Extendability of Continuous Maps Is Undecidable

Martin Čadek · Marek Krčál · Jiří Matoušek ·
Lukáš Vokřínek · Uli Wagner

Received: 26 February 2013 / Revised: 18 September 2013 / Accepted: 24 September 2013 /
Published online: 8 November 2013
© Springer Science+Business Media New York 2013

Abstract We consider two basic problems of algebraic topology: the *extension problem* and the *computation of higher homotopy groups*, from the point of view of computability and computational complexity.

The *extension problem* is the following: Given topological spaces X and Y , a subspace $A \subseteq X$, and a (continuous) map $f: A \rightarrow Y$, decide whether f can be extended to a continuous map $\tilde{f}: X \rightarrow Y$. All spaces are given as finite simplicial complexes, and the map f is simplicial.

Recent positive algorithmic results, proved in a series of companion papers, show that for $(k-1)$ -connected Y , $k \geq 2$, the extension problem is algorithmically solvable if the dimension of X is at most $2k-1$, and even in polynomial time when k is fixed.

Here we show that the condition $\dim X \leq 2k-1$ cannot be relaxed: for $\dim X = 2k$, the extension problem with $(k-1)$ -connected Y becomes undecidable. Moreover, either the target space Y or the pair (X, A) can be fixed in such a way that the problem remains undecidable.

M. Čadek · L. Vokřínek

Department of Mathematics and Statistics, Masaryk University, Kotlářská 2, 611 37 Brno,
Czech Republic

M. Krčál (✉) · J. Matoušek

Department of Applied Mathematics, Charles University, Malostranské nám. 25, 118 00 Praha 1,
Czech Republic
e-mail: krcal@kam.mff.cuni.cz

J. Matoušek

Institute of Theoretical Computer Science, ETH Zurich, 8092 Zurich, Switzerland

U. Wagner

Institut de Mathématiques de Géométrie et Applications, École Polytechnique Fédérale de Lausanne,
EPFL SB MATHGEOM, MA C1 553, Station 8, 1015 Lausanne, Switzerland

Our second result, a strengthening of a result of Anick, says that the computation of $\pi_k(Y)$ of a 1-connected simplicial complex Y is $\#P$ -hard when k is considered as a part of the input.

1 Introduction

One of the central themes in algebraic topology is to understand the structure of all *continuous* maps $X \rightarrow Y$, for given topological spaces X and Y (all maps between topological spaces in this paper are assumed to be continuous). For topological purposes, two maps $f, g: X \rightarrow Y$ are usually considered equivalent if they are *homotopic*, i.e., if one can be continuously deformed into the other;¹ thus, the object of interest is $[X, Y]$, the set of all homotopy classes of maps $X \rightarrow Y$.

Many of the celebrated results throughout the history of topology can be cast as information about $[X, Y]$ for particular spaces X and Y . In particular, one of the important challenges propelling the research in algebraic topology has been the computation of the *homotopy groups of spheres*² $\pi_k(S^n)$, for which only partial results have been obtained in spite of an enormous effort (see, e.g., [13, 22]).

A closely related question is the *extension problem*: given $A \subset X$ and a map $f: A \rightarrow Y$, can it be extended to a map $X \rightarrow Y$? For example, the famous *Brouwer fixed point theorem* can be restated as non-extendability of the identity map $S^n \rightarrow S^n$ to the ball D^{n+1} bounded by the sphere S^n . See [33] for a very clear and accessible introduction to the extension problem, including further examples and applications and covering the earlier developments until the late 1950s.

Computational Homotopy Theory In this paper, we consider the (theoretical) *computational complexity* of homotopy-theoretic questions such as the extension problem, the homotopy classification of maps, and the computation of homotopy groups. More precisely, we prove hardness and undecidability results that complement recent positive algorithmic results obtained in a series of companion papers [5, 6, 12]. To put our results into context, we first give more background.

By classical uncomputability results in topology (see, e.g., the survey [31]), most of these problems are *algorithmically unsolvable* if we place no restriction on the space Y . Indeed, by a result of Adjan and of Rabin, it is undecidable whether the fundamental group $\pi_1(Y)$ of a given finite simplicial complex Y is trivial, even if Y is assumed to be 2-dimensional. The triviality of $\pi_1(Y)$ is equivalent to $[S^1, Y]$ having

¹More precisely, f and g are homotopic, in symbols $f \sim g$, if there is a map $F: X \times [0, 1] \rightarrow Y$ such that $F(\cdot, 0) = f$ and $F(\cdot, 1) = g$. With this notation, $[X, Y] = \{[f] : f: X \rightarrow Y\}$, where $[f] = \{g : g \sim f\}$ is the *homotopy class* of f .

