in the Sofia University, Stephan Zhechev joined IST Austria in 2014 in the research group of Uli Wagner. His main research interests include computational aspects of algebraic topology and homotopy theory, constructive homological algebra and their applications to topics such as the study of embeddability of simplicial complexes in Euclidean spaces.

List of Publications

1. Filakovský, M., Franek, P., Wagner, U., Zhechev, S. Computing simplicial representatives of homotopy group elements. J. Appl. Comput. Topol. 2 (2018), no. 3-4, 177-231.

   An extended abstract appeared in the Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, 1135-1151, SIAM, Philadelphia, PA, 2018.

2. Filakovský, P., Wagner, U., Zhechev, S. Embeddability of simplicial complexes is undecidable. Preprint.

Table of Contents

1 Introduction

Given a finite graph G, a natural question one could ask is whether G is planar, i.e. whether it embeds in the plane $\mathbb{R}^2$. A criterion for planarity was provided by Kuratowski [44], which asserts that G is planar if and only if it does not contain K_5, the complete graph on five vertices or $K_{3,3}$, the complete bipartite graph as *topological minors*.[1] In other words, G is planar if and only if it does not contain a subdivided copy of K_5 or $K_{3,3}$. It is important to notice that Kuratowski's theorem considers continuous embeddings. However, Fáry [21] proved that if the graph G embeds continuously in the plane, it also embeds linearly, i.e. all the edges of G are drawn as straight line segments. Considering the computational aspect of the question, there exist many algorithms for planarity testing. One particular algorithm, devised by Robertson and Seymour, utilises a generalisation of Kuratowski's criterion due to Wagner [76] and runs in polynomial time on the size of G (for details, we refer to [17]). In a celebrated result, Hopcroft and Tarjan [35] devised a linear-time algorithm for testing graph planarity.[2] Finally, in the case when G is planar, under mild conditions, Tutte's embedding theorem [73] provides an algorithm that produces an embedding of G in $\mathbb{R}^2$.

The well rounded picture for graphs changes completely when we replace $(G, \mathbb{R}^2)$ by $(K, \mathbb{R}^d)$, where K is a finite simplicial complex of dimension $k \geq 2$. The first thing to notice is that Fáry's theorem is not valid any longer, e.g. for $k \geq 2$ there are contractible and even collapsible complexes, which do not embed *linearly* into $\mathbb{R}^{2k}$ (see [1]), but on the other hand, every contractible k-complex admits a *piecewise-linear embedding*[3] into $\mathbb{R}^{2k}$ (see [78]).

The difference between topological and PL embeddability is more subtle. Bryant [9] proved that they coincide in codimension three, namely when $d - k \geq 3$. The same is true also for $(k, d) = (2, 3)$, which follows from a combination of a result by Bing [7] and the classical *Hauptvermutung* for 2-dimensional polyhedra,[4] proved by Papakyriakopoulos [56]. Since we are mostly interested in the computational aspect, we are going to restrict our attention to PL embeddings. For the sake of simplicity of the exposition, we introduce the following notation.

---

[1] A graph H is called a topological minor of a graph G if G contains a subgraph, which is isomorphic to a subdivision of H.

[2] In fact, this algorithm was devised earlier than the algorithm by Robertson and Seymour. There are several other earlier and less efficient algorithms.

[3] A map $f\colon |K| \to |L|$ between the polyhedra of simplicial complexes is called *picewise-linear (or PL)*, if there exist subdivisions K' and L', so that the induced map $f\colon |K'| \to |L'|$ is linear on every simplex.

[4] The Hauptvermutung in dimension 2 asserts that any two 2-dimensional polyhedra, which are homeomorphic, are also PL homeomorphic.

Definition 1.1. *For a fixed pair of integers* $1 \leq k \leq d$, $\mathrm{EMBED}_{k\to d}$ *is the following algorithmic decision problem: Given a finite* k*-dimensional simplicial complex* K*, does there exists a piecewise-linear embedding* $f\colon K \hookrightarrow \mathbb{R}^d$*.*

The question $\mathrm{EMBED}_{k\to d}$ strongly depends on k and d. For instance, when $d \geq 2k+1$, K always embeds into $\mathbb{R}^d$ by general position and when $d < k$, K never embeds for dimensional reasons. Thus, the interesting range of dimensions is $k \leq d \leq 2k$, which is divided into two main subranges.

1.1 Deleted products and the meta-stable range

Let K be a simplicial complex, or more generally a topological space. We define the *deleted product* of K as $K^2_\Delta := (K\times K)\setminus\{(x,x) : x \in K\}$. Every embedding $f\colon K \hookrightarrow \mathbb{R}^d$ induces a continuous map $\widetilde{f}\colon K^2_\Delta \to S^{d-1}$, given by

$$\widetilde{f}(x,y) := \frac{f(x) - f(y)}{||f(x) - f(y)||}$$

The map $\widetilde{f}$ is well-defined because f is an embedding. Moreover, it has the property $\widetilde{f}(y,x) = -\widetilde{f}(x,y)$, i.e. it is equivariant with respect to the $\mathbb{Z}_2$-action on K^2_Δ, given by $(x,y) \mapsto (y,x)$ and the antipodal action of $\mathbb{Z}_2$ on S^{d-1}. This implies that the existence of such an equivariant map is a necessary condition for embeddability. When K is a finite simplicial complex of dimension k and $\frac{3}{2}(k+1) \leq d$, the famous theorem of Haefliger and Weber [78; 32] implies that this condition is also sufficient.

Theorem 1.2 (Haefliger–Weber)**.** *Let* K *be a finite simplicial complex with* $\dim K = k$*. When* $\frac{3}{2}(k+1) \leq d$*, there exists a PL embedding* $K \hookrightarrow \mathbb{R}^d$ *, if and only if there exists a* $\mathbb{Z}_2$*-equivariant map* $F : K^2_\Delta \to S^{d-1}$*.*

Notice that the question of embeddability, which is geometric in its nature, reduces to the question of existence of a particular type of symmetry preserving continuous functions, which is homotopy theoretic.

Algorithmic aspects of embeddings On the algorithmic side, following Theorem 1.2, in order to decide whether a finite simplicial complex K embeds into $\mathbb{R}^d$ when the pair (k,d) is in the meta-stable range, it suffices to devise an algorithm, which decides whether the set $[K^2_\Delta, S^{d-1}]_{\mathbb{Z}_2}$ of homotopy classes of $\mathbb{Z}_2$-equivariant maps is empty or not. In fact, when (k,d) are in the meta-stable range, this set is either empty or has the structure of a finitely generated abelian group (see [14]) and in a series of papers [13; 11; 14], Čadek et al have devised an algorithm, which computes the isomorphism type of $[K^2_\Delta, S^{d-1}]_{\mathbb{Z}_2}$ (see Theorem 1.4). The algorithm is the product of a build-up of three separate but related results, which we shortly outline.

Computing $[X,Y]$ Given a topological space Y, we can obtain a lot of information about it by studying the sets $[X,Y]$ of homotopy classes of continuous maps from a space X into Y, for different for X. For general spaces X,Y, the set $[X,Y]$ does not have any particular structure, nor it is finite. However, if in addition we ask that Y is d-connected ($d \geq 1$) and $\dim X \leq 2d$, we obtain a group structure on $[X, Y]$ (see [11]). Furthermore, when X and Y are both finite simplicial complexes, then $[X, Y]$ is a finitely generated abelian group, so it's isomorphism type can be presented by a finite set of generators and relations between them. In [11] the authors prove the following—:

Theorem 1.3 (Čadek et al.). *Let X, Y be finite simplicial complexes, such that Y is $(d-1)$-connected and $\dim X \leq 2d-2$ for some integer $d \geq 2$. Then, there exists an algorithm, which computes the isomorphism type of the finitely generated abelian group $[X, Y]$. When d is fixed, the algorithm runs in polynomial time in $\mathrm{size}(X) + \mathrm{size}(Y)$.*

Here, by definition, the size of a finite simplicial complex X, or $\mathrm{size}(X)$, is the number of simplices of X.

Computing $[X, Y]_G$ A natural generalisation of Theorem 1.3 is to introduce a free action of a finite group G on X and Y. The computational side of this question also fits into the context of Theorem 1.2. If we apply the same conditions as in Theorem 1.3 and as long as the action of G on X and Y is free, the set of G-equivariant homotopy classes of maps $[X, Y]_G$ is either empty or has the structure of a finitely generated abelian group (see [14] for details). In [14], the authors devise an algorithm, which computes the isomorphism type of this group.

Theorem 1.4 (Čadek et al.). *Let X, Y be finite simplicial complexes, such that Y is $(d-1)$-connected and $\dim X \leq 2d-2$ for some integer $d \geq 2$. Let G be a finite group, which acts* freely *on X and Y. Then, there exists an algorithm, which computes the isomorphism type $[X, Y]_G$. When d and G are fixed, the algorithm runs in polynomial time in $\mathrm{size}(X) + \mathrm{size}(Y)$.*

Combining this theorem with Theorem 1.2, the authors obtain the following corollary, which ensures that embedadbility is decidable in the meta-stable range.

Corollary 1.5 (Čadek et al.). *Let (k, d) be a pair with $\frac{3}{2}(k+1) \leq d$, then $\mathrm{EMBED}_{k \to d}$ is decidable. If d is fixed, the algorithm runs in polynomial time in $\mathrm{size}(K)$.*

Computing homotopy groups The d-th homotopy group $\pi_d(X)$ of a pointed topological space X is defined as the set of pointed homotopy classes of basepoint preserving continuous maps from S^d into X. Similar to the *homology groups* $H_d(X)$, the homotopy groups $\pi_d(X)$ provide a mathematically precise way of measuring the "d-dimensional holes" in X, but the latter are significantly more subtle and computationally much less tractable than the former. Computing and understanding homotopy groups has been one of the driving forces of algebraic topology in the last century with only partial results so far despite an enormous effort (see, e.g., [60; 42]); the amazing complexity of the problem is illustrated by the fact that even for the 2-dimensional sphere S^2, the higher homotopy groups $\pi_d(S^2)$ are nontrivial for infinitely many d and *known* only for a few dozen values of d.

The first algorithm that computes the homotopy groups of simply connected finite simplicial complexes, was given by Brown [8] in the late 50's. The algorithm is very inefficient and relies on exhaustive searches, but it has focused attention on the algorithmic point of view on homotopy groups.

The condition of simple connectivity is essential, since triviality of the fundamental group $\pi_1(Y)$, when Y is a finite simplicial complex, is undecidable. This follows via a standard reduction from a result of Adjan[2] and Rabin [59] on the algorithmic unsolvability of the triviality problem of a group given in terms of generators and relations. The undecidability result is true even if we restrict Y to be 2-dimensional.

Several more refined algorithms computing homotopy groups of simply connected finite simplicial complexes have been obtained as a part of general computational frameworks in algebraic topology; in particular, an algorithm based on the methods of Sergeraert et al. [71; 65] was described by Real [61]. More recently, Čadek et al. [13] proved that, for any fixed d, the homotopy group $\pi_d(X)$ of a given 1-connected finite simplicial complex can be computed in polynomial time on the number of simplices of X.

Theorem 1.6 (Čadek et al.)**.** *Let Y be a simply connected finite simplicial complex and $d \geq 2$ be an integer. Then, there exists an algorithm, which computes the isomorphism type of $\pi_d(Y)$. When d is fixed, the algorithm runs in polynomial time in $\mathrm{size}(Y)$.*

On the negative side, computing $\pi_d(X)$ is #P-hard if d is part of the input [4; 12] (and, moreover, W[1]-hard with respect to the parameter d [51]), even if X is restricted to be 4-dimensional.

Constructing explicit maps Having algorithms that compute $\pi_d(X)$, $[X, Y]$ and $[X, Y]_G$, a natural question is to devise algorithms, which represent their elements as simplicial maps. For instance, in the case of $\pi_d(X)$, that would mean an algorithm, which represents a set of generators $g_1, \ldots, g_k$ of $\pi_d(X)$ as simplicial maps $\Sigma_j^d \to X$ from suitable triangulations Σ_j^d of the d-sphere S^d. Further motivation to study those problems is provided by the algorithmic study of embeddability, where the goal would be, given a complex K, which we know embeds into $\mathbb{R}^d$, to algorithmically construct an embedding $\widetilde{K} \hookrightarrow \mathbb{R}^d$ from a suitable subdivision of $\widetilde{K}$ of K. Unlike the case for graphs, little is known about this question in higher dimensions.

Computing representatives for homotopy group elements Let X be a simply connected finite simplicial complex. Theorem 1.6 provides an algorithm, which computes the isomorphism type of $\pi_d(X)$ for a fixed $d \geq 2$. The output of this algorithm is a string of the form $(0, \ldots, 0, c_1, \ldots, c_p)$, where $c_i \in \mathbb{N}, 1 \leq i \leq p$ and $\pi_d(X) \cong \mathbb{Z} \oplus \ldots \oplus \mathbb{Z} \oplus \mathbb{Z}_{c_1} \oplus \ldots \oplus \mathbb{Z}_{c_p}$, where the number of copies of the integers is equal to the number of zeros in the string. In fact, the algorithm would also output a set of generators $\alpha_1, \ldots, \alpha_l$, but each of them given with some algebraic representation. In Chapter 2 we prove the following theorem.

Theorem 1.7. *There exists an algorithm that, given $d \geq 2$ and a 1-connected finite simplicial complex X, provided with a certificate for simple connectivity as described in Section 2.2.2, computes a set of generators $g_1, \ldots, g_k$ of $\pi_d(X)$ as simplicial maps $\Sigma_j^d \to X$. Here Σ_j^d ($j = 1, \ldots, k$) are suitable triangulations of S^d.*

For fixed d, the time complexity is exponential in the size *(number of simplices) of X; more precisely, it is $O(2^{P(\mathrm{size}(X))})$ where $P = P_d$ is a polynomial depending only on d.*

In addition, we also prove that the exponential-time complexity is *optimal*. That means that any other algorithm, which computes a set of generators of $\pi_d(X)$ as simplicial maps from suitably subdivided spheres, must have at least an exponential-time complexity. The details are given in Section 2.4.

1.2 Outside the meta-stable range

When we are outside the meta-stable range, namely when $k \leq d < \frac{3}{2}(k+1)$, the situation becomes more complicated and most of the questions remain unanswered. A first difference is that Theorem 1.2 does not provide a sufficient condition for embeddability anymore and in general, there is no known criterion for embeddability. As a result, except for the cases $\mathrm{EMBED}_{1\to 2}$ and $\mathrm{EMBED}_{2\to 2}$, outside the meta-stable range there are no known algorithms for deciding embeddability. Moreover, in contrast with the polynomial-time algorithm in the meta-stable range, according to [50; 18], $\mathrm{EMBED}_{k\to d}$ is NP-hard outside the stable range. In addition, in [50] it is also proved that $\mathrm{EMBED}_{(d-1)\to d}$ and $\mathrm{EMBED}_{d\to d}$ are undecidable for $d \geq 5$. We visualise these results in Figure 1.1.

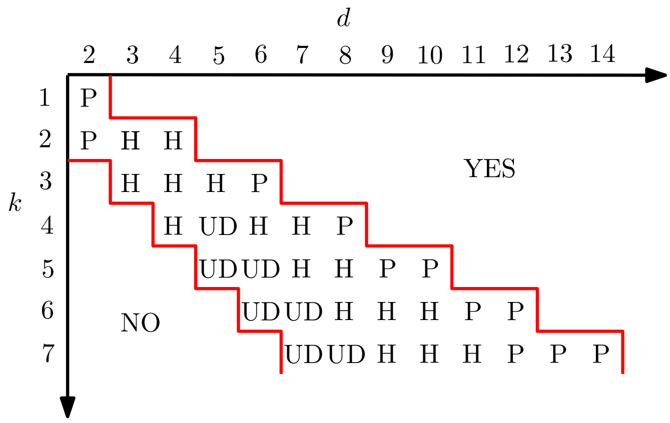

Figure 1.1: P polynomial-time decidable; H NP-hard; UD undecidable.

Immersibility is partially undecidable In related work, Manin and Weinberger [47] have shown that *immersibility*, a question related to embeddability, is undecidable for a large portion of what lies outside the meta-stable range. A PL map $f\colon K \to \mathbb{R}^d$ is said to be a *PL immersion* if it is locally an embedding, i.e. if every point $x \in K$ has a neighbourhood $x \in U \subset K$, such that $f|_U\colon U \to \mathbb{R}^d$ is a PL embedding. Observe that, in general, an immersion is not an embedding, since it allows the images under f of distant parts of K to intersect. Moreover, if we assume K to be a smooth manifold and replace PL by smooth, we obtain the definition of a *smooth immersion*. Focusing on PL[5] and smooth manifolds, in [47] the authors obtain the following result.

Theorem 1.8 (Manin–Weinberger)**.** *Let (k,d) be positive integers.*

- *If $\frac{4}{5}d \leq k \leq d-3$, smooth immersibility is undecidable when $(d-k)$ is even.*

[5] PL manifolds are a class of topological spaces, which admit particularly nice triangulations as simplicial complexes. We postpone the proper definition of a PL manifold untill Section 3.4

- *If $(d-k)=2$ and $d \geq 10$, both smooth immersibility and PL immersibility are undecidable.*

For more details about the specific finite encoding for smooth manifolds, as well as the details of the proof, we refer the reader to [47]. In codimension two, there are two different types of immersions of PL manifolds and the theorem considers the *locally flat* immersions. For the definition and details about why in codimension more than 2 those two types coincide, we refer to Chapter 8 in [79].

As a corollary to the Theorem 1.8, the authors also prove that for a narrow range, smooth *embeddability* is also undecidable.

Theorem 1.9 (Manin–Weinberger)**.** *When $\frac{10}{11}d < k < d-2$ and $d-k$ is even, smooth embeddability of a smooth k-manifold with boundary in $\mathbb{R}^d$ is undecidable.*

The question $\mathbf{EXTEMBED}_{k\to d}$ In Section 3 we consider the following question closely related to embeddability.

Definition 1.10. $\mathrm{EXTEMBED}_{k\to d}$ *is the question, given a k-dimensional simplicial complex K, a subcomplex L and an embedding $f\colon L \hookrightarrow \mathbb{R}^d$, does there exist an embedding $F\colon K \hookrightarrow \mathbb{R}^d$, such that $F|_L = f$.*

Observe that $\mathrm{EMBED}_{k\to d}$ is a special case of $\mathrm{EXTEMBED}_{k\to d}$, namely, when we set $L=\emptyset$.

There are several instances of the problem, which have already been studied. $\mathrm{EXTEMBED}_{1\to 2}$ corresponds to the question whether a planar embedding of a subgraph can be extended to a planar embedding of the entire graph. In [3] it is shown that this can be solved in linear time, which is in accordance with the linear-time decision algorithm for graph planarity [35]. More generally, in higher dimensions, the problem $\mathrm{EXTEMBED}_{k\to d}$ can be solved in polynomial-time for every fixed pair (k,d) of integers in the meta-stable range. This follows from Theorem 1.2 (Haefliger–Weber), together with results of Čadek et al. in [14].

On the other hand, since the problem $\mathrm{EMBED}_{k\to d}$ is a special case of $\mathrm{EXTEMBED}_{k\to d}$, the result by Matoušek et al. [50], which we discussed above, that $\mathrm{EMBED}_{(d-1)\to d}$ and $\mathrm{EMBED}_{d\to d}$ are undecidable for $d \geq 5$, also implies that $\mathrm{EXTEMBED}_{(d-1)\to d}$ and $\mathrm{EXTEMBED}_{d\to d}$ are undecidable for $d \geq 5$. The undecidability of $\mathrm{EXTEMBED}_{(d-1)\to d}$ for $d \geq 5$ follows also from the following result by Nabutovsky and Weinberger [55].

Theorem 1.11 (Theorem 1 in [55])**.** *For any fixed $d>3$ there is no algorithm deciding whether or not a given knot f: $S^d \hookrightarrow \mathbb{R}^{d+2}$ is trivial. Here f is a PL-embedding of the boundary of the standard $(d+1)$-simplex into $\mathbb{R}^{d+2}$.*

To see how this implies the undecidability of $\mathrm{EXTEMBED}_{(d-1)\to d}$, we need to make use of the following theorem.

Theorem 1.12 (Theorem 1 in [80])**.** *A knot $f: S^{d-2} \hookrightarrow S^d$ is a* trivial knot *if and only if f extends to an embedding $F\colon D^{d-1} \hookrightarrow S^d$, such that the following diagram commutes:*

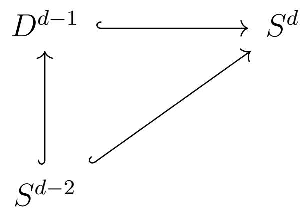

where $i\colon S^{d-2} \hookrightarrow D^{d-1}$ *is the inclusion of* $S^{d-2} = \partial D^{d-1}$.

Combining those two theorems, we obtain the following corollary, which ensures that $\mathrm{EXTEMBED}_{(d-1)\to d}$ is algorithmically undecidable for $d \geq 5$.

Theorem 1.13. *There is no algorithm, which decides whether a given embedding* f : $S^d \hookrightarrow S^{d+2}$ *extends to an embedding* $F\colon D^{d+1} \hookrightarrow S^{d+2}$.

Undecidability of $\mathbf{EMBED}_{k\to d}$ and $\mathbf{EXTEMBED}_{k\to d}$ In Chapter 3 we prove that both $\mathrm{EMBED}_{k\to d}$ and $\mathrm{EXTEMBED}_{k\to d}$ are undecidable for most pairs (k, d) outside the meta-stable range.

Theorem 1.14. $\mathrm{EXTEMBED}_{k\to d}$ *is* undecidable *for* $k \leq d < \left\lfloor \frac{3(k+1)}{2} \right\rfloor$, $k \geq 5$ *and* $(k, d) \neq (5, 7)$.

Theorem 1.15. $\mathrm{EMBED}_{k\to d}$ *is* undecidable *for* $k \leq d < \left\lfloor \frac{3(k+1)}{2} \right\rfloor$, $k \geq 5$ *and* $(k, d) \neq (5, 7)$.

Since $\mathrm{EMBED}_{k\to d}$ is a special case of $\mathrm{EXTEMBED}_{k\to d}$, Theorem 1.15 trivially implies Theorem 1.14. However, we state them separately because in our approach, we first prove the former and then use it as the base for the proof of the latter.

Following the discussion above, we only prove the theorems for the case $d - k \geq 2$, since the codimension 0 and codimension 1 cases have already been proved. Our methods are insufficient when considering some sporadic pairs (k, d), which remain open. We postpone the technical discussion of this issue until Chapter 3. We illustrate the results of Theorem 1.14 and Theorem 1.15 on Figure 1.2.

Remark 1.15.1. In essence, the undecidability in Theorem 1.14 follows from the undecidability of the halting problem, but the reduction from the halting problem takes different paths in different cases. More precisely, the undecidability of $\mathrm{EXTEMBED}_{(d-1)\to d}$ and $\mathrm{EXTEMBED}_{d\to d}$ ($d \geq 5$), as proven in [38], follows from the celebrated result of Novikov [75] on the algorithmic unsolvability of recognizing the 5-sphere, which is related to the algorithmic unsolvability of the word problem. In a similar manner, the alternative proof of the undecidability of $\mathrm{EXTEMBED}_{(d-1)\to d}$ ($d \geq 5$) follows from the algorithmic unsolvability of the word problem (see [55]). On the other hand, as will be shown in Chapter 3, the undecidability of $\mathrm{EXTEMBED}_{k\to d}$ ($k \leq d-2$, $d < \left\lfloor \frac{3(k+1)}{2} \right\rfloor$, $(k, d) \neq (5, 7)$) follows from Matiyasevich's result implying the undecidability of Hilbert's tenth problem [48] (algorithmically deciding the solvability of Diophantine equations). [6]

[6] We briefly discuss Hilbert's tenth problem in Chapter 3.

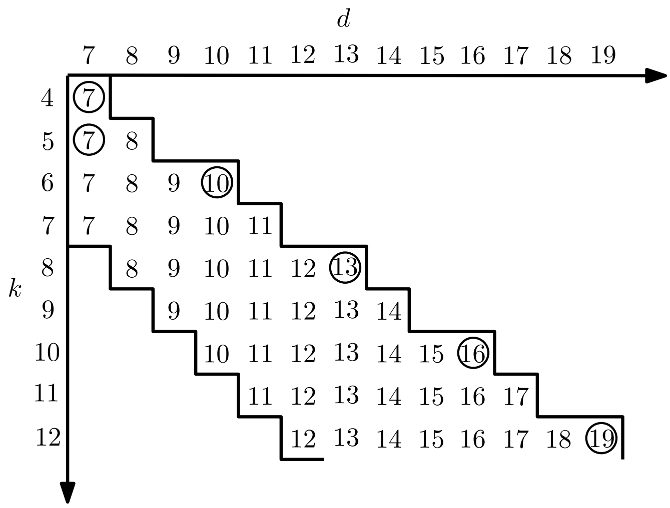

Figure 1.2: The question is still open for the encircled pairs.

1.3 Related work and open problems

The results in the present thesis fall into the broader area of *computational topology*, a field lying within the intersection of topology and computer science. A primary goal of computational topology is to study the algorithmic properties of various invariants, such as homology groups and homotopy groups, of topological spaces admitting some finite presentation. Typical problems often are devising efficient algorithms for computing invariants or studying the computational complexity of the problem of computing invariants. For reference to different flavours of computational topology, we refer the reader to [19; 81; 52].

Computational homotopy theory and applications.

The problem, considered in Chapter 2 forms part of a general effort to understand the computational complexity of problems in *homotopy theory*, both because of the intrinsic importance of these problems in topology and because of applications in other areas, such as the algorithmic study of embeddability of simplicial complexes. This is the central topic of the present thesis and its connection to computational homotopy theory has been pointed out above (for further reading we refer to [50; 27]). Other interesting and fruitful applications of computational homotopy theory are to questions in topological combinatorics (see, e.g., [46]), to the robust satisfiability of equations [26], or to quantitative questions in homotopy theory [31].

Homotopy-theoretic questions have been at the heart of the development of algebraic topology since the 1940's. In the 1990s, three independent groups of researchers proposed general frameworks to make various more advanced methods of algebraic topology (such as spectral sequences) *effective* (algorithmic): Schön [69], Smith [72], and Sergeraert, Rubio, Dousson, Romero, and coworkers (e.g., [71; 65; 62; 66]; also see [67] for an exposition). These frameworks yielded general *computability* results for homotopy-theoretic questions (including new algorithms for the computation of higher homotopy groups [61]), and in the case of Sergeraert et al., also a *practical implementation* in form of the Kenzo software package [34].

Building on the framework of *objects with effective homology* by Sergeraert et al., in recent years a variety of new results in computational homotopy theory were obtained [43; 12; 74; 24; 63; 64]. Here we can also include the results in [11; 13; 14], which we already discussed above in Theorems 1.3 1.4 1.6. They provide the first *polynomial-time algorithms* for the considered problems, by using a refined framework of *objects with polynomial-time homology* [43; 13] that allows for a computational complexity analysis. For an introduction to this area from a theoretical computer science perspective and an overview of some of these results, see, e.g., [10] and the references therein.

Explicit maps. The above algorithms often work with rather *implicit* representations of the homotopy classes in $\pi_d(X)$ (or, more generally, in $[X, Y]$) but do not yield explicit maps representing these homotopy classes.

For instance, the algorithm in [61] computes $\pi_d(X)$ as the *homology group* $H_d(F)$ of an auxiliary space $F = F_d(X)$ constructed from X in such a way that $\pi_d(X)$ and $H_d(F)$ are isomorphic as groups.[7]

More recently, Romero and Sergeraert [64] devised an algorithm that, given a 1-reduced (and hence simply connected) simplicial set[8] X and $d \geq 2$, computes the homotopy group $\pi_d(X)$ as the homotopy group $\pi_d(K)$ of an auxiliary simplicial set K (a so-called *Kan completion* of X) with $\pi_d(X) \cong \pi_d(K)$. Moreover, given an element of this group, the algorithm can compute an explicit simplicial map $\Sigma^d \to K$ from a suitable triangulation of S^d to K representing the given homotopy class. In this way, homotopy classes are represented by explicit maps, but as maps to the auxiliary space K, which is homotopy equivalent to but not homeomorphic to the given space X.

By contrast, our general goal in Chapter 2 is to is represent the elements of $\pi_d(X)$ by maps into the given space.

Quantitative homotopy theory. Another motivation for the result in Chapter 2, namely representing homotopy classes by simplicial maps and complexity bounds for such algorithms, is the connection to *quantitative questions* in homotopy theory [31; 22] and in the theory of embeddings [27]. Given a suitable measure of *complexity* for the maps in question, typical questions are: What is the relation between the complexity of a given null-homotopic map $f : X \to Y$ and the minimum complexity of a nullhomotopy witnessing this? What is the minimum complexity of an embedding of a simplicial complex K into $\mathbb{R}^d$? In quantitative homotopy theory, complexity is often quantified by assuming that the spaces are metric spaces and by considering Lipschitz constants (which are closely related to the sizes of the simplicial representatives of maps and homotopies [22]). For embeddings, the connection is even more direct: a typical measure is the smallest number of simplices in a subdivision $\widetilde{K}$ or K such that there exists a simplexwise linear-embedding $\widetilde{K} \hookrightarrow \mathbb{R}^d$.

Constructing simplicial representatives of elements of $[X, Y]$ As we mentioned above, a natural continuation of the result in Chapter 2 would be a constructive version of Theorem 1.3. Let X and Y be spaces, satisfying the conditions of Theorem 1.3, so the set $[X, Y]$ of homotopy classes of maps has the structure of a finitely generated

[7] Similarly, the algorithm in [13] constructs an auxiliary chain complex C such that $\pi_d(X)$ is isomorphic to the homology group $H_{d+1}(C)$ and computes the latter.

[8] Simplicial sets are topological spaces, which admit a combinatorial structure very similar to simplicial complexes but more flexible. We introduce them formally in Chapter 2.2

abelian group, therefore admitting a finite presentation. We would like to obtain an efficient algorithm, which for an element $\alpha \in [X, Y]$, i.d. a homotopy class, constructs a suitable subdivision $\widetilde{X}$ of X and a simplicial map $f\colon \widetilde{X} \to Y$, which represents α. More precisely, since the geometric realisations $|X|$ and $|\widetilde{X}|$ are homeomorphic, there is an isomorphism $\phi\colon [X, Y] \cong [\widetilde{X}, Y]$ and we require that $\phi(\alpha) = [f]$. In this section we briefly outline our partial progress on this problem.

Along with the isomorphism type of the abelian group $[X, Y]$, for any given element $\alpha \in [X, Y]$, Theorem 1.3 also outputs an implicit algebraic representation. For the sake of presentation, we assume that $\dim X = 2d$, where d is the connectivity of Y, i.e. that X is of the maximal permitted dimension. In fact, the element $\alpha \in [X, Y]$ will be represented as a map into an auxiliary space $P_{2d}(Y)$, called the $(2d)$-th *Postnikov stage* of Y, which approximates Y homotopically up to dimension $2d$. Because of the model of P_{2d} used by the authors, a map $X \to P_{2d}$ can be represented in the form $(0, \ldots, 0, c^{d+1}, \ldots, c^{2d})$, where $c^i \in C^i(X; \pi_i(Y))$. We refer to [11] for further details.

Starting with an element $\alpha \in [X, Y]$ with a representation $(0, \ldots, 0, c^{d+1}, \ldots, c^{2d})$, our strategy for constructing a map $\widetilde{X} \to Y$, representing α is to use this data and define the map by induction on the skeleta of X. We briefly outline the steps without details.

1. Define $\alpha\colon X^d \to Y$ to be the constant map, sending the whole d-skeleton of X to the basepoint $* \in Y$.

2. Every $(d+1)$-simplex $\sigma \in X$ is labelled by an element $c^{d+1}(\sigma) \in \pi_{d+1}(Y)$. Since $\alpha(\partial\sigma) \mapsto * \in Y$, the desired map $\alpha\colon X^{(d+1)} \to Y$ can be factored through a map $\widetilde{\alpha}\colon \bigvee S^{d+1} \to Y$, where we have one copy of S^{d+1} for each $(d+1)$-simplex of X. Moreover, if S^{d+1}_σ is the sphere corresponding to σ, then $[\widetilde{\alpha}|_{S^{(d+1)}_\sigma}] = c^{d+1}(\sigma) \in \pi_{d+1}(Y)$.[9] Next, we use a slightly modified version of Theorem 1.7 to represent the homotopy class $c^{d+1}(\sigma)$ as a map $f_\sigma\colon \widetilde{\sigma} \to Y$ from a subdivision of σ. Performing this for every $(d+1)$-simplex of X yields a map $\widetilde{X^{(d+1)}} \to Y$ from a subdivision of the $(d+1)$-skeleton of X.

3. Let $\tau \in X^{(d+2)}$ be a $(d+2)$-simplex. In the previous step we constructed a map $\widetilde{\partial\tau} \to Y$, from a subdivision of the boundary of τ. By construction, we know that this map is nullhomotopic, so we construct a map $\widetilde{\tau} \to Y$ from a subdibision of τ, extending the map from the boundary. Performing this procedure for every $(d+2)$-simplex of X produces a map $\widetilde{X^{(d+2)}} \to Y$.

4. The extension in the previous step is arbitrary and in general, we would not obtain a map from the correct homotopy class $\alpha \in [X, Y]$. In this step, we use the data given by $c^{d+2}(\sigma) \in \pi_{d+2}(Y)$ in order to correct the map obtained in the previous step.

5. We repeat Step 3 and Step 4 inductively for $d+2, d+3, \ldots, 2d$.

The first step is trivial and Step 2 is just an iterated application of the results in Chapter 2. Step 4 can be made precise and can be achieved using mostly the techniques developed in [11]. The bottleneck of the approach is Step 3, i.e. constructing

[9] This follows from the construction in [11], but it is also intuitively clear.

explicit nullhomotopies for nullhomotopic maps from a sphere, which is related to the following question in quantitative topology, posed by Gromov (see [31]):

Question 1.15.1. Let X, Y be metric spaces and $f\colon X \to Y$ be an L-Lipschitz function, which is nullhomotopic. What is the minimal Lipschitz constant for a nullhomotopy $F\colon X \times [0,1] \to Y$ for f. More generally, given L-Lipschitz functions $f, g\colon X \to Y$, what is the minimal Lipschitz constant for a homotopy $H\colon X \times [0,1] \to Y$ between f and g.

For details about the relation between the two problems and results for particular classes of metric spaces we refer the reader to [22; 15].

While we were unable to find a general solution for the problem of constructing explicit nullhomotopies, by making use of particular combinatorial properties of the spheres produced by Theorem 2.1, we managed to solve it for a sufficiently large class of triangulations of spheres. We obtained the following theorem.

Theorem 1.16. *Let $d \geq 2$ be a fixed integer. Let X and Y be finite simplicial complexes such that Y is d-connected and $\dim X \leq 2d$. Assume also that Y has a certificate for simple connectivity as described in Section 2.2.2. Then there exists an algorithm that, computes the generators $g_1, \ldots, g_k$ of $[X, Y]$ as simplicial maps $\widetilde{X}_j \to Y$, for suitable triangulations $\widetilde{X}_j$ of X, $j = 1, \ldots, k$.*

The time complexity of the algorithm provided by this theorem is a tower of exponentials of hight at least d. To a large extent this is due to the way we are solving the problem of constructing explicit nullhomotopies. This complexity is also in contrast with the singly exponential time complexity of the algorithm in Theorem 2.1, as well as the general belief that the problem should be solvable in singly exponential time. We were also not able to produce any convincing evidence against this expectation. That is the reason why we do not consider our solution to be satisfactory and continue working on the problem.

Constructing explicit embeddings In a subsequent step, we hope to generalize this further to the *equivariant* setting $[X, Y]_G$ of [14], in which a finite group G of symmetries acts freely on the spaces X, Y and all maps and homotopies are required to be *equivariant*, i.e., to preserve the symmetries. That would be a crucial step in a strategy to construct explicit embeddings in the meta-stable range. Given a finite k-dimensional simplicial complex, for $d \geq \frac{3(k+1)}{2}$, Theorem 1.2 (Haefliger–Weber) provides a necessary and sufficient condition for the existence of an embedding $K \hookrightarrow \mathbb{R}^d$. Namely, an embedding exists if and only if there exists a $\mathbb{Z}_2$-equivariant map $F\colon K^2_\Delta \to S^{d-1}$. The proof of the theorem is, in principle, constructive, but in order to turn it into an algorithm that computes an embedding, it would be necessary to have a map F given explicitly. A constructive version of Theorem 1.4, similar to Theorem 1.7, would compute $F\colon \widetilde{K^2_\Delta} \to S^{d-1}$ as a simplicial map from a suitable subdivision of K^2_Δ. The next step would be, starting with map $f\colon K \to \mathbb{R}^d$, which is not an embedding, to use the arguments in the proof of Theorem 1.2 and the map F, to resolve the self-intersections of f, thus constructing an embedding $g\colon \widetilde{K} \to \mathbb{R}^d$ from a suitable subdivision $\widetilde{K}$ of K.

Such an algorithm is still a far reaching goal, which poses hard questions to be answered. While it provides a strong motivation for obtaining constructive versions of Theorem 1.6, 1.3 and 1.4, those generalisations are interesting in their own right.

2 Constructing simplicial representatives of homotopy group elements

This chapter is a joint work with Marek Filakovský, Peter Franek and Uli Wagner and has appeared as [25].

In this chapter, we provide a complete proof of Theorem 1.7. For computational purposes, we consider spaces that have a combinatorial description as *simplicial sets* and maps between them as *simplicial maps*. These are very similar to simplicial complexes and in particular have a combinatorial description, but are much more flexible. Every ordered simplicial complex naturally gives rise to a unique simplicial set. We introduce simplicial sets formally in Section 2.2. We assume such an object to be encoded as a list of its nondegenerate simplices and boundary operators given via finite tables. Similar to the definition for a simplicial complex, we define the *size* of a finite simplicial set X to be the number of its non-degenerate simplices.

Since our input space X is *simply connected*, i.e., that it is connected and has trivial fundamental group $\pi_1(X)$, we will further assume that X is given as a 1-reduced simplicial set. That means that it has only one vertex and no edges, which the additional flexibility of simplicial sets allows. This will significantly simplify the presentation of the proofs. In Section 2.3 we outline the additional steps needed when the input is a simplicial complex, as in the statement of Theorem 1.7.

Theorem 2.1. *There exists an algorithm that, given $d \geq 2$ and a finite 1-reduced simplicial set X, computes the generators $g_1, \ldots, g_k$ of $\pi_d(X)$ as simplicial maps $\Sigma_j^d \to X$, for suitable triangulations Σ_j^d of S^d, $j = 1, \ldots, k$.*

For fixed d, the time complexity is exponential in the size *(number of simplices) of X; more precisely, it is $O(2^{P(\mathrm{size}(X))})$ where $P = P_d$ is a polynomial depending only on d.*

Any element of $\pi_d(X)$ can be expressed as a sum of generators, and expressing the sum of two explicit maps from spheres into X as another explicit map is a simple operation. Hence, the algorithm in Theorem 2.1 can convert *any* element of $\pi_d(X)$ into an explicit simplicial map.