²We recall that the k th homotopy group $\pi_k(Y)$ of a space Y is defined as the set of all homotopy classes of *pointed* maps $f: S^k \rightarrow Y$, i.e., maps f that send a distinguished *basepoint* $s_0 \in S^k$ to a distinguished basepoint $y_0 \in Y$ (and the homotopies F also satisfy $F(s_0, t) = y_0$ for all $t \in [0, 1]$). Strictly speaking, one should write $\pi_k(Y, y_0)$ but for a path-connected Y , the choice of y_0 does not matter. Moreover, if Y is *simply connected*, i.e., if $\pi_1(Y)$ is trivial, then the pointedness of the maps does not matter either and one can identify $\pi_k(Y)$ with $[S^k, Y]$. For $k \geq 1$, each $\pi_k(Y)$ is a group, which for $k \geq 2$ is Abelian; the definition of the group operation will be reviewed in Sect. 4.

only one element, represented by the constant map, and so $[S^1, Y]$ is uncomputable in general. Moreover, by the Boone–Novikov theorem, it is undecidable whether a given pointed map $f: S^1 \rightarrow Y$ is homotopic to a constant map, and this homotopic triviality is equivalent to the extendability of f to the 2-dimensional ball D^2 . Therefore, the extension problem is undecidable as well.³

In these results, the difficulty stems from the intractability of the fundamental group of Y . Thus, a reasonable restriction is to assume that $\pi_1(Y)$ is trivial (which in general cannot be tested, but in many cases of interest it is known), or more generally, that Y is k -connected, meaning that $\pi_i(Y)$ is trivial for all $i \leq k$ (equivalently, every map $S^i \rightarrow Y$, $i \leq k$, can be extended to D^{i+1}). A basic and important example of a $(k-1)$ -connected space is the sphere S^k .

For a long time, the only positive result concerning the computation of $[X, Y]$ was that of Brown [4], who showed that $[X, Y]$ is computable under the assumption that Y is 1-connected and that all the higher homotopy groups $\pi_k(Y)$, $2 \leq k \leq \dim X$, are finite (the second assumption is rather strong and not satisfied if Y is a sphere, for example). Brown also gave an algorithm that, given $k \geq 2$ and a finite 1-connected simplicial complex Y , computes $\pi_k(Y)$.

In the 1990s, three independent collections of works appeared with the goal of making various more advanced methods of algebraic topology *effective* (algorithmic): by Schön [28], by Smith [30], and by Sergeraert, Rubio, Dousson, and Romero (e.g., [24–26, 29]; also see [27] for an exposition). New algorithms for computing higher homotopy groups follow from these methods; see Real [23] for an algorithm based on Sergeraert et al.

An algorithm that computes $\pi_k(Y)$ for a given 1-connected simplicial complex Y in polynomial time for every fixed $k \geq 2$ was recently presented in [6], also relying on [12] and on the methods of effective homology developed earlier by Sergeraert et al.

The problem of computing $[X, Y]$ was addressed in [5], where it was shown that its structure is computable assuming that Y is $(k-1)$ -connected and $\dim(X) \leq 2k-2$, for some integer $k \geq 2$. These assumptions are sometimes summarized by saying that X and Y are *in the stable range*.

As observed in [6], the methods of [5] can also be used to obtain an algorithmic solution of the extension problem. Here $\dim X$ can even be 1 beyond the stable range;⁴ thus, given finite simplicial complexes $A \subseteq X$ and Y and a simplicial map $f: A \rightarrow Y$, where Y is $(k-1)$ -connected and $\dim X \leq 2k-1$, $k \geq 2$, it can be decided algorithmically whether f can be extended to a continuous map $X \rightarrow Y$. The algorithm again runs in polynomial time for k fixed, and the same holds for the algorithm mentioned above for computing $[X, Y]$ in the stable range.

New Undecidability Results For the algorithms for homotopy classification and extendability, we have two types of assumptions: The first is that the dimension of X

³For undecidability results concerning numerous more loosely related topological problems we refer to [20, 21, 31] and references therein.

⁴In the border case $\dim X = 2k-1$, the algorithm just decides the existence of an extension, while for $\dim X \leq 2k-2$ it also yields a classification of *all* possible extensions up to homotopy.

is suitably bounded in terms of the connectivity of Y (in the stable range or at most one more). This is essential for the algorithms to work at all.⁵ The second assumption is that the relevant dimensional parameter k is fixed, which guarantees that the algorithm runs in polynomial time.

Our main result is that, for the extension problem, the first assumption is necessary and sharp.

Theorem 1.1 *Let $k \geq 2$ be fixed.*

- (a) (Fixed Target) *There is a fixed $(k - 1)$ -connected finite simplicial complex $Y = Y_k$ such that the following problem is algorithmically unsolvable: Given finite simplicial complexes $A \subseteq X$ with $\dim X = 2k$ and a simplicial map $f: A \rightarrow Y$, decide whether there exists a continuous map $X \rightarrow Y$ extending f . For k even, we can take Y_k to be the sphere S^k .*
- (b) (Fixed Source) *There exist fixed finite simplicial complexes $A = A_k$ and $X = X_k$ with $A \subseteq X$ and $\dim X = 2k$ such that the following problem is algorithmically unsolvable: Given a $(k - 1)$ -connected finite simplicial complex Y and a simplicial map $f: A \rightarrow Y$, decide whether there exists a continuous map $X \rightarrow Y$ extending f .*

The theorem is stated in terms of simplicial complexes since these are a standard input model for topological spaces in computational topology that we assume may be most familiar to most readers. For the purposes of our reductions, we actually work with *simplicial sets* (see Sect. 3.2), which offer a more flexible, but still purely combinatorial, way of representing topological spaces. The simplicial sets are then converted into simplicial complexes by a suitable subdivision.

When constructing A , X , and Y as simplicial sets, we can furthermore ensure that Y has a certain additional property, namely that it is $(k - 1)$ -reduced, which provides an immediate certificate that Y is $(k - 1)$ -connected; this is proved in Appendix B. Thus, in particular, the difficulty of the extension problem does not lie in verifying the $(k - 1)$ -connectedness of Y .

While most of the previous undecidability results in topology rely on the word problem in groups and its relatives, our proof of Theorem 1.1 relies on the undecidability of *Hilbert's tenth problem*, which is the solvability of a system of polynomial Diophantine equations, i.e., the existence of an integral solution of a system of the form

$$p_i(x_1, \dots, x_r) = 0, \quad i = 1, 2, \dots, s, \quad (1)$$

where $p_1, \dots, p_s$ are r -variate polynomials with integer coefficients. This problem is undecidable by a celebrated result of Matiyasevich [14], building on earlier work by Davis, Putnam, and Robinson; also see [15, 17] for additional background and further references.

⁵We remark that the stable range assumption guarantees that $[X, Y]$ has a canonical Abelian group structure, which we exploit heavily (for instance, it means that $[X, Y]$ has a finite description even when it is an infinite set). In the special case $\pi_k(Y) \cong [S^k, Y]$, by contrast, the group structure has a different origin and is available for all dimensions k .

On the Hardness of Computing $[X, Y]$ When $\dim X = 2k$ and Y is $(k - 1)$ -connected, we can no longer equip $[X, Y]$ with the group structure of the stable range. Thus, it is not clear in what sense the potentially infinite set $[X, Y]$ could be computed in general. A natural computational problem in this setting is to decide whether $|[X, Y]| > 1$; in other words, whether there is a homotopically nontrivial map $X \rightarrow Y$ for given simplicial complexes X and Y as above.

We can prove that this problem is NP-hard for every *even* $k \geq 2$; in order to keep this paper reasonably concise, the proof is to be presented in the Ph.D. thesis of the second author. The reduction is very similar to that of Theorem 1.1. We can show that the problem is at least as hard as deciding the existence of a nonzero integral solution of the quadratic system (Q-SYM) defined in Sect. 2 below with all the constant terms b_q equal to zero. This problem may well be undecidable, but as far as we know, the best known lower bound is that of NP-hardness.