Theorem 2.1 also has the following *quantitative* consequence: Fix some standard triangulation Σ of the sphere S^d, e.g., as the boundary of a $(d+1)$-simplex. By the classical *Simplicial Approximation Theorem* [33, 2.C], for any continuous map $f\colon S^d \to X$, there is a subdivision Σ' of Σ and a simplicial map $f'\colon \Sigma' \to X$ that is homotopic to f. Theorem 2.1 implies that if f represents a generator of $\pi_d(X)$, then the size of Σ' can be bounded by an exponential function of the number of simplices of X.

Furthermore, we can show that the exponential dependence on the number of simplices in X is inevitable:

Theorem 2.2. *Let $d \geq 2$ be fixed. Then there is an infinite family of d-dimensional 0-reduced 1-connected simplicial sets X such that for any simplicial map $\Sigma \to X$ representing a generator of $\pi_d(X)$, the triangulation Σ of S^d on which f is defined has size at least $2^{\Omega(\mathrm{size}(X))}$. If $d \geq 3$, we may even assume that X are 1-reduced.*

Consequently, any algorithm for computing simplicial representatives of the generators of $\pi_d(X)$ for 1-reduced simplicial set X has time complexity at least $2^{\Omega(\mathrm{size}(X))}$.

In the boundary case of 1-reduced simplicial sets for $d = 2$, we don't know whether the lower complexity bound is sub-exponential or not. However, we can show that the algorithm from Theorem 2.1 is optimal in that case as well, see a discussion in Section 2.4, page 30.

In Section 2.3 and 2.4, we state and prove generalizations of Theorem 2.1 and 2.2 denoted as Theorem 2.11 and 2.17. They remove the 1-reduceness assumption and replace it by a more flexible certificate of simple connectivity, allowing the input space X to be a more flexible simplicial set or simplicial complex.

Structure of the chapter. In Section 2.1, we give a high-level description of the main ingredients of the algorithm from Theorem 2.1. In Section 2.2, we review a number of necessary technical definitions regarding simplicial sets and the frameworks of effective and polynomial-time homology, in particular Kan's simplicial version of loop spaces and polynomial-time loop contractions for infinite simplicial sets. In Section 2.3, we formally describe the algorithm from Theorem 2.1 and give a high level proof based on a number of lemmas which are proved in subsequent chapters. Section 2.4 contains the proof of Theorem 2.2. The rest of the paper contains several technical parts needed for the proof of Theorem 2.1: in Section 2.5, we describe Berger's effective Hurewicz inverse and analyze its running time (Theorem 2.13), in Section 2.6, we prove that the stages of the Whitehead tower have polynomial-time contractible loops (Lemma 2.14). Finally, in Section 2.7, we show how to reduce the case when the input is a simplicial complex X^{sc}, as presented in the formulation of Theorem 1.7, to the case of an associated simplicial set X and convert a map $\Sigma \to X$ into a map from a subdivision $Sd(\Sigma)$ into X^{sc} (Lemma 2.16).

2.1 Outline of the Algorithm

In this section we present a high-level description of the main steps and ingredients involved in the algorithm from Theorem 2.1.

The algorithm in a nutshell.

1. In the simplest case when the space X is $(d-1)$-connected (i.e., $\pi_i(X) = 0$ for all $i \leq d-1$.), the classical Hurewicz Theorem [33, Sec. 4.2] yields an isomorphism $\pi_d(X) \cong H_d(X)$ between the dth homotopy group and the dth homology group of X. Computing generators of the homology group is known to be a computationally easy task (it amounts to solving a linear system of equations over the integers). The key is then converting the homology generators into the corresponding homotopy generators, i.e., to compute an inverse of the Hurewicz

isomorphism. This was described in the work of Berger [5; 6]. We analyze the complexity of Berger's algorithm in detail and show that it runs in exponential time in the size of X (assuming that the dimension d is fixed).

2. For the general case, we construct an auxiliary simplicial set F_d together with a simplicial map $\psi_d : F_d \to X$ that has the following properties:

   - F_d is a simplicial set that is $d-1$ connected, and
   - $\psi_d \colon F_d \to X$ induces an isomorphism $\psi_{d*} \colon \pi_d(F_d) \to \pi_d(X)$.

   Our construction of F_d is based on computing stages of the Whitehead tower of X [33, p. 356]; this is similar to Real's algorithm, which computes $\pi_d(X)$ as $H_d(F_d)$ as an abstract abelian group.

   The overall strategy is to use Berger's algorithm on the space F_d and compute generators of $\pi_d(F_d)$ as simplicial maps. Then we use the simplicial map ψ_d to convert each generator of $\pi_d(F_d)$ into a map $\Sigma^d \to X$, and these maps generate $\pi_d(X)$. The main technical task for this step is to show that Berger's algorithm can be applied to F_d. For this, we need to construct a polynomial algorithm for explicit contractions of loops in F_d (this space is 1-connected but not 1-reduced in general).

Our contributions. The main ingredients of the algorithm outlined above are the computability of stages of the Whitehead tower [61] as simplicial sets with polynomial-time homology and Berger's algorithmization of the inverse Hurewicz isomorphism [5; 6].

The idea that these two tools can be combined to compute explicit representatives of $\pi_d(X)$ is rather natural and is also mentioned, for the special case of 1-reduced simplicial sets, in [64, p. 3]; however, there are a number of technical challenges to overcome in order to carry out this program (as remarked in [64, p. 3]: "Clemens Berger's algorithm, quite complex, has never been implemented, severely limiting the current scope of this approach, same comment with respect to the theoretical complexity of such an algorithm."). On a technical level, our main contributions are as follows:

- We give a complexity analysis of Berger's algorithm to compute the inverse of the Hurewicz isomorphism (Theorem 2.13).

- We show that the homology generators of the Whitehead stage F_d can be computed in polynomial time (Lemma 2.12).

- Berger's algorithm requires an explicit algorithm for loop contraction—a certificate of 1-connectedness of the space F_d. While F_d is not 1-reduced in general, we describe an explicit algorithm for contracting its loop and show that Berger's algorithm can be applied.

We remark that the Whitehead tower stages are simplicial sets with infinitely many simplices, and we need the machinery of objects with polynomial-time homology to carry out the last two steps.

2.2 Definitions and Preliminaries

In this section, we give the necessary technical definitions that will be used throughout this chapter. In the first part, we recall the standard definitions for simplicial sets and the toolbox of effective homology.

Afterwards, we present Kan's definiton of a loop space and further formalize our definition of (polynomial-time) loop contractions.

2.2.1 Simplicial Sets and Polynomial-Time Effective Homology

Simplicial sets and their computer representation. A simplicial set X is a graded set X indexed by the non-negative integers together with a collection of mappings $d_i\colon X_n \to X_{n-1}$ and $s_i\colon X_n \to X_{n+1}$, $0 \le i \le n$ called the *face* and *degeneracy* operators. They satisfy the following identities:

$$
\begin{aligned}
d_i d_j &= d_{j-1} d_i && \text{for } i < j,\\
d_i s_i &= d_{i+1} s_i = \mathrm{id} && \text{for } 0 \le i < n,\\
d_i s_j &= s_j d_{i-1} && \text{for } i > j+1,\\
d_i s_j &= s_{j-1} d_i && \text{for } i < j,\\
s_i s_j &= s_{j+1} s_i && \text{for } i \le j.
\end{aligned}
$$

More details on simplicial sets and the motivation behind these formulas can be found in [53; 30].

Simplicial maps between simplicial sets are maps of graded sets which commute with the face and degeneracy operators. The elements of X_n are called *n-simplices*. We say that a simplex $x \in X_n$ is *(non-)degenerate* if it can(not) be expressed as $x = s_i y$ for some $y \in X_{n-1}$. If a simplicial set X is also a graded (Abelian) group and face and degeneracy operators are group homomorphisms, we say that X is a simplicial (Abelian) group.

A simplicial set is called k-reduced for $k \ge 0$, if it has a single i-simplex for each $i \le k$.

For a simplicial set X, we define the chain complex $C_*(X)$ to be a free Abelian group enerated by the elements of X_n with differential

$$\partial(c) = \sum_{i=0}^{n} (-1)^i d_i(c).$$

A simplicial set is *locally effective*, if its simplices have a specified finite encoding and algorithms are given that compute the face and degeneracy operators. A simplicial map f between locally effective simplicial sets X and Y is *locally effective*, if an algorithm is given that for the encoding of any given $x \in X$ computes the encoding of $f(x) \in Y$.

We define a simplicial set to be *finite* if it has finitely many non-degenerate simplices. Such simplicial set can be algorithmically represented in the following way. The encoding of non-degenerate simplices can be given via a finite list and the encoding of a degenerate simplex $s_{i_k} \dots s_{i_1} y$ for $i_1 < i_2 < \dots < i_k$ and a non-degenerate y can be assumed to be a pair consisting of the sequence $(i_1, \dots, i_k)$ and the encoding of y. The face operators are fully described by their action on non-degenerate simplices and can be given via finite tables. In this way, any simplicial set with finitely many non-degenerate simplices is naturally locally effective. Any choice of an implementation of the encoding and face operators is called a *representation* of the simplicial set. The *size of a representation* is the overall memory space one needs to store the data which represent the simplicial set.

Geometric realization. To each simplicial set X we assign a topological space $|X|$ called its geometric realization. The construction is similar to that of simplicial complexes. Let Δ_j be the geometric realization of a standard j-simplex for each $j \geq 0$. For each k, we define $D_i : \Delta_{k-1} \hookrightarrow \Delta_k$ to be the inclusion of a $(k-1)$-simplex into the i'th face of a k-simplex and $S_i : \Delta_k \to \Delta_{k-1}$ be the geometric realization of a simplicial map that sends the vertices $(0, 1, \dots, k)$ of Δ_k to the vertices $(0, 1, \dots, i, i, i+1, \dots, k-1)$. The geometric realization $|X|$ is then defined to be a disjoint union of all simplices X factored by the relation $\sim$

$$|X| := (\bigsqcup_{n=0}^{\infty} X_n \times \Delta_n)/\sim$$

where $\sim$ is the equivalence relation generated by the relations $(x, D_i(p)) \sim (d_i(x), p)$ for $x \in X_{n+1}$, $p \in \Delta_n$ and the relations $(x, S_i(p)) \sim (s_i(x), p)$ for $x \in X_{n-1}$, $p \in \Delta_n$.

Similarly, a simplicial map between simplicial complexes naturally induces a continuous map between their geometric realizations.

Simplicial complexes and simplicial sets. In any simplicial complex X^{sc}, we can choose an ordering of vertices and define a simplicial sets X^{ss} that consists of all non-decrasing sequences of points in X^{sc}: the dimension of $(V_0, \dots, V_d)$ equals d. The face operator is d_i omits the i'th coordinate and the degeneracy s_j doubles the j'th coordinate. Moreover, choosing a maximal tree T in the 1-skeleton of X enables us to construct a simplicial set $X := X^{ss}/T$ in which all vertices and edges in the tree, as well as their degeneracies, are considered to be a base-point (or its degeneracies). The geometric realizations of X^{sc} and X are homotopy equivalent and X is 0-reduced, i.e. it has one vertex only.

Homotopy groups. Let (X, x_0) be a pointed topological space. The k-th homotopy group $\pi_k(X, x_0)$ of (X, x_0) is defined as the set of pointed homotopy[1] classes of pointed continuous maps $(S^k, *) \to (X, x_0)$, where $* \in S^k$ is a distinguished point. In particular, the 0-th homotopy group has one element for each path connected component of X. For $k = 1$, $\pi_1(X, x_0)$ is the fundamental group of X, once we endow it with the group operation that concatenates loops starting and ending in x_0. The group operation on $\pi_k(X, x_0)$ for $k > 1$ assigns to $[f], [g]$ the homotopy class of the composition $S^k \xrightarrow{\pi} S^k \vee S^k \xrightarrow{f \vee g} X$ where π factors an equatorial $(k-1)$-sphere containing x_0 into a point. Homotopy groups π_k are commutative for $k > 1$.

If the choice of base-points is understood from the context or unimportant, we will use the shorter notation $\pi_k(X)$. For a simplicial set X, we will use the notation $\pi_k(X)$ for the k'th homotopy group of its geometric realization $|X|$.

[1] A homotopy $F : S^k \times I \to X$ is pointed if $F(*, t) = x_0$ for all $t \in I$.

An important tool for computing homotopy groups is the *Hurewicz theorem*. It says that whenever X is $(d-1)$-connected, then there is an isomorphism $\pi_d(X) \to H_d(X)$. Moreover, if the element of $\pi_d(X)$ is represented by a simplicial map $f : \Sigma^d \to X$ and $\sum_j k_j \sigma_j$ represents a homology generator of $H_d(\Sigma^d)$, then the Hurewicz isomorphism maps $[f]$ to the homology class of the formal sum $\sum_j k_j f(\sigma_j)$ of d-simplices in X.

Effective homology. We call a chain complex C_* *locally effective* if the elements $c \in C_*$ have finite (agreed upon) encoding and there are algorithms computing the addition, zero, inverse and differential for the elements of C_*.

A locally effective chain complex C_* is called *effective* if there is an algorithm that for given $n \in \mathbb{N}$ generates a finite basis $c_\alpha \in C_n$ and an algorithm that for every $c \in C_*$ outputs the unique decomposition of c into a linear combination of c_α's.

Let C_* and D_* be chain complexes. A *reduction* $C_* \Rrightarrow D_*$ is a triple (f, g, h) of maps such that $f : C_* \to D_*$ and $g : D_* \to C_*$ are chain homomorphisms, $h : C_* \to C_*$ has degree 1, $fg = \mathrm{id}$ and $fg - \mathrm{id} = h\partial + \partial h$, and further $hh = hg = fh = 0$.

A locally effective chain complex C_* has *effective homology* (C_* is a *chain complex with effective homology*) if there is a locally effective chain complex $\tilde{C}_*$, reductions $C_* \Lleftarrow \tilde{C}_* \Rrightarrow C_*^{\mathrm{ef}}$ where C_*^{ef} is an effective chain complex, and all the reduction maps are computable.

Eilenberg-MacLane spaces. Let $d \geq 1$ and π be an Abelian group. An Eilenberg-MacLane space $K(\pi, d)$ is a topological space with the properties $\pi_d(K(\pi, d)) \simeq \pi$ and $\pi_j(K(\pi, d)) = 0$ for $0 < j \neq d$. It can be shown that such space $K(\pi, d)$ exists and, under certain natural restrictions, has a unique homotopy type. If π is finitely generated, then $K(\pi, d)$ has a locally effective simplicial model [43].

Globally polynomial-time homology and related notions. In many auxiliary steps of the algorithm, we will construct various spaces and maps. To analyse the overall time complexity, we need to parametrize all these objects by the very initial input, which is in our case an encoding of a finite 1-reduced simplicial set (or, in Theorem 2.11, a more general space endowed with certain explicit certificate of 1-connectedness).

More generally, let $\mathcal{I}$ be a parameter set so that for each $I \in \mathcal{I}$ an integer $\mathrm{size}(I)$ is defined. We say that F is a parametrized simplicial set (group, chain group, …), if for each $I \in \mathcal{I}$, a locally effective simplicial set (group, chain group, …) $F(I)$ is given. The simplicial set F is *locally polynomial-time*, if there exists a locally effective model of $F(I)$ such that for each $k \in \mathbb{N}$ and an encoding of a k-simplex $x \in F(I)$, the encoding of $d_i(x)$ and $s_j(x)$ can be computed in time polynomial in $\mathrm{size}(\mathsf{enc}(x)) + \mathrm{size}(I)$. The polynomial, however, may depend on k. A polynomial-time map between parametrized simplicial sets F and G is an algorithm that for each $k \in \mathbb{N}$, $I \in \mathcal{I}$ and an encoding of an k-simplex x in $F(I)$ computes the encoding of $f(x)$ in time polynomial in $\mathrm{size}(\mathsf{enc}(x)) + \mathrm{size}(I)$: again, the polynomial may depend on k.

Similarly, a locally polynomial-time (parametrized) chain complex is an assignment of a computer representation $C_*(I)$ of a chain complex with a distinguished basis in each gradation, such that all these basis elements have some agreed-upon encoding. A chain $\sum_j k_j \sigma_j$ is assumed to be represented as a list of pairs $(k_j, \mathsf{enc}(\sigma_j))_j$ and has size $\sum_j (\mathrm{size}(k_j) + \mathrm{size}(\mathsf{enc}(\sigma_j)))$, where we assume that the size of an integer k_j is its bit-size. Further, an algorithm is given that computes the differential of a chain

$z \in C_k(I)$ in time polynomial in $\mathrm{size}(z) + \mathrm{size}(I)$, the polynomial depending on k. The notion of a polynomial-time chain map is straight-forward.

A *globally polynomial-time chain complex* is a locally polynomial-time chain complex EC that in addition has all chain groups $EC(I)_k$ finitely generated and an additional algorithm is given that for each k computes the encoding of the generators of $EC(I)_k$ in time polynomial in $\mathrm{size}(I)$. Finally, we define a *simplicial set with globally polynomial-time homology* to be a locally polynomial-time parametrized simplicial set F together with reductions $C_*(F) \Lleftarrow \tilde{C} \Rrightarrow EC$ where $\tilde{C}, EC$ are locally polynomial-time chain complexes, EC is a globally polynomial-time chain complex and the reduction data are all polynomial-time maps, as usual the polynomials depending on the grading k.

The name "polynomial-time homology" is motivated by the following:

Lemma 2.3. *Let F be a parametrized simplicial set with polynomial-time homology and $k \geq 0$ be fixed. Then all generators of $H_k(F(I))$ can be computed in time polynomial in $\mathrm{size}(I)$.*

Proof. For the globally polynomial-time chain complex EF and each fixed j, we can compute the matrix of the differentials $d_j : EF(I)_j \to EF(I)_{j-1}$ with respect to the distinguished bases in time polynomial in $\mathrm{size}(I)$: we just evaluate d_k on each element of the distinguished basis of $EF(I)_k$. Then the homology generators of $H_k(EC)$ can be computed using a Smith normal form algorithm applied to the matrices of d_k and d_{k+1}, as is explained in standard textbooks (such as [54]). Polynomial-time algorithms for the Smith normal form are nontrivial but known [41].

Let $x_1, \ldots, x_m$ be the cycles generating $H_k(EF(I))$. We assume that reductions

$$C_*(F) \overset{(f,g,h)}{\Lleftarrow} \tilde{F} \overset{(f',g',h')}{\Rrightarrow} EF$$

are given and all the reduction maps are polynomial. Thus we can compute the chains

$$fg'(x_1), fg'(x_2), \ldots, fg'(x_m)$$

in polynomial time and it is a matter of elementary computation to verify that they constitute a set of homology generators for $H_k(F(I))$. □

2.2.2 Loop Spaces and Polynomial-Time Loop Contraction

Principal bundles and loop group complexes. In the text we will frequently deal with principal twisted Cartesian products: these are simplicial analogues of principal fiber bundles. The definitions in this section come from Kan's article [40].

We first define the Cartesian product $X \times Y$ of simplicial sets X, Y: The set of n-simplices $(X \times Y)_n$ consists of tuples (x, y), where $x \in X_n, x \in Y_n$. The face and degeneracy operators on $X \times Y$ are given by $d_i(x, y) = (d_i x, d_i y)$, $s_i(x, y) = (s_i x, s_i y)$.

Definition 2.4 (Principal Twisted Cartesian product)**.** *Let B be a simplicial set with a basepoint $b_0 \in B_0$ and G be a simplicial group. We call a graded map (of degree -1) $\tau \colon B_{n+1} \to G_n, n \geq 0$ a* twisting operator *if the following conditions are satisfied:*

- $d_n\tau(\beta) = \tau(d_{n+1}b)^{-1}\tau(d_nb)$
- $d_i\tau(\beta) = \tau(d_ib)$ *for* $0 \le i < n$
- $s_i\tau(b) = \tau(s_ib)$*,* $i < n$*, and*
- $\tau(s_nb) = 1_n$ *for all* $b \in B_n$ *where* 1_n *is the unit element of* G_n*.*

Let B*,* G*,* τ *be as above. We will define a* twisted Cartesian product $B \times_\tau G$ *to be a simplicial set* E *with* $E_n = B_n \times G_n$*, and the face and degeneracy operators are also as in the Cartesian product, i.e.* $d_i(b, g) = (d_ib, d_ig)$ *, with the sole exception of* d_n*, which is given by*

$$d_n(b, g) := (d_nb, \tau(b)d_n(g)), \quad (b, g) \in B_n \times G_n.$$

It is not trivial to see why this should be the right way of representing fiber bundles simplicially, but for us, it is only important that it works, and we will have explicit formulas available for the twisting operator for all the specific applications.

We remark that in the literature one can find multiple definitions of twisted operator and twisted product [53; 40; 5] and that they, in essence differ from each other based on the decision whether the twisting “compresses” the first two or the last two face operators. Here, we follow the same notation as in [5].

Definition 2.5. *Let* X *be a* 0*-reduced simplicial set. Then we define* GX *to be a (non-commutative) simplicial group such that*

- GX_n *has a generator* $\overline{\sigma}$ *for each* $(n+1)$*-simplex* $\sigma \in X$ *and a relation* $\overline{s_ny} = 1$ *for each simplex in the image of the last degeneracy* s_n*.*
- *The face operators are given by* $d_i\overline{\sigma} := \overline{d_i\sigma}$ *for* $i < n$ *and* $d_n\overline{\sigma} := (\overline{d_{n+1}\sigma})^{-1}\overline{d_n\sigma}$
- *The degeneracy operators are* $s_i\overline{\sigma} := \overline{s_i\sigma}$*.*

We use the multiplicative notation, with 1 being the neutral element. It is shown in [40] that GX is a discrete simplicial analog of the loop space of X.

For algorithmic puroposes, we assume that an elements $\prod_j \overline{\sigma}_j^{k_j}$ of GX is represented as a list of pairs (σ_j, k_j) and has size $\sum_j \mathrm{size}(\sigma_j) + \mathrm{size}(k_j)$.

Definition 2.6. *Let* X *be a* 0*-reduced simplicial set. We say that a map* $c_0 : GX_0 \to GX_1$ *is a* contraction *of loops in* X*, if* $d_0c_0(x) = x$ *and* $d_1c_0(x) = 1$ *for each* $x \in GX_0$*.*

In case where X *has finitely many nondegenerate* 1*-simplices, we define the size* $\mathrm{size}(c_0)$ *to be the sum*

$$\sum_{\gamma \in X_1} \mathrm{size}(c_0(\gamma)).$$

Loop contraction for simplicial complexes. Let X^{sc} be a simplicial complex. Let T be a spanning tree in the 1-skeleton of X^{sc} and R a chosen vertex. For each oriented edge $e = (v_1v_2)$ we define a formal inverse to be $e^{-1} := (v_2v_1)$ and we also consider degenerate edges (v, v). A *loop* is defined as a sequence $e_1, \ldots, e_k$ of oriented edges in X^{sc} such that

- The end vertex of e_i equals the initial vertex of e_{i+1}, and
- The initial vertex of e_1 and the end vertex of e_k equal R.

Every edge e that is not contained in T gives rise to a unique loop l_e. Further, every loop in X^{sc} is either a concatenation of such l_e's, or can be derived from such concatenation by inserting and deleting consecutive pairs (e, e^{-1}) and degenerate edges. Before we formally define our combinatorial version of loop contraction, we need the following definition.

Definition 2.7. *Let S be a set, $U \subseteq S$, $F(S)$ and $F(U)$ be free groups generated by S, U, respectively.*[2] *Let $h_U : F(S) \to F(S)$ be a homomorphism that sends each $u \in U$ to 1 and each $s \in S \setminus U$ to itself. We say that an element x of $F(S)$ equals y modulo U, if $h_U(x) = y$.*

An example of an element that is trivial modulo U is the word $s\,u\,s^{-1}$, where $s \in S$ and $u \in U$.

Definition 2.8. *Let S be the set of all oriented edges and oriented degenerate edges in X^{sc} and assume that a spanning tree T is chosen. Let U be the set of all oriented edges in T, including all degenerate edges. A* contraction of an edge α *is a sequence of vertices $A_0, A_1, \dots, A_s$ and $B_1, \dots, B_s$ such that*

- *for each i, $\{A_i, A_{i+1}, B_{i+1}\}$ is a simplex of X^{sc}, and*
- *the element of $F(S)$*

$$(A_0B_1)(B_1A_1)(A_1B_2)(B_2A_2)\dots(B_sA_s)(A_sA_{s-1})(A_{s-1}A_{s-2})\dots(A_1A_0) \qquad (2.1)$$

  equals α modulo U.

  A loop contraction in a simplicial complex *is the choice of a contraction of α for each edge $\alpha \in X^{sc} \setminus T$.*

The size of the contraction of α is defined to be the number of vertices in (2.1) and the size $\mathrm{size}(c)$ of the loop contraction on X^{sc} is the sum of the sizes over all $\alpha \in X^{sc} \setminus T$.

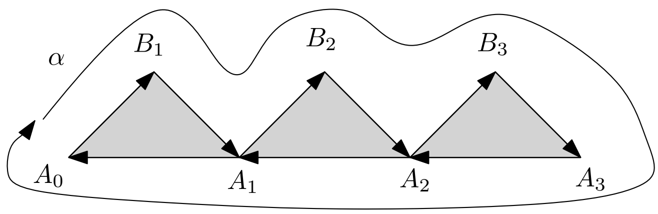

Figure 2.1: The loop ranging over the boundary of this geometric shape equals α, after ignoring edges in the maximal tree and canceling pairs (e, e^{-1}). The interior of the triangles gives rise to a contraction.

[2]Formally, elements of $F(S)$ are sequences of symbols s^ϵ for $\epsilon \in \{1, -1\}$ and $s \in S$ with the relation $s^1 s^{-1} = 1$, where 1 represents the empty sequence. The group operation is concatenation.

The geometry behind this definition is displayed in Figure 2.1. The sequence of A_i's and B_j's gives rise to a map from the sequence of (full) triangles into X^{sc}. The big loop around the boundary is combinatorially described by (2.1). We can continuously contract all of its parts that are in the tree T to a chosen basepoint, as the tree is contractible. Further, we can continuously contract all pairs of edges (e, e^{-1}) and what remains is the original edge α: with all the tree contracted to a point, it will be transformed into a loop that geometrically corresponds to l_α. The interior of the full triangles then constitutes its "filler", hence a certificate of the contractibility of l_α.

A loop contraction in the sense of Definition 2.1 exists if and only if the space X^{sc} is simply connected. One could choose different notions of loop contraction. For instance, we could provide, for each α, a simplicial map from a triangulated 2-disc into X^{sc} such that the oriented boundary of the disc would be mapped exactly to l_α. The description from Definition 2.8 could easily be converted into such map. We chose the current definition because of its canonical and algebraic nature. The connection between Definitions 2.6 and 2.8 is the content of the following lemma.

Lemma 2.9. *Let X^{sc} be a 1-connected simplicial complex with a chosen orientation of all simplices, X^{ss} the induced simplicial set, T a maximal tree in X^{sc}, and $X := X^{ss}/T$ the corresponding 0-reduced simplicial set. Assume that a loop contraction in the simplicial complex X^{sc} is given, such as described in Definition 2.8. Then we can algorithmically compute $c_0(\alpha) \in GX_1$ such that $d_0 c_0(\alpha) = \alpha$ and $d_1 c_0(\alpha) = 1$, for every generator α of GX_0. Moreover, the computation of $c_0(\alpha)$ is linear in the size of X^{sc} and the size of the simplicial complex contraction data.*

Proof. For each i, the triangle $\{A_i, A_{i+1}, B_{i+1}\}$ from Def. 2.8 is in the simplicial complex X^{sc}. There is a unique oriented 2-simplex in X^{ss} of the form (V_0, V_1, V_2) (possibly degenerate) such that $\{V_0, V_1, V_2\} = \{A_i, A_{i+1}, B_{i+1}\}$. Let us denote such oriented simplex by σ_i, and its image in GX_1 by $\overline{\sigma}_i$. We will define an element $g_i \in GX_1$ such that it satisfies

$$d_0 g_i \simeq \overline{(A_i, A_{i+1})} \quad \text{and} \quad d_1 g_i \simeq \overline{(A_i, B_{i+1})}\ \overline{(B_{i+1}, A_{i+1})} \tag{2.2}$$

where $\simeq$ is an equivalence relation that identifies any element $\overline{(U, V)} \in GX_1$ with $\overline{(V, U)}^{-1}$ (note that only one of the symbols (U, V) and (V, U) is well defined in X^{ss}, resp. X.) Explicitly, we can define g_i with these properties as follows:

- If $\sigma = (B_{i+1}, A_i, A_{i+1})$, then $g_i := \overline{\sigma}_i$,
- If $\sigma = (A_i, A_{i+1}, B_{i+1})$, then $g_i := s_0 \overline{(d_2 \sigma)}\, \overline{\sigma}_i\, s_0 d_0 (\overline{\sigma}_i)^{-1}$
- If $\sigma = (A_{i+1}, B_{i+1}, A_i)$, then $g_i = s_0 d_0 \overline{\sigma_i}^{-1}\, \overline{\sigma_i}\, s_0 (\overline{d_1 \sigma_i})^{-1}$
- If $\sigma = (B_{i+1}, A_{i+1}, A_i)$, then $g_i := \overline{\sigma_i}^{-1}$
- If $\sigma = (A_{i+1}, A_i, B_{i+1})$, then $g_i := s_0 d_0 \overline{\sigma_i}\ \overline{\sigma_i}^{-1}\, s_0 (\overline{d_2 \sigma_i})^{-1}$
- If $\sigma = (A_i, B_{i+1}, A_{i+1})$, then $g_i := s_0 (\overline{d_1 \sigma_i})\, \overline{\sigma_i}^{-1} s_0 d_0 \overline{\sigma_i}$.

Let $g := g_0 \ldots, g_s$. The assumption (2.1) together with equation (2.2) immediately implies that $d_1 g (d_0 g)^{-1} = \overline{\alpha}$. Thus we define $c_0(\overline{\alpha}) := s_0 d_1(g)\, g^{-1}$. Algorithmically, to construct g amounts to going over all the triples (A_i, A_{i+1}, B_{i+1}) from a given sequence of $A_i' s$ and B_j's, checking the orientation and computing g_i for every i. □

Polynomial-time loop contraction. Let F be a parametrized simplicial set such that each $F(I)$ is 0-reduced. Using constructions analogous to those defined above, GF is a parametrized locally-polynomial simplicial group whereas we assume a simple encoding of elements of GF_i as follows. If $x = \prod_j \overline{\sigma_j}^{k_j} \in GF(I)_k$ where σ_j are $(k+1)$-simplices in $F(I)$, not in the image of s_k, then we assume that x is stored in the memory as a list of pairs $(k_j, \text{enc}(\sigma_j))$ and has size $\sum_j(\text{size}(k_j) + \text{size}(\sigma_j))$ where some σ_i may be equal to σ_j for $i \neq j$. Face and degeneracy operators are defined in Definition (2.5) and it is easy to see that for any locally polynomial-time simplicial set F, GF is a locally polynomial-time simplicial group.

Definition 2.10. *Let F be a locally polynomial simplicial set. We say that F has* polynomially contractible loops*, if there exists an algorithm that for a 0-simplex $x \in GF(I)$ computes a 1-simplex $c_0(x) \in GF(I)$ such that $d_0 x = x$, $d_1 x = 1 \in GF(I)_0$, and the running-time is polynomial in $\text{size}(x) + \text{size}(I)$.*

2.3 Proof of Theorem 2.1

We will prove a stronger statement of Theorem 2.1 formulated as follows.

Theorem 2.11. *There exists an algorithm that, given $d \geq 2$ and a finite 0-reduced simplicial set X (alternatively, a finite simplicial complex) with an explicit loop contraction c_0 (such as in Definition 2.6 or 2.8) computes the generators $g_1, \dots, g_k$ of $\pi_d(X)$ as simplicial maps $\Sigma_j^d \to X$, for suitable triangulations Σ_j^d of S^d, $j = 1, \dots, k$.*

For fixed d, the time complexity is exponential in the size *of X and the size of the loop contraction c_0; more precisely, it is $O(2^{P(\text{size}(X)+\text{size}(c_0))})$ where $P = P_d$ is a polynomial depending only on d.*

This immediately implies Theorem 2.1, as for a 1-reduced simplicial set, the contraction c_0 is trivial, given by $c_0(1) = 1$.

The proof of Theorem 2.11 is based on a combination of four statements presented here as Lemma 2.12, Theorem 2.13, Lemma 2.14 and Lemma 2.16. Each of them is relatively independent and their proofs are delegated to further sections.

First we present an algorithm that, given a 1-connected finite simplicial set X and a positive integer d, outputs a simplicial set F_d and a simplicial map ψ_d such that

- the simplicial set F_d is $d-1$ connected, it has polynomial-time effective homology and polynomially contractible loops.
- the simplicial map $\psi_d \colon F_d \to X$ is polynomial-time and induces an isomorphism $\psi_{d*} \colon \pi_d(F_d) \to \pi_d(X)$.

Whitehead tower. We construct simplicial sets F_d as stages of a so-called *Whitehead tower* for the simplicial set X. It is a sequence of simplicial sets and maps

$$\cdots \longrightarrow F_d \xrightarrow{f_d} F_{d-1} \xrightarrow{f_{d-1}} \cdots \xrightarrow{f_4} F_3 \xrightarrow{f_3} F_2 = X.$$

where f_i induces an isomorphism $\pi_j(F_{i+1}) \to \pi_j(F_i)$ for $j > i$ and $\pi_j(F_i) = 0$ for $j < i$. We define $\psi_d = f_d f_{d-1} \dots f_3$. One can see that F_d, ψ_d satisfy the desired properties.

Lemma 2.12. *Let $d \geq 2$ be a fixed integer. Then there exists a polynomial-time algorithm that, for a given 1-connected finite simplicial set X, constructs the stages $F_2, \ldots, F_d$ of the Whitehead tower of X.*

The simplicial sets $F_k(X)$, parametrized by 1-connected finite simplicial sets X, have polynomial-time homology and the maps f_k are polynomial-time simplicial maps.

Proof. The proof is by induction. The basic step is trivial as $F_2 = X$. We describe how to obtain F_{k+1}, f_{k+1} assuming that we have computed F_k, $2 \leq k < d$.

1. We compute simplicial map $\varphi_k \colon F_k \to K(\pi_k(X), k) = K(\pi_k(F_k), k)$ that induces an isomorphism $\varphi_{k*} \colon \pi_k(F_k) \to \pi_k(K(\pi_k(X), k)) \cong \pi_k(X)$. This is done using the algorithm in [13], as $K(\pi_k(X), k)$ is the first nontrivial stage of the Postnikov tower for the simplicial set F_k.

   For the simplicial set $K(\pi_k(X), k)$ and for such simplicial sets there is a classical principal bundle (twisted Cartesian product) (see [53]):

$$\begin{array}{c} K(\pi_k(X), k-1) \\ \downarrow \\ E(\pi_k(X), k-1) = K(\pi_k(X), k) \times_\tau K(\pi_k(X), k-1) \\ \Big\downarrow \delta \\ K(\pi_k(X), k) \end{array}$$

2. We construct F_{k+1} and f_{k+1} as a pullback of the twisted Cartesian product:

$$\begin{array}{ccc} K(\pi_k(X), k-1) & \xrightarrow{\cong} & K(\pi_k(X), k-1) \\ \downarrow & & \downarrow \\ F_{k+1} := F_k \times_{\tau'} K(\pi_k(X), k-1) & \dashrightarrow & K(\pi_k(X), k) \times_\tau K(\pi_k(X), k-1) \\ \Big\downarrow f_{k+1} & & \Big\downarrow \delta \\ F_k & \xrightarrow{\varphi_k} & K(\pi_k(X), k). \end{array}$$

It can be shown that the pullback, i.e. simplicial subset of pairs $(x, y) \in F_k \times E(\pi_k(X), k-1)$ such that $\delta(y) = \varphi_k(x)$, can be identified with the twisted product as above [53], where the twisting operator τ' is defined as $\tau\varphi_k$.

To show correctness of the algorithm, we assume inductively, that F_k has polynomial-time effective homology. According to [13, Section 3.8], the simplicial sets $K(\pi_k(X), k-1)$, $E(\pi_k(X), k-1)$, $K(\pi_k(X), k)$ have polynomial-time effective homology and maps φ_k, δ are polynomial-time. Further, they are all obtained by an algorithm that runs in polynomial time.

As F_{k+1} is constructed as a twisted product of F_k with $K(\pi_k(X), k)$, Corollary 3.18 of [13] implies that F_{k+1} has polynomial-time effective homology and f_{k+1} is a polynomial-time map.[3]

[3] We remark that the paper [13] uses a different formalization of twised cartesian product than the one employed by us. However, the paper [23], on which the Corollary 3.18 of [13] is based, can be reformulated in context of the definition used here. We do not provide full details, only remark that one has to make a choice of *Eilenberg-Zilber reduction data* that corresponds to the definition of twisted cartesian product.

The sequence of simplicial sets $F_{k+1} \xrightarrow{f_{k+1}} F_k \xrightarrow{\varphi_k} K(\pi_k(X), k)$ induces the long exact sequence of homotopy groups

$$\cdots \longrightarrow \pi_i(F_{k+1}) \xrightarrow{f_{k+1*}} \pi_i(F_k) \xrightarrow{\varphi_{k*}} \pi_i(K(\pi_k(X), k)) \longrightarrow \pi_{i-1}(F_{k+1}) \longrightarrow \cdots$$

The reason why this is the case follows from a rather technical argument that identifies the simplicial set F_{k+1} with a so called *homotopy fiber* of the map $\varphi_k \colon F_k \to K(\pi_k(X), k)$. In more detail, the category of simplicial sets is right proper [30, II.8.6–7] and map δ is a so-called Kan fibration [53, § 23]. This makes the pullback F_{k+1} coincide with so-called homotopy pullback. Further, the simplicial set $E(\pi_k(X), k-1)$ is contractible, hence the homotopy pullback is a homotopy fiber. The induced exact sequence is due to [58, chapter I.3].

The inductive assumption, together with the fact that φ_k induces an isomorphism $\varphi_{k*} \colon \pi_k(F_k) \to \pi_k(K(\pi_k(X), k))$ imply that f_k induces an isomorphism $\pi_j(F_{k+1}) \to \pi_j(F_k)$ for $j > k$ and $\pi_j(F_{k+1}) = 0$ for $j \leq k$. □

The lemma implies that the simplicial sets F_k have polynomial-time effective homology and maps $\psi_k = f_k f_{k-1} \dots f_3$ are polynomial-time as they are defined as a composition of polynomial-time maps f_i.

The following theorem is a key ingredient of our algorithm.

Theorem 2.13 (Effective Hurewicz Inverse)**.** *Let $d > 1$ be fixed and F be an $(d-1)$-connected 0-reduced simplicial set parametrized by a set $\mathcal{I}$, with polynomial-time homology and polynomially contractible loops.*

Then there exists an algorithm that, for a given d-cycle $z \in Z_d(F(I))$, outputs a simplicial model Σ^d of the d-sphere and a simplicial map $\Sigma^d \to F(I)$ whose homotopy class is the Hurewicz inverse of $[z] \in H_d(F(I))$.

Moreover, the time complexity is bounded by an exponential of a polynomial function in $\mathrm{size}(I) + \mathrm{size}(z)$.

The construction of an effective Hurewicz inverse is the main result of [5] and further details are provided in Section 2.5. It exploits a combinatorial version of Hurewicz theorem given by Kan in [39] where $\pi_d(F)$ is described in terms of $\pi_{d-1}(\widetilde{GF})$ where $\widetilde{GF}$ is a non-commutative simplicial group that models the loop space of F. Kan showed that the Hurewicz isomorphism can be identified with a map $H_{d-1}(\widetilde{GF}) \to H_{d-1}(\widetilde{AF})$ induced by Abelianization. Berger then describes the inverse of the Hurewicz homomorphism as a composition of the maps $1, 2, 3$ in the diagram

$$\begin{array}{ccc} \pi_d(F) & \xleftarrow{\;h^{-1}\;} & H_d(F) \\ {\scriptstyle 3}\uparrow & & \downarrow{\scriptstyle 1} \\ H_{d-1}(\widetilde{GF}) & \xleftarrow[2]{} & H_{d-1}(\widetilde{AF}). \end{array}$$

Arrow 1 is induced by a chain homotopy equivalence and arrow 3 by Berger's explicit geometric model of the loop space. To algorithmize arrow 2, we need an algebraic machinery that includes an explicit contraction of k-loops in $\widetilde{GF}$ for all $k < d-1$. Those are based partially on linear computations in the Abelian group $\widetilde{AF}$ and partially on explicit inductive formulas dealing with commutators. The lowest-dimensional contraction operation, however, cannot be algorithmized, without some external input. The possibility of providing it is is the content of the following claim:

Lemma 2.14. *Let $d \geq 2$ be a fixed integer and $\mathcal{I}$ be the set of all 1-connected 0-reduced finite simplicial sets with an explicit loop contraction c_0. Then the simplicial set F_d from Lemma 2.12, parametrized by $\mathcal{I}$, has polynomial-time contractible loops.*

The proof is constructive, based on explicit formulas in our model of F_d: the details are in Section 2.6.

The core of the algorithm we will describe works with simplicial sets and simplicial maps between them. If our input is a simplicial complex, we need tools to convert them into maps between simplicial complexes. The next two lemmas address this.

Lemma 2.15. *Let Y be a finite simplicial set. Then there exists a polynomial-time algorithm that computes a simplicial complex Y^{sc} with a given orientation of each simplex, and a map $\gamma : Y^{sc} \to Y$ (still understood to be a map between simplicial sets) such that the geometric realization of γ is homotopic to a homeomorphism.*

This construction is given in [12, Appendix B].[4] Explicitly, the simplicial complex Y^{sc} is defined to be $Y^{sc} := B_*(Sd(Y))$, where Sd is the barycentric subdivision functor and B_* a functor introduced in [37]: Y^{sc} can be constructed recursively by adding a vertex v_σ for each nondegenerate simplex $\sigma \in Sd(Y)$ and replacing σ by the cone with apex v_σ over $B_*(\partial\sigma)$. The subdivision $Sd(Y)$ is a regular simplicial set and $B_*(Sd(Y))$ coincides with the flag simplicial complex of the poset of nondegenerate simplices of $Sd(Y)$. It follows that the geometric realisations $|Y^{sc}|$ is homeomorphic[5] to $|Y|$. Simplices of Y^{sc} are naturally oriented and the explicit description of γ is given in [12, p. 61] and the references therein.

In our main algoritm, $Y = \Sigma^d$ will be a triangulation of the d-sphere and X a simplicial set derived from a simplicial complex X^{sc} by contracting its spanning tree into a point. The following lemma shows that we can convert a map $\Sigma^{sc} \to X$ into a map $(\Sigma^{sc})' \to X^{sc}$ between simplicial complexes.

Lemma 2.16. *Let $d > 0$ be fixed. Assume that X^{sc} is a given simplicial complex with a chosen ordering of vertices and a maximal spanning tree T; we denote the underlying simplicial set by X^{ss}. Let $p : X^{ss} \to X := X^{ss}/T$ be the projection to the associated 0-reduced simplicial set. Let Σ be a given d-dimensional simplicial complex with a chosen orientation of each simplex, Σ^{ss} the induced simplicial set, and $f : \Sigma^{ss} \to X$ a simplicial map.*

Then there exists a subdivision $\mathrm{Sd}(\Sigma)$ and a simplicial map $f' : \mathrm{Sd}(\Sigma) \to X^{sc}$ between simplicial complexes[6] *such that*

$$|\Sigma| = |\mathrm{Sd}(\Sigma)| \stackrel{|f'|}{\to} |X^{sc}| \stackrel{|p|}{\to} |X|$$

[4] A version of this lemma is given as [12, Proposition 3.5]. However, we also need the fact that $|Y^{sc}|$ is homeomorphic to $|Y|$, which is not explicitly mentioned in this reference, but follows easily from the construction.

[5] The subdivision $Sd(Y)$ has geometric realisation homeomorphic to $|Y|$ by [28, Thm 4.6.4]. The realisation of $Sd(X)$ is a regular CW complex and $B_*(Sd(Y))$ coincides with the first derived subdivision of this regular CW complex, as defined in [29, p. 137]. The geometric realisation of the resulting simplicial complex is still homeomorphic to $|Y|$ and $|Sd(Y)|$ by [29, Prop. 5.3.8].

[6] The constructed map f does not necessarily preserves orientations: it only maps simplices to simplices.

is homotopic to $|\Sigma^{ss}| \stackrel{|f|}{\to} |X|$. *Moreover,* f' *can be computed in polynomial time, assuming an encoding of the input* f, Σ, X^{sc}, X *and* T.

Thus if Σ is a sphere and f corresponds to a homotopy generator, f' is the corresponding homotopy generator represented as a simplicial map between simplicial complexes. We remark that the algorithm we describe works even if d is a part of the input, but the time complexity would be exponential in general, as the number of vertices in our subdivision $\mathrm{Sd}(\Sigma)$ would grow exponentially with d.

The proof of Lemma 2.16 is given in Section 2.7.

Proof of Theorem 2.11. First assume that a finite simplicial complex X^{sc} is given together with a loop contraction. Then the algorithm goes as follows.

1. We choose an ordering of vertices and convert X^{sc} into a simplicial set. Choosing a spanning tree and contracting it to a point creates a 0-reduced simplicial set X homotopy equivalent to X^{sc}. By Lemma 2.9, we can convert the input data into a list $c_0(\alpha)$ for all generators α of GX_0 in polynomial time.

2. We construct the simplicial set F_d from Lemma 2.12 as simplicial set with polynomial-time effective homology. Hence by Lemma 2.3 we can compute the generators of $H_d(F_d)$ in time polynomial in $\mathrm{size}(X)$. Due to Lemma 2.14 and Theorem 2.13, we can convert these homology generators to homotopy generators $\Sigma^d_j \to F_d$ in time exponential in $P(\mathrm{size}(X) + \mathrm{size}(c_0))$ where P is a polynomial.

3. We compose the representatives of $\pi_d(F_d)$ with ψ_d to obtain representatives $\Sigma^d_j \to X$ of the generators of $\pi_d(X)$, another polynomial-time operation. This way, we compute explicit homotopy generators as maps into the simplicial set X.

4. We use Lemma 2.15 to compute simplicial complexes Σ^{sc}_j and maps $\Sigma^{sc}_j \to \Sigma^d$ homotopic to homeomorphisms. The compositions $\Sigma^{sc}_j \to \Sigma^d_j \to X$ still represent a set of homotopy generators. Finally, by Lemma 2.16, we can compute, for each j, a subdivision of the sphere Σ^{sc}_j and a simplicial map from this subdivision into the simplicial *complex* X^{sc}, in time polynomial in the size of the representatives $\Sigma^{sc}_j \to X$.

In case when the input is a 0-reduced simplicial set X with a loop contraction c_0, only steps 2 and 3 are performed. In either case, the overall exponential complexity bound comes from Berger's Effective Hurewicz inverse theorem. □

2.4 Proof of Theorem 2.2

Similarly as in the proof of Theorem 2.1, we prove a slightly stronger version of Theorem 2.2 that also includes finite simplicial complexes.

Theorem 2.17. *Let* $d \geq 2$ *be fixed. Then*

1. *there is an infinite family of* d*-dimensional* 1*-connected finite simplicial complexes* X *such that for any simplicial map* $\Sigma \to X$ *representing a generator of* $\pi_d(X)$, *the triangulation* Σ *of* S^d *on which* f *is defined has size at least* $2^{\Omega(\mathrm{size}(X))}$.

2. *there is an infinite family of d-dimensional $(d-1)$-connected and $(d-2)$-reduced simplicial sets X such that for any simplicial map $\Sigma \to X$ representing a generator of $\pi_d(X)$, the triangulation Σ of S^d on which f is defined has size at least $2^{\Omega(\mathrm{size}(X))}$.*

Consequently, any algorithm for computing simplicial representatives of the generators of $\pi_d(X)$ has time complexity at least $2^{\Omega(\mathrm{size}(X))}$. [7]

The second item immediately implies Theorem 2.2.

In the first item, we don't assume any certificate for 1-connectedness. However, we suspect that any algorithm that computes representatives of $\pi_d(X)$ for simplicial complexes X *must* necessarily use some explicit certificate of simple connectivity, but so far we have not been able to verify this.

Lemma 2.18. *Let $d \geq 2$.*

1. *There exists a sequence $\{X_k\}_{k\geq 1}$ of d-dimensional $(d-1)$-connected simplicial complexes, such that $H_d(X_k) \simeq \mathbb{Z}$ for all k and for any choice of a cycle $z_k \in Z_d(X_k)$ generating the homology group, the largest coefficient in z_k grows exponentially in $\mathrm{size}(X_k)$.*

2. *There exists a sequence $\{X_k\}_{k\geq 1}$ of d-dimensional $(d-1)$-connected and $(d-2)$-reduced simplicial sets, such that $H_d(X_k) \simeq \mathbb{Z}$ for all k and for any choice of cycles $z_k \in Z_d(X_k)$ generating the homology, the largest coefficient in z_k grows exponentially[8] in $\mathrm{size}(X_k)$.*

Proof of Theorem 2 based on Lemma 2.18. Let $\{X_k\}_{k\geq 1}$ be the sequence of simplicial sets or simplicial complexes from Lemma 2.18. Since they are $(d-1)$-connected, by the theorem of Hurewicz, $\pi_d(X_k) \simeq H_d(X_k) \simeq \mathbb{Z}$. For each k, let Σ_k be a simplicial set or simplicial complex with $|\Sigma_k| = S^d$, and $f_k : \Sigma_k \to X_k$ a simplicial map representing a generator of $\pi_d(X_k)$. The generator of $H_d(\Sigma_d)$ contains each non-degenerate d-simplex with a coefficient ± 1 (this follows from the fact that Σ_k is a triangulation of the d-sphere and the d-homology of the d-sphere is generated by its fundamental class). The Hurewicz isomorphism $\pi_d(X_k) \to H_d(X_k)$ maps such a representative to the formal sum of simplices

$$f_k \mapsto \sum_{\sigma \text{ is a } d-\text{simplex in } (\Sigma_k)} \pm f_k(\sigma) \in C_d(X_k)\,,$$

which represents a generator of $H_d(X_k)$. It follows from Lemma 2.18 that the number of d-simplices in Σ_k grows exponentially in $\mathrm{size}(X_k)$. Moreover, the complexity of any algorithm that computes $f_k : \Sigma_k \to X_k$ is at least the size of Σ_k, which completes the proof. □

It remains to define the sequence from Lemma 2.18:

Proof of Lemma 2.18.

[7] We write $f(x) = \Omega(g(x))$ if $\limsup_{x\to\infty} |\frac{f(x)}{g(x)}| \geq 0$.

[8] With a slight abuse of language, we assume that each X_k not only a simplicial set but also its algorithmic representation with a specified size such as explained in Section 2.2.

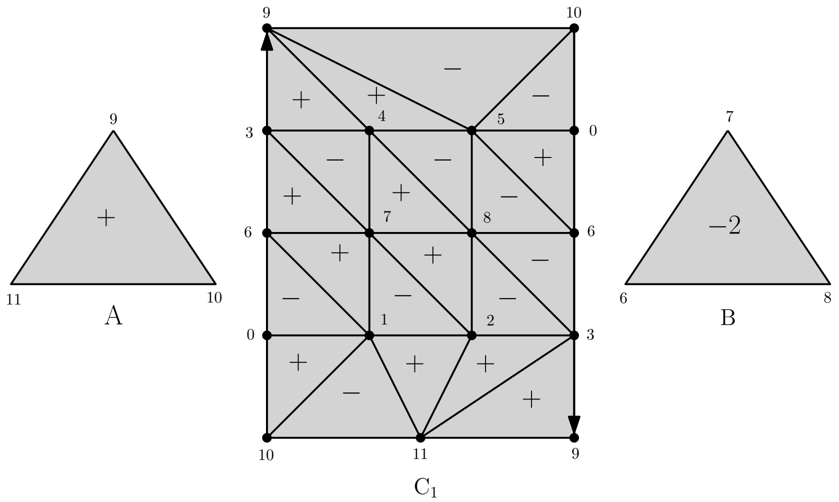

Figure 2.2: The Möbius band is the mapping cylinder of a degree 2 map $S^1 \to S^1$. The triangulation has four layers because starting from the boundary, which is a triangle, we first need to pass to a hexagon in order to cover the middle triangle twice, obtaining the desired degree 2 map. Connecting k copies of the Möbius band creates a mapping cylinder of a degree 2^k map, using only linearly (in k) many simplices. Gluing the full triangles A and B to the ends of this mapping cylinder finishes the construction of X_k. The red coefficients exhibit a generator ξ of $H_2(X_1) = Z_2(X_1) \simeq \mathbb{Z}$ given as a formal sum of 2-simplices.

1. We begin by constructing for every $d \geq 2$, a sequence of $\{X_k\}_{k\geq 1}$ of $(d-1)$-connected simplicial complexes, such that $H_d(X_k) \simeq \mathbb{Z}$ for all k, and for any choice of a cycle $z_k \in Z_d(X_k)$ generating the homology group, the largest coefficient in z_k grows exponentially in $\mathrm{size}(X_k)$.

   We start with $d = 2$. The idea is to glue X_k out of k copies of a triangulated mapping cylinders of a degree 2 map $S^1 \to S^1$, i.e. k Möbius bands, and then fill in the two open ends with one triangle each (A and B in Figure 2.2). The case $k = 1$ is shown in Figure 2.2. For $k \geq 2$, we take k copies of the triangulated Möbius band and identify the middle circle of each one to the boundary of the next one.

   We observe that, up to homotopy equivalence, X_k consists of a 2-disc with another 2-disc which is attached to it via the boundary map $S^1 \to S^1$ of degree 2^k. Therefore, X_k is simply connected and has $H_2(X_k) \simeq \mathbb{Z}$ and any homology generator will contain the 2-simplex A with coefficient ± 1 and B with coefficient $\pm 2^k$.

   Similarly for $d > 2$, the simplicial complex X_k is obtained by glueing k copies of a triangulated mapping cylinder of a degree 2 map $S^{d-1} \to S^{d-1}$, and the two open ends are filled in with two triangulated d-balls.

2. For every $k \geq 1$ we define the simplicial sets X_k to have one vertex $*$, no non-degenerate simplices up to dimension $d-2$, k non-degenerate $(d-1)$-simplices $\sigma_1, \ldots, \sigma_k$ that are all spherical (that is, for all i, j, $d_i\sigma_j = *$ is the degeneracy of the only vertex of X_k), and $k+1$ d-simplices $A, C_1, C_2, \ldots, C_{k-1}, B$ such that

- $d_0A = \sigma_1$, $d_jA = *$ for $j > 0$,
- $d_0C_i = \sigma_i$, $d_1C_i = \sigma_{i+1}$, $d_2C_i = \sigma_i$ and $d_jC_i = *$ for $j > 2$, and
- $d_0B = \sigma_k$, $d_jB = *$ for $j > 0$.

X_k does not have any non-degenerate simplices of dimension larger than d. The relations of a simplicial set are satisfied, because d_id_j is trivial in all cases.

The boundary operator in the associated normalised chain complex $C_*(X_i)$ acts on basis elements as

- $\partial A = \sigma_1$
- $\partial C_i = 2\sigma_i - \sigma_{i+1}$, and
- $\partial B = \sigma_k$.

To see that X_k is $(d-1)$-connected for $d > 2$, it is enough to prove that $H_{d-1}(X_k)$ is trivial (by 1-reduceness and Hurewicz theorem). This is true, because σ_1 is the boundary of A and for $i > 1$, σ_i is the boundary of the chain

$$2^{i-1}A - 2^{i-2}C_1 - \ldots - 2C_{i-2} - C_{i-1}.$$

In the case $d = 2$, X_k is not 1-reduced, but we can show 1-connectedness similarly as in the proof of the first part: up to homotopy, X_k consists of two discs with boundaries together via a map of degree 2^{k-1}.

There are no non-degenerate $(d+1)$-simplices, so $H_d(X_k) \simeq Z_d(X_k)$ and a simple computation shows that every cycle is a multiple of

$$2^{k-1}A - 2^{k-2}C_1 - 2^{k-3}C_2 - \ldots - C_{k-1} - B.$$

The computer representation of X_k has size that grows linearly with k, but the coefficients of homology generators grow exponentially with k, so they grow exponentially with $\mathrm{size}(X_k)$.

□

Discussion on optimality. If $d = 2$ and X is a 1-reduced simplicial set, then generators of $H_2(X)$ can be computed via the Smith normal form of the differential $\partial_3 : C_3(X) \to C_2(X)$. Using canonical bases, the matrix of $\partial_3 = d_0 - d_1 + d_2 - d_3$ satisfies that the sum of absolute values over each column is at most 4. We were not able to find any infinite family of such matrices so that the smallest coefficient in any set of homology generating cycles grows exponentially with the size of X (that is, the size of the matrix). However, if a set of homology-generating cycles with subexponential coefficients always exists and can be found algorithmically in polynomial time, our main algorithm given as Theorem 2.1 is optimal in this case as well. This is because the exponential complexity of the algorithm only appears in the geometric realisation of an element of GX_1^{sph} with large (exponential) exponents (see “Arrow 3” in Section 2.5), and the only source of such exponents is the homology $H_1(AX) \simeq H_2(X)$.

2.5 Effective Hurewicz Inverse

Here we will prove Theorem 2.13 by directly describing the algorithm proposed in [5] and analysing its running time.

Definition 2.19. *Let G be a simplicial group. Then the Moore complex $\widetilde{G}$ is a (possibly non-abelian) chain complex defined by $\widetilde{G}_i := G_i \cap (\bigcap_{j>0} \ker d_j)$ endowed with the differential $d_0 : \widetilde{G}_i \to \widetilde{G}_{i-1}$.*

It can be shown that $d_0 d_0 = 1$ in $\widetilde{G}$ and that $\mathrm{Im}(d_0)$ is a normal subgroup of $\ker d_0$ so that the homology $H_*(\widetilde{G})$ is well defined.

Definition 2.20. *Let F be a 0-reduced simplicial set, GF the associated simplicial group from Def. 2.5, and $\widetilde{GF}$ its Moore complex. We define AF to be the Abelianization of GF and $\widetilde{AF}$ to be the Moore complex of AF. The simplicial group AF is also endowed with a chain group structure via $\partial = \sum_j (-1)^j d_j$. If $\sigma \in F_k$, we will denote by $\overline{\sigma}$ the corresponding simplex in GF_{i-1}, resp. AF_{i-1}.*

Note that, following Def. 2.5, the "last" differential $d_k \overline{\sigma}$ in AF_k equals $\overline{d_k \sigma} - \overline{d_{k+1}\sigma}$. Clearly, the Abelianization map $p : GF \to GF/[GF, GF] = AF$ takes $\widetilde{GF}$ into $\widetilde{AF}$.

Kan showed in [39] that for $d > 1$ and a $(d-1)$-connected simplicial set F, the Hurewicz isomorphism can be identified with the map $H_{d-1}(\widetilde{GF}) \to H_{d-1}(\widetilde{AF})$ induced by Abelianization, whereas these groups are naturally isomorphic to $\pi_d(F)$ and $H_d(F)$, respectively. Our strategy is to construct maps representing the isomorphisms $1, 2, 3$ in the commutative diagram

$$\begin{array}{ccc} \pi_d(F) & \xleftarrow{\;h^{-1}\;} & H_d(F) \\ \uparrow{\scriptstyle 3} & & \downarrow{\scriptstyle 1} \\ H_{d-1}(\widetilde{GF}) & \xleftarrow[\;2\;]{} & H_{d-1}(\widetilde{AF}). \end{array} \tag{2.3}$$

Here h stands for the Hurewicz isomorphism, 1 is induced by a homotopy equivalence of chain complexes, 2 is the inverse of $H_{d-1}(p)$ where p is the Abelianization, and 3 represents an isomorphism between the $(d-1)$'th homology of $\widetilde{GF}$ (that models the loop space of F) and $\pi_d(F)$. The algorithms representing $1, 2, 3$ will act on representatives, that is, 1 and 2 will convert cycles to cycles and 3 will convert a cycle to a simplicial map $\Sigma^d \to F$ where $|\Sigma^d| = S^d$. In what follows, we will explicitly describe the effective versions of $1, 2, 3$ and show that the underlying algorithms are polynomial for arrows 1,2 and exponential for arrow 3.

Arrow 1.

Let F be a 0-reduced simplicial set, $C_*(F)$ be the (unreduced) chain complex of F and AF_{*-1} the shifted chain complex of AF defined by $(AF_{*-1})_i := AX_{i-1}$. As a chain complex, AF_{*-1} is a subcomplex of $C_*(F)$ generated by all simplices that are not in the image of the last degeneracy. Let $\widetilde{AF}_{*-1}$ be the Moore complex of AF_{*-1}.

Lemma 2.21. *There exists a polynomial-time strong chain deformation retraction* $(f, g, h) : C_*(F) \to \widetilde{AF}_{*-1}$. *That is,* $f : C_*(F) \to \widetilde{AF}_{*-1}$, $g : \widetilde{AF}_{*-1} \to C_*(F)$ *are polynomial-time chain-maps and* $h : C_*(F) \to C_{*+1}(F)$ *is a polynomial map such that* $fg = \mathrm{id}$ *and* $gf - \mathrm{id} = h\partial + \partial h$.

Proof. First we will define a chain deformation retraction from $C_*(F)$ to AF_{*-1} represented by $f_0 : C_*(F) \to AF_{*-1}$, $g_0 : AF_{*-1} \to C_*(F)$ and $h_0 : C_*(F) \to C_{*+1}(F)$.

The chain complex AF_{*-1} consists of Abelian groups AF_{k-1} freely generated by k-simplices in F that are not in the image of the last degeneracy s_{k-1}. On generators, we define $f_0(\sigma) = \overline{\sigma}$ whenever σ is a k-simplex not in $\mathrm{Im}(s_{k-1})$ and $f_0(x) = 0$ otherwise. Deciding whether σ is in the image of s_{k-1} amounts to deciding $\sigma = s_{k-1}d_k\sigma$ which can be done in time polynomial in $\mathrm{size}(I) + \mathrm{size}(\sigma)$, the polynomial depending on k. It follows that f_0 is a locally polynomial map.

The remaining maps are defined by $g_0(\overline{\sigma}) \coloneqq \sigma - s_{k-1}d_k\sigma$ and $h_0(\sigma) \coloneqq (-1)^k s_k\sigma$. These maps are locally polynomial as well and it is a matter of straight-forward computations to check that f_0 and g_0 are chain maps, $f_0g_0 = \mathrm{id}$ and $g_0f_0 - \mathrm{id} = h_0\partial + \partial h_0$.

Further, we define another chain deformation retraction from AF to $\widetilde{AF}$. For each $p \geq 0$, let A^p be a chain subcomplex of AF defined by

$$(A^p)_k \coloneqq \{x \in AF_k : d_i x = 0 \quad \text{for} \quad i > \max\{k - p, 0\}\}$$

that is, the kernel of the p last face operators, not including d_0 (d_i refers here to the face operators in AF). Then A^{p+1} is a chain subcomplex of A^p and we define the maps $f_{p+1} : (A^p)_k \to (A^{p+1})_k$ by $f_{p+1}(x) = x - s_{k-p-1}d_{k-p}x$ whenever $k - p > 0$, and $f_{p+1}(x) = x$ otherwise; $g_{p+1} : A^{p+1} \to A^p$ will be an inclusion, and $h_{p+1} : (A^p)_k \to (A^p)_{k+1}$ via $h_{p+1}(x) = (-1)^{k-p}s_{k-p}x$ if $k - p > 0$ and 0 otherwise. A simple calculation shows that f_{p+1}, g_{p+1} are chain maps, $f_{p+1}g_{p+1} = \mathrm{id}$, $g_{p+1}f_{p+1} - \mathrm{id} = h_{p+1}\partial + \partial h_{p+1}$ and it is clear that $f_{p+1}, g_{p+1}, h_{p+1}$ are polynomial-time maps.

By definition, the Moore complex $\widetilde{AF} = \cap_{p>0} A^p$. The strong chain deformation retraction (f, g, h) from $C_*(F)$ to $\tilde{A}F_{*-1}$ is then defined by the infinite compositions

$$f \coloneqq \dots f_{k+1}f_k \dots f_1 f_0$$
$$g \coloneqq g_0 g_1 \dots g_k g_{k+1} \dots$$

and the infinite sum

$$h = h_0 + g_1h_1f_1 + (g_1g_2)h_2(f_2f_1) + \dots$$

which are all well-defined, because when applying them to an element x, only finitely many of f_j, g_j differ from the identity map and only finitely many h_j are nonzero. These are the maps f, g, h from the lemma and we need to show that if the degree k is fixed, then we can evaluate f, g, h on $C_k(F)$ resp. $\tilde{A}F_{k-1}$ in time polynomial in the input size. However, for fixed k , the definition of f, g, h includes only f_i, g_i, h_i for $i < k$. Then f, g are composed of k polynomial-time maps and h is a sum of k polynomial-time maps. □

The polynomial-time version of arrow 1 is then induced by applying the map f from Lemma 2.21.

Arrow 2.

Lemma 2.22 (Boundary certificate)**.** *Let $d > 1$ be fixed and let F be a $(d-1)$-connected simplicial set with polynomial-time homology. There is an algorithm that, for $j < d-1$ and a cycle $z \in Z_j(\widetilde{AF})$, computes an element $c^A(z) \in \widetilde{AF}_{j+1}$ such that $d_0 c^A(z) = z$. The running time is polynomial in $\operatorname{size}(z) + \operatorname{size}(I)$.*

Proof. First note that the $(d-1)$-connectedness of F implies that $H_{j+1}(F) \simeq H_j(\widetilde{AF})$ are trivial for $j < d-1$, so each cycle in these dimensions is a boundary.

By assumption, F has a polynomial-time homology, which means that there exists a globally polynomial-time chain complex E_*F, a locally polynomial-time chain complex Y and polynomial-time reductions from Y to $C_*(F)$ and E_*F

$$E_*F \overset{\mathrm{P}}{\Lleftarrow} Y \overset{\mathrm{P}}{\Rrightarrow} C_*(F).$$

Let (f', g', h') be chain homotopy equivalence of Y and $\widetilde{AF}_{*-1}$ defined as the composition of $Y \Rrightarrow C_*(F)$ and the chain homotopy equivalence of $C_*(F)$ and $\widetilde{AF}_{*-1}$ described in Lemma 2.21. Further, let f'', g'', h'' be the maps defining the reduction $Y \Rrightarrow E_*F$: all of these maps are polynomial-time.

Let $j < d-1$ and $z \in Z_j(\widetilde{AF})$, $z = \sum_j k_j y_j$. Then the element $f''g'(z)$ is a cycle in $E_{j+1}F$ and can be computed in time polynomial in $\operatorname{size}(z) + \operatorname{size}(I)$. In particular, the size of $f''g'(z)$ is bounded by such polynomial. The number of generators of $E_{j+2}F$ and $E_{j+1}F$ is polynomial in $\operatorname{size}(I)$ and we can compute, in time polynomial in $\operatorname{size}(I)$, the boundary matrix of $\partial : E_{j+2}F \to E_{j+1}F$ with respect to the generators.

Next we want to find an element $t \in E_{j+2}F$ such that $\partial t = f''g'(z)$. Using generating sets for $E_{j+2}F$, $E_{j+1}F$, this reduces to a linear system of Diophantine equations and can be solved in time polynomial in the size of the ∂-matrix and the right hand side $f''g'(z)$ [41].

Finally, we claim that $c^A(z) := f'g''(t) - f'h''g'(z)$ is the desired element mapped to z by the differential in $\widetilde{AF}$. This follows from a direct computation

$$\begin{aligned}
\partial c^A(z) :=& \partial f'g''(t) - \partial f'h''g'(z) = \\
=& f'g''(\partial t) - \partial f'h''g'(z) = \\
=& f'g''f''g'(z) - \partial f'h''g'(z) = \\
=& f'(h''\partial + \partial h'' + \mathrm{id})g'(z) - \partial f'h''g'(z) = \\
=& f'h''g'\partial z + \partial f'h''g'(z) + f'g'(z) - \partial f'h''g'(z) = \\
=& 0 + f'g'(z) = z
\end{aligned}$$

The computation of t as well as all involved maps are polynomial-time, hence the computation of $c^A(z)$ is polynomial too. □

The next lemma will be needed as an auxiliary tool later.

Lemma 2.23. *Let S be a countable set with a given encoding, G be the free (non-abelian) group generated by S, and define $\operatorname{size}(\prod_j s_j^{k_j}) := \sum_j (\operatorname{size}(s_j) + \operatorname{size}(k_j))$. Let $G' := [G, G]$ be its commutator subgroup.*

Then there exists a polynomial-time algorithm that for an element $g = \prod_j s_j^{k_j}$ in $G' \subseteq G$, computes elements $a_i, b_i \in G$ such that $g = \prod_j [a_j, b_j]$.

In other words, we can decompose commutator elements into simple commutators in polynomial-time at most.

Proof. Let us choose a linear ordering on S and let $g = \prod_j s_j^{k_j}$ be in G': that is, for each j, the exponents $\{k_{j'} : s_{j'} = s_j\}$ sum up to zero. We will present a bubble-sort type algorithm for sorting elements in g. Going from the left to right, we will always swap $s_j^{k_j}$ and s_{j+1}^{k+1} whenever $s_{j+1} < s_j$. Such swap always creates a commutator, but that will immediately be moved to the initial segment of commutators.

More precisely, assume that Init is the initial segment, $x = s_j^{k_j}$ and $y = s_{j+1}^{k_{j+1}}$ should be swapped, Rest represent the segment behind y, and Commutators is a final segment of commutators. The swapping will consists of these steps:

$$\begin{aligned} & \text{Init } x \ y \text{ Rest Commutators} \\ \mapsto\ & \text{Init } y \ x \ [x^{-1}, y^{-1}] \text{ Rest Commutators} \\ \mapsto\ & \text{Init } y \ x \text{ Rest } \big([x^{-1}, y^{-1}]\ [[y^{-1}, x^{-1}], \text{Rest}^{-1}] \text{ Commutators}\big) \end{aligned}$$

where the parenthesis enclose a new segment of commutators. Before the parenthesis, x and y are swapped. Each such swap requires enhancing the commutator section by two new commutators of size at most $\mathrm{size}(g)$, hence each such swap has complexity linear in $\mathrm{size}(g)$.

Let as call everything before the commutator section a “regular section”. Going from left to right and performing these swaps will ensure that the largest element will be at the end of the regular section. But no later then that, the largest element y_{largest} disappears from the regular section completely, because all of its exponents add up to 0. Again, starting from the left and performing another round of swaps will ensure that the second-largest elements disappear from the regular section; repeating this, all the regular section will eventually disappear which will happen in at most $\mathrm{size}(g)^2$ swaps in total. Each swap has complexity linear in $\mathrm{size}(g)$ and the overall time complexity is not worse than cubic. □

Lemma 2.24. *Assume that F is a parametrized simplicial set with polynomially contractible loops. Let $k > 0$, $\gamma \in GF_k$ be spherical and $\alpha \in GF_k$ is arbitrary. There is a polynomial-time algorithm that computes $\delta \in GF'_{k+1}$ such that $d_0\delta = [\alpha, \gamma]$ and $d_i\delta = 1$ for all $i > 0$.*

In other words, a simple commutator of a spherical element with another element can always be “contracted” in GF' in polynomial time. Our proof roughly follows the construction in Kan [39, Sec. 8]

Proof. For $x \in GF_0$, we will denote by c_0x the element of $\widetilde{GF}_1$ such that $d_0c_0x = x$: this can be computed in polynomial-time by the assumption on polynomial loop contractions. For the simplex $\alpha \in GF_k$, we define $(k+1)$-simplices $\beta_0, \dots, \beta_k$ by $\beta_k := s_0^k c_0 d_0^k \alpha$ and inductively $\beta_{j-1} := (s_j d_j \beta_j) \cdot (s_j \alpha^{-1}) \cdot (s_{j-1}\alpha)$ for $j < k$. Then the following relations hold:[9]

- $d_0\beta_0 = \alpha.$

[9] Kan uses a slightly different convention in [39] but the resulting properties are the same. The sequence $\beta_0, \dots, \beta_k$ can be interpreted as a discrete path from α to the identity element.

- $d_j\beta_j = d_j\beta_{j-1}$, $1 \le j \le k$
- $d_{k+1}\beta_k = 1$.

The second and third equations are a matter of direct computation, while the first follows from the more general relation $d_0^{j+1}\beta_j = d_0^j\alpha$ which can be proved by induction. If k is fixed, then all $\beta_0, \dots, \beta_k$ can be computed in polynomial time.

The desired element δ is then the alternating product

$$\delta := [\beta_0, s_0\gamma]\,[\beta_1, s_1\gamma]^{-1} \,\dots\, [\beta_k, s_k\gamma]^{\pm 1}.$$

□

Lemma 2.25. *Under the assumptions of Theorem 2.13, there exist homomorphisms $c_j : GF_j \to GF_{j+1}$ for $0 \le j < d-1$ such that*

1. $d_0c_j = \mathrm{id}$,
2. $d_ic_j = c_{j-1}d_{i-1}$, $0 < i \le j+1$, *and*
3. $c_js_i = s_{i+1}c_{j-1}$ *for* $0 < j < d-1$ *and* $0 \le i < j$,
4. $d_1c_0(x) = 1$ *for all* $x \in GF_0$.

If d is fixed and $x \in GF_j$, $j < d-1$, then $c_j(x)$ can be computed in polynomial time.

Proof. The homomorphism c_0 can be constructed directly from the assumption on polynomial contractibility of loops. We have a canonical basis of GF_0 consisting of all non-degenerate 1-simplices of F. For $\sigma \in F_1$, we denote by $\overline{\sigma}$ the corresponding generator of GF_0. The we define $c_0(\prod \overline{\sigma}_j^{k_j})$ to be $\prod b_j^{k_j}$ where b_j is the element of GF_1 such that $d_0b_j = \overline{\sigma}_j$ and $d_1b_j = 1$.

In what follows, assume that $1 \le k < d-1$ and c_i have been defined for all $i < k$. We will define c_k in the following steps.

Step 1. Contractible elements.

Let $x \in GF_k$. We will say that x is *contractible* and $y \in GF_{k+1}$ is a *contraction of* x, if $d_0y = x$ and $d_iy = c_{k-1}d_{i-1}x$ for all $i > 0$.

The general strategy for defining c_k will be to find a contraction h for each basis element ($(k+1)$-simplex) $g \in GF_k$ and define $c_k(g) := h$. This will enforce properties 1 and 2. Moreover, in case when g is degenerate, the contraction will be chosen in such a way that property 3 holds too; otherwise it holds vacuously. Property 4 only deals with c_0 and is satisfied by the definition of loop contraction (a polynomial-time c_0 is given as an input in Theorem 2.13).

Step 2. Contraction of degenerate elements.

Let $g = s_iy$ be a basis element of GF_k, $y \in GF_{k-1}$. Then g can be uniquely expressed as s_jz where j is the maximal i such that $g \in \mathrm{Im}(s_i)$. We then define $c_k(g) := s_{j+1}c_{k-1}(z)$. Note that

$$d_0c_k(g) = d_0s_{j+1}c_{k-1}(z) = s_jd_0c_{k-1}(z) = s_jz = g,$$

so property 1 is satisfied. To verify property 2, first assume that $i \in \{j+1, j+2\}$. Then we have

$$d_i c_k(g) = d_i s_{j+1} c_{k-1}(z) = c_{k-1}(z) = c_{k-1} d_{i-1} s_j z = c_{k-1} d_{i-1} g.$$

This fully covers the case $k = 1$, because then the only possibility is $j = 0$ and $i \in \{1, 2\}$. Further, let $k > 1$. If $i \leq j$, then we have

$$\begin{aligned} d_i c_k g =& d_i c_k s_j z = d_i s_{j+1} c_{k-1}(z) = s_j d_i c_{k-1}(z) = s_j c_{k-2} d_{i-1} z = \\ =& c_{k-1} s_{j-1} d_{i-1} z = c_{k-1} d_{i-1} s_j z = c_{k-1} d_{i-1} g \end{aligned}$$

and if $i > j + 2$, then the computation is completely analogous, using the relation $d_i s_{j+1} = s_{j+1} d_{i-1}$ instead.

So far, we have shown that $c_k(g) \coloneqq s_{j+1} c_{k-1} g$ is a contraction of g. It remains to show property 3. That is, we have to show that if $g = s_j z$ can also be expressed as $s_i y$, then $c_k(s_i y) = s_{i+1} c_{k-1} y$.

The degenerate element g has a unique expression $g = s_{i_u} \ldots s_{i_1} s_{i_0} v$ where $i_0 < i_1 < \ldots < i_u = j$ and is expressible as $s_i x$ if and only if $i = i_j$ for some $j = 0, 1, \ldots, u$. Choosing such $i < j$, we can rewrite g as $g = s_j s_i v$ for some v and then $g = s_i s_{j-1} v$, so that $y = s_{j-1} v$ and $z = s_i v$. Then we again use induction to show

$$\begin{aligned} c_k(s_i y) =& s_{j+1} c_{k-1}(z) = s_{j+1} c_{k-1} s_i v = s_{j+1} s_{i+1} c_{k-2} v = \\ =& s_{i+1} s_j c_{k-2} v = s_{i+1} c_{k-1} s_{j-1} v = s_{i+1} c_{k-1} y \end{aligned}$$

as required.

Step 3. Decomposition into spherical and conical parts.

We will call an element $\hat{x} \in GF_k$ to be *conical*, if it is a product of elements that are either degenerate or in the image of c_{k-1}. Let $x \in GF_k$ be arbitrary. We define $x_k \coloneqq x$ and inductively $x_{i-1} \coloneqq x_i (s_{i-1} d_i x_i)^{-1}$. In this way we obtain $x_0, \ldots, x_n$ such that x_i is in the kernel of d_j for $j > i$ and $x = x_0 y$ where y is a product of degenerate simplices. Further, let $x^s \coloneqq x_0 (c_{k-1} d_0 x_0)^{-1}$. A simple computation shows that x^s is *spherical*, that is, $d_i x^s = 1$ for all i. We obtain an equation $x = x^s \hat{x}$ where $\hat{x} = (c_{k-1}(d_0 x_0) y$; this is a decomposition of x into a spherical part x^s and a conical element $\hat{x}$.

We will define c_k on non-degenerate basis elements $g = \overline{\sigma}$ by first decomposing $g = g^S \hat{g}$ into a spherical and conical part, finding contractions h_1 of g^S and h_2 of $\hat{g}$, and defining $c_k(g) \coloneqq h_1 h_2$. Then $c_k(g)$ is a contraction of g and hence satisfies properties 1 and 2: property 3 is vacuously true once g is non-degenerate.

Step 4. Contraction of the conical part.