#P-Hardness Our second result concerns the problem of computing the higher homotopy groups $\pi_n(Y) \cong [S^n, Y]$ for a simply connected space Y , if n is not considered fixed but part of the input (n is given in unary encoding).⁶ Anick [2] proved that this problem is #P-hard,⁷ where Y can even be assumed to be a 4-dimensional space.⁸

However, Anick's hardness result has the following caveat: It assumes that the input space Y is given in a very concise form, as a *cell complex* with the degrees of the *attaching maps* encoded in binary (see Sect. 5.3 for a review of the construction). A straightforward way of converting this cell complex to a simplicial complex yields a 4-dimensional simplicial complex with an exponential number of simplices, which renders the hardness result meaningless for simplicial complexes. In Sect. 6.3, we provide a different way of converting Anick's concise encoding of the input space Y into a homotopy equivalent⁹ simplicial complex that can be constructed in polynomial time, and in particular, has only polynomially many simplices. This yields the following result.

Theorem 1.2 *It is #P-hard to compute the rank of $\pi_n(Y)$ (i.e., the number of summands of $\pi_n(Y)$ isomorphic to $\mathbb{Z}$) for a given number $n \in \mathbb{N}$ (encoded in unary) and a given simply connected 4-dimensional simplicial complex Y .*

⁶Note that with a unary encoding of n , the size of the input is significantly larger than with a binary (or decimal encoding), and hence the hardness result is correspondingly stronger.

⁷Somewhat informally, the class of #P-hard problems consists of computational problems that should return a natural number (as opposed to YES/NO problems) and are at least as hard as counting the number of all Hamiltonian cycles in a given graph, or counting the number of subsets with zero sum for a given set of integers, etc. These problems are clearly at least as hard as NP-complete problems, and most likely even less tractable.

⁸Actually, the hardness already applies to the potentially easier problem of computing the *rational homotopy groups* $\pi_n(Y) \otimes \mathbb{Q}$; practically speaking, one asks only for the rank of $\pi_n(Y)$, i.e., the number of direct summands isomorphic to $\mathbb{Z}$.

⁹Spaces X and Y are homotopy equivalent if there are maps $f: X \rightarrow Y$ and $g: Y \rightarrow X$ such that the compositions fg and gf are homotopic to identities. From the point of view of homotopy theory, such X and Y are indistinguishable and, in particular, $\pi_k(X) = \pi_k(Y)$ for all $k \geq 0$.

Outline of the Proof of Theorem 1.1 We aim to present our results in a way that makes the statements of the results and the main steps and ideas accessible while assuming only a moderate knowledge of topology on the side of the reader.

We thus review a number of basic topological concepts and provide proofs for various assertions and facts that may be rather elementary for topologists. On the other hand, some of the proofs assume a slightly stronger topological background, since reviewing every single notion and fact would make the paper too lengthy.

For proving Theorem 1.1, we present an algorithm that converts a given system of Diophantine equations into an instance of the extension problem; i.e., it constructs simplicial complexes A , X , and Y and a map $f: A \rightarrow Y$ such that there is an extension of f to all of X iff the given system of equations is solvable. Moreover, as stated in the theorem, there are actually two versions of the reduction: The first uses a fixed target space $Y = Y_k$ and encodes the equations into A , X , and f . The second uses a fixed pair (X_k, A_k) of source complexes and encodes the equations into f and Y .

We will actually work only with *quadratic* Diophantine equations of a slightly special form (which is sufficient; see Sect. 2). The unknowns are represented by the degrees of restrictions of the desired extension $\tilde{f}$ to suitable k -dimensional spheres. The quadratic terms in the equations are obtained using the *Whitehead product*, which is a binary operation that, for a space Z , assigns to elements $\alpha \in \pi_k(Z)$ and $\beta \in \pi_\ell(Z)$ an element $[\alpha, \beta] \in \pi_{k+\ell-1}(Z)$; see Sect. 4.2.

Here is a rough outline of the proof strategy. First we focus on Theorem 1.1(a) (fixed target) with k even, which is the simplest among our constructions.