Let $\hat{x} \coloneqq c_{k-1}(d_0 x_0)\, y$ be the conical part defined in the previous step. By construction, $x_0 \in \tilde{GF}_k$ and y is a product of degenerate elements $s_{i_1} u_1 \ldots s_{i_l} u_l$. We define the contraction of $c_{k-1}(d_0 x_0)$ to be

$$\tilde{c}_k(c_{k-1}(d_0 x_0)) \coloneqq s_0 c_{k-1}(d_0 x_0).$$

Note that this satisfies property 1 as $d_0 \tilde{c}_k c_{k-1}(d_0 x_0) = c_{k-1}(d_0 x_0)$. For property 2, we first verify

$$d_1 \tilde{c}_k c_{k-1}(d_0 x_0) = d_1 s_0 c_{k-1}(d_0 x_0) = c_{k-1}(d_0 x_0) = c_{k-1} d_0 c_{k-1}(d_0 x_0).$$

Not let $i \geq 2$. If $k = 1$, then the remaining face operator is d_2 and we have

$$d_2 \tilde{c}_1 c_0(d_0 x_0) = d_2 s_0 c_0(d_0 x_0) = s_0 d_1 c_0(d_0 x_0) = 1 = c_0 d_1 c_0(d_0 x_0)$$

using axiom 4 for c_0. Finally, if $i \geq 2$ and $k \geq 2$, we have

$$\begin{aligned} d_i \tilde{c}_k c_{k-1}(d_0 x_0) =& d_i s_0 c_{k-1}(d_0 x_0) = s_0 d_{i-1} c_{k-1}(d_0 x_0) = s_0 c_{k-1} d_{i-2} d_0 x_0 = \\ =& s_0 c_{k-1} d_0 d_{i-1} x_0 = s_0 c_{k-1} d_0 1 = 1 = c_{k-1} c_{k-2} d_0 d_{i-1} x_0 = \\ =& c_{k-1} c_{k-2} d_{i-2} d_0 x_0 = c_{k-1} d_{i-1} c_{k-1}(d_0 x_0), \end{aligned}$$

where we exploited the fact that $x_0 \in \widetilde{GF}_k$ and hence $d_u x_0 = 1$ for $u \geq 2$.

The contraction of degenerate elements y has already been defined in Step 2, so we can define a contraction of $c_{k-1}(d_0 x_0) y$ to be $s_0 c_{k-1}(d_0 x_0)\, c_k(y)$.

Step 5. Contraction of commutators.

Let $g' \in GF'_k$ be an element of the commutator subgroup. By Lemma 2.23, we can algorithmically decompose g' into a product of simple commutators, so to find a contraction of g', it is sufficient to find a contraction of each simple commutator $[x, y]$ in this decomposition.

Let $x = x^S\, \hat{x}$ and $y = y^S\, \hat{y}$ be the decompositions into spherical and conical parts described in Step 3. Using the notation ${}^b a := bab^{-1}$, we can decompose $[x, y]$ as follows [5, p. 60]:

$$[x, y] = \left([x, y][\hat{y}, x]\right)\left([x, \hat{y}][\hat{y}, \hat{x}]\right)[\hat{x}, \hat{y}] = [{}^{xy}x^{-1}, {}^{xy}(y^{-1}\hat{y})]\,[{}^{x}\hat{y}, {}^{x}(x^{-1}\hat{x})]\,[\hat{x}, \hat{y}]. \quad (2.4)$$

Both $x^{-1}\hat{x}$ and $y^{-1}\hat{y}$ are spherical simplices and so are their conjugations. It follows that equation (2.4) can be rewritten to $[x, y] = [\alpha_1, \gamma_1]\,[\alpha_2, \gamma_2]\,[\hat{x}, \hat{y}]$ where γ_1 and γ_2 are spherical. All of these decompositions are done by elementary formulas and are polynomial-time in the size of x and y.

By Lemma 2.24 we can find an elements $\lambda_i \in \widetilde{GF}_{k+1}$ such that $d_0 \lambda_i = [\alpha_i, \gamma_i]$, $i = 1, 2$, in polynomial time. Further, both $\tilde{x}$ and $\tilde{y}$ are conical and they are in the form $\tilde{x} = c_0(d_0 x_0) x_{deg}$ where $x_0 \in \widetilde{GF}_k$ and x_{deg} is degenerate; similar decomposition holds for y. In Step 4 we showed how to compute elements c^x and c^y such that c^x, c^y is a contraction of $\hat{x}$, $\hat{y}$, respectively. Then $[c^x, c^y]$ is a contraction of $[\hat{x}, \hat{y}]$ and $\lambda_1 \lambda_2 [c^x, c^y]$ is a contraction of $[x, y]$.

Step 6. Contraction of spherical elements.

The last missing step is to compute a contraction of the spherical element g^S where g^S is the spherical part of a basis element $g \in GF_k$.

Let us denote by p the projection $GF \xrightarrow{p} AF$. The projection $z := p(g^S)$ is in the kernel of all face operators and hence a cycle in $\widetilde{AF}_k$. By Lemma 2.22, we can compute $t := c_k^A(z) \in \widetilde{AF}_{k+1}$ such that $d_0 t = z$, in polynomial time. Let $h \in GF_{k+1}$ be any p-preimage[10] of t. Let $h_k := h$ and inductively define $h_{j-1} := h_j (s_{j-1} d_j h_j)^{-1}$ for $j < k$. Then h_0 is in the kernel of all faces except d_0, that is, $h_0 \in \widetilde{GF}_{k+1}$. It follows that $p(h_0) \in \widetilde{AF}_{k+1}$ is in the kernel of all faces except d_0. We claim that $p(h_0) = t$.This can be shown as follows: assume that $p(h_j) = t$, then $p(h_{j-1}) = p(h_j) + p(s_{j-1} d_j h_j^{-1}) = t + s_{j-1} d_j t = t + 0 = t$.

[10]For $t = \sum_j k_j \overline{\sigma}_j$, we may choose $h = \prod_j \overline{\sigma}_j^{k_j}$ (choosing any order of the simplices).

We have the following commutative diagram:

$$\begin{array}{ccccc} & & h_0 & \longmapsto & t \\ \widetilde{GF}'_{k+1} & \hookrightarrow & \widetilde{GF}_{k+1} & \xrightarrow{p} & \widetilde{AF}_{k+1} \\ \downarrow d_0 & & \downarrow d_0 & & \downarrow d_0 \\ \widetilde{GF}'_k & \hookrightarrow & \widetilde{GF}_k & \xrightarrow{p} & \widetilde{AF}_k \\ & & g^S & \longmapsto & z \end{array}$$

Both g^S and $d_0 h_0$ are mapped by p to the same element z: it follows that $g^S(d_0 h_0)^{-1}$ is mapped by p to zero and hence is an element of the commutator subgroup. Let $\tilde{h}$ be the contraction of $g^S(d_0 h_0)^{-1}$, computed in Step 5, and finally let $h := \tilde{h} h_0$. Then h is an element of $\widetilde{GF}_{k+1}$ and a direct computation shows that $d_0 h = g^S$ as desired.

This completes the construction of c_k: for each non-degenerate basis element g of GF_k, $c_k(g)$ is defined to be the product of the contraction of g^S and the contraction[11] of $\hat{g}$.

All the subroutines described in the above steps are polynomial-time. Thus we showed that if there exists a polynomial-time algorithm for c_{k-1}, then there also exists a polynomial-time algorithm for c_k. The existence of a polynomial-time c_0 follows from the assumption on polynomial loop contractibility and d is fixed, thus there exists a polynomial-time algorithm that for $x \in GF_j$ computes $c_j(x)$ for each $j < d-1$. □

Lemma 2.26 (Construction of arrow 2)**.** *Under the assumption of Theorem 2.13, let $z \in Z_{d-1}(\widetilde{AF})$ be a cycle. Then there exists a polynomial-time algorithm that computes a cycle $x \in Z_{d-1}(\widetilde{GF})$ such that the Abelianization of x is z.*

The assignment $z \mapsto x$ is hence an effective inverse of the isomorphism

$$H_{d-1}(\widetilde{GF}) \to H_{d-1}(\widetilde{AF})$$

on the level of representatives.

Proof. Let c_{d-2} be the contraction from Lemma 2.25 and $z \in Z_{d-1}(\widetilde{AF})$ be a cycle. First choose $y \in GF_{d-1}$ such that $p(y) = z$. Creating the sequence $y_n := y$, $y_{j-1} := y_j s_{j-1} d_j y_j^{-1}$ for decreasing j, yields an element $y_0 \in \widetilde{GF}_{d-1}$ that is still mapped to z by p, similarly as in Step 4 of Lemma 2.25. The equation $pd_0(y_0) = d_0 p(y_0) = d_0 z = 0$ shows that $d_0 y_0$ is in the commutator subgroup $\widetilde{GF}'_{d-2}$. We define $x := y_0 c_{d-2}(d_0 y_0)^{-1}$: this is already a cycle in $\widetilde{GF}_{d-1}$ and $p(x) = p(y_0) = z$. □

Arrow 3.

The construction of map 3 is one of the main results from [6] and involves further definitions. Here, we describe the main points of the construction only while details are given in later sections.

[11] The connectivity assumption on F was exploited in the existence of the contraction c_j^A on the Abelian part.

Given a 0-reduced simplicial set F, there exists a simplicial group $\overline{\Omega}F$ that is a discrete analog of a loopspace of F i.e. $\pi_{d-1}(\overline{\Omega}F) \cong \pi_d(F)$. Further, there is a homomorphism of simplicial groups $t\colon GF \to \overline{\Omega}F$ that induces an isomorphism on the level of homotopy groups. This is described in [6, Proposition 3.3].

The homomorphism t is given later by formula (2.6) and the simplicial set $\overline{\Omega}F$ is described in the next section. Here, we remark that the size of $t(g)$ is exponential in size of g.

Finally, Lemma 2.31 describes an algorithm that for a spherical element $\gamma \in \overline{\Omega}F_{d-1}$ constructs a simplicial map $\gamma_{\mathsf{sph}}\colon \Sigma^d(\gamma) \to F$ such that $\pi_{d-1}(\overline{\Omega}F) \ni [\gamma] \simeq [\gamma_{\mathsf{sph}}] \in \pi_d(F)$.

The size of γ_{sph} is polynomial in $\mathrm{size}(\gamma)$. Hence, given a spherical $g \in \widetilde{GF}_{d-1}$, the algorithm produces $t(g)_{\mathsf{sph}}\colon \Sigma^d(t(g)) \to F$ that is exponential with respect to $\mathrm{size}(g)$.

Berger's model of the loop space.

Definition 2.27 (Oriented multigraph on X_n)**.** *Let X be a 0-reduced simplicial set. We define a directed multigraph $MX_n = (V_n, E_n)$, where the set of vertices $V_n = X_n$ and the set of edges E_n is given by*

$$E_n = \{[x,i]^\epsilon \mid x \in X_{n+1}, 0 \le i \le n, \epsilon \in \{1,-1\}\}.$$

We define maps $\mathsf{source}, \mathsf{target} : E_n \to V_n$ *by setting* $\mathsf{source}\,[x,i] = d_{i+1}x$, $\mathsf{target}\,[x,i] = d_i x$ *and* $\mathsf{source}\,[x,i]^{-1} = \mathsf{target}\,[x,i]$ *and* $\mathsf{target}\,[x,i]^{-1} = \mathsf{source}\,[x,i]$.

An edge $[x,i]^\epsilon \in E_n$ is called compressible*, if $x = s_i x'$ for some $x' \in X_n$.*

Definition 2.28 (Paths)**.** *Let $X \in \mathsf{sSet}$. A sequence of edges in MX_n*

$$\gamma = [x_1, i_1]^{\epsilon_1}[x_2, i_2]^{\epsilon_2} \cdots [x_k, i_k]^{\epsilon_k} \tag{2.5}$$

*is called an n-*path*, if* $\mathsf{target}\,[x_j, i_j]^{\epsilon_j} = \mathsf{source}\,[x_{j+1}, i_{j+1}]^{\epsilon_{j+1}}$, $1 \le j < k$.

Moreover, for every $x \in V_n = X_n$ we define a path of length zero 1_x with the property $\mathsf{source}\,1_x = x = \mathsf{target}\,1_x$ *and relations $a1_x = a$ whenever* $\mathsf{target}\,a = x$ *and $1_x b = b$ whenever* $\mathsf{source}\,b = x$.

The set of paths on MX_n is denoted by IX_n. Let $\gamma \in IX_n$ by as in (2.5). We define $\mathsf{source}\,\gamma = \mathsf{source}\,[x_1,i_1]^{\epsilon_1}$ *and* $\mathsf{target}\,\gamma = \mathsf{target}\,[x_k,i_k]^{\epsilon_k}$. *The* inverse *of γ, denoted γ^{-1}, is defined as*

$$\gamma^{-1} = [x_k, i_k]^{-\epsilon_k} \cdots [x_1, i_1]^{-\epsilon_1}.$$

if $\gamma = 1_x$, then $\gamma^{-1} = \gamma$. Note that each path is either equal to 1_x for some x or can be represented in a form such as (2.5), without any units.

For algorithmic purposes, we assume that a path $\gamma = [x_1, i_1]^{\epsilon_1}[x_2, i_2]^{\epsilon_2} \cdots [x_k, i_k]^{\epsilon_k}$ is represented as a list of triples (x_j, i_j, ϵ_j) and has size

$$\mathrm{size}(\gamma) := \sum_j \mathrm{size}(x_j) + \mathrm{size}(i_j) + \mathrm{size}(\epsilon_j),$$

which is bounded by a linear function in $\sum_j \mathrm{size}(x_j)$.

Given an edge $[x,i]^\epsilon \in MX_n$, we define operators

$$d_0, \dots d_n \colon E_n \to IX_{n-1} \text{ and } s_0, \dots, s_n \colon E_n \to IX_{n+1}$$

called *face* and *degeneracy* operators, respectively. These are given as follows

$$d_j[x,i]^\epsilon = \begin{cases} [d_j x, i-1]^\epsilon, & j<i; \\ 1_{d_i d_{i+1} x}, & i=j; \\ [d_{j+1}x, i]^\epsilon, & j>i. \end{cases} \qquad s_j[x,i]^\epsilon = \begin{cases} [s_j x, i+1]^\epsilon, & j<i; \\ [s_i x, i+1][s_{i+1}x, i])^\epsilon, & i=j; \\ [s_{j+1}x, i]^\epsilon, & j>i. \end{cases}$$

One can now extend the definition of face and degeneracy operators to paths, i.e. we define operators $d_0, \dots d_n \colon IX_n \to IX_{n-1}$ and $s_0, \dots, s_n \colon IX_n \to IX_{n+1}$

$$d_j\gamma = \begin{cases} d_j([x_1,i_1]^{\epsilon_1})d_j([x_2,i_2]^{\epsilon_2})\cdots d_j([x_k,i_k]^{\epsilon_k}) & \text{if} \quad \gamma = [x_1,i_1]^{\epsilon_1}[x_2,i_2]^{\epsilon_2}\cdots[x_k,i_k]^{\epsilon_k}, \\ 1_{d_j x} & \text{if} \quad \gamma = 1_x, x \in X_n. \end{cases}$$

$$s_j\gamma = \begin{cases} s_j([x_1,i_1]^{\epsilon_1})s_j([x_2,i_2]^{\epsilon_2})\cdots s_j([x_k,i_k]^{\epsilon_k}) & \text{if} \quad \gamma = [x_1,i_1]^{\epsilon_1}[x_2,i_2]^{\epsilon_2}\cdots[x_k,i_k]^{\epsilon_k} \\ 1_{s_j x} & \text{if} \quad \gamma = 1_x, x \in X_n. \end{cases}$$

With the operators defined above, one can see that IX is in fact a simplicial set.

For any $\gamma, \gamma' \in IX$ such that $\mathsf{target}\,\gamma = \mathsf{source}\,\gamma'$, we define a composition $\gamma \cdot \gamma'$ in an obvious way.

If the simplicial set X is 0-reduced, we denote the unique basepoint $* \in X_0$. Abusing the notation, we denote the iterated degeneracy of the basepoint $\underbrace{s_0 \cdots s_0 *}_{k-\text{times}} \in X_k$ by $*$ as well. With that in mind, we define simplicial subsets PX, ΩX of IX as follows:

$$PX = \{\gamma \in IX \mid \mathsf{target}\,\gamma = *\} \quad \Omega X = \{\gamma \in IX \mid \mathsf{source}\,\gamma = * = \mathsf{target}\,\gamma\}.$$

We remark that simplicial sets $PX, \Omega X$ intuitively capture the idea of pathspace and loopspace in a simplicial setting.

Definition 2.29. *A path $\gamma = [x_1,i_1]^{\epsilon_1}[x_2,i_2]^{\epsilon_2}\cdots[x_k,i_k]^{\epsilon_k} \in IX$ is called* reduced*, if for every $1 \le j < k$ the following condition holds:*

$$(x_j = x_{j+1} \,\&\, i_j = i_{j+1}) \Rightarrow \epsilon_j = \epsilon_{j+1}.$$

e.g. an edge in the path γ is never followed by its inverse.

An edge $[x,i]^\epsilon \in E_n$ is called compressible*, if $x = s_i x'$ for some $x' \in X_n$. A path is* compressed *if it does not contain any compressible edge.*

We define relation $\sim_R$ on IX (or rather on each IX_n) as a relation generated by

$$[x,i]^\epsilon[x,i]^{-\epsilon} \sim_R 1_{\mathsf{source}\,([x,i]^\epsilon)}, \quad n \in \mathbb{N}_0, [x,i]^\epsilon \in E_n.$$

Similarly, we define $\sim_C$ on IX as a relation generated by

$$[x,i]^\epsilon \sim_C 1_{\mathsf{source}\,([x,i]^\epsilon)}, \quad \text{if } [x,i]^\epsilon \in E_n \text{ is compressible.}$$

We finally define $\overline{I}X = (IX/\sim_C)/\sim_R$. Similarly, one defines $\overline{P}X, \overline{\Omega}X$.

For $\gamma, \gamma' \in IX_n$, we write $\gamma \sim \gamma'$ if they represent the same element in $\overline{I}X_n$. The symbol $\overline{\gamma}$, denotes the (unique) compressed and reduced path such that $\gamma \sim \overline{\gamma}$. One can see $\overline{I}X$ ($\overline{P}X, \overline{\Omega}X$) as the set of reduced and compressed paths in $IX(PX, \Omega X)$.

In a natural way, we can extend the definition of face and degeneracy operators d_i, s_i on sets $\overline{I}X(\overline{P}X,\overline{\Omega}X)$ by setting $d_i\gamma = \overline{d_i\gamma}$ and $s_i\gamma = \overline{s_i\gamma}$. One can check that this turns $\overline{I}X$, $\overline{P}X$ and $\overline{\Omega}X$ into simplicial sets.

Similarly, we define operation $\cdot\colon \overline{\Omega}X_n \times \overline{\Omega}X_n \to \overline{\Omega}X_n$ by $\gamma \cdot \gamma' \mapsto \overline{\gamma\gamma'}$, i.e. we first compose the loops and then assign the appropriate compressed and reduced representative. With the operation defined as above, $\overline{\Omega}X$ is a simplicial group.

Homomorphism $t\colon GX \to \overline{\Omega}X$. We first describe how to any given $x \in X_n$ assign a path $\gamma_x \in \overline{P}X_n$ with the property $\mathsf{source}\,\gamma_x = x$ and $\mathsf{target}\,\gamma_x = *$:

For $x \in X_n$, $n > 0$, the 0-reducedness of X gives us $d_{i_1}d_{i_2}\cdots d_{i_n}x = *$, here $i_j \in \{0,\ldots,j\}$, $0 < j \leq n$. In particular, $d_0d_1\cdots d_{n-1}x = *$. Using this, we define

$$\gamma_x = [s_nx, n-1][s_ns_{n-1}d_{n-1}x, n-2]\cdots[s_ns_{n-1}\cdots s_1d_1d_2\cdots d_{n-1}x, 0].$$

Ignoring the degeneracies, one can see the sequence of edges as a path

$$x \to d_{n-1}x \to d_{n-2}d_{n-1}x \to \cdots \to d_0d_1\cdots d_{n-1}x.$$

We define the homomorphism t on the generators of GX_n, i.e. on the elements $\overline{x}$, where $x \in X_{n+1}$ as follows:

$$t(\overline{x}) = \overline{\gamma^{-1}_{d_{n+1}x}[x,n]\gamma_{d_nx}}. \tag{2.6}$$

This is an element of $\overline{\Omega}X_n$.

The algorithm representing the map t has *exponential time complexity* due to the fact that an element $\overline{\sigma}^k$ with size $\mathrm{size}(\sigma) + \mathrm{size}(k)$ is mapped to

$$\overline{\underbrace{\gamma^{-1}_{d_{n+1}x}[x,n]\gamma_{d_nx} \;\;\ldots\;\; \gamma^{-1}_{d_{n+1}x}[x,n]\gamma_{d_nx}}_{k \text{ times}}}$$

which in general can have size proportional to k. Assuming an encoding of integers such that $\mathrm{size}(k) \simeq \ln(k)$, this amounts to an exponential increase.

Universal preimage of a path. Intuitively, one can think of the simplicial set IX of paths as of a discretized version of space of continuous maps $|X|^{[0,1]}$. In particular, $\gamma \in IX_{d-1}$ is a walk through a sequence of d-simplices in X that connect $\mathsf{source}\,\gamma$ with $\mathsf{target}\,\gamma$. However, in the continuous case an element $\mu \in |X|^{[0,1]}$ corresponds to a continuous map $\mu\colon [0,1] \to |X|$. We want to push the parallels further, namely, given any nontrivial[12] $\gamma \in IX_{d-1}$, we aim to define a simplicial set $\mathsf{Dom}(\gamma)$ and a simplicial map $\gamma_{\mathsf{map}} : \mathsf{Dom}(\gamma) \to X$ with the following properties:

1. $|\mathsf{Dom}(\gamma)| = D^d$.

2. γ_{map} maps $\mathsf{Dom}(\gamma)$ to the set of simplices contained in the path γ.

We will utilize the following construction given in [6].

Definition 2.30. *Let $\gamma \in IX_{d-1}$. We define $\mathsf{Dom}(\gamma)$ and γ_{map} as follows. Suppose, that $\gamma = [y_1,i_1]^{\epsilon_1}[y_2,i_2]^{\epsilon_2}\cdots[y_k,i_k]^{\epsilon_k}$. For every edge $[y_j,i_j]^{\epsilon_j}$, let α_j be the simplicial map $\Delta^d \to y_j$ sending the nondegenerate d simplex in Δ^d to y_j.*

[12] By nontrivial we mean that $\gamma \neq 1_x$ for any $x \in X_{d-1}$.

We define Dom(γ) *as a quotient of the disjoint union of* k *copies of* Δ^d*:*

$$\mathsf{Dom}(\gamma) = \bigsqcup_{i=1}^{k} \Delta^d / \sim$$

where each copy of Δ^d *corresponds to a domain of a unique* α_j *and the relation is given by*

$$(\alpha_j)^{-1}\mathsf{target}\left([y_j, i_j]^{\epsilon_j}\right) \sim (\alpha_{j+1})^{-1}\mathsf{source}\left([y_{j+1}, i_{j+1}]^{\epsilon_{j+1}}\right).$$

The map γ_{map} *is induced by the collection of maps* $\alpha_1, \ldots, \alpha_k$*:*

$$\begin{array}{ccc} \bigsqcup_{i=1}^{k} \Delta^d & & \\ \downarrow & \searrow^{\alpha_1,\ldots,\alpha_k} & \\ \mathsf{Dom}(\gamma) & \xrightarrow{\gamma_{\mathsf{map}}} & X. \end{array}$$

We recall that simplicial set $\overline{I}X$ was defined as the set of "reduced and compressed" paths in IX. Similarly, one introduces a reduced and compressed versions of the construction Dom. As a final step we then get

Lemma 2.31 (Section 2.4 in [6]). *Let* $\gamma \in \overline{\Omega}X_{d-1}$ *such that* $d_i\gamma = 1 \in \overline{\Omega}X$ *for all* i*. Then the map* $\gamma_{\mathsf{map}}\colon \mathsf{Dom}(\gamma) \to X$ *factorizes through a simplicial set model of the sphere* $\Sigma^d(\gamma)$ *as follows:*

$$\begin{array}{ccc} \mathsf{Dom}(\gamma) & & \\ \downarrow & \searrow^{\gamma_{\mathsf{map}}} & \\ \Sigma^d(\gamma) & \xrightarrow{\gamma_{\mathsf{sph}}} & X. \end{array}$$

Further, $\pi_{d-1}(\overline{\Omega}X) \ni [\gamma] \simeq [\gamma_{\mathsf{sph}}] \in \pi_d(X)$*.*

We will not give the proof of correctness of Lemma 2.31 (it can be found in [6]). Instead, in the next section, we only describe the algorithmic construction of $\gamma_{\mathsf{sph}}\colon \Sigma^d(\gamma) \to X$ and give a running time estimate.

Algorithm from Lemma 2.31.

The algorithm accepts an element $\gamma \in \overline{\Omega}X_{d-1}$ such that $d_i\gamma = 1 \in \overline{\Omega}X$ for all i, a spherical element. We divide the algorithm into four steps that correspond to the four step factorization in the following diagram:

$$\begin{array}{ccc} \mathsf{Dom}(\gamma) & & \\ \downarrow & \searrow^{\gamma_{\mathsf{map}}} & \\ \overline{\mathsf{Dom}}(\gamma) & \searrow^{\gamma_{\mathsf{c}}} & \\ \downarrow & & \\ \overline{\overline{\mathsf{Dom}}}(\gamma) & \xrightarrow{\gamma_{\mathsf{cr}}} & X \\ \downarrow & \nearrow_{\gamma_{\mathsf{sph}}} & \\ \Sigma^d(\gamma) & & \end{array}$$

$\mathsf{Dom}(\gamma)$: We interpret γ as an element in IX and construct $\gamma_{\mathsf{map}}\colon \mathsf{Dom}(\gamma) \to X$. This is clearly linear in the size of γ.

$\overline{\mathsf{Dom}}(\gamma)$: The algorithm checks, whether an edge $[y, j]^{\epsilon}$ in $d_{i_1} d_{i_2} \dots d_{i_\ell} \gamma$, where $0 \leq i_1 < i_2 < \dots < i_\ell < (d - \ell - 2)$ is *compressible*, i.e. $y = s_j d_j y$. If this is the case, add a corresponding relation on the preimages: $\alpha^{-1}(y) \sim s_j d_j \alpha^{-1}(y)$. Factoring out the relations, we get a map $\gamma_{\mathsf{c}}\colon \overline{\mathsf{Dom}}(\gamma) \to X$.

Although the number of faces we have to go through is exponential in d, this is not a problem, since d is deemed as a constant in the algorithm and so is 2^d. Hence the number of operations is again linear in the size of γ.

$\overline{\overline{\mathsf{Dom}}}(\gamma)$: Let $k < d$. We know that $\overline{d_k \gamma} = 1_*$, so after removing all compressible elements from the path $d_k \gamma$, it will contain a sequence of pairs $([y_i, j_i]^{\epsilon_i}, [y_i, j_i]^{-\epsilon_i})$ such that, after removing all $[y_u, j_u]^{\pm 1}$ for all $u < v$, then $[y_v, j_v]^{\epsilon_v}$ and $[y_v, j_v]^{-\epsilon_v}$ are next to each other.[13] Each such pair $([y_i, j_i]^{\epsilon_i}, [y_i, j_i]^{-\epsilon_i})$ corresponds to a pair of indices (l_i, m_i) corresponding to the positions of those edges in $d_k \gamma$. These sequences are not unique, but can be easily found in time linear in $\mathrm{length}(\gamma)$. Then we glue $\alpha_{l_i}^{-1}(y_i)$ with $\alpha_{m_i}^{-1}(y_i)$ for all i. Performing such identifications for all k defines the new simplicial set $\overline{\overline{\mathsf{Dom}}}(\gamma)$.

$\Sigma^d(\gamma)$: It remains to identify $\alpha^{-1}(\mathsf{source}\,\gamma)$ and $\alpha^{-1}(\mathsf{target}\,\gamma)$ with the appropriate degeneracy of the (unique) basepoint. The resulting space $|\Sigma^d(\gamma)|$ is a d-sphere.

2.6 Polynomial-Time Loop Contraction in F_d

In this section, we show that simplicial sets F_k, $2 \leq k \leq d$ constructed algorithmically in Section 2.3 have polynomial-time contractible loops, thus proving Lemma 2.14. We first give the contraction on F_2 and show that the contraction $F_i, i > 3$ follows from the contraction on F_3. The majority of the effort in this section is then concentrated on the description of the contraction c_0 on F_3.

Notation. We will further use the following shorthand notation: For a 0-reduced simplicial set X we will denote the iterated degeneracy $s_0 \cdots s_0 *$ of its unique basepoint $*$ by $*$ and we set $\pi_i = \pi_i(X)$. For any Eilenberg-Maclane space $K(\pi_i, i-1)$, $i \geq 2$, we denote its basepoint and its degeneracies by 0. From the context, it will always be clear which simplicial set we refer to.

Loop contraction on F_2. Assuming that X is a 0-reduced, 1-connected simplicial set with a given algorithm that computes the contraction on loops $c_0\colon (GX)_0 \to (GX)_1$, the contraction c_0 on F_2 is automatically defined, as $X = F_2$.

Loop contraction on F_i, $i > 3$. Suppose we have defined the contraction on the generators of $G_0(F_3)$. i.e. for any $(x, k) \in (X \times_{\tau'} K(\pi_2, 1))_1$ we have

$$c_0(\overline{(x,k)}) = \overline{(x_1, k_1)}^{\epsilon_1} \cdots \overline{(x_n, k_n)}^{\epsilon_n} \qquad (x_j, k_j) \in (F_3)_2, \epsilon_j \in \mathbb{Z}, 1 \leq j \leq n$$

[13]For example, $[a,1][b,2][b,2]^{-1}[a,1]^{-1}$ can be split into a sequence $([b,2],[b,2]^{-1}), ([a,1],[a,1]^{-1})$.

such that $d_0c_0(\overline{(x,k)}) = \overline{(x,k)}$ and $d_1c_0(\overline{(x,k)}) = 1$. In detail, we get the following:

$$\overline{(x,k)} = d_0c_0(\overline{(x,k)}) = \overline{(d_0x_1, d_0k_1)}^{\epsilon_1} \cdots \overline{(d_0x_n, d_0k_n)}^{\epsilon_n} \tag{2.7}$$

$$1 = d_1c_0(\overline{(x,k)}) = \left(\overline{(d_2x_1, \tau'(x_1)d_2k_1)}^{-1} \cdot \overline{(d_1x_1, d_1k_1)}\right)^{\epsilon_1} \cdots \tag{2.8}$$
$$\left(\overline{(d_2x_n, \tau'(x_n)d_2k_n)}^{-1} \cdot \overline{(d_1x_n, d_1k_n)}\right)^{\epsilon_n}$$

We now aim to give a reduction on the generators of $G_0(F_i)$, $i > 3$. Simplicial set F_i is an iterated twisted product of the form

$$\big(((X \times_{\tau'} K(\pi_2, 1)) \times_{\tau'} K(\pi_3, 2)) \times_{\tau'} \cdots \times_{\tau'} K(\pi_{i-2}, i-3)\big) \times_{\tau'} K(\pi_{i-1}, i-2)$$

As simplicial sets $K(\pi_{i-1}, i-2)$ are 1-reduced for $i > 3$, we can identify elements of $(F_i)_1$ with vectors $(x, k, 0, \dots, 0)$, where $k \in K(\pi_2, 1)_1, x \in X_1$. We further shorthand the series of $i-3$ zeros in the vector with $\mathbf{0}$. Hence generators $G_0(F_i)$ are of the form $\overline{(x, k, \mathbf{0})}$. The 1-reducedness also implies that $\tau'(\alpha) = 0$ whenever $\alpha \in (F_i)_2$, $i > 2$.

Finally, we set

$$c_0(\overline{(x,k,\mathbf{0})}) = \overline{(x_1, k_1, \mathbf{0})}^{\epsilon_1} \cdots \overline{(x_n, k_n, \mathbf{0})}^{\epsilon_n} \qquad (x_j, k_j, \mathbf{0}) \in (F_i)_2, \epsilon_j \in \mathbb{Z}, 1 \le j \le n$$

The (almost) freeness of $G_0(F_i)$, the fact that $K(\pi_{i-1}, i-2)$ are 1-reduced for $i > 3$ and equations (2.7), (2.8) give that $d_0c_0(\overline{(x,k,\mathbf{0})}) = \overline{(x,k,\mathbf{0})}$ and $d_1c_0(\overline{(x,k,\mathbf{0})}) = 1$.

Before the definition of contraction on simplicial set F_3, we remind the basic facts involving the simplicial model of Eilenberg-MacLane spaces we are using.

Eilenberg–MacLane spaces. As noted in Section 2.2, given a group π and an integer $i \ge 0$ an Eilenberg–MacLane space $K(\pi, i)$ is a space satisfying

$$\pi_j(K(\pi, i)) = \begin{cases} \pi & \text{for } j = i, \\ 0 & \text{else.} \end{cases}$$

In the rest of this section, by $K(\pi, i)$ we will always mean the simplicial model which is defined in [53, page 101]

$$K(\pi, i)_q = Z^i(\Delta^q; \pi),$$

where $\Delta^q \in \mathsf{sSet}$ is the standard q-simplex and Z^i denotes the cocycles. This means that each q–simplex is regarded as a labeling of the i–dimensional faces of Δ^q by elements of π such that they add up to $0 \in \pi$ on the boundary of every $(i+1)$-simplex in Δ^q, hence elements of $K(\pi, q)_q$ are in bijection with elements of π. The boundary and degeneracy operators in $K(\pi, k)$ are given as follows: For any $\sigma \in K(\pi, i)_q$, $d_j(\sigma) \in K(\pi, k)_{q-1}$ is given by a restriction of $\sigma \in K(\pi, i)$ to the j-th face of Δ^q. To define the degeneracy we first introduce mapping $\eta_j \colon \{0, 1, \dots, q+1\} \to \{0, 1, \dots, q\}$ given by

$$\eta_j(\ell) = \begin{cases} \ell & \text{for } \ell \le j, \\ \ell - 1 & \text{for } \ell > j. \end{cases}$$

Every mapping η_j defines a map $C^*(\eta_j) \colon C^*(\Delta^q) \to C^*(\Delta^{q+1})$.The degeneracy $s_j\sigma$ is now defined to be $C^*(\eta_j)(\sigma)$ (see [53, § 23]).

It follows from our model of Eilenberg-MacLane space, that elements of $K(\pi_2, 1)_2$ can be identified with labelings of 1-faces of a 2-simplex by elements of π_2 that sum up to zero.

As π_2 is an Abelian group, we use the additive notation for π_2. We identify the elements of $K(\pi_2,1)_2$ with triples (k_0,k_1,k_2), $k_i\in\pi_2$, $0\le i\le 2$, such that $k_0-k_1+k_2=0\in\pi_2$.

Loop contraction on F_3. Let X be a 0-reduced, 1-connected simplicial set with a given algorithm that computes the contraction on loops $c_0\colon (GX)_0\to (GX)_1$.

In the rest of the section, we will assume $x\in X_1$. Then by our assumptions $c_0\overline{x}=\overline{y_1}^{\epsilon_1}\cdots\overline{y_n}^{\epsilon_n}$, where $y_i\in X_2, \epsilon_i\in\mathbb{Z}$, $1\le i\le n$. Let $k_i=\tau'(y_i)$.

We first show that in order to give a contraction on elements of the form $\overline{(x,0)}$ and $\overline{(x,k)}$, it suffices to have the contraction on elements of the form $\overline{(*,k)}$:

Contraction on element $(x,0)$. Let $\overline{(x,0)}\in G_0(F_3)$. We define

$$c_0\overline{(x,0)}=\prod_{i=1}^{n}\left(c_0\overline{(*,k_i)}^{-1}\overline{(s_1d_2y_i,(k_i,k_i,0))}\cdot\overline{(y_i,0)}\right)^{\epsilon_i}.$$

Contraction on element (x,k). Suppose $\overline{(x,k)}\in(GF_3)_0$. The formula for the contraction is given using the formulae on contraction on $\overline{(x,0)}$ and $\overline{(*,k)}$ as follows

$$c_0\overline{(x,k)}=\overline{(s_0x,(k,0,-k))}\cdot s_0\overline{(x,0)}^{-1}\cdot s_0\overline{(*,-k)}\cdot c_0(\overline{(*,-k)})^{-1}\cdot c_0(\overline{(x,0)})$$

Contraction on element $(*,k)$. We formalize the existence of the contraction as Proposition 2.35 given at the end of this section. Due to the fact that the proof is rather technical, we need to define and prove some preliminary results first:

Definition 2.32. *Let $Z=\{z\in(GF_3)_1\mid d_0z=1\}$ and let $W=\{d_1z\mid z\in Z\}$ We define an equivalence relation $\sim$ on the elements of W in the following way: We say that $w\sim w'$ if there exists $z\in Z$, $\alpha,\beta\in(GF_3)_1$ such that $d_1z=w$, $\alpha z\beta\in Z$ and $d_1(\alpha z\beta)=w'$.*

Lemma 2.33. *Let $w\in W$ such that*

1. *$w=\overline{(x,k)}^{\epsilon}\cdot\alpha$, where $\alpha\in(GF_3)_1$ Then $w=\overline{(x,k)}^{\epsilon}\cdot\alpha\sim\alpha\cdot(x,k)^{\epsilon}=w'$.*
2. *$w=\overline{(*,k)}^{\epsilon}\cdot\alpha$, where $\alpha\in(GF_3)_0$. Then $w\sim w'=\overline{(*,-k)}^{-\epsilon}\cdot\alpha$.*
3. *$w=\overline{(*,-k)}^{-1}(x,0)\cdot\alpha$, where $\alpha\in(GF_3)_0$. Then $w\sim w'=\overline{(x,k)}\cdot\alpha$.*
4. *$w=\overline{(x,0)}^{-1}\overline{(x,k)}\cdot\alpha$, where $\alpha\in(GF_3)_0$. Then $w\sim w'=\overline{(*,k)}\cdot\alpha$.*
5. *$w=\overline{(*,-l)}^{-1}\overline{(*,k)}\cdot\alpha$, where $\alpha\in(GF_3)_0$. Then $w\sim w'=\overline{(*,k+l)}\cdot\alpha$.*

Proof. In all cases, we assume $z\in Z$ such that $d_1z=w$ and we give a formula for $z'\in Z$ with $d_1z'=w'$:

1. $z'=s_0\overline{(x.k)}^{-\epsilon}\cdot z\cdot s_0\overline{(x,k)}^{\epsilon}$.
2. $z'=\overline{(*,(k,0,-k))}^{\epsilon}\cdot(s_0\overline{(*,k)})^{-\epsilon}\cdot z$.

3. $z' = (s_0\overline{(x,k)}) \cdot \overline{(s_0x,(k,0,-k))}^{-1} \cdot z.$

4. $z' = (s_0\overline{(*,k)})\overline{(s_1x,(k,k,0))}^{-1} \cdot z.$

5. $z' = \overline{(s_0(*,k+l))(*,(k+l,k,-l))}^{-1} \cdot z.$

□

Lemma 2.34. *Let $z \in (GF_3)_1$, $z \in Z$ with*

$$d_1z = w = \overline{(*,-k_1)}^{-1} \cdot \overline{(x_1,0)}^{\epsilon_1} \cdots \overline{(*,-k_n)}^{-1} \cdot \overline{(x_n,0)}^{\epsilon_n}$$

where $\overline{x_1}^{\epsilon_1} \cdots \overline{x_n}^{\epsilon_n} = 1$ in GX_0, $x_i \in X$, $k_i \in \pi_2(X)$, $\epsilon_i \in \{1,-1\}$, $1 \le i \le n$. Then $w \sim \overline{(\sum_{i=1}^n k_i, *)}$.