- The spaces X and A are simplest to describe as cell complexes. The subcomplex A is a union of r spheres S^{2k-1} , which intersect only at a single common point. This union is called a wedge sum and denoted by $A = S^{2k-1} \vee \dots \vee S^{2k-1}$. The space X is homotopy equivalent to another wedge sum, of s spheres S^k ; i.e., $X \simeq S^k \vee \dots \vee S^k$.
- The fixed $(k-1)$ -connected target space Y is the k -sphere S^k .
- Maps $X \rightarrow S^k$ can be described completely by their restrictions to the k -spheres in the wedge sum. Each such restriction is characterized, uniquely up to homotopy, by its degree—this can be an arbitrary integer. Thus, a potential extension $\tilde{f}$ can be encoded into a vector $\mathbf{x} = (x_1, \dots, x_r)$ of integers.
- Similarly, the map $f: A \rightarrow S^k$ can be described by its restrictions to the $(2k-1)$ -spheres in the wedge sum. Crucially for our construction, the homotopy group $\pi_{2k-1}(S^k)$ has an element of infinite order, namely, the Whitehead square $[\iota, \iota]$, where ι is the identity $S^k \rightarrow S^k$ (we still assume k even). We will work with maps f whose restriction to the q th sphere is (homotopic to) an integral multiple $b_q[\iota, \iota]$, for some (unique) integer b_q . Thus, f is specified by the vector $\mathbf{b} = (b_1, \dots, b_s)$ of these integers.
- Given arbitrary integers $a_{ij}^{(q)}$, $1 \leq i < j \leq r$, $q = 1, 2, \dots, s$, we construct the pair (X, A) in such a way that, taking $\tilde{f}: X \rightarrow Y$ specified by $\mathbf{x}$ as above, the restriction of $\tilde{f}$ to the q th sphere of A is homotopic to $\sum_{i < j} a_{ij}^{(q)} x_i x_j [\iota, \iota]$ (here the addition and multiplication by integers are performed in $\pi_{2k-1}(S^k)$). Since $[\iota, \iota]$ is an element of $\pi_{2k-1}(S^k)$ of infinite order, $\tilde{f}$ is an extension of f iff $\sum_{i < j} a_{ij}^{(q)} x_i x_j = b_q$ for all $q = 1, 2, \dots, s$.

- In this way, we can simulate an arbitrary system of quadratic equations by an extension problem. Some more work is still needed to describe X and A as finite simplicial complexes and f as a simplicial map.
- For Theorem 1.1(a) with k odd, the Whitehead square $[\iota, \iota]$ as above no longer has infinite order. Instead, we use $Y = S^k \vee S^k$ and replace $[\iota, \iota]$ by the Whitehead product $[\iota_1, \iota_2]$ of the inclusions of the two spheres into Y . This leads to skew-symmetric systems of quadratic equations, and showing that these are still undecidable needs some work (see Sect. 2).

In Theorem 1.1(b) with k even, the (fixed) source space X is homotopy equivalent to S^k and $A = S^{2k-1}$. Under the homotopy equivalence $X \simeq S^k$, the inclusion $A \hookrightarrow X$ becomes the Whitehead square $[\iota, \iota]$, and for k odd it is replaced by $[\iota_1, \iota_2]$. In both cases, the system of quadratic equations is encoded into the structure of the cell complex Y and the map $f: A \rightarrow Y$.

2 Diophantine Equations and Undecidability

We will need to work with quadratic Diophantine equations of two special forms:

$$\sum_{1 \leq i < j \leq r} a_{ij}^{(q)} x_i x_j = b_q, \quad q = 1, 2, \dots, s, \quad (\text{Q-SYM})$$

where $a_{ij}^{(q)}, b_q \in \mathbb{Z}$, and $x_1, \dots, x_r$ are the unknowns (i.e., the left-hand sides are quadratic forms with no square terms), and

$$\sum_{1 \leq i < j \leq r} a_{ij}^{(q)} (x_i y_j - x_j y_i) = b_q, \quad q = 1, 2, \dots, s, \quad (\text{Q-SKEW})$$

with $a_{ij}^{(q)}, b_q \in \mathbb{Z}$, and unknowns $x_1, \dots, x_r, y_1, \dots, y_r$ (so here we deal with skew-symmetric bilinear forms).