Proof. We achieve the proof using a sequence of equivalences given in Lemma 2.33. Without loss of generality we can assume that $x_1 = x_2^{-1}$ and $\epsilon_1, \epsilon_2 = 1$ (If this is not the case, we can use rule (1) and/or relabel the elements). Using (1) gives us

$$\begin{aligned} w =&\overline{(*,-k_1)}^{-1} \cdot \overline{(x_2,0)}^{-1} \cdot \overline{(*,-k_2)}^{-1} \cdot \overline{(x_2,0)} \cdots \overline{(*,-k_n)}^{-1} \cdot \overline{(x_n,0)}^{\epsilon_n} \\ \sim&\overline{(*,-k_2)}^{-1} \cdot \overline{(x_2,0)} \cdots \overline{(*,-k_n)}^{-1} \cdot \overline{(x_n,0)}^{\epsilon_n} \cdot \overline{(*,-k_1)}^{-1} \cdot \overline{(x_2,0)}^{-1}. \end{aligned}$$

Then successive use of (3),(1),(4), (1) and finally (5) gives us

$$\begin{aligned} w \sim&\overline{(x_2,k_2)} \cdots \overline{(*,-k_n)}^{-1} \cdot \overline{(x_n,0)}^{\epsilon_n} \cdot \overline{(*,-k_1)}^{-1} \cdot \overline{(x_2,0)}^{-1}. \\ \sim&\overline{(x_2,0)}^{-1} \cdot \overline{(x_2,k_2)} \cdots \overline{(*,-k_n)}^{-1} \cdot \overline{(x_n,0)}^{\epsilon_n} \cdot \overline{(*,-k_1)}^{-1} \\ \sim&\overline{(*,k_2)} \cdots \overline{(*,-k_n)}^{-1} \cdot \overline{(x_n,0)}^{\epsilon_n} \cdot \overline{(*,-k_1)}^{-1} \\ \sim&\overline{(*,k_1+k_2)} \cdot \overline{(*,-k_3)}^{-1} \cdot \overline{(x_3,0)} \cdots \overline{(*,-k_n)}^{-1} \cdot \overline{(x_n,0)}^{\epsilon_n} \end{aligned}$$

multiple use or rules (2) and (1) and gives us

$$w \sim \overline{(*,-k_1-k_2-k_3)}^{-1} \cdot \overline{(x_3,0)} \cdots \overline{(*,-k_n)}^{-1} \cdot \overline{(x_n,0)}^{\epsilon_n}$$

So far, we have produced some element $z' \in Z \subseteq (GF_3)_1$ such that $d_0z' = 1$,

$$d_1z' = \overline{(*,-k_1-k_2-k_3)}^{-1} \cdot \overline{(x_3,0)} \cdots \overline{(*,-k_n)}^{-1} \cdot \overline{(x_n,0)}^{\epsilon_n}$$

and further $\overline{x_3}^{\epsilon_3} \cdots \overline{x_n}^{\epsilon_n} = 1$ in GX_0.

It follows that the construction described above can be applied iteratively until all elements $\overline{(x_i,0)}$ are removed and we obtain $w \sim \overline{(-\sum_{i=1}^n k_i, *)}^{-1} \sim \overline{(\sum_{i=1}^n k_i, *)}$. □

Proposition 2.35. *Let $k \in \pi_2(X)$. Then there is an algorithm that computes an element $z \in (GF_3)_1$ such that $d_0z = \overline{(*,k)}$ and $d_1z = 1$.*

Proof. Given an element $k \in \pi_2 \cong H_2(X)$, one can compute a cycle $\gamma \in Z_2(X)$ such that

$$[\gamma] = k \in \pi_2(X) \cong H_2(X) \cong H_2(K(\pi_2,2)) \cong \pi_2(K(\pi_2,2)),$$

were the middle isomorphism is induced by φ_2 and the other isomorphisms follow from Hurewicz theorem.

If one considers $\gamma \in \widetilde{AX}_1$ then by Lemma 2.26 one can algorithmically compute a spherical element $\gamma' = \overline{y_1}^{\epsilon_1} \cdots \overline{y_n}^{\epsilon_n} \in \widetilde{GX}_1$ where $y_i \in X_2$ and $\tau' y_i = k_i \in \pi_2(X)$, such that $d_0\gamma' = 1 = d_1\gamma'$ and $\sum_{i=1}^n \epsilon_i \cdot k_i = k$.

We define $z' \in (GF_3)_1$ by

$$z' = (\prod_{i=1}^n \overline{(s_0 d_0 y_i, (k_i, 0, -k_i))}^{\epsilon_i}) \cdot (\prod_{i=1}^n \overline{(y_i, (k_i, 0, -k_i))}^{\epsilon_i})^{-1}.$$

Observe that $d_0(z') = 1$ and

$$d_1 z' = \left(\overline{(*, -k_1)}^{-1} \cdot \overline{(d_0 y_1, 0)}\right)^{\epsilon_1} \cdots \left(\overline{(*, -k_n)}^{-1} \cdot \overline{(d_0 y_n, 0)}\right)^{\epsilon_n}.$$

We apply Lemma 2.34 on z' and get an element $z'' \in (GF_3)_1$ with the property $d_0 z'' = 1$ and $d_1 z'' = \overline{(*, k)}$. We define $z = s_0\overline{(*, k)} \cdot (z'')^{-1}$. Thus $d_0 z = \overline{(*, k)}$ and $d_1 z = 1$. □

Computational complexity. We first observe that that formulas for c_0 on a general element $\overline{(x, k)}$ depend polynomially on the size of $c_0(\overline{x})$ and the size of contractions on $\overline{(*, k)}$. Hence it is enough to analyse the complexity of the algorithm described in Proposition 2.35:

The computation of γ' is obtained by the polynomial-time Smith normal form algorithm presented in [41] and the polynomial-time algorithm in Lemma 2.26. The size of z' depends polynomially (in fact linearly) on size of γ'. The algorithm described in Lemma 2.34 runs in a linear time in the size of z'.

To sum up, the algorithm computes the formula for contraction on the elements of GF_i in time polynomial in the input ($\operatorname{size} X + \operatorname{size} c_0(GX)$).

2.7 Reconstructing a Map to the Original Simplicial Complex

This section contains the proof of Lemma 2.16.

Edgewise subdivision of simplicial complexes. In [20], the authors present, for $k \in \mathbb{N}$, the *edgewise subdivision* $\operatorname{Esd}_k(\Delta^m)$ of an m-simplex Δ^m that generalizes the two-dimensional sketch displayed in Figure 2.3. This subdivision has several nice properties: in particular, the number of simplices of $\operatorname{Esd}_k(\Delta^m)$ grows polynomially with k. Explicitly, the subdivision can be represented as follows.

- The vertices of $\operatorname{Esd}_k(\Delta^m)$ are labeled by coordinates $(a_0, \ldots, a_m)$ such that $a_j \geq 0$ and $\sum_j a_j = k$.
- Two vertices $(a_0, \ldots a_m)$ and $(b_0, \ldots, b_m)$ are *adjacent*, if there is a pair $j < k$ such that $|b_j - a_j| = |b_k - a_k| = 1$ and $a_i = b_i$ for $i \neq j, k$.
- Simplices of $\operatorname{Esd}_k(\Delta^m)$ are given by tuples of vertices such that each vertex of a simplex is adjacent to each other vertex.

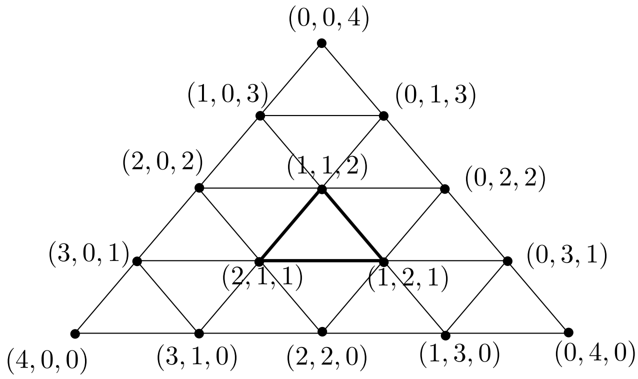

Figure 2.3: Edgewise subdivision of a 2-simplex for $k = 4$. In this case, there exists a copy of the 2-simplex completely in the "interior", defined by vertices $(2, 1, 1)$, $(1, 2, 1)$ and $(1, 1, 2)$. All other vertices are at the "boundary": more formally, their coordinatates contain a zero.

We define the *distance* of two vertices to be the minimal number of edges between them. An edgewise k-subdivision of Δ^m induces an edgewise k-subdivision of all faces, hence we may naturally define an edgewise subdivision of any simplicial complex.

Constructing the map $\mathrm{Esd}_k(\Sigma) \to X^{sc}$. Let R be a chosen root in the tree T. We denote the tree-distance of a vertex W from R by $\mathrm{dist}_T(W)$. Let

$$l := \max\{\mathrm{dist}_T(V) : \ V \text{ is a vertex of } X^{sc}\}$$

be the maximal tree-distance of some vertex from R. For each vertex V of X^{sc}, there is a unique path in the spanning tree that goes from R into V. Further, we define the maps $M(j) : (X^{sc})^{(0)} \to (X^{sc})^{(0)}$ from vertices of X^{sc} into vertices of X^{sc} such that

- $M(j)(V) := V$ if $j \geq \mathrm{dist}_T(V)$, and
- $M(j)(V)$ is the vertex on the unique tree-path from R to V that has tree-distance j from R, if $j < \mathrm{dist}_T(V)$.

If, for example, $R - U - V - W$ is a path in the tree, then $M(0)(W) = R$, $M(1)(W) = U$ etc. Clearly, $M(l) = M(l+1) = \dots$ is the identity map, as l equals the longest possible tree-distance of some vertex.

Assume that d is the dimension of Σ and $k := l(d+1)+1$. We will define $f' : \mathrm{Esd}_k(\Sigma) \to X^{sc}$ simplexwise. Let $\tau \in \Sigma$ be an m-simplex and $f(\tau) = \tilde{\sigma} \in X$ be its image in the simplicial set X. If σ is the degeneracy of the base-point $* \in X$, then we define $f'(x) := R$ for all vertices x of $\mathrm{Esd}_k(\tau)$: in other words, f' will be constant on the subdivision of τ. Otherwise, $\tilde{\sigma}$ is not the degeneracy of a point and has a unique lift $\sigma \in X^{ss}$. (Recall that $X := X^{ss}/T$.) Let $(V_0, \dots, V_m)$ be the vertices of σ (order given by orientation): these vertices are not necessarily different, as σ may be degenerate.

In the algorithm, we will need to know which faces of σ are in the tree T. We formalize this as follows: let $S \subseteq 2^m$ be the family of all subsets of $\{0, 1, \dots, m\}$ such that

- For each $\{i_0, \dots, i_j\} \in S$, $\{V_{i_0}, \dots, V_{i_j}\}$ is in the tree (that is, it is either an edge or a single vertex),

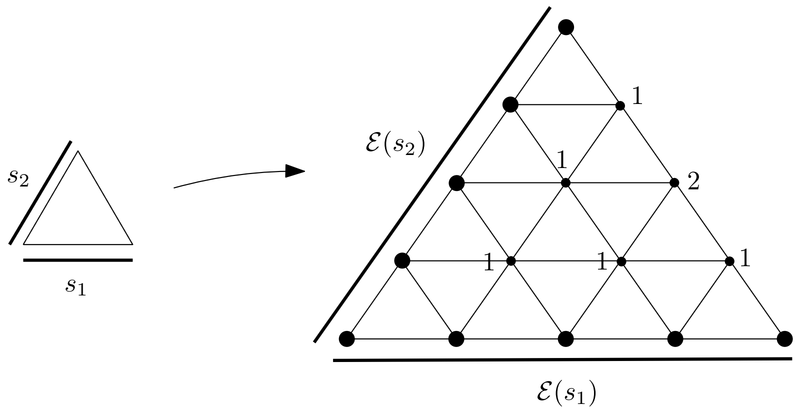

Figure 2.4: Illustration of extended faces. Here $S = \{s_1, s_2\}$ corresponds to the lower- and left-face of a 2-simplex. The extended faces $\mathcal{E}(s_1)$ and $\mathcal{E}(s_2)$ are sets of vertices of $\mathrm{Esd}_k(\Delta^2)$ that are on the lower- and left- boundary. The corresponding extended tree $\mathcal{E}(T)$ is the union of all these vertices. The integers indicate edge-distances dist_{ET} of vertices in $\mathrm{Esd}_k(\Delta^2)$ from $\mathcal{E}(T)$.

- Each set in S is maximal wrt. inclusion.

Elements of S correspond to maximal faces of σ that are in the tree, in other words, to faces of $\tilde{\sigma}$ that are degeneracies of the base-point.

Definition 2.36. *Let Δ^m be an oriented m-simplex, represented as a sequence of vertices $(e_0, \ldots, e_m)$. For any face $s \subseteq \{e_0, \ldots, e_m\}$, we define the* extended face $\mathcal{E}(s)$ in $\mathrm{Esd}_k(\Delta^m)$ *to be the set of vertices $(x_0, \ldots, x_m)$ in* $\mathrm{Esd}_k(\Delta^m)$ *that have nonzero coordinates only on positions i such that $e_i \in S$.*

The geometric meaning of this is illustrated by Figure 2.4.

Definition 2.37. *For $S \subseteq 2^m$, we define the* extended tree $\mathcal{E}(T)$ *to be the union of the extended faces $\mathcal{E}(s)$ in $\mathrm{Esd}_k(\Delta^m)$ for all $s \in S$. The edge-distance of a vertex x in $\mathrm{Esd}_k(\Delta^m)$ from $\mathcal{E}(T)$ will be denoted by $\mathrm{dist}_{ET}(x)$.*

In words, $\mathcal{E}(T)$ it is the union of all vertices in parts of the boundary of $\mathrm{Esd}_k(\Delta^m)$ that correspond to the faces of σ that are in the tree, see Fig. 2.4. The number $\mathrm{dist}_{ET}(x)$ is the distance to x from those boundary parts that correspond to faces of σ that are in the tree.

To define a simplicial map from $\mathrm{Esd}_k(\tau)$ to X^{sc}, we need to label vertices of $\mathrm{Esd}_k(\tau)$ by vertices of X^{sc} such that the induced map takes simplices in $\mathrm{Esd}_k(\tau)$ to simplices in X^{sc}. Recall that $V_0, \ldots, V_m$ are the vertices of σ. For $x = (x_0, \ldots, x_m)$, we denote by $\arg\max x$ the smallest index of a coordinate of x among those with maximal value (for instance, $\arg\max(4, 2, 1, 4, 0) = 0$, as the first 4 is on position 0). The geometric meaning of $V_{\arg\max x}$ is illustrated by Figure 2.5.

Now we are ready to define the map $f' : \mathrm{Esd}_k(\tau) \to X^{sc}$. It is defined on vertices x with coordinates $(x_0, \ldots, x_m)$ by

$$f'(x_0, \ldots, x_m) := M(\mathrm{dist}_{ET}(x))(V_{\arg\max\ x}). \tag{2.9}$$

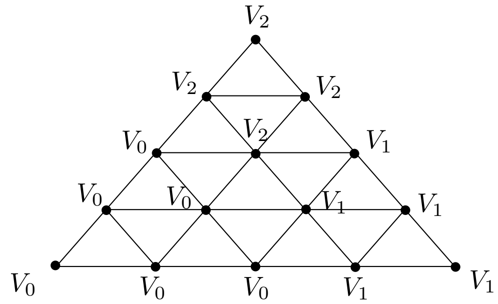

Figure 2.5: Labelling vertices of $\mathrm{Esd}_k(\Delta^2)$ by $V_{\arg\max x}$.

Geometrically, most vertices x will be simply mapped to V_j for which the j'th coordinate of x is dominant. In particular, a unique m-simplex "most in the interior of $\mathrm{Esd}_k(\tau)$" with coordinates

$$\begin{pmatrix} j+1 \\ j \\ \cdots \\ j \\ \hline j+1 \\ \cdots \\ j+1 \end{pmatrix}^T, \begin{pmatrix} j \\ j+1 \\ \cdots \\ j \\ \hline j+1 \\ \cdots \\ j+1 \end{pmatrix}^T, \ldots, \begin{pmatrix} j \\ j \\ \cdots \\ j+1 \\ \hline j+1 \\ \cdots \\ j+1 \end{pmatrix}^T, \begin{pmatrix} j \\ j \\ \cdots \\ j \\ \hline j+2 \\ \cdots \\ j+1 \end{pmatrix}^T, \ldots, \begin{pmatrix} j \\ j \\ \cdots \\ j \\ \hline j+1 \\ \cdots \\ j+2 \end{pmatrix}^T \tag{2.10}$$

for suitable j will be labeled by $V_0, V_1, \ldots, V_m$; in other words, it will be mapped to σ.[14]

However, vertices x close to those boundary parts of $\mathrm{Esd}_k(\tau)$ that correspond to the tree-parts of σ, will be mapped closer to the root R and all the extended tree $\mathcal{E}(T)$ will be mapped to R. One illustration is in Figure 2.6.

Computational complexity. Assuming that we have a given encoding of Σ, f, X, X^{sc} and a choice of T and R, defining a simplicial map $f' : \mathrm{Esd}_k(\Sigma) \to X^{sc}$ is equivalent to labeling vertices of $\mathrm{Esd}_k(\Sigma)$ by vertices of X^{sc}. Clearly, the maximal tree-distance l of some vertex depends only polynomially on the size of X^{sc} and can be computed in polynomial time, as well as the maps $M(0), \ldots, M(l)$. Whenever $j > l$, we can use the formula $M(j) = \mathrm{id}$. Further, $k = l(d+1)+1$ is linear in l, assuming the dimension d is fixed. If $\tau \in \Sigma$ is an m-simplex, then the number of vertices in $\mathrm{Esd}_k(\tau)$ is polynomial[15] in k, and their coordinates can be computed in polynomial time. Finding the lift σ of $f(\tau) = \tilde{\sigma}$ is at most a linear operation in $\mathrm{size}(X^{sc}) + \mathrm{size}(\tilde{\sigma})$. Converting $\sigma \in X^{ss}$ into an ordered sequence $(V_0, V_1, \ldots, V_m)$ amounts to computing its vertices $d_0 d_1 \ldots \hat{d_i} \ldots, d_m \sigma$, where d_i is omitted. Collecting information on faces of σ that are in the tree and the set of vertices $\mathcal{E}(T)$ is straight-forward: note that assuming fixed dimensions, there are only constantly many faces of each simplex to be checked. If $s = \{i_0, \ldots, i_j\}$ is a face, then the edge-distance of a vertex x from $\mathcal{E}(s)$ equals to $\sum_u x_{i_u}$. Applying formula (2.9) to x requires to compute the edge-distance of x from $\mathcal{E}(T)$: this equals to the minimum of the edge-distances of x from $\mathcal{E}(s)$ for all faces s of σ that are in the tree. Computing $\arg\max x$ is a trivial operation. Finally, the number of simplices τ of Σ is bounded by

[14] If $\dim(\tau) = d$ is maximal, then $j = l$ and this most-middle simplex has particularly nice coordinates $(l+1, l, \ldots, l), \ldots, (l, \ldots, l, l+1)$.

[15] Here the assumption on the fixed dimension d is crucial.

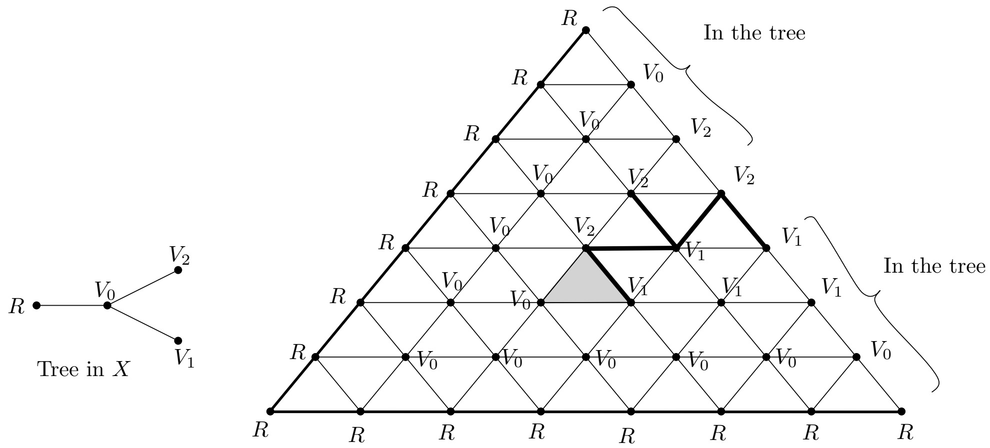

Figure 2.6: Example of the labeling induced by formula (2.9). We assume that $f(\tau) = \tilde{\sigma}$ where σ is a simplex of X^{sc} with three different vertices $V_0V_1V_2$. In this example, the tree connects $R - V_0 - V_1$ as well as $R - V_0 - V_2$ and the edge V_1V_2 is *not* in the tree. On the right, we give the induced labeling of vertices of $\mathrm{Esd}_k(\tau)$ which determines a simplicial map to X^{sc}. The bottom and left faces of σ are in the tree, hence the bottom and left extended faces in $\mathrm{Esd}_k(\tau)$ are all mapped into R. The right face of σ is the edge V_1V_2 that is not in the tree: the corresponding right extended face in $\mathrm{Esd}_k(\tau)$ is mapped to a loop $R - V_0 - V_1 - V_2 - V_0 - R$, where V_1V_2 is the only part that is *not* in the tree. The most interior simplex in $\mathrm{Esd}_k(\tau)$ is highlighted and is the only one mapped to σ.

the size of Σ, so applying (2.9) to each vertex of $\mathrm{Esd}_k(\Sigma)$ only requires polynomially many steps in $\mathrm{size}(\Sigma, f, X^{sc}, T, X)$.

Correctness. What remains is to prove that formula (2.9) defines a well-defined simplicial map and that $|\mathrm{Esd}_k(\Sigma)| \to |X^{sc}| \to |X|$ is homotopic to $|\Sigma| \to |X|$.

Lemma 2.38. *The above algorithm determines a well-defined simplicial map* $\mathrm{Esd}_k(\Sigma) \to X^{sc}$.

Proof. First we claim that formula (2.9) defines a global labeling of vertices of $\mathrm{Esd}_k(\Sigma)$ by vertices of X^{sc}. For this we need to check that if τ' is a face of τ, then (2.9) maps vertices of $\mathrm{Esd}_k(\tau')$ compatibly. This follows from the following facts, each of them easily checkable:

- If τ' is spanned by vertices of τ corresponding to $s \subseteq \{0, \dots, m\}$, then a vertex $x' := (x_0, \dots, x_j)$ in $\mathrm{Esd}_k(\tau')$ has coordinates x in $\mathrm{Esd}_k(\tau)$ equal to zero on positions $\{0 \dots, m\} \setminus s$ and to $x_0, \dots, x_j$ on other positions, successively.

- If $V'_k := V_{i_k}$ for $s = (i_0, \dots, i_j)$ are the vertices of the corresponding face of σ, then

$$V'_{\arg\max x'} = V_{\arg\max x}$$

- The extended tree $\mathcal{E}'(T)$ in $\mathrm{Esd}_k(\tau')$ equals the intersection of the extended tree in $\mathrm{Esd}_k(\tau)$ with $\mathcal{E}(\tau')$

- The distance $\mathrm{dist}_{ET}(x')$ in $\mathrm{Esd}_k(\tau')$ equals $\mathrm{dist}_{ET}(x)$ in $\mathrm{Esd}_k(\tau)$.

Further, we need to show that this labeling defines a well-defined simplicial map, that is, it maps simplices to simplices. We claim that each simplex in $\mathrm{Esd}_k(\tau)$ is mapped either to some subset of $\{V_0, \ldots, V_m\}$ or to some edge in the tree T, or to a single vertex.

We will show the last claim by contradiction. Assume that some simplex is *not* mapped to a subset of $\{V_0, \ldots, V_m\}$, and also it is *not* mapped to an edge of the tree and *not* mapped to a single vertex. Then there exist two vertices x and y in this simplex that are labeled by U and W in X^{sc}, such that either U or W is not in $\{V_0, \ldots, V_m\}$, UW is not in the tree, and $U \neq W$.

The fact that at least one of $\{U, W\}$ does not belong to $\{V_0 \ldots, V_m\}$, implies that $\mathrm{dist}_{ET}(x) < l$ or $\mathrm{dist}_{ET}(y) < l$ (as $M(j)$ maps each $V_{\arg\max x}$ to itself for $j \geq l$).

Without loss of generality, assume that $\arg\max x = 0$ and $\arg\max y = 1$. Then the coordinates of x and y are either

$$x = (j+1, j, x_3, \ldots, x_m), \quad y = (j, j+1, x_3, \ldots, x_m)$$

such that $x_i \leq j+1$ for all $i \geq 3$, or

$$x = (j, j, x_3, \ldots, x_m), \quad y = (j-1, j+1, x_3, \ldots, x_m)$$

for some j such that $x_i \leq j$ for all $i \geq 3$.

We claim that $V_0 \neq V_1$ and that the edge V_0V_1 is *not* in the tree. This is because there exists a tree-path from R via U to V_0 and also a tree-path from R via W to V_1 (and $U \neq W$): both $V_0 = V_1$ as well as a tree-edge V_0V_1 would create a circle in the tree. In coordinates, this means that vertices $(*, *, 0, 0, \ldots, 0)$ are not contained in $\mathcal{E}(T)$, apart of $(k, 0, 0, \ldots, 0)$ and $(0, k, 0, \ldots, 0)$. So, any vertex in $\mathcal{E}(T)$ has a zero on either the zeroth or the first coordinate. This immediately implies that $\mathrm{dist}_{ET}(x) \geq j$ and $\mathrm{dist}_{ET}(y) \geq j$. Keeping in mind that coordinates of x (and y) has to sum up to $k = l(d+1)+1$, the smallest possible value of j is $j = l$ (if $m = d$ is maximal), in which case $x = (l+1, l, l, \ldots, l)$ and $y = (l, l+1, \ldots, l)$. This choice, however, would contradict the fact that either $\mathrm{dist}_{ET}(x) < l$ or $\mathrm{dist}_{ET}(y) < l$. Therefore we have a strict inequality $j > l$. Finally, we derive a contradiction having either $\mathrm{dist}_{ET}(x) \geq j > l > \mathrm{dist}_{ET}(x)$, or a similar inequality for y.

This completes the proof that each simplex is either mapped to a subset of $\{V_0, \ldots, V_m\}$ or to an edge in the tree or to a single vertex: the image is a simplex in X^{sc} in either case. □

Lemma 2.39. *The geometric realisations of* $pf' : \mathrm{Esd}_k(\Sigma) \to X$ *and* $f : \Sigma \to X$ *are homotopic.*

Proof. First we reduce the general case to the case when all maximal simplices in Σ (wrt. inclusion) have the same dimension d. If this were not the case, we could enrich any lower-dimensional maximal simplex $\tau = \{x_0, \ldots, x_j\} \in \Sigma$ by new vertices $y^\tau_{j+1}, \ldots, y^\tau_d$ and produce a maximal d-simplex

$$\tilde{\tau} = \{x_0, \ldots, x_j, y^\tau_{j+1}, \ldots, y^\tau_d\}.$$

Thus we produce a simplicial complex $\tilde{\Sigma} \supseteq \Sigma$ with the required property. Whenever $f(\tau)$ is mapped to $\tilde{\sigma}$ where $\sigma = (V_0, \ldots, V_j)$, we define $f(\tilde{\tau})$ to be $s_j^{d-j}\tilde{\sigma}$, a degenerate simplex with lift $(V_0, \ldots, V_j, V_j, \ldots, V_j)$. The map $f' : \tilde{\Sigma} \to X^{sc}$ is constructed from

$f : \tilde{\Sigma} \to X$ as above and if we prove that $|f|$ is homotopic to $|pf'|$ as maps $|\tilde{\Sigma}| \to |X|$, it immediately follows that their restrictions are homotopic as maps $|\Sigma| \to |X|$ as well.

Further, assume that all maximal simplices have dimension d. Let $\tau \in \Sigma$ be a d-dimensional simplex and let τ^{int} be the simplex in $\mathrm{Esd}_k(\tau)$ spanned by the vertices

$$(l+1, l, \ldots, l), \ldots, (l, \ldots, l, l+1),$$

that is, the simplex in the interior of τ that is mapped by pf' to $\tilde{\sigma}$. Let $H_\tau(\cdot, 1) : |\tau| \to |\tau|$ be a linear map that takes $|\tau|$ linearly to $|\tau^{int}|$ via mapping the i'th vertex to $(l, \ldots, l+1, 1 \ldots, l)$ where the $l+1$ is on position i. Further, let H_τ be a linear homotopy $|\tau| \times [0,1] \to |\tau|$ between the identity $H_\tau(\cdot, 0) = \mathrm{id}$ and $H_\tau(\cdot, 1)$. The composition $|pf'|H_\tau$ then gives a homotopy $|\tau| \times [0,1] \to |X|$ between the restrictions $(|pf'|)|_{|\tau|}$ and $(|f|)|_{|\tau|}$. For a general $x \in |\Sigma|$, there exists a maximal d-simplex $|\tau|$ such that $x \in |\tau|$ and we define a homotopy

$$(x,t) \mapsto |pf'|H_\tau(x,t).$$

It remains to show that this map is independent on the choice of τ.

Let us denote the (ordered) vertices of τ by $\{v_0, v_1, \ldots, v_d\}$ and let $\delta \subseteq \tau$ be one of its faces: further, let w_i be the vertex of τ^{int} with barycentric coordinates $(l, \ldots, l, l+1, l, \ldots, l)/k$ in $|\tau|$ such that the $l+1$ is in position i. The homotopy H_τ sends points in $|\delta|$ onto the span of points w_i for which $v_i \in \delta$. For $y \in |\delta|$, the j-th barycentric coordinate of $H_\tau(y,t)$ is equal to $t\,(l/k)$ for each $j \notin \delta$. In particular, the j-th coordinate of $H_\tau(y,t)$ is between 0 and l/k for $j \notin \delta$, and hence it is not the "dominant" coordinate. It follows that each $z := H_\tau(x,t)$ is contained in the interior of a unique simplex Δ of $\mathrm{Esd}_k(\tau)$ such that $v_{\arg\max x} \in \delta$ for all vertices x of Δ.

Let $i_0 < i_1 \ldots < i_k$ be the indices such that $v_{i_j} \in \delta$ and $j_1 < \ldots < j_{d-k}$ be the remaining indices. Let $\tau' = (v'_0, \ldots, v'_d)$ be another d-simplex containing δ as a face. Assume, for simplicity, that the vertices of τ' are ordered so that vertices of δ have orders $i_0, \ldots, i_k$—such as it is in τ. Let σ, σ' be the lift of $f(\tau)$, $f(\tau')$ respectively, and V_i, V'_i the i-th vertex of σ, σ' respectively.

We define a "mirror" map $m : |\tau| \to |\tau'|$, which to a point with barycentric coordinates $(x_0, \ldots, x_d)$ with respect to τ assigns a point in $|\tau'|$ with the same barycentric coordinates with respect to τ'. Clearly, $H_{\tau'}(y,t) = m(H_\tau(y,t))$ for $y \in |\tau|$ and whenever z is in the interior of a simplex $\Delta \in \mathrm{Esd}_k(\tau)$, then $m(z)$ is in the interior of $m(\Delta)$, where vertices of Δ and $m(\Delta)$ have the same barycentric coordinates with respect to τ and τ', respectively. If, moreover, Δ is such that each of its vertices r have coordinates $\leq l/k$ on positions $j_1, \ldots, j_{d-k}$, then $V_{\arg\max r} = V'_{\arg\max m(r)}$.

To summarize these properties, $H_\tau(y,t)$ and $H_{\tau'}(y,t)$ satisfy that[16]

- they have the same coordinates wrt. τ, τ', respectively,
- they are in the interior of simplices $\Delta \in \mathrm{Esd}_k(\tau)$, $\Delta' \in \mathrm{Esd}_k(\tau')$ whose vertices have the same coordinates wrt. τ, τ', respectively,
- the $\arg\max$ labeling induces the same labeling of vertices of Δ, Δ' by vertices of δ, respectively.

[16] In general, vertices of δ may have different order in τ and τ' and the assumption on compatible ordering was chosen only to increase readability. If $i'_0 < \ldots < i'_k$ are such that $v'_{i'_j} = v_{i_j}$ (orders of δ-vertices wrt. τ') and $j'_1 < \ldots < j'_{d-k}$ are positions of the remaining vertices in τ', then m is defined so that it maps $x \in |\tau|$ with τ-coordinates $(x_0, \ldots, x_d)$ into $x' \in |\tau'|$ with coordinates $x'_{i'_j} = x_{i_j}$ and $x'_{j_k} = x_{j_k}$.

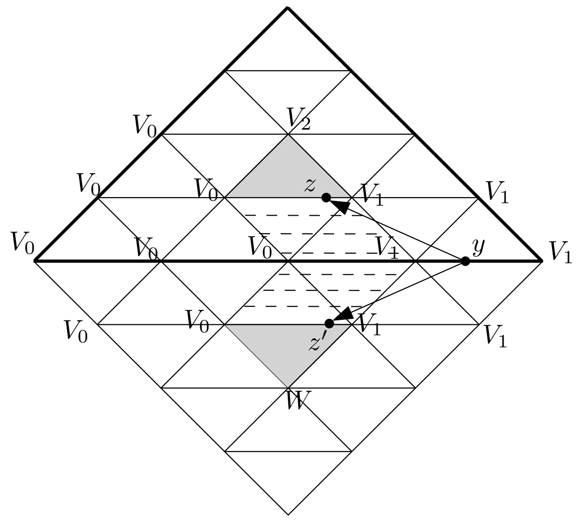

Figure 2.7: The homotopy H_τ takes y linearly into z and $H_{\tau'}$ takes y into z'. Due to the symmetry represented by the horizontal line, $|pf'|$ maps $H_\tau(y,t)$ into the same point of X as $|pf'|H_{\tau'}(y,t)$.

The map pf' takes each m-simplex Δ in $\mathrm{Esd}_k(\tau)$ with vertices t_u labeled by $V_{\arg\max t_u}$ onto $p(V_{\arg\max t_0},\ldots,V_{\arg\max t_m})$ and it follows from the above properties that $m(\Delta)$ is mapped to the same simplex. We conclude that $|pf'|H_\tau(y,t) = |pf'|H_{\tau'}(y,t)$ for each $y \in |\delta|$ and $t \in [0,1]$. $\square$

3 Embeddability of simplicial complexes is undecidable outside the meta-stable range

This chapter is a joint work with Marek Filakovský and Uli Wagner

In this chapter we present the complete proofs of Theorem 1.14 and Theorem 1.15. First, we recall their statements.

Theorem 1.14. $\mathrm{EXTEMBED}_{k\to d}$ *is* undecidable *for* $k \le d < \left\lfloor \frac{3(k+1)}{2} \right\rfloor$*,* $k \ge 5$ *and* $(k, d) \neq (5, 7)$*.*

Theorem 1.15. $\mathrm{EMBED}_{k\to d}$ *is* undecidable *for* $k \le d < \left\lfloor \frac{3(k+1)}{2} \right\rfloor$*,* $k \ge 5$ *and* $(k, d) \neq (5, 7)$*.*

As discussed in Chapter 1, the codimension 0 and 1 cases have already been proved [50; 55]. Thus, we will focus our attention to the case when $(d - k) \ge 2$.

Structure of the Chapter We begin by introducing the *extension problem* as defined in [12]. We then focus on a specific instance of this problem, which the authors prove is undecidable and provide a family of reductions from such instances to instance of $\mathrm{EXTEMBED}_{k\to d}$ for all pairs (k, d), satisfying the conditions of Theorem 1.14 for $(d - k) \ge 3$. By extending these arguments, we provide a similar proof of Theorem 1.15 for $(d - k) \ge 3$. Finally, we provide an additional construction to cover the cases when $(d - k) = 2$, thus completing the two proofs. In Section 3.4 we provide all the necessary background from PL topology and related topics and in Section 3.7, we prove the main technical statement, which is at the heart of the proofs of Theorem 1.14 and Theorem 1.15. Once we have introduced the techniques we use, we will also explain why our methods fail for some pairs (k, d), which lie at the boundary just outside of the meta-stable range.

The undecidable problem EXT_m

Definition 3.1 (EXT_m)**.** *Given finite simplicial complexes* $A \subset X$ *and* Y*, such that* $\dim X = 2m$ *and* Y *is* $(m-1)$*-connected, and a simplicial map* $f : A \to Y$*, decide whether there exists a continuous map* $X \to Y$ *extending* f*.*

The authors of [12] provide a construction, which translates a system of Diophantine equations into an extension problem of this type, so that an extension exists if and only if the system has an integer solution. Thus, they reduce a version of Hilbert's tenth problem to an extension problem. On the other hand, Hilbert's tenth problem is undecidable by a celebrated result of Matiyasevich [48], building on earlier work by Robinson, Davis and Putnam(we refer to [49] for further details), which yields the undecidability of the extension problem.

Theorem 3.2 (Theorem 1 (a) in [12])**.** *Let $m \geq 2$ be fixed. There is a fixed $(m-1)$-connected finite simplicial complex $Y = Y_m$ such that the following problem is algorithmically unsolvable: Given finite simplicial complexes $A \subseteq X$ with $\dim X = 2m$ and a simplicial map $f\colon A \to Y$, decide whether there exists a continuous map $F\colon X \to Y$ extending f. For m even, we can take Y_m to be the sphere S^m.*

It is important to note here, that the property of whether the map f extends or not depends solely on the homotopy class of f.

We will outline the construction in [12] without providing any details about its undecidability. Let $m \geq 2$ be the dimension parameter of the extension problem, as in Theorem 3.2. Since there is a slight difference between m even and odd, we will consider them separately.

m **even:** Following the notation in Theorem 3.2, we set $A = S^{2m-1} \vee \ldots \vee S^{2m-1}$, $Y = S^m$ and $X = \mathrm{Cyl}(\phi)$ to be the mapping cylinder of a map $\phi\colon A \to S^m \vee \ldots \vee S^m$. We will denote by $W := S^m \vee \ldots \vee S^m$ the target wedge sum . The two wedge sums are always finite and typically consist of different number of spheres. We don't explicitly enumerate them, since this will not be of any importance for our further arguments. The extension problem in this case is the following:

$$
\begin{array}{ccc}
S^{2m-1} \vee \ldots \vee S^{2m-1} & \longrightarrow & S^m \\
\big\downarrow & \nearrow & \\
\mathrm{Cyl}(\phi) & &
\end{array}
\qquad (3.1)
$$

m **odd:** In the odd case we only modify the target, by changing it from S^m to $S^m \vee S^m$.

Remark 3.2.1. In Chapter 6 in [12], the authors provide an algorithm, which for every instance of EXT_m, computes a representation of the corresponding topological spaces and maps between them as finite simplicial complexes and simplicial maps.

Formulation using a double mapping cylinder

For our purposes, it will prove convenient to reformulate the constructions above using *double mapping cylinders*, which we now introduce.

Definition 3.3. *Let A, W and Y be topological spaces and let $\phi\colon A \to W$ and $\psi\colon A \to Y$ be continuous maps. Then, the* double mapping cylinder *of the maps (ϕ, ψ), denoted by* DCyl(ϕ, ψ)*, is the topological space $((A \times I) \sqcup Y \sqcup W)/\sim$, where $(a, 0) \sim \psi(a)$ and $(a, 1) \sim \phi(a)$.*

Alternatively, one can see the double mapping cylinder as being obtained by glueing the mapping cylinders of ϕ and ψ along their base, i.e. DCyl$(\phi, \psi) =$ Cyl$(\phi) \cup_A$ Cyl(ψ).

Assume that $m \geq 2$ is even. Then, we can reformulate the extension problem 3.1 in the following way.

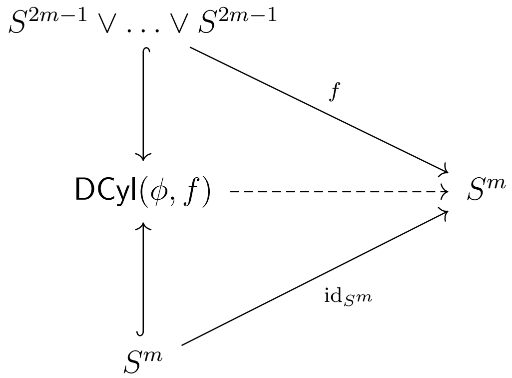

(3.2)

Proposition 3.4. *The extensions problems 3.1 and 3.2 are equivalent, i.e. there exists an extension in one of them if and only if there exists an extension in the other one.*

Proof. We prove the two directions of the equivalence separately.

- (3.2 $\Rightarrow$ 3.1) This direction is trivial, since we can consider the mapping cylinder Cyl(ϕ) as being a subspace of the double mapping cylinder DCyl$(\phi; f)$.
- (3.1 $\Rightarrow$ 3.2) If we think about the double mapping cylinder DCyl$((\phi; f)$ as being glued out of the mapping cylinders Cyl(ϕ) and Cyl(f), then the existence of an extension in 3.1 covers the upper half of DCyl$((\phi; f)$. The lower half looks in the following way.

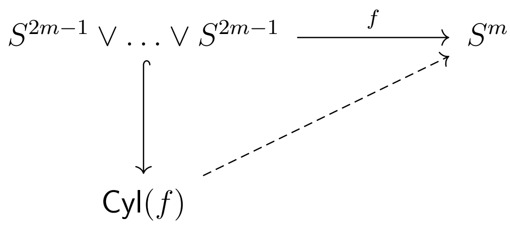

Here an extension always exists, since Cyl(f) deformation retracts onto S^m. Putting both extensions together, provides the desired one DCyl$(\phi; f) \to S^m$.

□

The same argumentation works for the case when m is odd, where we just need to replace the target S^m with $S^m \vee S^m$.

Remark 3.4.1. Following Remark 3.2.1 and using the same argumentations as in Chapter 6 in [12], we can obtain an algorithm, which, for every instance of EXT_m with a double mapping cylinder reformulation, computes a representation of the corresponding topological spaces and maps between them as finite simplicial complexes and simplicial maps.

This instance of the extension problem is going to serve as a base for our undecidability results. For the sake of presentation, we introduce the following notation.

Definition 3.5 ($\widetilde{\mathrm{EXT}}_m$)**.** *Given finite simplicial complexes $A = S^{2m+1} \vee \ldots \vee S^{2m+1}$, $W = S^m \vee \ldots \vee S^m$ and simplicial maps $\phi\colon A \to W$ and $f\colon A \to Y$, decide whether there exists a continuous map $F\colon$ DCyl$(\phi; f) \to Y$, which extends $(f, \mathrm{id}_Y)\colon A \sqcup Y \to Y$, where $Y \subset$ DCyl$(\phi; f)$ is included as the lower rim of the double mapping cylinder. Here Y is either S^m or $S^m \vee S^m$, depending on whether m is even or odd.*

3.1 A reduction $\widetilde{\mathrm{EXT}}_m \leq \mathrm{EXTEMBED}_{2m+1\to 3m+2}$

We first prove Theorem 1.14, since it contains the main arguments required also by Theorem 1.15.

We begin by showing how to construct a reduction $\widetilde{\mathrm{EXT}}_m \leq \mathrm{EXTEMBED}_{2m+1\to 3m+2}$, for a particular instance of the extension problem with $m \geq 2$. The proof of Theorem 1.14 will then consist of a systematic application of such basic reductions. There is an inherited difference between the even and odd cases for the parameter m, so we consider them separately.

m **even:** Let $A = S^{2m-1} \vee \ldots \vee S^{2m-1}$, $W = S^m \vee \ldots \vee S^m$, $\phi\colon A \to W$, $f\colon A \to S^m$ be our input extension problem. Denote by $X :=$ DCyl$(\phi; f)$ the double mapping cylinder of the maps (ϕ, f). By Remark 3.4.1, we assume that those spaces and maps are given as finite simplicial complexes and simplicial maps. For dimensional reasons, it is apparent that we cannot embed X into S^m or even turn the map f into an embedding, therefore we would need to introduce changes in the setting. We reinterpret the target sphere S^m in the extension problem in the following way. Consider the spheres S^{3m+2} and S^{2m+1}. By the discussion before Proposition 3.33, since $(3m + 2) - (2m + 1) \geq 3$, all embeddings of S^{2m+1} into S^{3m+2} are ambient isotopic.[1] Let $s\colon S^{2m+1} \hookrightarrow S^{3m+2}$ be the standard embedding, as defined in Section 3.4, i.e. s is the inclusion of S^{2m+1} into S^{3m+2} as the standard sphere in the first $2m + 2$ coordinates. Let $S^m \subset S^{3m+2}$ be *complementary* to S^{2m+1}, i.e. the standard sphere on the last $m + 1$ coordinates of S^{3m+2}, shifted slightly, so it does not intersect S^{2m+1}. The easiest case to imagine is a Hopf link in S^3, consisting of two circles, lying in perpendicular planes.

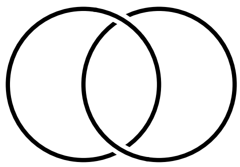

Now, we can think of the map $f\colon A \to S^m \subset S^{3m+2} \setminus \mathrm{Int}(N(S^{2m+1}))$, as having S^{3m+2} as its target and missing a small open neighbourhood $\mathrm{Int}(N(S^{2m+1}))$ of $S^{2m+1} \subset S^{3m+2}$. We put this map in general position and use Remark 3.55.2 to homotope it to an embedding $g\colon A \hookrightarrow S^{3m+2}$.

Proposition 3.6. *The map $g\colon A \hookrightarrow S^{3m+2} \setminus \mathrm{Int}(N(S^{2m+1}))$ extends to a map $F\colon X =$ DCyl$(\phi; f) \to S^{3m+2} \setminus \mathrm{Int}(N(S^{2m+1}))$ if and only if the original extension problem has a solution.*

[1] *Ambient isotopy* is an equivalence relation between embeddings, which reflects their geometric properties alongside their topological ones. We give a formal definition in Section 3.4.

Proof. As we remarked above, the existence of a solution for the extension problem is of homotopy theoretic nature. To prove the proposition, we need to check that we only changed the setting up to homotopy. Indeed, by Proposition 3.33, we have the homotopy equivalence $S^{3m+2} \setminus \mathrm{Int}(N(S^{2m+1})) \sim S^m$, so homotopically, our target is nothing but S^m. On the other hand, the maps $f, g\colon A \to S^{3m+2} \setminus \mathrm{Int}(N(S^{2m+1}))$ are homotopic by construction. That concludes the proof. □

Observe, that even if there exists an extension for the map $g\colon A \hookrightarrow S^{3m+2}$, we still need to prove that we can modify it to be an embedding. This is granted by a more technical theorem. Before we state it, we need to introduce some further notation. Let Σ be either S^m or $S^m \vee S^m$ and let $s\colon \Sigma \hookrightarrow S^{3m+2}$ be the standard embeddeding, as introduced in Section 3.4. Then, denote by $Q := S^{3m+2} \setminus \mathrm{Int}(N(\Sigma))$, where $\mathrm{Int}(N(\Sigma))$ is a small open neighbourhood of Σ.[2]

Theorem 3.7. *Let $A \subset X$ and Q be as introduced above, and $f\colon A \hookrightarrow Q$ be an embedding of A. Then, if $F\colon X \to Q$ is any PL extension of f, there exists an embedding $G\colon X \hookrightarrow Q$, which is homotopic to F and $F|_A = G|_A = f$.*

The reason why we introduce Σ is so, that the theorem also covers the case when m is odd, which we explain below. The proof of this theorem will occupy much of the later sections.

Finally, in order to obtain the desired $\mathrm{EXTEMBED}_{2m+1\to 3m+2}$ problem, we need to change our perspective once again, which we illustrate in the following diagrams. Observe, that $\dim A = 2m-1$ and $\dim X = 2m$, so in the right diagram, the complexes on the left side have dimension $(2m+1)$.

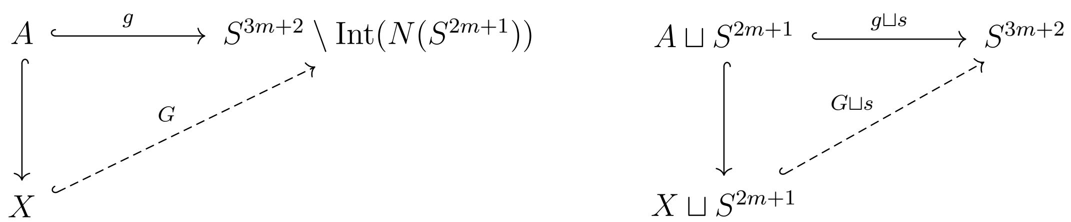

In other words, rather than thinking about our embeddings $g\colon A \hookrightarrow S^{3m+2} \setminus \mathrm{Int}(N(S^{2m+1}))$ and $G\colon X \hookrightarrow S^{3m+2} \setminus \mathrm{Int}(N(S^{2m+1}))$ being into $S^{3m+2} \setminus \mathrm{Int}(N(S^{2m+1}))$, we think of them as being embeddings of the disjoint unions $A \sqcup S^{2m+1}$ and $X \sqcup S^{2m+1}$ into S^{3m+2}. By compactness, since the image of X avoids the image of S^{2m+1}, it also avoids a small neighbourhood $N(S^{2m+1})$ of it. This concludes the step.

m **odd:** In this case, the only difference is that instead of S^m, the target in the extension problem is $S^m \vee S^m$. As we already mentioned, Proposition 3.33 applies here as well, ensuring the homotopy equivalence $S^{3m+2} \setminus \mathrm{Int}(N(S^{2m+1} \vee S^{2m+1})) \sim S^m \vee S^m$, where $\mathrm{Int}(N(S^{2m+1} \vee S^{2m+1}))$ is a small open neighbourhood of $S^{2m+1} \vee S^{2m+1}$ inside S^{3m+2}. Thus, from the initial extension problem, we can easily obtain an embedding $g\colon A \hookrightarrow S^{3m+2} \setminus \mathrm{Int}(N(S^{2m+1} \vee S^{2m+1}))$, which extends to a map $F\colon \mathsf{DCyl}(\phi; f) \to S^{3m+2} \setminus \mathrm{Int}(N(S^{2m+1} \vee S^{2m+1}))$ if and only if the original problem has a solution. Finally,

[2]Recall, that we have also fixed a *standard* embedding of S^{2m+1} into S^{3m+2}.

if an extension exists, once again we make use of Theorem 3.7 to turn it into an embedding $G\colon \text{DCyl}(\phi; f) \hookrightarrow S^{3m+2} \setminus \text{Int}(N(S^{2m+1} \vee S^{2m+1}))$. Switching our point of view, we obtain the following $\text{EXTEMBED}_{2m+1\to 3m+2}$ problem, to which the original EXT_m problem reduces.

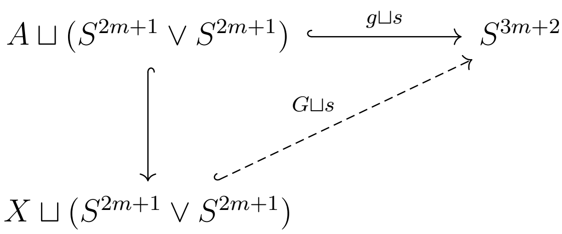

Proof of Theorem 1.14 for $(d-k) \geq 3$. We are finally ready to show how to cover all the pairs (k, d), which satisfy the conditions of the theorem and such that $(d-k) \geq 3$. We begin with a simple observation, which plays a key role in the proof.

Observation 3.8. *A reduction $\widetilde{\text{EXT}}_m \leq \text{EXTEMBED}_{k\to d}$, as constructed above, yields a family of reductions $\widetilde{\text{EXT}}_m \leq \text{EXTEMBED}_{k+\ell\to d+\ell}$, $\ell \geq 0$. Indeed, in the construction we picked S^{2m+1} and S^{3m+2} and used the fact that $S^{3m+2} \setminus N(S^{2m+1}) \sim S^m$. However, the same is true if we raise the dimensions of the two spheres equally, i.e. $S^{3m+2+\ell} \setminus N(S^{2m+1+\ell}) \sim S^m$, $\ell \geq 0$. Moreover, all the tools we use for the original reduction still work in this case.*

However, in order to use Theorem 3.7, the dimensions of these additional spheres have to be at least $(2m+1)$ and $(3m+2)$. To see, why, consider an instance of $\widetilde{\text{EXT}}_m$, say for m even. This problem is represented in the following diagram.

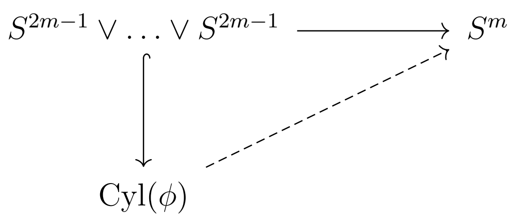

In order to obtain an EXTEMBED problem, we pick S^{3m+1} and S^{2m} and observe that $S^{3m+1} \setminus N(S^{2m}) \sim S^m$. Once again, we can interpret the S^m, which is the target in this extension problem, as the complimentary sphere to the standardly embedded S^{2m} inside S^{3m+1}. We then put the map $F\colon \text{DCyl}(\phi; f) \to S^m \subset S^{3m+1}$, which we get from the extension problem, in general position, and observe that the set of self-intersections of F has dimension $\dim S(F) = 2 * 2m - (3m+1) = m-1$.[3] On the other hand, $S^{3m+1} \setminus N(S^{2m}) \sim S^m$ is $(m-1)$-connected and in order to use Theorem 3.7, we would need its connectivity to be at least one more than $\dim S(f) = m-1$, i.e. it has to be at least m-connected. By raising the dimension of the additional spheres, we can solve this problem, because, while not changing the connectivity of the target, we

[3] We will thoroughly explain how to compute the dimension of the set of self-intersections of a PL map in general position in Section 3.4.

decrease the dimension of $S(F)$ by at least one. For that reason, our methods do not cover the cases $(k, \frac{3k}{2}+1)$ for k even.

Finally, following Observation 3.8, in order to cover all the other pairs, which fall into the assumptions of Theorem 1.14, it suffices to cover the boundary cases for k odd. The way to do that is by considering the even cases we cannot solve and shifting them up with one dimension. We illustrate the idea with the pair $(7, 11)$. Starting with the problem $\widetilde{\mathrm{EXT}}_3$ and picking S^6 and S^{10} as additional spheres, we would obtain $\mathrm{EXTEMBED}_{6\to 10}$, which we know we cannot solve. Therefore, we raise the dimensions of the spheres to S^7 and S^{11}. This removes the obstruction for using Theorem 3.7 and we can obtain a reduction $\widetilde{\mathrm{EXT}}_3 \leq \mathrm{EXTEMBED}_{7\to 11}$. Doing the same for the boundary case for every even k and combining it with Observation 3.8 gives us a sequence of diagonals, which cover the whole range of pairs (k, d), that fall into the assumption of Theorem 1.14 such that $(d-k) \geq 3$. That concludes the proof.

In Section 3.3 we will prove Theorem 1.14 for the cases $(d-k) = 2$, $(k, d) \neq (5, 7)$. We recall Figure 1.2, which illustrates the combined result.

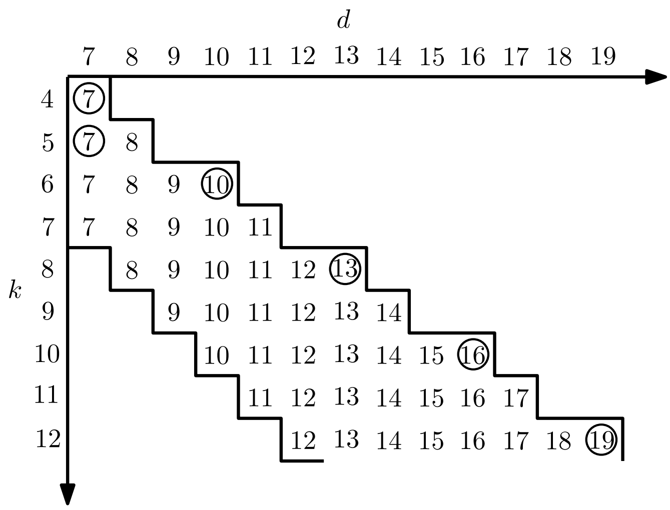

Figure 3.1: Undecidability of $\mathrm{EXTEMBED}_{k\to d}$. The pairs that are not covered by our result are encircled.

3.2 A reduction $\widetilde{\mathrm{EXT}}_m \leq \mathrm{EMBED}_{2m+1\to 3m+2}$

In this section, we provide a proof of Theorem 1.15 when $(d-k) \geq 3$. We will construct, for every pair (k, d), which falls under these assumptions, a reduction $\widetilde{\mathrm{EXT}}_m \leq \mathrm{EMBED}_{k\to d}$ from a suitable instance of the extension problem. In fact, we will first use the construction in the previous section, to obtain an instance of $\mathrm{EXTEMBED}_{k\to d}$, that we will then turn into an instance of $\mathrm{EMBED}_{k\to d}$. According to Theorem 1.14, and since the two theorems cover the same ranges of pairs (k, d), that will ensure the undecidability of the problem $\mathrm{EMBED}_{k\to d}$ for the relevant pairs with $(d-k) \geq 3$. We will prove the case when $(d-k) = 2$ in Section 3.3. We consider the even and odd cases separately.

m even We begin by providing some geometric intuition for the construction. Let the following be an embedding extension problem, as constructed in the previous section.

$$\begin{array}{ccc} A = S^{2m-1}\vee\ldots\vee S^{2m-1}\sqcup S^{2m+1} & \xhookrightarrow{g\sqcup s} & S^{3m+2} \\ \big\downarrow & \overset{G\sqcup s}{\dashrightarrow} & \\ X = \mathsf{DCyl}(\phi; f)\sqcup S^{2m+1} & & \end{array}$$

Observation 3.9. *From the way we constructed the* $\mathrm{EXTEMBED}_{2m+1\to 3m+2}$ *problem above, the embedded copies* $G(S^m)$ *and* $s(S^{2m+1})$ *inside* S^{3m+2} *have homological linking number*[4] ± 1. *Indeed,* s *sends* S^{2m+1} *onto the standard* $(2m+1)$ *sphere on the first* $(2m+2)$ *coordinates of* S^{3m+2} *and* S^m *onto a shifted copy of the standard sphere on the remaining* $(m+1)$ *coordinates, which might be slightly altered up to an ambient isotopy, which does not change the linking number.*

The general idea of the reduction is to take the space $X \sqcup S^{2m+1}$ and turn it from a disjoint union into a single space, by adding a scaffold that attaches S^{2m+1} to $S^m \subset X$, where we regard this copy of S^m as one of the *rims* of the double mapping cylinder X. This scaffold will be a finite simplicial complex, which contains a copy of S^m and S^{2m+1} and has the property that no matter how we embed it into S^{3m+2}, the embedded spheres S^m and S^{2m+1} are linked with linking number ± 1, depending on the chosen orientations. More precisely, we introduce the following definition of a *linking scaffold*, which we construct explicitly in Section 3.4.3.

Definition 3.10. *Let* k, ℓ *be nonnegative integers. We call a simplicial complex* L *a* (k,ℓ)-linking scaffold *if the following conditions hold:*

- $S^k \sqcup S^\ell \subseteq L$.
- *There exists a PL embedding* $L \hookrightarrow S^{k+\ell+1}$.
- *For any PL embedding* $f\colon L \hookrightarrow S^{k+\ell+1}$, *the spheres* $f(S^k)$ *and* $f(S^\ell)$ *are linked with linking number* ± 1.

We are now ready to construct the $\mathrm{EMBED}_{2m+1\to 3m+2}$ problem, associated to the given $\mathrm{EXTEMBED}_{2m+1\to 3m+2}$ problem. Let $\mathsf{link}(2m+1; m)$ be the linking scaffold constructed in Section 3.4.3, which has $\dim(\mathsf{link}(2m+1;m)) = 2m+1$. Then, glueing together the double mapping cylinder $X = \mathsf{DCyl}(\phi; f)$ and $\mathsf{link}(2m+1;m)$, we obtain the connected simplicial complex

$$LX := \mathsf{link}(2m+1; m) \cup_{S^m} \mathsf{DCyl}(\phi; f). \tag{3.3}$$

[4]For a general precise definition, we refer to Chapter 2.5 in [57]. In our situation, one way to think about linking numbers is the following. Let $f\colon S^p \sqcup S^q \hookrightarrow S^{p+q+1}$ be an embedding. We say that $f(S^p)$ and $f(S^q)$ are *linked* inside S^{p+q+1} with *(homological) linking number* ℓ if the map $f\colon S^p \to S^{p+q+1}\setminus S^q \sim S^p$ has degree $\deg(f) = \ell$. The linking number does not depend on the choice of the sphere we remove.

Here the identification on the side of X, is along the lower rim S^m of X and the copy of S^{2m+1}, which is originally disjoint with X and on the side of the linking scaffold, it is along the distinguished copies $S^m, S^{2m+1} \subset \mathsf{link}(2m+1;m)$, which are mentioned in Definition 3.10.

Lemma 3.11. *If there exists an embedding $G \sqcup s\colon X \sqcup S^{2m+1} \hookrightarrow S^{3m+2}$, then there exists a PL embedding $h\colon LX \hookrightarrow S^{3m+2}$.*

Proof. Lemma 3.48 ensured that there is an embedding $h\colon \mathsf{link}(2m+1;m) \hookrightarrow S^{3m+2}$ such that $h(\mathsf{link}(2m+1;m)) \subset S^{2m+1} \times B^{m+1}$. In other words, the image of the embedding h is contained in a regular neighbourhood of S^{2m+1} inside S^{3m+2}. On the other hand, the image $G(X)$ avoids a small regular neighbourhood of S^{2m+1}, so by an ambient isotopy, we can obtain the desired embedding. □

Proof of Theorem 1.15 for m even and $(d-k) \geq 3$. Consider an instance of $\mathrm{EXTEMBED}_{2m+1\to 3m+2}$ and the associated instance of $\mathrm{EMBED}_{2m+1\to 3m+2}$ as just constructed. We will prove that one has a solution if and only if the other one has a solution.

- If $\mathrm{EXTEMBED}_{2m+1\to 3m+2}$ has a solution, then Lemma 3.11 guarantees that $\mathrm{EMBED}_{2m+1\to 3m+2}$ also has a solution.
- If $\mathrm{EMBED}_{2m+1\to 3m+2}$ has a solution, then there exists an embedding $h\colon LX \hookrightarrow S^{3m+2}$. In particular, this give an embedding $h_{X\sqcup S^{2m+1}}\colon X \sqcup S^{2m+1} \hookrightarrow S^{3m+2}$. Finally, since $\mathsf{link}(2m+1;m)$ is a linking scaffold, $h(S^{2m+1})$ and $h(S^m)$ are linked in S^{3m+2} with linking number ± 1. That means, that we can apply an ambient isotopy and obtain an embedding $G\colon X \sqcup S^{2m+1} \hookrightarrow S^{3m+2}$, where $G(S^{2m+1})$ is standardly embedded, which implies that $\mathrm{EXTEMBED}_{2m+1\to 3m+2}$.

That concludes the proof.

□

m odd The odd case works in the same way. We only need to make use of the wedge linking scaffold $\mathsf{link}_w(2m+1,m)$, which serves the same purpose as a linking scaffold, but for a pair $(S^{2m+1} \vee S^{2m+1}, S^m \vee S^m)$.

3.3 The cases $(d-k)=2$, $(d,k) \neq (5,7)$

In this section we complete the proofs of Theorem 1.14 and Theorem 1.15. More precisely, we prove the following proposition.

Proposition 3.12. *$\mathrm{EMBED}_{k\to k+2}$ is undecidable for $k > 5$.*

From the proposition, it trivially follows that $\mathrm{EXTEMBED}_{k\to k+2}$ is undecidable for $k > 5$. Our proof does not cover the case $(5,7)$, which is left open.

The strategy of the proof of Proposition 3.12 is to show that for every $k > 5$ the undecidability of $\mathrm{EMBED}_{k\to k+2}$ follows from the undecidability of $\mathrm{EMBED}_{k-1\to k+2}$, which we already proved. Our proof makes use of the following technical lemma, which we prove in Section 3.5.

Lemma 3.13. *Let K be a finite simplicial complex, $\tau \in K$ a maximal* [5] *ℓ-simplex of K and $d \geq \ell + 3$ an integer. Let $\widetilde{K} := K \cup_\tau \Delta^{\ell+1}$ be the complex, obtained from K by coning over τ. Then, there is an embedding $\widetilde{K} \hookrightarrow S^d$ if and only if there is an embedding $K \hookrightarrow S^d$.*

Proof of Proposition 3.12. Let $(\ell, \ell+3)$ be a pair of integers with $\ell \geq 5$. Our main theorem implies that $\mathrm{EMBED}_{\ell \to \ell+3}$ is undecidable, by constructing a family of ℓ-dimensional finite simplicial complexes, which encode an undecidable instance of the extension problem. Let K be any of these complexes. By construction, K is not pure and by an inductive application of Lemma 3.13, we can construct a complex $\widetilde{K}$, which is $(\ell+1)$-dimensional, pure and there exists an embedding $\widetilde{K} \hookrightarrow S^{\ell+3}$ if and only if there exists an embedding $K \hookrightarrow S^{\ell+3}$. This operation translates the undecidable instance of the problem $\mathrm{EMBED}_{\ell \to \ell+3}$, which we constructed in our main theorem, into an instance of the problem $\mathrm{EMBED}_{\ell+1 \to \ell+3}$ for any $\ell \geq 5$. This concludes the proof.

□

The rest of the exposition will be dedicated to the proof of Theorem 3.7. We begin by outlining the necessary preliminaries from PL topology.

3.4 Preliminaries

In this section we introduce the basic notions and constructions we need for the proof of Theorem 3.7. All topological spaces will be assumed compact unless otherwise stated.

Definition 3.14. *Let K, L be simplicial complexes. A continuous map $f\colon |K| \to |L|$ is a piecewise linear (PL) map if there exist subdivisions $\widetilde{K}, \widetilde{L}$ of K and L and a simplicial map that $\widetilde{f}\colon \widetilde{K} \to \widetilde{L}$ such that $|\widetilde{f}| = f$.*

Polyhedra and PL maps Let X be a compact topological space. A *triangulation* of X is a homeomorphism $t\colon |K| \to X$, where $|K|$ is the geometric realization of some finite simplicial complex.

Definition 3.15 (Polyhedron)**.** *A compact* polyhedron *is a compact topological space X, provided with a family Tr_X of triangulations of X, such that the following conditions are satisfied:*

1. *If $t\colon |K| \to X$ is a triangulation from Tr_X and $\phi\colon |L| \to |K|$ is a PL homeomorphism, then $t \circ \phi\colon |L| \to X$ also belongs to Tr_X.*
2. *If $t_1, t_2 \in \mathrm{Tr}_X$, then $t_2^{-1} \circ t_1$ is a PL homeomorphism.*

Polyhedra are the main objects of interest in this Chapter. They are quite flexible, since they grant us the possibility to pick different triangulations, depending on their requried properties.

[5]A simplex $\tau \in K$ is called *maximal* if τ is not a face of any other simplex $\sigma \in K$. The simplicial complex K is called *pure* if all the maximal simplices of K have the same dimension.

Example 3.16. *The following are classical examples of polyhedra.*

1. *Any finite simplicial complex* K *with* Tr_K *consisting of all PL homeomorphisms onto* K.
2. *A PL* k*-sphere is a polyhedron* S*, which admits a triangulation by the boundary of the standard* $(k+1)$*-simplex* Δ^{k+1}.
3. *A PL* $(k+1)$*-ball is a polyhedron* B*, which admits a triangulation by the standard* $(k+1)$*-simplex* Δ^{k+1}.

Next, we define the appropriate type of maps between polyhedra.

Definition 3.17. *Let* X, Y *be polyhedra and* $f\colon X \to Y$ *be a continuous map. We say that* f *is a* PL *map if there exist triangulations* $t_X\colon |K| \to X$ *in* Tr_X *and* $t_Y\colon |L| \to Y$ *in* Tr_Y*, such that the map* $t_Y^{-1} \circ f \circ t_X\colon |K| \to |L|$ *is a PL map as defined in Definition 3.14.*

Let X be a polyhedron and $A \subset X$ be a subspace. We say that A is a subpolyhedron of X if the inclusion map $i\colon A \hookrightarrow X$ is a PL map.

For a PL map $f\colon X \to Y$, we define its *singular set* of the map f as $S(f) := \mathrm{Cl}(\{x \in X | f^{-1}(f(x)) \neq \{x\}\})$, i.e. the closure of the set of self-intersections of f. Observe that $S(f) = \emptyset$ if and only if the map f is an embedding. We further define, for every $i \geq 2$, $S_i(f) := \{x \in X || f^{-1}(f(x))| \geq i\}$. That means that $S_i(f)$ is the set of i-fold intersection points of f. Observe that $S(f) = \mathrm{Cl}(S_2(f))$ and for every $i \geq 2$, $S_i(f) \supseteq S_{i+1}(f)$.

PL manifolds Piecewise-linear (or PL) manifolds are a class of polyhedra, which, similar to smooth manifolds, have a particularly nice structure. Before we introduce them, we define a class of simplicial complexes that would serve as triangulations for PL manifolds.

Definition 3.18 (Combinatorial manifold)**.** *Let* K *be a finite simplicial complex. We say that* K *is a* k*-dimensional* combinatorial manifold *if for every vertex* $v \in K$*, the link* $\mathrm{lk}(v, K)$ *is PL homeomorphic either to a PL* k*-sphere* S^k *or to a PL* k*-ball* B^k*. Vertices, whose links are PL balls belong to the boundary of the combinatorial manifold* K.

Definition 3.19 (PL manifold)**.** *Let* Q *be a polyhedron. We say that* Q *is a* k*-dimensional* PL manifold *if, for one, therefore for every, triangulation* $t\colon |K| \to X$*,* K *is a combinatorial* k*-manifold.*

Definition 3.20. *Let* M, Q *be PL manifolds with boundary. We say that an embedding* $f\colon M \hookrightarrow Q$ *is a* proper *embedding if* $f(\partial M) \subset \partial Q$*. In particular, if* $\partial M = \emptyset$*, every embedding* $M \hookrightarrow Q$ *is proper.*

Unknotting of PL spheres and balls In this paragraph we outline several classical results about unknotting of spheres and balls inside PL manifolds, which will be necessary for our main proofs. Since we will be working a lot with embeddings, we first define right type of equivalence between two embeddings.

Definition 3.21. *Let* P *be a polyhedron (not necessarily a manifold). An* ambient isotopy *of* P *is a PL map* $H\colon P \times [0,1] \to P$*, such that.*

1. $H(\bullet, t)$ *is a homeomorphism for every* $t \in [0,1]$.
2. $H(\bullet, 0)$ *is the identity of* P.

Definition 3.22 (Unknotting)**.** *Let* X *be a polyhedron and* M *be a PL manifold with* $\dim X < \dim M$. *We say that* X unknots *in* M *if every two embeddings* $f, g\colon X \hookrightarrow M$, *which are homotopic, are also ambient PL isotopic.*

Theorem 3.23 ([79] Theorem 9)**.** S^p *unknots in* S^q*, provided* $p \leq q-3$.

In fact, the theorem is also true for properly embedded balls, again in codimension at least three, but we will not need that. A similar fact is proved in [77] in the discussion after the proof of Theorem 8. We state it here, since we will make use of it later.

Proposition 3.24 (Lickorish)**.** B^r *unknots in* S^p*, provided* $r \leq p-3$.

Let $S^{p_1} \cup_{B^r} \ldots \cup_{B^r} S^{p_l}$ be the polyhedron, consisting of $S^{p_1}, \ldots, S^{p_l}$, identified along a copy of B^r. By Proposition 3.24, B^r unknots in S^p, provided $r \leq p-3$. Hence $S^{p_1} \cup_{B^r} \ldots \cup_{B^r} S^{p_l}$ is well defined up to homeomorphism and does not depend on the particular choice of the embeddings of B^r into the different spheres, as long as $r \leq p_i - 3\,,\ 1 \leq i \leq l$.

The following proposition follows from a result by Lickorish. The original statement was formulated for a wedge of two spheres, but the arguments apply for the wedge of any finite number of spheres.

Proposition 3.25 (Theorems 8 and 9 in [77])**.** $S^{p_1} \vee \ldots \vee S^{p_l}$ *unknot in* S^q*, given* $p_i \leq q-3\,,\ 1 \leq i \leq l$. *Moreover,* $S^{p_1} \cup_{B^r} \ldots \cup_{B^r} S^{p_l}$ *unknots in* S^q*, provided* $p_i \leq q-3\,,\ 1 \leq i \leq l$.

General position General position is tool, which allows one to minimize the intersection between two polyhedra inside a PL manifold. We state here some results, which we will extensively use in our exposition. For further details we refer the reader to Chapter 6 of [79].

Definition 3.26. *Let* $P_1, P_2 \subset M$ *be two polyhedra inside a PL* m*-manifold* M. *We say that* P_1 *is in* general position *with respect to* P_2 *if* $\dim(P_1 \cap P_2) \leq \dim P_1 + \dim P_2 - m$.

Theorem 3.27. *Let* $P_1, P_2 \subset \mathrm{Int}(M)$ *be two polyhedra inside the PL manifold* M. *Then, there is an ambient isotopy of* M*, keeping* ∂M *fixed, such that the image of* P_1 *is in general position with respect the image of* P_2.

There is a similar notion for PL maps.

Definition 3.28. *Let* $f\colon X \to M$ *be a PL map from a polyhedron* X *to a PL* m*-manifold* M. *Denote by* $c := m - \dim X$. *We say that* f *is in* general position *if for every* $i \geq 2\,,\ \dim S_i(f) \leq \dim X - (i-1)c$*, where* $S_i(f)$ *is the set of* i*-fold intersection points of the map* f.

Theorem 3.29. *Let* $f\colon X \to M$ *be a PL map from a polyhedron to a PL manifold* M. *Then,* f *is homotopic to a map* $g\colon X \to M$*, which is in general position.*

As long as we are inside an ambient PL manifold, we can always put polyhedra in general position to each other, as well as PL maps. In what follows, we will always be assuming that all given maps are always in general position.

Collapses and collapsible polyhedra Let $B^{q-1} \subset B^q$ be an $(q-1)$-ball inside the q-ball B^q. We say that B^{q-1} is a *face* of the q-ball B^q and denote it by $B^{q-1} \prec B^q$ if there exists a triangulation $t\colon \Delta^q \to B^q$ of B^q, such that for some $0 \leq i \leq q$, $t(d_i\Delta^q) = B^{q-1}$.

Let $A \subset X$ be polyhedra. We say that there is an *elementary collapse* from X to A if there exists a ball B^q and a face $B^{q-1} \prec B^q$ of B^q, such that

$$X = A \cup B^q$$
$$B^{q-1} = A \cap B^q$$

That means that we can obtain A from X by erasing a PL ball.

Definition 3.30. *Let $A \subset X$ be a subpolyhedron of the polyhedron X. We say that X* collapses *onto A, denoted by $X \searrow A$, if there exists a finite sequence $A \subset Z_1 \subset Z_2 \subset \ldots \subset Z_l = X$, such that for every $1 \leq i \leq l-1$, there is an elementary collapse from Z_{i+1} to Z_i. In particular, if A is a point, we say that X is* collapsible*, denoted by $X \searrow 0$.*

Example 3.31.

1. *If B is a PL ball, then $B \searrow 0$.*
2. *Let X be a polyhedron. Then $CX \searrow 0$, where CX is the cone over X. Moreover, if $A \subset X$ is any subpolyhedron of X, then $CX \searrow CA$, i.e. a cone collapses onto any subcone.*

The mapping cylinder of a map Let $f\colon X \to Y$ be a continuous map between the topological spaces X and Y. We define the *mapping cylinder* $\mathrm{Cyl}(f) := (X \times [0,1] \sqcup Y)/((x,1) \sim f(x))$. If X and Y are polyhedra and f is a PL map, then the mapping cylinder $\mathrm{Cyl}(f)$ is naturally a subpolyhedron of the join $X * Y$ of X and Y. For more details, we refer the reader to Chapter 2 of [79].

Lemma 3.32. *Let $m \geq 2$, $A = S^{2m-1} \vee \ldots \vee S^{2m-1}$ and $W = S^m \vee \ldots \vee S^m$. Then the mapping cylinder of any PL map $\phi\colon A \to W$ embeds in S^{3m+2}.*

Proof. The mapping cylinder of every PL map $A \to W$ is naturally a subset of their join $A * W$. We have that A embeds in $\mathbb{R}^{2m}$ and W embeds in $\mathbb{R}^{m+1}$. Therefore, $A * W$ embeds in $\mathbb{R}^{3m+2}$, and also in S^{3m+2}. □

Removing S^p from S^q For $p \leq q$, define the *standard* embedding of S^p into S^q as follows. First, embed $S^p \hookrightarrow \mathbb{R}^{p+1}$ as the unit sphere, and then embed $\mathbb{R}^{p+1} \hookrightarrow \mathbb{R}^q \subset S^q$, where the first map is the inclusion of $\mathbb{R}^{p+1}$ into $\mathbb{R}^q$ as the first $p+1$ coordinates. Similarly, define the *standard* embedding of $S^p \vee \ldots \vee S^p$ into S^q as follows. First, embed $S^p \vee \ldots \vee S^p$ into $\mathbb{R}^{p+1}$ as in Figure 3.4. Then, embed $\mathbb{R}^{p+1} \hookrightarrow \mathbb{R}^q \subset S^q$ in the same way as in the case a single copy of S^p.

It is well known that if we remove a standardly embedded copy of S^p from S^q $(p < q)$, the resulting space is homotopy equivalent to S^{q-p-1}. In fact, a similar statement is also true if we replace S^p by a wedge $S^p \vee S^p$.