Lemma 2.1 *The solvability of the system (Q-SYM), as well as that of (Q-SKEW), in the integers is algorithmically undecidable.*

Proof First, it is well known and easy to see that the solvability of a general quadratic system of Diophantine equations is no easier than the solvability of an arbitrary Diophantine system (1), and thus undecidable.¹⁰

First we show undecidability for (Q-SYM); this system differs from a general quadratic system only by the lack of linear terms and squares. Given a general quadratic system

$$\sum_{1 \leq i, j \leq r} a_{ij}^{(q)} x_i x_j + \sum_{1 \leq i \leq r} b_i^{(q)} x_i = c_q, \quad q = 1, \dots, s, \quad (4)$$

¹⁰The idea is to represent higher degree monomials in the general system using new variables; e.g., for the monomial $x^3 y$ we can introduce new variables t_1, t_2 , new quadratic equations $t_1 = x^2$ and $t_2 = xy$, and replace $x^3 y$ by $t_1 t_2$.

we add new variables x_0, x'_0 and $x'_1, \dots, x'_r$, and we replace the terms $x_i x_j$ with $x_i x'_j$ and x_i with $x_i x'_0$. We also add the following equations:

$$x_0 x'_0 = 1; \quad x_i x'_0 - x_0 x'_i = 0, \quad i = 1, \dots, s.$$

The resulting system is of the form (Q-SYM) (assuming an indexing of the variables such that the x_i precede the x'_i) and it forces $x_0 = x'_0 = \pm 1$, and $x_i = x'_i$. Thus, each of its solutions corresponds either to a solution of the original system (4) (when $x_0 = x'_0 = 1$), or to a solution of the system obtained from (4) by changing the sign of all the linear terms (when $x_0 = x'_0 = -1$). Since there is an obvious bijection $x_i \mapsto -x_i$ between the solutions of (4) and those of the system with negated linear terms, the solvability of the constructed system (Q-SYM) is equivalent to the solvability of (4).

Next, we show that (Q-SKEW) is no easier than (Q-SYM). Given a general system (Q-SYM), we add new variables x_0, y_0, x'_0, y'_0 and, for each $i = 1, \dots, r$, also x'_i, y_i , and y'_i . We replace each term $x_i x_j$ in the original system (Q-SYM) by the antisymmetric expression $x_i y'_j - x'_j y_i$, and we add the following equations (for $i = 1, 2, \dots, r$):

$$\begin{aligned} x_0 y'_0 - x'_0 y_0 &= 1, & x_0 y_i - x_i y_0 &= 0, & x'_0 y'_i - x'_i y'_0 &= 0, \\ (x_0 y'_i - x'_i y_0) - (x_i y'_0 - x'_0 y_i) &= 0. \end{aligned}$$

This gives a system of the form (Q-SKEW), which we call the *new system*.

It is clear that each solution of (Q-SYM) yields a solution of the new system. Conversely, supposing that the new system has a solution, we claim that it also has a solution with $x_0 = y'_0 = 1$ and $y_0 = x'_0 = 0$. Once we have a solution satisfying these additional conditions, it is easy to check that $x_1, \dots, x_r$ form a solution of the original system.

To verify the claim, for notational convenience, let us index the x and y variables in the new system by the set $I = \{0, 1, \dots, r, 0', 1', \dots, r'\}$, where $x_{i'} = x'_i$ and $y_{i'} = y'_i$. We suppose that $(x_i, y_i : i \in I)$ form a solution of the new system. Since $x_0 y'_0 - x'_0 y_0 = 1$, the 2×2 matrix $\begin{pmatrix} x_0 & x'_0 \\ y_0 & y'_0 \end{pmatrix}$ has determinant 1 and thus an integral inverse matrix, which we denote by T .

Let us define new values $(\bar{x}_i, \bar{y}_i : i \in I)$ by $\begin{pmatrix} \bar{x}_i \\ \bar{y}_i \end{pmatrix} = T \cdot \begin{pmatrix} x_i \\ y_i \end{pmatrix}$, $i \in I$. We have $\bar{x}_0 = \bar{y}'_0 = 1$ and $\bar{y}_0 = \bar{x}'_0 = 0$, and it remains to show that the $\bar{x}_i$ and $\bar{y}_i$ satisfy the new system. This happens because, for every $i, j \in I$, we have

$$\begin{aligned} \bar{x}_i \bar{y}_j - \bar{x}_j \bar{y}_i &= \det \begin{pmatrix} \bar{x}_i & \bar{x}_j \\ \bar{y}_i & \bar{y}_j \end{pmatrix} = \det \left(T \cdot \begin{pmatrix} x_i & x_j \\ y_i & y_j \end{pmatrix} \right) \\ &= \det T \cdot \det \begin{pmatrix} x_i & x_j \\ y_i & y_j \end{pmatrix} = x_i y_j - x_j y_i. \end{aligned} \quad \square$$

3 Cell Complexes and Simplicial Sets

This section and the next one primarily present known material from topology. In several cases we need to adapt results from the literature to our needs, which is sometimes best done by re-proving them. Readers may want to skim these two sections quickly and return to them later when needed.