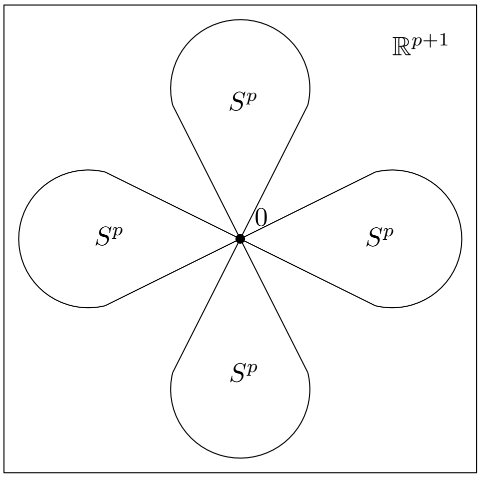

Figure 3.2: Standard embedding of $S^p \vee S^p \vee S^p \vee S^p$ into $\mathbb{R}^{p+1}$

Proposition 3.33. *Let $f\colon S^p \vee S^p \hookrightarrow S^q$ ($p < q$) be the standard embedding. Then $S^q \setminus f(S^p \vee S^p) \sim S^{q-p-1} \vee S^{q-p-1}$.*

Proof. Throughout the proof, we will identify the sphere S^q to the one-point compactification $\mathbb{R}^q \cup \{\infty\}$ of $\mathbb{R}^q$. The standard embedding f is ambient isotopic to the embedding $g\colon S^p \vee S^p \hookrightarrow S^q$, which sends $S^p \vee S^p$ to two parallel affine p-spaces $A_1^p, A_2^p \subset \mathbb{R}^q$, which intersect at infinity. If we remove $g(S^p \vee S^p)$ from S^q, what remains is $\mathbb{R}^q$ with two parallel affine p-spaces removed. Let $\mathbb{R}^p$ be the unique linear p-subspace of $\mathbb{R}^q$, which is parallel to A_1 and A_2, and let $\mathbb{R}^{q-p}$ be its complement. We construct the desired homotopy equivalence in two steps. First, we project $\mathbb{R}^q \setminus A_1^p \cup A_2^p$ to $\mathbb{R}^{q-p}$. The resulting space is $\mathbb{R}^{q-p}$ with two points removed, which is homotopy equivalent to $S^{q-p-1} \vee S^{q-p-1}$. This completes the proof. $\square$

Observe that, since $p < q$, any two embeddings of $S^p \vee S^p$ into S^q are homotopic. Moreover, since $p \leq q-3$, by Proposition 3.25, any two PL embeddings of $S^p \vee S^p$ into S^q are also ambient isotopic. Therefore, for $p \leq q-3$, Proposition 3.33 holds for any embedding $S^p \vee S^p \hookrightarrow S^q$.

3.4.1 Regular and relative regular neighbourhoods

In this section we we provide the necessary background on regular and relative regular neighbourhoods. We first define them in PL manifolds and then give a generalization to simplicial complexes due to Cohen [16]. For our purposes, all PL manifolds will be assumed compact.

The way we will usually construct neighbourhoods is by first refining the triangulation by subdivision. However, it will often be the case that the barycentric subdivision would not be the most convenient choice. We would rather use a more flexible subdivision, called the *derived* subdivision. We obtain it in the same way as the barycentric subdivision, but we don't necessarily pick the barycenters of the simplices when we are coning over their boundaries, but rather give ourselves the freedom to choose any interior point instead. The barycentric subdivision is an example of a derived subdivision. We can also iterated it, obtaining second, third etc. derived subdivisions. We refer to Chapter I of [79] for a precise definition.

Definition 3.34. *Let J be a simplicial complex and K be a subcomplex. We say that K is* full *in J if no simplex of $J \setminus K$ has all its vertices in K.*

Given a polyhedron P in a manifold Q, it will be important for the construction of regular neighbourhood of P in Q that we pick a triangulation of Q, such that P is a full subcomplex. However, this is not a problem due to the following remark.

Remark 3.34.1. Let J be a simplicial complex and K be a subcomplex. Then K' is a full subcomplex of J', where J' is the first derived subdivision of J.

Remark 3.34.1 permits us to always pick triangulations, so that the polyhedra we work with are full subcomplexes of the ambient manifold.

Let J be a simplicial complex and X a subset (not necessarily a simplicial complex). The simplicial neighbourhood $N(X, J)$ of X in J is the minimal subcomplex of J containing all simplices of J, which have non-empty intersection with X.

Suppose P is a subpolyhedron of the polyhedron Z. Pick a triangulation J of Z and K of P such that K is a full subcomplex of J. The polyhedron $|N(K, J')|$ is called a *derived neighbourhood* of P in J. From the remark above it follows, that if we pick any triangulation J of Z and K of P, such that K is a subcomplex of J, then $|N(K, J'')|$ is a derived neighbourhood of K in J, where J'' is the second derived subdivision of the triangulation of J.

Regular neighbourhoods in manifolds

Definition 3.35 (Regular neighbourhood)**.** *Let P be a polyhedron in a PL m-manifold M. A* regular neighbourhood *of P in M is a polyhedron N, such that:*

1. *N is a closed topological neighbourhood of P in M.*
2. *N is a PL m-manifold.*
3. *$N \searrow P$.*

The following theorem, due to Whitehead, ensures the existence and uniqueness of regular neighbourhoods in PL manifolds.

Theorem 3.36 ([68] Theorem 1.6.4)**.** *If P is a polyhedron in the PL manifold M , then*

1. *(Existence) Any derived neighbourhood of P in M is a regular neighbourhood of P in M.*
2. *(Uniqueness) If N_1 and N_2 are any two regular neighbourhoods of P in M , then there is a PL homeomorphism h of N_1 onto N_2, keeping P fixed.*
3. *(Uniqueness) If $P \subset$ Int(M), then any two regular neighbourhoods of P in $\mathrm{Int}(M)$ are ambient isotopic leaving $P \cup \partial M$ fixed.*

An important fact about regular neighbourhoods, which we will be using extensively, is the following.

Theorem 3.37 ([79] Theorem 5)**.** *Let P be a polyhedron in an m-manifold M, such that $P \searrow 0$. Then, any derived neighbourhood of P in M is an m-ball.*

It follows from Theorem 3.36 and Theorem 3.37 that any regular neighbourhood of a collapsible polyhedron in a manifold is a ball of maximal dimension.

Relative regular neighbourhoods In this section we give an extension to the definition of a regular neighbourhood for the case when the ambient space is not a manifold. We also introduce the notion of a relative regular neighbourhood following the exposition in [16].

Definition 3.38 (Relative regular neighbourhood)**.** *Let P, R, V be subpolyhedra of the polyhedron Z. We say that V is a* relative regular neighbourhood *of $P \bmod R$ in Z if there exist a triangulations of (K, L, M, J) of (P, R, V, Z), such that K,L and M are full subcomplexes of J and $M = N(K \setminus L, J')$ in that triangulation.*

We can distinguish the following special cases of this definition:

- When J is a manifold: relative regular neighbourhood in a manifold.
- When $L = \emptyset$: regular neighbourhood in a general polyhedron.
- When J is a manifold and $L = \emptyset$: it can be shown that the definition coincides with definition 3.35.

It is clear that relative regular neighbourhoods always exist. Indeed, given $P, R \subset Z$, pick triangulations (K, L, J) of (P, R, Z), such that K and L are subcomplexes. Then $N(K \setminus L, J'')$ is a regular neighbourhood of $P \bmod R$ in Z. Moreover, every two regular neighbourhoods of $P \bmod R$ in Z are ambient isotopic, as ensured by the following theorem.

Theorem 3.39 ([16] Theorem 3.1)**.** *Let V and W be two regular neighbourhoods of $P \bmod R$ in Z. Then, there is an ambient isotopy $H \colon Z \times [0, 1] \to Z$ such that:*

1. *$H(\bullet, 0)$ is the identity on Z.*
2. *H keeps $P \cup R$ fixed.*
3. *$H(V, 1) = W$.*

Let $X, Y \subset Z$ be subpolyhedra of the polyhedron Z. We introduce the following notation

$$X_R := \mathrm{Cl}(X \setminus Y) \quad ; \quad Y_R := \mathrm{Cl}(X \setminus Y) \cap Y\,.$$

Similarly, if K and L are subcomplexes of simplicial complex J, we denote

$$K_R := N(K \setminus L, K) \quad ; \quad L_R := L \cap K_R\,.$$

The following lemma connects the two notations.

Lemma 3.40 ([16] lemma 2.3). *If $X = |K|$ and $Y = |L|$ then $X_R = |K_R|$ and $Y_R = |L_R|$.*

The next lemma ensures that if $f\colon Z_1 \to Z_2$ is a PL map between polyhedra, and $Y \subset X \subset Z_2$ are subpolyhedra of Z_2, then, under mild conditions on f, the preimage of a relative regular neighbourhood of $X \bmod Y$ in Z_2 is a relative regular neighbourhood of $f^{-1}(X) \bmod f^{-1}(Y)$ in Z_1. The lemma is stated for simplicial complexes, but can be applied for polyhedra simply by picking appropriate triangulations.

Lemma 3.41 ([16] lemma 2.14). *Let J_1, J_2 be simplicial complexes and $f\colon J_1 \to J_2$ be a simplicial map. Let J_1', J_2' be first derived subdivisions of J_1 and J_2, such that $f\colon J_1' \to J_2'$ is also simplicial. If K and L are subcomplexes of J_2, such that $L \subset K$ and $f^{-1}(L_R) = (f^{-1}(L))_R$. Then*

1. $f^{-1}N(K \setminus L, J_2') = N(f^{-1}(K) \setminus f^{-1}(L), J_1')$
2. $f^{-1}\partial N(K \setminus L, J_2') = \partial N(f^{-1}(K) \setminus f^{-1}(L), J_1')$

Observe that if $L = \emptyset$, i.e. when we have absolute regular neighbourhood as opposed to relative, the condition $f^{-1}(L_R) = (f^{-1}(L))_R$ becomes vacuous.

3.4.2 Mapping cylinder neighbourhoods

In this section, we prove the following propositions, which will be important for the proof of Theorem 3.7.

Proposition 3.42. *Let $m \geq 2$. Let $A = S^{2m-1} \vee \ldots \vee S^{2m-1}$, $W = S^m \vee \ldots \vee S^m$, $\phi\colon A \to W$ be a PL map and $Z :=$ Cyl(ϕ) be the mapping cylinder of ϕ. Then, Z is a regular neighbourhood of W.*

From Proposition 3.42 and the uniqueness theorem for regular neighbourhoods in polyhedra, it will follow that any regular neighbourhood N of W in Z, is PL homeomorphic to Z. While Z is trivially a topological neighbourhood of W, it is not clear from definition 3.38 that it is a *regular* neighbourhood. In order to prove Proposition 3.42, we need to introduce some further notions.

Definition 3.43. *A closed subset $U \subset Z$ is called a* mapping cylinder neighbourhood *of W in Z if there exists a map $f\colon \partial U \to A$* [6] *and a homeomorphism $h\colon U \to \mathrm{Cyl}(f)$, such that $h|_{\partial U \cup W} = 1$.*

In other words, a mapping cylinder neighbourhoods of W is a subpolyhedron $U \subset Z$, which is a topological neighbourhood of W and is homeomorphic to the mapping cylinder of some map. The following theorem ensures the uniqueness of such neighbourhoods.

Theorem 3.44 (Theorem 1 in [45]). *Let U, W be mapping cylinder neighbourhoods of W in Z. Then, there is a homeomorphism $h\colon U \to W$, which leaves a neighbourhood of A fixed.*

[6] For a subset $U \subset Z$, when no confusion can occure, we will denote by ∂U the topological boundary of U in Z.

Proof of Proposition 3.42. Consider a regular neighbourhood $M \subset Z$ of W in Z. By definition $M \searrow W$, and in particular, the collapse gives us a map $\psi\colon \partial M \to W$, so that M is the mapping cylinder $\mathrm{Cyl}(\psi)$. By Definition 3.43, M is a mapping cylinder neighbourhood of W in Z. Since Z is the mapping cylinder of a map with target W, the same is true for Z as well. By Theorem 3.44, for any two mapping cylinder neighbourhoods M_1, M_2 of W in Z, there exists a homeomorphism $h\colon M_1 \to M_2$, which keeps a small neighbourhood of W fixed. In particular, that means that $M \cong Z$, i.e. M is smaller copy of $\mathrm{Cyl}(\phi)$ inside Z, and so is any other regular neighbourhood of W in Z. □

Proposition 3.45 (Proposition 7.5 in [16])**.** *Let M_1 and M_2 be regular neighbourhoods of W in Z, such that $M_2 \subset \mathrm{Int}(M_1)$.*[7] *Then $M_1 \setminus \mathrm{Int}(M_2) \cong \partial M_1 \times [0,1]$.*

The proposition implies that if M is any regular neighbourhood of W in Z, then $Z \setminus \mathrm{Int}(M) \cong \partial M \times [0,1] \cong A \times [0,1]$.

3.4.3 Linking scaffolds

In this section, we give a construction for the linking scaffolds, which we use in the proof of Theorem 1.15. We first recall their definition.

Definition 3.10. *Let k, ℓ be nonnegative integers. We call a simplicial complex L a (k,ℓ)*-linking scaffold *if the following conditions hold:*

- $S^k \sqcup S^\ell \subseteq L$.
- *There exists a PL embedding $L \hookrightarrow S^{k+\ell+1}$.*
- *For any PL embedding $f\colon L \hookrightarrow S^{k+\ell+1}$, the spheres $f(S^k)$ and $f(S^\ell)$ are linked with linking number ± 1.*

For our purposes, the difference between the positive and negative linking number is not essential, as it amounts to a change in the orientations of the copies of S^k and S^ℓ inside L.

In [70] the authors present a simplicial complex P, which has similar properties to a linking scaffold.

The construction in [70] Let $k, \ell \in \mathbb{Z}$ be such that $0 \leq \ell < k$ and let $m = k + \ell + 1$. Let Δ^{m+1} be the $(m+1)$-simplex, spanned by the vertices $p_0, \dots, p_{m+1}$, by Δ^{k+1} the $(k+1)$-simplex, spanned by the vertices $p_0, \dots, p_{k+1}$ and by $\Delta^{\ell+1}$, the $(\ell+1)$-simplex spanned by the vertices $p_{k+2}, \dots p_{m+1}$.

Let $\mathsf{sk}_k\Delta^{m+1}$ denote the k-skeleton of Δ^{m+1} and let $C(\mathsf{sk}_\ell\Delta^{m+1}; p_{m+2})$ denote the cone over the ℓ-skeleton of Δ^{m+1} with apex p_{m+2}. We define $P(k,\ell) := \mathsf{sk}_k(\Delta^{m+1}) \cup_{\mathsf{sk}_\ell\Delta^{m+1}} C(\mathsf{sk}_\ell\Delta^{m+1}; p_{m+2})$. In other words, $P(k,\ell)$ is the simplicial complex, obtained from the k-skeleton of Δ^{m+1} together with the cone over the ℓ-skeleton of Δ^{m+1}.

Lemma 3.46 (Lemma 1.1 in [70])**.** *There exists a PL map $f\colon P(k,\ell) \to \mathbb{R}^m$ and distinguished spheres $S^k_P, S^\ell_P \subset P(k,\ell)$ such that $f_{S^k_P \sqcup S^\ell_P}$ is an embedding and $f(S^k_P)$ and $f(S^\ell_P)$ are linked with linking number ± 1 inside $\mathbb{R}^m$.*

[7] In the case when $M_1 = Z$, by $M_2 \subset \mathrm{Int}(M_1)$ we mean that $M_2 \subset Z$.

We briefly outline the construction of the map $f\colon P(k,\ell)\to\mathbb{R}^m$, without proving any of its properties. Consider an m-simplex $\Delta^m_q\subset\mathbb{R}^m$ with vertices $q_1,\dots,q_{m+2}$. Let q_0 be the barycenter of Δ^m_q and $C(\mathsf{sk}_{k-1}(\Delta^m_q);q_0)$ be the cone in $\mathbb{R}^m$ with vertex q_0 over the $(k-1)$-skeleton of Δ^m_q. Let $h_1\colon \mathsf{sk}_k(\Delta^{m+1})\to\mathsf{sk}_k(\Delta^m_q)\cup C(\mathsf{sk}_{k-1}(\Delta^m_q);q_o)$ be the unique PL homeomorphism, which is linear on each simplex of $\mathsf{sk}_k(\Delta^{m+1})$ and such that $h_1(p_i)=q_i$, $i=0,\dots,m+1$.

Let b be the barycenter of $(k+1)$-simplex in $\mathbb{R}^m$ with vertices $q_0,\dots,q_{k+1}$. Let $h_2\colon C(\mathsf{sk}_\ell(\Delta^{m+1});p_{m+2})\to\mathbb{R}^m$ be the map, which is linear on each simplex of the cone $C(\mathsf{sk}_\ell(\Delta^{m+1});p_{m+2})$ and such that $h_2|_{\mathsf{sk}_\ell(\Delta^{m+1})}=h_1|_{\mathsf{sk}_\ell(\Delta^{m+1})}$ and $h_2(p_{m+2}=b$. In [70] it is shown that the map $h_2(C(\mathsf{sk}_\ell(\Delta^{m+1}));p_{m+2})\setminus\mathrm{Int}(\Delta^{\ell+1}))$ is a PL embedding and that $h_2(C(\mathsf{sk}_\ell(\Delta^{m+1});p_{m+2})\setminus\mathrm{Int}(\Delta^{\ell+1}))\cap h_1(\mathsf{sk}_k(\Delta^{m+1}))$ is a subset of $h_1(\mathsf{sk}_\ell(\Delta^{m+1}))$, which implies that the following map is a PL embedding.

$$h_1\cup h_2\colon \mathsf{sk}_k(\Delta^{m+1})\cup(C(\mathsf{sk}_\ell(\Delta^m+1);p_{m+2})\setminus\mathrm{Int}(\Delta^\ell+1)).$$

Next, denote by d the barycenter of the k-simplex in $\mathbb{R}^m$ with vertices $q_0,\dots,q_k$ and by $\Delta^{\ell+1}_q\subset\mathbb{R}^m$ the $(\ell+1)$-simplex with vertices $b,q_{k+2},\dots,q_{m+1}$. Let $g\colon\Delta^{\ell+1}\to C(\partial\Delta^{\ell+1}_q;d)$ be the PL homeomorphism, which maps an interior point $c\in\mathrm{Int}(\Delta^{\ell+1})$ onto d, $g(p_i)=q_i$, $i=k+2,\dots,m+1$ and which is linear on each simplex of the triangulation of $\Delta^{\ell+1}$ with vertices $c,p_{m+2},p_{k+2},\dots,p_{m+1}$.

Finally, we define the map $f\colon P(k,\ell)\to\mathbb{R}^m\hookrightarrow S^m$ as $f=h\cup g$. For the proof that it has the desired properties, we refer to [70].

Remark 3.46.1. One could check that $g(\mathrm{Int}(\Delta^{\ell+1}))\cap\mathrm{im}(h)=g(\Delta^{\ell+1})\cap h(\Delta^k)=\{d\}$. In fact, that is the only self-intersection of the map f, as we defined it. Therefore, the restriction $f\colon P(k,\ell)\setminus\mathrm{Int}(\Delta^{\ell+1})\to S^m$ will be an embedding. Moreover, Lemma 1.4 in [70] states that for any embedding $g\colon P(k,\ell)\setminus\mathrm{Int}(\Delta^{\ell+1})\to S^m$, the spheres $g(S^k)$ and $g(S^\ell)$ have an odd linking number inside S^m.

The linking scaffolds $L(k,\ell)$ and $\mathsf{link}(k,\ell)$

Consider the simplices $\Delta^k,\Delta^{\ell+1}\subset P(k,\ell)$. We obtain the simplicial complex $L(k,\ell)$ from $P(k,\ell)$ in two steps.

1. Subdivide Δ^k by coning over its boundary with vertex x_k, and $\Delta^{\ell+1}$ by coning over its boundary with vertex x_ℓ.
2. Identify the newly introduced vertices x_k and x_ℓ.

We define the distinguished spheres $S^k_L\subset L(k,\ell)$ to be the same as the distinguished spheres of $P(k,\ell)$ with the simplex $\Delta^k\subset S^k_P$ being subdivided. In order to define the simplicial complex $\mathsf{link}(k,\ell)$, we start with $L(k,\ell)$ and glue to S^ℓ_L a cylinder $S^\ell\times[0,1]$.

$$\mathsf{link}(k,\ell)=L(k,\ell)\cup_{S^\ell_L\sim(S^\ell\times\{0\})}(S^\ell\times[0,1]).$$

The first main result of this section is the following lemma.

Lemma 3.47. *Let $k,\ell\in\mathbb{Z}$, $0\le\ell<k$. Then, the simplicial complexes $L(k,\ell)$ and $\mathsf{link}(k,\ell)$ are (k,ℓ)-linking scaffolds.*

Proof. We first prove the lemma for $L(k,\ell)$. Recall that we obtained it from $P(k,\ell)$ by identifying the points x_k and x_ℓ, which are interior for the simplices Δ^k and $\Delta^{\ell+1}$. Starting with the restricted map $f\colon P(k,\ell) \setminus (\mathrm{Int}(\Delta^k) \cup \mathrm{Int}(\Delta^{\ell+1})) \to S^m$, we define a map $g\colon L(k,\ell) \to S^m$ by setting $g(x_k = x_l) = d$, where d is the single self-intersection point of the map $f\colon P(k,\ell) \to S^k$ and extending linearly on $\mathrm{Int}(\Delta^k) \cup \mathrm{Int}(\Delta^{\ell+1})$. It is easy to check, that the map $g\colon L(k,\ell) \to S^m$ is an embedding. Moreover, following Lemma 3.46, $g(S^k)$ and $g(S^\ell)$ are linked in S^m with linking number ± 1.

What remains to be seen is that for any other embedding embedding $h\colon L(k,\ell) \hookrightarrow \mathbb{R}^m$, the spheres $h(S^k)$ and $h(S^\ell)$ will also have linking number ± 1 inside S^m. Let $D^{\ell+1} \subset L(k,\ell)$ be the disc, which bound S^ℓ. By construction, $S^k \cap D^{\ell+1} = \{x_k = x_\ell\}$, therefore, for any embedding $h \hookrightarrow L(k,\ell) \hookrightarrow \mathbb{R}^m$, the images $h(S^k)$ and $h(D^{\ell+1})$ will intersect in a single point. Following Lemma 1.4 in [70], $h(S^k)$ and $h(S^\ell)$ will be linked with an odd linking number, which means that $h(S^k)$ and $h(D^{\ell+1})$ have to intersect transversally, i.e. the following picture is not allowed.

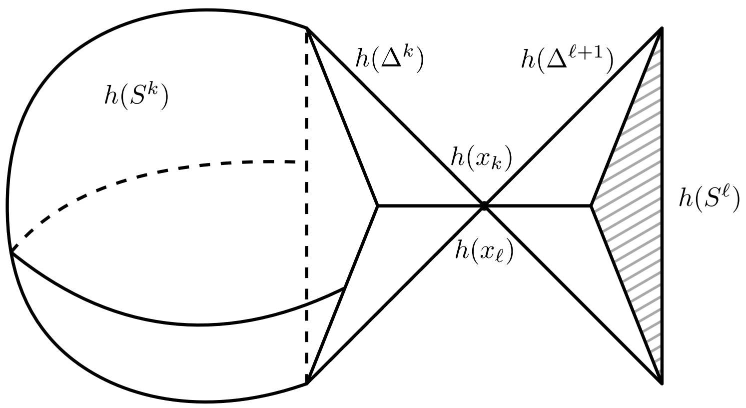

On the other hand, the linking number cannot be more than 1 (in absolute value), since $h(S^k)$ and $h(D^{\ell+1})$ only intersect once. This concludes the proof for $L(k,\ell)$.

The proof that $\mathsf{link}(k,\ell)$ is also a linking scaffold is a simple extension of the above arguments. ☐

For the proof of Theorem 1.15 we would need a stronger statement, which requires further assumptions on k and l.

Lemma 3.48. *Let $\ell \geq 2$ and $k = 2\ell + 1$. Let S_L^k, S_L^ℓ be the distinguished spheres inside $\mathsf{link}(k,\ell)$. Then, there exists an embedding $f\colon \mathsf{link}(k,\ell) \hookrightarrow T \subset \mathbb{R}^{k+\ell+1}$, where T is a simplicial complex, which is PL homeomorphic to $S^k \times \Delta^{\ell+1}$. Moreover, it has the following properties:*

- $f(S_L^\ell) \subset \partial T$
- $f(\mathsf{link}(k,\ell)) \cap \partial T = f(S_L^\ell)$

Proof. We first construct an embedding of $L(k,\ell)$ into $S^k \times \Delta^{\ell+1} \subset \mathbb{R}^{k+\ell+1}$, where the copy of S^k is standardly embedded into $\mathbb{R}^{k+\ell+1}$. Let $\widetilde{g}\colon L(k,\ell) \hookrightarrow \mathbb{R}^{k+\ell+1}$ be any embedding. By Theorem 3.23, there is an ambient isotopy ξ, taking $\widetilde{g}(S_L^k)$ to the standard embedding $S^k \hookrightarrow \mathbb{R}^{k+\ell+1}$. Let $\widetilde{\xi}$ be the final homeomorphism $\mathbb{R}^{k+\ell+1}$ of ν. Then, $\widetilde{\nu} \circ \widetilde{g}\colon L(k,\ell) \hookrightarrow \mathbb{R}^{k+\ell+1}$ is an embedding of $L(k,\ell)$, which embeds S_L^k standardly.

Next, by compactness, there exists a ball $B^{k+\ell+1} \subset \mathbb{R}^{k+\ell+1}$, such that $\widetilde{\nu} \circ \widetilde{g}(L(k,\ell) \subset \mathrm{Int}(B^{k+\ell+1})$. By choosing an appropriate triangulation of $\Delta^{\ell+1}$, we can use the inclusion $B^{k+\ell+1} \cong B^k \times \Delta^{\ell+1} \subset S^k \times \Delta^{\ell+1}$, which gives us the desired embedding, which we denote by $g\colon L(k,\ell) \hookrightarrow T$.

Let $h\colon S^\ell_L \to \partial T$ be an embedding, so that the spheres $h(S^\ell)$ and $g(S^k_L)$ are linked in $\mathbb{R}^{k+\ell+1}$ with the same linking number as the spheres $g(S^\ell_L)$ and $g(S^k_L)$. By connectivity, the maps $g|_{S^\ell_L}, h\colon S^\ell_g \to T$ are homotopic. Let $H\colon S^\ell_L \times [0,1] \to T$ be a homotopy between them, which, by general position, we may assume to be an embedding of $S^\ell_L \times [0,1]$ into T.

The image $H(S^\ell_L \times [0,1])$ will intersect $g(X)$ in a finite set of isolated points, which will be disjoint from $H(S^\ell_L \times \partial[0,1])$. Let $\epsilon > 0$ be small enough, so that $H(S^\ell_L \times [\epsilon,1]) \cap g(X) = \emptyset$. The embeddings $g|_{S^\ell_L}, H|_{S^\ell_L \times \{\epsilon\}}\colon S^\ell \to T$ are isotopic and by the result of Hudson and Zeeman [36], since the codimension is at least three, they are also ambient isotopic. Let μ be an ambient isotopy of T between $g|_{S^\ell_L}$ and $H|_{S^\ell_L \times \{\epsilon\}}$ and let $\widetilde{\mu}$ be the final homeomorphism of μ. We define the desired embedding $f\colon \mathrm{link}(k,\ell) \to T$ as $f_{L(k,\ell)} := \widetilde{\mu}(g(L(k,\ell)))$ and $f_{S^\ell \times [0,1]} := \widetilde{\mu}(H(S^\ell_L \times [\epsilon,1]))$. □

Wedge linking scaffolds

We now present a generalisation of the linking scaffolds. Let $Y := S^\ell_1 \vee S^\ell_2$ and $\Sigma = S^k_1 \vee S^k_2$.

Definition 3.49. *Let k,ℓ be nonnegative integers. We call a simplicial complex L a (k,ℓ)-*wedge linking scaffold *if the following conditions hold:*

- $Y \sqcup \Sigma \subset L$.
- *There exists a PL embedding $L \hookrightarrow S^m$, where $m = k+\ell+1$.*
- *If $f\colon L \hookrightarrow S^m$ is any PL embedding, then the spheres $f(S^k_1)$ and $f(S^\ell_1)$, are linked in S^m with linking number ± 1, the spheres $f(S^k_2)$ and $f(S^\ell_2)$, are linked in S^m with linking number ± 1 and $f(S^k_i)$ is not linked with $f(S^\ell_j)$ if $i \neq j$.*

We now present an explicit construction of a wedge linking scaffold for a pair (k,ℓ), based on the construction of a linking scaffold. In fact, we will simply take two linking scaffolds $L_1(k,\ell)$ and $L_1(k,\ell)$, one for each pair (S^k_1, S^ℓ_1) and (S^k_2, S^ℓ_2) and glue them together. More precisely, we define the wedge linking scaffold $L_w(k,\ell)$ as

$$L_w(k,\ell) = L(k,\ell)_1 \cup_{p^1_0 = p^2_0,\, p^1_{m+1} = p^2_{m+1}} L(k,\ell)_2 \tag{3.4}$$

Here p^1_0, p^2_0, p^1_{m+1} and p^2_{m+1} are vertices from $L(k,\ell)_1$ and $L(k,\ell)_2$ such that $p^1_0 \in S^k_{L_1}, p^1_{m+1} \in S^\ell_{L_1}, p^2_0 \in S^k_{L_2}, p^2_{m+1} \in S^\ell_{L_2}$. In a similar way, we define the wedge linking scaffold $\mathrm{link}_w(k,\ell)$.

$$\mathrm{link}_w(k,\ell) = L_w(k,\ell) \cup_{S^\ell_{L_1} \vee S^\ell_{L_2} \sim (S^\ell \vee S^\ell \times \{0\})} ((S^\ell \vee S^\ell) \times [0,1]).$$

The simplicial complexes $L_w(k,\ell)$ and $\mathrm{link}_w(k,\ell)$ have very similar properties to $L(k,\ell)$ and $\mathrm{link}(k,\ell)$.

Lemma 3.50. *Let $k, \ell \in \mathbb{Z}$, $0 \leq \ell < k$. Then, the simplicial complex $L_w(k,\ell)$ is a (k,ℓ)-wedge linking scaffolds.*

Proof. The lemma follows from Lemma 3.47. More precisely, first it is clear that $L_w(k,\ell)$ embeds into S^m. Moreover, observe that the only points the two copies $X_1(k,\ell), X_2(k,\ell) \subset L_w(k,\ell)$ share, are the the base points of the two wedges $S_1^k \vee S_2^k$ and S_1^ℓ, S_2^ℓ. That implies that, by construction, for any embedding $f\colon L_w(k,\ell) \to S^m$, the sphere $f(S_1^k)$ and $f(S_1^\ell)$ are linked inside S^m with linking number ± 1 and the same is true for the spheres $f(S_2^k)$ and $f(S_2^\ell)$. On the other hand, if we consider any other pair of spheres, say S_1^k and S_2^ℓ, if Δ_1^{k+1} and $\Delta_2^{\ell+1}$ are the discs they bound inside $L_w(k,\ell)$, then $f(\Delta_1^{k+1}) \cap \Delta_2^{\ell+1} = \emptyset$, which implies that $f(S_1^k)$ and $f(S_2^\ell)$ cannot be linked. That concludes the proof. $\square$

As a consequence of Lemma 3.50 and Lemma 3.48, we also obtain the following.

Lemma 3.51. *Let $\ell \geq 2$ and $k = 2\ell + 1$. Let $(S^k \vee S^k)_L, (S^\ell \vee S^\ell)_L$ be the distinguished wedges inside $\mathsf{link}_w(k,\ell)$. Then, there exists an embedding $f\colon \mathsf{link}_w(k,\ell) \hookrightarrow T \subset \mathbb{R}^{k+\ell+1}$, where T is a simplicial complex, which is PL homeomorphic to $(S^k \vee S^k) \times \Delta^{\ell+1}$. Moreover, it has the following properties:*

- $f((S^\ell \vee S^\ell)_L) \subset \partial T$
- $f(\mathsf{link}_w(k,\ell)) \cap \partial T = f((S^\ell \vee S^\ell)_L)$

3.5 The proof of Lemma 3.13

In this section we prove the following technical lemma, which was crucial for the proofs of the codimension two cases of Theorem 1.14 and Theorem 1.15.

Lemma 3.13. *Let K be a finite simplicial complex, $\tau \in K$ a maximal ℓ-simplex of K and $d \geq \ell + 3$ an integer. Let $\widetilde{K} := K \cup_\tau \Delta^{\ell+1}$ be the complex, obtained from K by coning over τ. Then, there is an embedding $\widetilde{K} \hookrightarrow S^d$ if and only if there is an embedding $K \hookrightarrow S^d$.*

Proof. In one direction the lemma is trivial, so we only consider the opposite direction. Let $f\colon K \hookrightarrow S^d$ be an embedding. We are going to construct an embedding $F\colon \widetilde{K} \hookrightarrow S^d$. Let $N = N(f(\tau), S^d) \mathrm{mod} f(\partial\tau)$ be a relative regular neighbourhood of $f(\tau)$ modulo $f(\partial\tau)$ in S^d. Since, $f(\tau)$ is a PL ℓ-ball in S^d, N is a d-ball, with the property that $N \cap f(K \setminus \mathrm{Int}(\tau)) = \partial N \cap f(K \setminus \mathrm{Int}(\tau)) = f(\partial\tau)$, i.e. $f(\tau)$ is a standardly embedded ℓ-ball inside the d-ball N. Since $d \geq \ell + 3$. By Theorem 9 in Chapter 4 of [79], $f(\tau)$ unknots in N. That implies that there is an ambient isotopy $H\colon N \times [0,1] \to N$, keeping ∂N fixed, which takes $f(\tau)$ to the cone $H_1(f(\tau)) = C(f(\partial\tau)) \subset N$. Let $p \in \mathrm{Int}(N)$ be a point, disjoint from $H_1(f(\tau))$. Then, the cone $C(H_1(f(\tau), p)$ with apex p is PL homeomorphic to the standard $(\ell+1)$-simplex. We define the desired embedding $F\colon \widetilde{K} \hookrightarrow S^d$ by $F|_{K \setminus \mathrm{Int}(\tau)} := f$ and $F(\Delta^{\ell+1}) := C(H_1(f(\tau), p)$. □

3.6 The proof of theorem 3.7

First, we remind our notation. Let $A := S^{2m-1} \vee \ldots \vee S^{2m-1}$, $W := S^m \vee \ldots \vee S^m$ and $X :=$ DCyl$(\phi; f)$ be the double mapping cylinder of the maps $\phi\colon A \to W$ and $f\colon A \to Y$, where Y is either S^m or $S^m \vee S^m$, depending on whether m is even or odd. All the proofs we present work in the same way for both cases. We consider A, W and Y as being included into X as the *central section*, the *lower rim* and the *upper rim* of the double mapping cylinder. Let Σ be either S^{2k+1} or $S^{2k+1} \vee S^{2k+1}$, depending on whether m is even or odd. Once again all the proofs will work for both cases. We are also given a *standard* embedding $\Sigma \hookrightarrow S^{3k+2}$, as constructed in Section 3.4, so we consider Σ as a subset of S^{3k+2} under this embedding. Let Q denote S^{3m+2} with a small open neighbourhood $\mathrm{Int}(N(\Sigma))$ of Σ removed.

We start with a map $f\colon X \to Q$, which is an embedding on $A \subset X$, and we want to homotope it to an embedding, such that the homotopy keeps A fixed. Moreover, by general position we can assume $f|_W$ and $f|_Y$ to be embeddings as well.

The proof of the theorem is essentially a combination of Proposition 3.52 and Proposition 3.55. The idea is to first construct small regular neighbourhoods of the W and Y in X, and homotope f to a map $g\colon X \to Q$, which embeds those neighbourhoods in Q. Thus, we make sure that $S(g)$ will be disjoint from the upper and lower rim of the double mapping cylinder X. Most importantly, being away from W and Y, $S(g)$ will be contained in a part of X, which is homeomorphic to $A \times [0,1]$. Once we have that, we can use the particularly nice structure of $A \times [0,1]$ and resolve all the self-intersections of the map g, thus ending up with the desired embedding.

Proof of theorem 3.7. Let M_Y, M_W be the regular neighbourhoods of Y and W in X, given by proposition 3.52. Let N_Y, N_W be regular neighbourhoods of $f(Y)$ and $f(W)$ in Q such that $M_Y = f^{-1}(N_Y)$ and $M_W = f^{-1}(N_W)$. From the proposition, we know that f is homotopic to a map $g\colon X \to Q$, such that, $g|_{M_Y}$ is a proper embedding into N_Y, and $g|_{M_W}$ is a proper embedding into N_W. Let $\widetilde{X} := X \setminus (\mathrm{Int}(M_Y) \cup \mathrm{Int}(M_W))$. By remark 3.53.2, we can assume that $S(g) \cap \partial\widetilde{X} = \emptyset$, i.e. the self-intersections of the map g are away from $\partial\widetilde{X}$.

Define $Z := Q \setminus (\mathrm{Int}(N_Y) \cup \mathrm{Int}(N_W))$ and observe, that $g|_{\widetilde{X}}$ can be regarded as a map $g|_{\widetilde{X}}\colon \widetilde{X} \to Z$. By Lemma 3.54, Z is a $(3m+2)$-dimensional $(m-1)$-connected PL manifold. Therefore, we can apply Proposition 3.55 to the map $g|_{\widetilde{X}}$ and homotope it to an embedding $g_1\colon \widetilde{X} \hookrightarrow Z$. Observe that, by using Z as a target, we ensure that the homotopy $H : g \sim g_1$ does not affect the neighbourhoods M_Y and M_W, which were already embedded. Finally, we define the embedding $X \hookrightarrow Q$ to be equal to g on $M_Y \cup M_W$ and to g_1 on $\widetilde{X}$. This is consistent, since the homotopy H preserves small neighbourhoods of $\partial\widetilde{X}$, so g and g_1 coincide on $\partial\widetilde{X}$. We illustrate the construction on Figure 3.6.

□

3.6.1 Pushing $S(f)$ away from the boundary of X

We first show how to homotope f to a map g, which is an embedding when restricted to small neighbourhoods of W and Y in X. Once we have that, in the next section

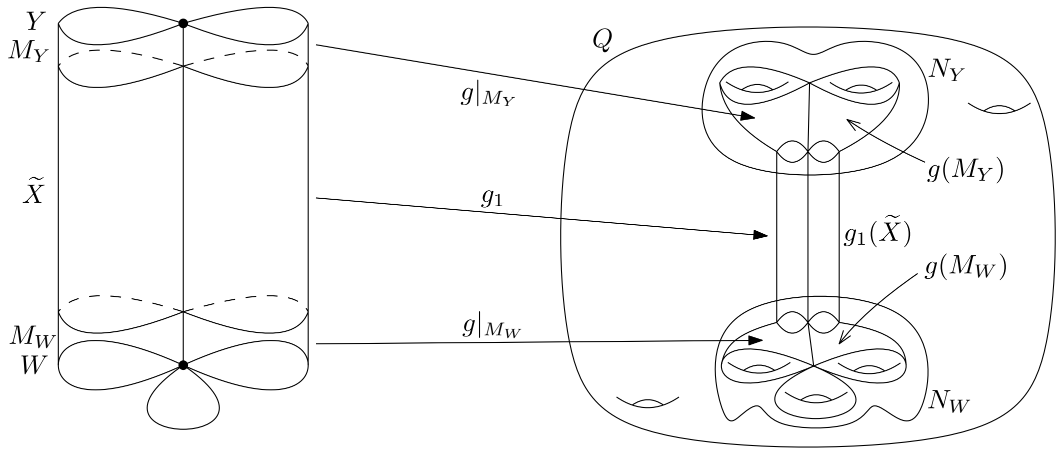

Figure 3.3: The final embedding

we carefully resolve the remaining self-intersections of g. We begin with the following proposition.

Proposition 3.52. *There exists a PL map $g\colon X \to Q$ and regular neighbourhoods M_W, M_Y of W and Y in X, with the following properties:*

- *g is homotopic to f.*
- *$g|_W = f|_W$ and $g|_Y = f|_Y$.*
- *$g|_{M_W}$ is an embedding.*
- *$g|_{M_Y}$ is an embedding.*
- *$S(g) \cap (M_Y \cup M_W) = \emptyset$.*

We will construct the map g in several steps. First, we need the following lemma.

Lemma 3.53. *There exist a regular neighborhood $\overline{N_W}$ of $f(W)$ in S^{3m+2}, which misses Σ, a small regular neighborhood $\overline{M_W}$ of W in X, and an embedding $m\colon \overline{M_W} \hookrightarrow \overline{N_W}$ that agrees with f on W (i.e., $m|_W = f|_W$) and such that $m^{-1}(\overline{N_W}) = \overline{M_W}, m^{-1}(\partial\overline{N_W}) = \partial\overline{M_W}$.*

Moreover, the restrictions $m|_{\partial M_W}$ and $f_{\partial\overline{M_W}}$ are homotopic as maps to $\partial\overline{N_W}$.

Proof. By lemma 3.32 there is an embedding $g\colon X \hookrightarrow S^{3m+2}$. Moreover, by proposition 3.25, we know that $g|_W$ and $f|_W$ are ambient isotopic. Let $h\colon S^{3m+2} \cong S^{3m+2}$ be the final homeomorphism of this isotopy; in particular $g|_W = h \circ f|_W$.

Since f is in general position, $f(W) \cap f(\Sigma) = \emptyset$. Therefore, $g(W) = h \circ f(W)$ misses $h \circ f(\Sigma)$. Pick a regular neighbourhood $\widetilde{N_W}$ of $g(W)$ in S^{3m+2} small enough, so it misses $h \circ f(\Sigma)$. Pick triangulations of X and S^{3m+2}, so that g is simplicial. Let $\widetilde{M_W} = g^{-1}(\widetilde{N_W})$. By Lemma 3.41 (for $L = \emptyset$), $\widetilde{M_W}$ is a regular neighbourhood of W in X.

Then $\overline{N_W} := h^{-1}(\widetilde{N_W})$ is a regular neighbourhood of $h^{-1}(g(W)) = f(W)$ in S^{3m+2} that misses
$h^{-1}(h(f(\Sigma))) = f(\Sigma)$, and $\widetilde{m} := h^{-1} \circ g|_{\widetilde{M_W}}$ is an embedding of $\widetilde{M_W}$ into $\overline{N_W}$ that agrees

with f on W. Moreover, by construction, it satisfies $\widetilde{m}^{-1}(\overline{N_W}) = \widetilde{M_W}$, $\widetilde{m}^{-1}(\partial\overline{N_W}) = \partial\widetilde{M_W}$.

Let $M := f^{-1}(\overline{N_W})$. By lemma 3.41, $f^{-1}(\partial\overline{N_W}) = \partial\overline{M_W}$,[8] and both $\overline{M_W}$ and $\widetilde{M_W}$ are regular neighborhoods of W in X. By the uniqueness theorem for regular neighbourhoods (theorem 3.39), $\overline{M_W}$ and $\widetilde{M_W}$ are ambient isotopic in X, keeping W fixed. Let $\ell\colon X \to X$ be the final homeomorphism of this isotopy, i.e., $\ell(\overline{M_W}) = \widetilde{M_W}$, with the restriction of ℓ to W being the identity. Set $m := \widetilde{m} \circ \ell$. Then m is an embedding of $\overline{M_W}$ into $\overline{N_W}$ that agrees with f on W, and $m^{-1}(\overline{N_W}) = \overline{M_W}$, $m^{-1}(\partial\overline{N_W}) = \partial\overline{M_W}$.

Finally, we claim that the restriction $m|_{\partial\overline{M_W}}$ is homotopic to $f|_{\partial\overline{M_W}}$ as a map into $\partial\overline{N_W}$. To see this, first note that m and f are homotopic as maps $\overline{M_W} \to \overline{N_W}$. This is true, since $\overline{M_W} \searrow W$, $\overline{N_W} \searrow f(W)$, and $f(W) = m(W)$. Let $H\colon \overline{M_W} \times [0,1] \to \overline{N_W}$ be a PL homotopy such that $H_W(\bullet, t) = f$, $0 \leq t \leq 1$. Consider its restriction $H|_{\partial\overline{M_W}}$, which gives a homotopy $f|_{\partial\overline{M_W}} \cong m|_{\partial\overline{M_W}} \to \overline{N_W}$. By general position, $H(\partial\overline{M_W} \times [0,1])$ misses $f(W)$, as well as a small regular neighbourhood N_1 of it. By Corollary 2 to Theorem 8 in [79], $\overline{N_W} \setminus \mathrm{Int}(N_1) \cong \partial\overline{N_W} \times [0,1]$. Let $r\colon \overline{N_W} \setminus \mathrm{Int}(N_1) \to \partial\overline{N_W}$ be a deformation retraction. Then, the composition $r \circ (H|_{\partial\overline{M_W} \times [0,1]})$ of the retraction with the original homotopy gives us the desired homotopy between $m|_{\partial\overline{M_W}}$ and $f|_{\partial\overline{M_W}}$, as maps to $\partial\overline{N_W}$. □

Remark 3.53.1. The same argumentation also works for the lower rim $Y \subset X$. Namely, there exist a regular neighborhood $\overline{N_Y}$ of $f(Y)$ in S^{3m+2}, which misses Σ, a small regular neighborhood $\overline{M_Y}$ of Y in X, and an embedding $m^Y\colon \overline{M_Y} \hookrightarrow \overline{N_Y}$ that agrees with f on Y (i.e., $m^Y|_Y = f|_Y$) and such that $m^{-1}(\overline{N_Y}) = \overline{M_Y}$, $(m^Y)^{-1}(\partial\overline{N_Y}) = \partial\overline{M_Y}$.

Proof of Proposition 3.52. First, observe that by proposition 3.45, $\widetilde{X} := X \setminus (\mathrm{Int}(\overline{M_Y}) \cup \mathrm{Int}(\overline{M_W})) \cong A \times [0,1]$. Let $\widetilde{f}\colon \widetilde{X} \to Q$ be the restriction of f to $\widetilde{X}$.

In order to define g, we will separate X in several pieces, and define g on each one, so that we can fit them all into one map with the desired properties. First, define g on $\overline{M_Y}$ and $\overline{M_W}$ to be equal to their embeddings into Q, as given by lemma 3.53 and remark 3.53.1.

Next, let $H_W\colon (\partial\overline{M_W} = A) \times [0,1] \to \partial\overline{N_W} \subset Q$ and $H_Y\colon (\partial\overline{M_Y} = A) \times [0,1] \to \partial\overline{N_Y} \subset Q$ be PL homotopies between $f|_{\partial\overline{M_W}} \sim g|_{\partial\overline{M_W}}$ and $f|_{\partial\overline{M_Y}} \sim g|_{\partial\overline{M_Y}}$, respectively, as provided by lemma 3.53 and remark 3.53.1.

Now, for some $\epsilon > 0$, divide the interval $[0,1] = [0,\epsilon] \cup [\epsilon, 1-\epsilon] \cup [1-\epsilon, 1]$, which also gives us a decomposition $\widetilde{X} = \widetilde{X}_1 \cup \widetilde{X}_2 \cup \overline{X}$, where each of the three pieces is a smaller copy of the cylinder $A \times [0,1]$ inside $\widetilde{X}$. In particular, each of those small cylinders is PL homeomorphic to $\widetilde{X}$. Denote by $\overline{H}\colon \overline{X} \cong \widetilde{X}$ the PL homeomorphism, which stretches the middle segment to the whole cylinder $\widetilde{X}$. Finally, we define $g|_{\widetilde{X}_1} := H_Y$, $g|_{\widetilde{X}_2} := H_W$ and $g|_{\overline{X}} := \widetilde{f} \circ \overline{H}$. This gives us a map $g\colon X \to Q$, which, by construction, is homotopic to f. We illustrate the map g on Figure 3.6.1.

Observe that the image of H_W is contained in $\partial\overline{N_W}$, therefore it is disjoint from $\mathrm{Int}(\overline{N_W})$, and the same is true for H_Y and $\partial\overline{N_Y}$. Moreover, from the way we constructed $\overline{M_Y}$ and $\overline{M_W}$, $S(\widetilde{f}) \cap (\mathrm{Int}(\overline{M_W}) \cup \mathrm{Int}(\overline{M_Y})) = \emptyset$). We thus see that $g|_{\mathrm{Int}(\overline{M_W}) \cup \mathrm{Int}(\overline{M_Y})}$ is

[8] Here $\partial\overline{M_W}$ denotes the topological boundary of $\overline{M_W}$ in X.

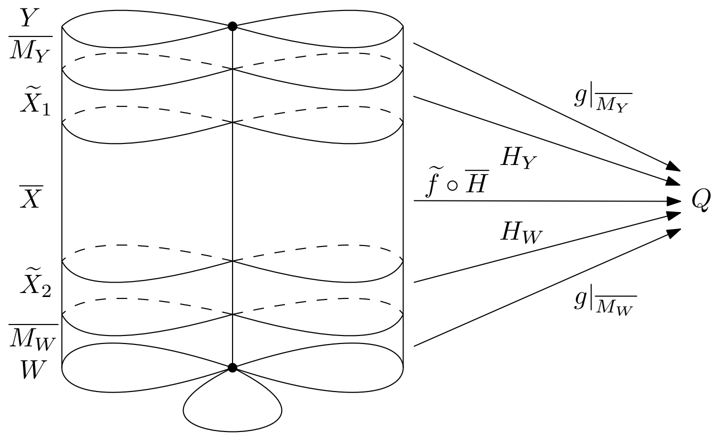

Figure 3.4: The map g

an embedding. Let $M_Y \subset \operatorname{Int}(\overline{M_Y})$ and $M_W \subset \operatorname{Int}(\overline{M_W})$ be smaller regular neighbourhoods of Y and W in X. Then, $g|_{M_Y \cup M_W}$ is an embedding. Observe also that, by construction, the property $S(g) \cap (M_W \cup M_Y) = \emptyset$ is satisfied. This concludes the proof. $\square$

Remark 3.53.2. Let M_Y, M_W be the regular neighbourhoods of Y and W in X, given by Proposition 3.25, which the map $g\colon X \to Q$ embeds into Q. Consider smaller regular neighbourhoods $\widetilde{M_Y} \subset \operatorname{Int}(M_Y), \widetilde{M_W} \subset \operatorname{Int}(M_W)$ of Y and W in X. They will have the same properties as the original ones. In particular there will be regular neighbourhoods $\widetilde{N_Y} \subset \operatorname{Int}(N_Y), \widetilde{N_W} \subset \operatorname{Int}(N_W)$ of $g(Y) = f(Y)$ and $g(W) = f(W)$ in Q such that $\widetilde{M_Y} = g^{-1}(\widetilde{N_Y})$ and $\widetilde{M_W} = g^{-1}(\widetilde{N_W})$ and $g|_{\widetilde{M_Y} \cap \widetilde{M_W}}$ will be an embedding. Now, if we denote $\widetilde{X} := X \setminus (\operatorname{Int}(\widetilde{M_Y}) \cup \operatorname{Int}(\widetilde{M_W}))$, by removing those smaller regular neighbourhoods, we ensure that the map g embeds also a small regular neighbourhood of $\partial\widetilde{X}$, namely $M_Y \setminus \operatorname{Int}(\widetilde{M_Y})$ and $M_W \setminus \operatorname{Int}(\widetilde{M_W})$.

Lemma 3.54. *Let N_Y and N_W be regular neighbourhoods of $g(Y) = f(Y)$ and $g(W) = f(W)$ in Q. Then $Z := Q \setminus (\operatorname{Int}(N_Y) \cup \operatorname{Int}(N_W))$ is a $(3m+2)$-dimensional $(m-1)$-connected PL manifold with boundary.*

Proof. The fact that Z is a PL manifold is guaranteed by Lemma 17 in [79].

In order to see that Z is $(m-1)$-connected, first observe that, since $N_Y \searrow f(Y)$ and $N_W \searrow f(W)$, Z is homotopy equivalent to $\widetilde{Z} := Q \setminus (f(Y) \cup f(W))$. Now, let $\psi\colon S^{m-1} \to \widetilde{Z}$ be a map. Since $\widetilde{Z} \subset Q$, we can think about ψ as a map $\psi\colon S^{m-1} \to Q$. Since Q is $(m-1)$-connected, there is a map $\Psi\colon D^m \to Q$ such that $\Psi|_{\partial D^m} = \psi$. By general position, $\Psi(D^m) \subset Q$ will miss $f(A) \cap f(W)$, so Ψ can also be seen as a map $\Psi\colon D^m \to \widetilde{Z}$ with $\Psi|_{\partial D^m} = \psi\colon S^{m-1} \to \widetilde{Z}$. Therefore ψ is nullhomotopic, $\widetilde{Z}$ is $(m-1)$-connected and so is Z. That concludes the proof. $\square$

3.6.2 Resolving the singularities of g

We first prove a slightly more general statement, which contains the core arguments we need. Following the notation we have already established, we set $A = S^{2m-1} \vee \ldots \vee S^{2m-1}$.

Proposition 3.55. *Let $g\colon A\times[0,1]\to Y$ be a PL map, where Y is a $(3m+2)$-dimensional $(m-1)$-connected PL manifold and $S(g)\cap A\times\partial[0,1]=\emptyset$. Then g is homotopic to an embedding. Moreover, the homotopy can be chosen so, that it keeps $A\times([0,\epsilon]\cup[1-\epsilon,1])$ fixed, for some sufficiently small $\epsilon>0$.*

Observe that, since $S(g)\cap A\times\partial[0,1]=\emptyset$, there exists a $\widetilde{I}\subset I$ such that $S(g)\cap I\subset\widetilde{I}$.

Our strategy for the proof of Proposition 3.55 is to use the fact that $A\times[0,1]$ has a particularly nice structure. It consists of a finite number of cylinders $S^{2m-1}\times[0,1]$, each of which is a PL manifold, glued along a common line segment. We divide the construction into two separate lemmas, but before we state and prove them, we begin with the following remark.

Remark 3.55.1. Since $\dim(A\times[0,1])=2m$ and $\dim Y=3m+2$ ($m\geq 5$), in general position, the map g does not have any triple intersection points. Therefore, if we pick triangulations of $A\times[0,1]$ and Y, so that g is simplicial and $S(g)$ is a subcomplex of $A\times[0,1]$, then $S(g)$ consists of pairs of simplices $\xi_1,\xi_2\in A\times[0,1]$, such that $g(\xi_1)=g(\xi_2)$. Thus, the singular set $S(g)$ of g decomposes into a union of the following:

- For every cylinder $\sigma=S^{2m-1}\times[0,1]$, the self-intersection set $S_\sigma:=S(g)\cap\sigma$ of $g|_\sigma$
- For every pair of cylinders σ,τ, the intersection set $S_{\sigma\tau}:=g^{-1}(g(\sigma)\cap g(\tau))\subset\sigma\cup\tau$.

Moreover, any two of those subsets can only meet on $I:=\{*\}\times[0,1]$, which by general position, we can assume to be embedded by the map g. That means that, while for any $p\in I$, $g^{-1}(g(p))=\{p\}$, it is possible that there is a sequence of double points of g, none of which belongs to I, which converges to p. Therefore, it will be important that all the homotopies we construct would be keeping I fixed.

Remark 3.55.2. Using the same argumentation as in the proof of Proposition 3.55, we can prove that any map $g\colon A\to Y$, where Y is a $(3m+2)$-dimensional $(m-1)$-connected PL manifold, is homotopic to an embedding.

The first lemma, we are going to prove, deals with the self-intersection of the map g, restricted to a single cylinder $S^{2m-1}\times[0,1]\subset A\times[0,1]$.

Lemma 3.56. *Let $\sigma:=S^{2m-1}\times[0,1]$ be one of the cylinders of $A\times[0,1]$. Then, we can homotope g to a map $g_1\colon A\times[0,1]\to Y$, such that $g_1|_\sigma$ is an embedding, $S(g_1)\subset S(g)$. Moreover, the homotopy between g and g_1 can be chosen so that it is constant outside σ and on $I\cup S^{2m-1}\times([0,\epsilon]\cup[1-\epsilon,1])$, for some sufficiently small $\epsilon>0$. That means we can resolve all self-intersections of $g|_\sigma$ without introducing any new self-intersections for the map g.*

Proof. Pick triangulations of $A\times[0,1]$ and Y, so that g is simplicial, and $S(g)$ is a subcomplex of $A\times[0,1]$. Since $\dim(S_\sigma\cup\widetilde{I})=m-2$, and σ is $(2m-2)$-connected, the inclusion $S_\sigma\cup\widetilde{I}\hookrightarrow\sigma$ is nullhomotopic. By general position, we can embed the cone $C(S_\sigma\cup\widetilde{I})$ over $S_\sigma\cup\widetilde{I}$ in σ, so that $C(S_\sigma\cup\widetilde{I})\cap(S(g)\setminus S_\sigma)\subset\widetilde{I}$. Consider $g(C(S_\sigma\cup\widetilde{I}))\subset Y$. Since Y is $(m-1)$-connected and $\dim g(C(S_\sigma\cup\widetilde{I}))\leq m-1$, the inclusion $g(C(S_\sigma\cup\widetilde{I}))\hookrightarrow Y$ is nullhomotopic. By general position, we can embed the cone $C:=Cg(C(S_\sigma\cup\widetilde{I}))$ in Y, so that $C\cap g(S(g)\setminus S_\sigma)\subset g(\widetilde{I})$ and $C\cap g(\mathrm{Cl}(A\times[0,1]\setminus\sigma))=g(\widetilde{I})$.[9]

[9]Here we also use the fact that by general position $g|_I$ is an embedding.

Let N be a relative regular neighbourhood of $C \bmod g(\widetilde{I})$ in Y, which is small enough, so that:

- $N \cap g(S(g) \setminus S_\sigma) \subset \widetilde{I}$
- $N \cap g(\mathrm{Cl}((A \times [0,1]) \setminus \sigma) \subset g(\widetilde{I})$
- $N \cap g(\partial\sigma) = \emptyset$

By the properties of relative regular neighbourhoods, $g(\widetilde{I}) \subset \partial N$ and by Theorem 5 in [79], since C is collapsible, N is a $(3m+2)$-ball.

Let $M := g^{-1}(N)$. Then, by lemma 3.41, M is a relative regular neighbourhood of $C(S_\sigma \cup \widetilde{I}) \bmod \widetilde{I}$ in σ, which is a $2m$-ball, since $C(S_\sigma \cup \widetilde{I})$ is collapsible. Moreover, by construction $M \cap (S(g) \setminus S_\sigma) \subset \widetilde{I} \subset \partial M$, $M \cap S^{2m-1} \times ([0,\epsilon] \cup [1-\epsilon, 1]) = \emptyset$, for a sufficiently small $\epsilon > 0$ and $g|_{\partial M}$ is an embedding of ∂M into ∂N. Finally, we modify g inside $\mathrm{Int}(M)$ by extending linearly from ∂M to a proper embedding of M into N, and denote the resulting map by $g_1 \colon A \times [0,1] \to Y$. Observe that, since $M \cap I \subset \partial M$, we do not modify g on I or outside σ. That completes the proof. □

It is important that the homotopy keeps I fixed, since we want to apply it for every cylinder separately, and the different cylinders share the line segment I.

The other ingredient in the proof of theorem 3.55 is a tool to resolve intersections between two different cylinders σ and τ within $A \times [0,1]$. This is ensured by the following lemma.

Lemma 3.57. *Let σ, τ be two different copies of $S^{2m-1} \times [0,1]$ in $A \times [0,1]$, for which $g(\sigma) \cap g(\tau) \neq \emptyset$. Then, we can homotope the map g to a map $g_1 \colon A \times [0,1] \to Y$, so that $g_1(\sigma) \cap g_1(\tau) = \emptyset$, and $S(g_1) \subset S(g)$. Moreover, the homotopy between g and g_1 can be chosen so that it is constant outside $\sigma \cup \tau$ and keeps I and $(S^{2m-1} \vee S^{2m-1}) \times ([0,\epsilon] \cup [1-\epsilon, 1])$ fixed, for some sufficiently small $\epsilon > 0$.*

Proof. First, by lemma 3.56, we assume that $g|_\sigma$ and $g|_\tau$ are embeddings.

By the assumptions on the map g (see the statement of Proposition 3.55), $S_{\sigma\tau} \cap \partial I = \emptyset$. We also recall that $\widetilde{I} \subset I$ was chosen such that $S(g) \cap I \subset \widetilde{I}$.

We have $\dim(S_\sigma \cup \widetilde{I}) \leq m-2$ and σ is $(2m-2)$-connected, so the inclusion map $S_\sigma \cup \widetilde{I} \hookrightarrow \sigma$ is nullhomotopic. Let $C_\sigma := C(S_\sigma \cup \widetilde{I})$ be the cone over $S_\sigma \cup \widetilde{I}$. By general position, C_σ embeds in σ, so that $C_\sigma \cap S(g) \subset S_\sigma \cup \widetilde{I}$ and C_σ misses the images under g of all cylinders of $A \times [0,1]$ other than σ and τ, except possibly at $\widetilde{I}$. Let C_τ be the embedded cone over $S_\tau \cup \widetilde{I}$ in τ with similar properties.

Next, consider $g(C_\sigma \cup C_\tau) \subset Y$. This is a $(m-1)$-dimensional polyhedron in the $(m-1)$-connected manifold Y, so the inclusion map $g(C_\sigma \cup C_\tau) \hookrightarrow Y$ is nullhomotopic. By general position, we can embed the cone $C := C(g(C_\sigma \cup C_\tau))$ in Y, so that $C \cap g(A \times [0,1]) = g(C_\sigma \cup C_\tau)$.

Pick triangulations of $A \times [0,1]$ and Y, so that C_σ and C_τ are full subcomplexes of $A \times [0,1]$, C is a full subcomplex of Y, and g is simplicial. Pick first derived subdivisions $(A \times [0,1])'$ and Y' of $A \times [0,1]$ and Y, so that g is still simplicial. Then, following definition 3.38, the first derived neighbourhood $N := N(C \setminus g(\widetilde{I}), Y')$ is a relative regular neighbourhood of $C \bmod g(\widetilde{I})$ in Y'.

Let g_σ, g_τ be the restrictions of g to σ and τ , respectively. It is easy to see that the pair $(C, \widetilde{I})$ satisfies the conditions of lemma 3.41, which then implies that $M_\sigma := g_\sigma^{-1}(N)$ and $M_\tau := g_\tau^{-1}(N)$ are relative regular neighbourhoods of $C_\sigma \bmod \widetilde{I}$ in σ and $C_\tau \bmod \widetilde{I}$ in τ. Moreover, again by lemma 3.41, $g_\sigma(M_\sigma)$ and $g_\tau(M_\tau)$ are proper embeddings of M_σ and M_τ into N. Since C_σ , C_τ and C are all collapsible, M_σ and M_τ are $2m$-balls and N is a $(3m+2)$-ball, so that $\widetilde{I} \subset \partial M_\sigma$, $\widetilde{I} \subset \partial M_\tau$ and $g(\widetilde{I}) \subset \partial N$.

According to proposition 3.24, $\widetilde{I} \cong B^1$ is unknotted in ∂M_σ and ∂M_τ. By corollary 3.25, $\partial M_\sigma \cup_{\widetilde{I}} \partial M_\tau$ [10] is unknotted in ∂N, therefore we can ambient isotope the embedding $g|_{\partial M_\sigma \cup_{\widetilde{I}} \partial M_\tau} \colon \partial M_\sigma \cup_{\widetilde{I}} \partial M_\tau \hookrightarrow \partial N$ to the standard one,[11] which we denote by $g_1 \colon \partial M_\sigma \cup_{\widetilde{I}} \partial M_\tau \hookrightarrow \partial N$. We can now extend g_1 to a proper embedding $g_1 \colon M_\sigma \cup_{\widetilde{I}} M_\tau \hookrightarrow N$.

We obtain our desired map by setting it to be equal to g on $(A \times [0,1]) \setminus \mathrm{Int}(M_\sigma) \cup \mathrm{Int}(M_\tau)$, and to g_1 on $\mathrm{Int}(M_\sigma) \cup \mathrm{Int}(M_\tau)$. The modified map clearly has the desired properties. In order to see that it is also homotopic to the original map g, observe, that we only modify the original map g inside the contractible subspace $M_\sigma \cup_{\widetilde{I}} M_\tau$ of $A \times [0,1]$, so that the image of this contractible subspace under both g and the modified map is contained in the ball N, which is also contractible. Therefore, the original map g and the modified map are homotopic, which concludes the proof. □

Proof of Proposition 3.55. First we apply lemma 3.56 for each of the cylinders of $A \times [0,1]$. Then, for every pair of cylinders, we apply lemma 3.57, thus obtaining the desired embedding. □

[10] The polyhedron $\partial M_\sigma \cup_{\widetilde{I}} \partial M_\tau$ consists of ∂M_σ and ∂M_τ, identified along $\widetilde{I}$.

[11] We define the *standard* embedding of $S^p \cup_{B^r} S^p$ into S^q in the same way as we defined it for $S^p \vee S^p$.

Bibliography

[1] Karim Adiprasito and Bruno Benedetti. Linear embeddings of contractible and collapsible complexes. Preprint, 2018.

[2] S. I. Adyan. Algorithmic unsolvability of problems of recognition of certain properties of groups. *Dokl. Akad. Nauk SSSR (N.S.)*, 103:533–535, 1955.

[3] Patrizio Angelini, Giuseppe Di Battista, Fabrizio Frati, Vít Jelínek, Jan Kratochvíl, Maurizio Patrignani, and Ignaz Rutter. Testing planarity of partially embedded graphs. *ACM Trans. Algorithms*, 11(4):32:1–32:42, April 2015. URL: http://doi.acm.org/10.1145/2629341, doi:10.1145/2629341.

[4] D. J. Anick. The computation of rational homotopy groups is $\#\wp$-hard. Computers in geometry and topology, Proc. Conf., Chicago/Ill. 1986, Lect. Notes Pure Appl. Math. 114, 1–56, 1989.

[5] C. Berger. *An effective version of the Hurewicz theorem*. Theses, Université Joseph-Fourier - Grenoble I, 1991. URL: https://tel.archives-ouvertes.fr/tel-00339314.

[6] Clemens Berger. Un groupoïde simplicial comme modèle de l'espace des chemins. *Bulletin de la Societé mathematique de France*, 123(1):1–32, 1995.

[7] R. H. Bing. An alternative proof that 3-manifolds can be triangulated. *Ann. Math. (2)*, 69:37–65, 1959.

[8] E. H. Brown (jr.). Finite computability of Postnikov complexes. *Ann. Math. (2)*, 65:1–20, 1957.

[9] J.L. Bryant. Approximating embeddings of polyhedra in codimension three. *Trans. Am. Math. Soc.*, 170:85–95, 1972.

[10] M. Cadek, M. Krcál, J. Matoušek, Lukás Vokřínek, and Uli Wagner. Extending continuous maps: polynomiality and undecidability. In *STOC*, pages 595–604, 2013.

[11] M. Čadek, M. Krčál, J. Matoušek, F. Sergeraert, L. Vokřínek, and U. Wagner. Computing all maps into a sphere. *J. ACM*, 61(3):17:1–17:44, June 2014.

[12] M. Čadek, M. Krčál, J. Matoušek, L. Vokřínek, and U. Wagner. Extendability of continuous maps is undecidable. *Discr. Comput. Geom.*, 51(1):24–66, 2013.

[13] M. Čadek, M. Krčál, J. Matoušek, L. Vokřínek, and U. Wagner. Polynomial-time computation of homotopy groups and Postnikov systems in fixed dimension. *Siam Journal on Computing*, 43(5):1728–1780, 2014.

[14] M. Čadek, M. Krčál, and L. Vokřínek. Algorithmic solvability of the lifting-extension problem. *Discrete Comput. Geom.*, 57(4):915–965, 2017.

[15] Gregory R. Chambers, Fedor Manin, and Shmuel Weinberger. Quantitative nullhomotopy and rational homotopy type. *Geometric and Functional Analysis*, 28(3):563–588, Jun 2018.

[16] Marshall M. Cohen. A general theory of relative regular neighborhoods. *Transactions of the American Mathematical Society*, 136:189–229, 1969. URL: `http://www.jstor.org/stable/1994710`.

[17] Marek Cygan, Fedor V. Fomin, Lukasz Kowalik, Daniel Lokshtanov, Daniel Marx, Marcin Pilipczuk, Michal Pilipczuk, and Saket Saurabh. *Parameterized Algorithms*. Springer Publishing Company, Incorporated, 1st edition, 2015.

[18] Arnaud de Mesmay, Yo'av Rieck, Eric Sedgwick, and Martin Tancer. Embeddability in $\mathbb{R}^3$ is np-hard. *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, 2018.

[19] H. Edelsbrunner and J. L. Harer. *Computational topology*. American Mathematical Society, Providence, RI, 2010.

[20] Herbert Edelsbrunner and Daniel R. Grayson. Edgewise subdivision of a simplex. In *Proceedings of the Fifteenth Annual Symposium on Computational Geometry*, SCG '99, pages 24–30, New York, NY, USA, 1999. ACM.

[21] I. Fáry. On straight line representation of planar graphs. *Acta Univ. Szeged., Acta Sci. Math.*, 11:229–233, 1948.

[22] Steve Ferry and Shmuel Weinberger. Quantitative algebraic topology and lipschitz homotopy. *Proceedings of the National Academy of Sciences*, 110(48):19246–19250, 2013.

[23] M. Filakovský. Effective chain complexes for twisted products. *Archivum Mathematicum*, 048:313–322, 2012.

[24] M. Filakovský and L. Vokřínek. Are two given maps homotopic? An algorithmic viewpoint. *Foundations of Computational Mathematics*, May 2019.

[25] Marek Filakovský, Peter Franek, Uli Wagner, and Stephan Zhechev. Computing simplicial representatives of homotopy group elements. *Journal of Applied and Computational Topology*, 2(3):177–231, Dec 2018.

[26] P. Franek and M. Krčál. Robust satisfiability of systems of equations. *J. ACM*, 62(4):26:1–26:19, 2015.

[27] Michael Freedman and Vyacheslav Krushkal. Geometric complexity of embeddings in $\mathbb{R}^d$. *Geometric and Functional Analysis*, 24(5):1406–1430, 2014.

[28] R. Fritsch and R. A. Piccinini. *Cellular structures in topology*, volume 19 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 1990. URL: `http://dx.doi.org/10.1017/CBO9780511983948`, `doi:10.1017/CBO9780511983948`.

[29] R. Geoghegan. *Topological Methods in Group Theory*. Graduate Texts in Mathematics. Springer New York, 2007. URL: `https://books.google.at/books?id=BwX6gblqV8MC`.

[30] P. G. Goerss and J. F. Jardine. *Simplicial homotopy theory*. Birkhäuser, Basel, 1999.

[31] M. Gromov. Quantitative homotopy theory. *Prospects in Mathematics: Invited Talks on the Occasion of the 250th Anniversary of Princeton University (H. Rossi, ed.)*, pages 45–49, 1999.

[32] A. Haefliger. Plongements différentiables dans le domaine stable. *Comment. Math. Helv.*, 37:155–176, 1962/1963.

[33] A. Hatcher. *Algebraic Topology*. Cambridge University Press, Cambridge, 2001.

[34] J. Heras, V. Pascual, J. Rubio, and F. Sergeraert. fKenzo: a user interface for computations in algebraic topology. *J. Symb. Comput.*, 46(6):685–698, 2011.

[35] J. Hopcroft and R. Tarjan. Efficient planarity testing. *J. Assoc. Comput. Mach.*, 21:549–568, 1974.

[36] John F. P. Hudson and Eric Christopher Zeeman. On combinatorial isotopy. *Publications Mathématiques de l'IHÉS*, 19:69–94, 1964.

[37] J. F. Jardine. Simplicial approximation. *Theory Appl. Categ.*, 12:No. 2, 34–72, 2004.

[38] Uli Wagner Jiří Matoušek, Martin Tancer. Hardness of embedding simplicial complexes in $\mathbb{R}^d$. *Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms*, 2009.

[39] D. Kan. The Hurewicz Theorem. In *International Symposium of Algebraic Topology, Autonomous University of Mexico and UNESCO*, pages 225–231, 1958.

[40] D. M. Kan. A combinatorial definition of homotopy groups. *Annals of Mathematics*, 67(2):282–312, 1958.

[41] R. Kannan and A. Bachem. Polynomial algorithms for computing the Smith and Hermite normal forms of an integer matrix. *SIAM J. Computing*, 8:499–507, 1981.

[42] S. O. Kochman. *Stable homotopy groups of spheres. A computer-assisted approach.* Lecture Notes in Mathematics 1423. Springer-Verlag, Berlin etc., 1990.

[43] M. Krčál, J. Matoušek, and F. Sergeraert. Polynomial-time homology for simplicial Eilenberg–MacLane spaces. *Foundat. of Comput. Mathematics*, 13:935–963, 2013.

[44] Casimir Kuratowski. Sur le problème des courbes gauches en topologie. *Fundamenta Mathematicae*, 15:271–283, 1930.

[45] Kyung Whan Kwun and Frank Raymond. Mapping cylinder neighborhoods. *Michigan Math. J. 10, no. 4, 353–357*, 1963.

[46] Isaac Mabillard and Uli Wagner. Eliminating higher-multiplicity intersections, II. The deleted product criterion in the r-metastable range. In *Proc. 32nd International Symposium on Computational Geometry (SoCG 2016)*, pages 51:1–51:12, 2016.

[47] Fedor Manin and Shmuel Weinberger. Algorithmic aspects of immersibility and embeddability. Preprint, `https://arxiv.org/abs/1812.09413`.

[48] Yu. Matiyasevich. The Diophantineness of enumerable sets. *Dokl. Akad. Nauk SSSR*, 191:279–282, 1970.

[49] Yu. V. Matiyasevich. *Hilbert's tenth problem*. Foundations of Computing Series. MIT Press, Cambridge, MA, 1993. Translated from the 1993 Russian original by the author, With a foreword by Martin Davis.

[50] J. Matoušek, M. Tancer, and U. Wagner. Hardness of embedding simplicial complexes in $\mathbb{R}^d$. *J. Eur. Math. Soc.*, 13(2):259–295, 2011.

[51] J. Matoušek. Computing higher homotopy groups is $W[1]$-hard. *Fundamenta Informaticae*, 2014.

[52] S. Matveev. *Algorithmic Topology and Classification of 3-Manifolds*. Springer, 2007.

[53] J. P. May. *Simplicial Objects in Algebraic Topology*. Chicago Lectures in Mathematics. University of Chicago Press, Chicago, IL, 1992. Reprint of the 1967 original.

[54] J. R. Munkres. *Elements of Algebraic Topology*. Addison-Wesley, Reading, MA, 1984.

[55] Alexander Nabutovsky and Shmuel Weinberger. Algorithmic unsolvability of the triviality problem for multidimensional knots. *Commentarii Mathematici Helvetici*, 71(1):426–434, Dec 1996.

[56] C. Papakyriakopoulos. A new proof for the invariance of the homology groups of a complex (in Greek). *Bull. Soc. Math. Grèce*, 22:1–154 (1946), 1943.

[57] V. V. Prasolov. *Elements of Homology Theory*. Graduate Studies in Mathematics. American Mathematical Society, 2007.

[58] D.G. Quillen. *Homotopical Algebra*. Lecture Notes in Mathematics. Springer Berlin Heidelberg, 1967.

[59] M. O. Rabin. Recursive unsolvability of group theoretic problems. *Ann. of Math. (2)*, 67:172–194, 1958.

[60] D. C. Ravenel. *Complex Cobordism and Stable Homotopy Groups of Spheres (2nd ed.)*. Amer. Math. Soc., 2004.

[61] P. Real. An algorithm computing homotopy groups. *Mathematics and Computers in Simulation*, 42:461—465, 1996.

[62] A. Romero, J. Rubio, and F. Sergeraert. Computing spectral sequences. *J. Symb. Comput.*, 41(10):1059–1079, 2006.

[63] A. Romero and F. Sergeraert. Effective homotopy of fibrations. *Applicable Algebra in Engineering, Communication and Computing*, 23(1-2):85–100, 2012.

[64] A. Romero and F. Sergeraert. A Bousfield–Kan algorithm for computing the effective homotopy of a space. *Foundations of Computational Mathematics*, pages 1–32, 2016.

[65] J. Rubio and F. Sergeraert. Constructive algebraic topology. *Bull. Sci. Math.*, 126(5):389–412, 2002.

[66] J. Rubio and F. Sergeraert. Algebraic models for homotopy types. *Homology, Homotopy and Applications*, 17:139–160, 2005.

[67] J. Rubio and F. Sergeraert. Constructive homological algebra and applications. Preprint, `arXiv:1208.3816`, 2012.

[68] T. Benny Rushing. Topological embeddings. *Academic Press*, 1973.

[69] R. Schön. *Effective Algebraic Topology*. Memoirs of the American Mathematical Society. American Mathematical Society, 1991.

[70] J. Segal and S. Spież. Quasi embeddings and embeddings of polyhedra in $\mathbb{R}^m$. *Topology Appl.*, 45(3):275–282, 1992.

[71] F. Sergeraert. The computability problem in algebraic topology. *Adv. Math.*, 104(1):1–29, 1994.

[72] J. R. Smith. m-Structures determine integral homotopy type. Preprint, `arXiv:math/9809151v1`, 1998.

[73] W. T. Tutte. How to draw a graph. *Proceedings of the London Mathematical Society*, s3-13(1):743–767, 1963.

[74] L. Vokřínek. Decidability of the extension problem for maps into odd-dimensional spheres. *Discrete & Computational Geometry*, 57(1):1–11, 2017.

[75] I. A. Volodin, V. E. Kuznetsov, and A. T. Fomenko. The problem of discriminating algorithmically the standard three-dimensional sphere. *Usp. Mat. Nauk*, 29(5):71–168, 1974. In Russian. English translation: Russ. Math. Surv. 29,5:71–172 (1974).

[76] K. Wagner. Über eine eigenschaft der ebenen komplexe. *Mathematische Annalen*, 114(1):570–590, Dec 1937.

[77] W.B.R.Lickorish. The piecewise linear unknotting of cones. *Topology Volume 4, Issue 1, Pages 67-91*, 1965.

[78] C. Weber. Plongements de polyèdres dans le domaine metastable. *Comment. Math. Helv.*, 42:1–27, 1967.

[79] E. Zeeman. Seminar on combinatorial topology. *Institut des hautes études scientifiques (Paris, France)*, 1963.

[80] E. C. Zeeman. Unknotting spheres. *Annals of Mathematics, Second Series, Vol. 72, No. 2, pp. 350-361*, 1960.

[81] A. J. Zomorodian. *Topology for Computing*, volume 16 of *Cambridge Monographs on Applied and Computational Mathematics*. Cambridge University Press, Cambridge, 2005.