Here we review two basic ways of building topological spaces from simple pieces: *cell complexes* and *simplicial sets*. Cell complexes, also known as *CW complexes*, are fairly standard in topology, and we will use them for a simple description of the various spaces in our proofs. Simplicial sets are perhaps less well known, and for us, they will mainly be a convenient device for converting cell complexes into simplicial complexes. Moreover, they are of crucial importance in the algorithmic results mentioned in the introduction. For a thorough discussion of simplicial complexes, simplicial sets, cell complexes, and the connections between the three, we refer to [7].

3.1 Cell Complexes

In the case of cell complexes, the building blocks are topological disks of various dimensions, called *cells*, which can be thought of as being completely “flexible” and which can be glued together in an almost arbitrary continuous fashion. Essentially the only condition is that each n -dimensional cell must be attached along its boundary to the $(n - 1)$ -skeleton of the space, i.e., to the part that has already been built, inductively, from lower dimensional cells. The formal definition is as follows.

We recall that if X and Y are topological spaces and if $f: A \rightarrow Y$ is a map defined on a subspace $A \subseteq X$, then the space $X \cup_f Y$ obtained by *attaching X to Y via f* is defined as the quotient of the disjoint union $X \sqcup Y$ under the equivalence relation generated by the identifications $a \sim f(a)$, $a \in A$.

A *closed* or *open n -cell* is a space homeomorphic to the closed n -dimensional unit disk D^n in n -dimensional Euclidean space or its interior $\mathring{D}^n$, respectively; a point is regarded as both a closed and an open 0-cell.

An *m -dimensional cell complex*¹¹ X is the last term of an inductively constructed sequence of spaces $X^{(0)} \subseteq X^{(1)} \subseteq X^{(2)} \subseteq \dots \subseteq X^{(m)} = X$, called the *skeletons* of X :

1. $X^{(0)}$ is a discrete set of points (possibly infinite) that are regarded as 0-cells.
2. Inductively, the n -skeleton $X^{(n)}$ is formed by attaching closed n -cells D_i^n (where i ranges over some arbitrary index set) to $X^{(n-1)}$ via *attaching maps* $\varphi_i: S_i^{n-1} = \partial D^n \rightarrow X^{(n-1)}$. Formally, we can consider all attaching maps together as defining a map $\varphi = \sqcup_i \varphi_i$ from the disjoint union $\bigsqcup_i S_i^{n-1}$ to $X^{(n-1)}$ and form $X^{(n)} = (\bigsqcup_i D_i^n) \cup_\varphi X^{(n-1)}$.

For every closed cell D_i^n , one has a *characteristic map*¹² $\Phi_i: D_i^n \rightarrow X^{(n)} \subseteq X$, which restricts to an embedding on the interior $\mathring{D}_i^n$. The image $\Phi_i(\mathring{D}_i^n)$ is commonly

¹¹Cell complexes can also be infinite dimensional, in which case some care must be taken in defining their topology, but we will deal with cell complexes that are made of finitely many cells, and are thus finite dimensional.

¹²The composition of the inclusion $D_i^n \hookrightarrow (\bigsqcup_i D_i^n) \sqcup X^{(n-1)}$ with the quotient map.

denoted by e_i^n , and it follows from the construction that every point of X is contained in a unique open cell (note that these are generally not open subsets of X , however).

As a basic example, the n -sphere is a cell complex with one n -cell and one 0-cell, obtained by attaching D^n to a point e^0 via the constant map that maps all of S^{n-1} to e_0 .

Subcomplexes A subcomplex $A \subseteq X$ is a subspace that is closed and a union of open cells of X . In particular, for each cell in A , the image of its attachment map is contained in A , so A is itself a cell complex (and its cell complex topology agrees with the subspace topology inherited from X).

The Homotopy Extension Property An important fact is that cell complexes have the *homotopy extension property*: Suppose that X is a cell complex and that $A \subseteq X$ is a subcomplex. If we are given a map $f_0: A \rightarrow Y$ into some space Y , an extension $\tilde{f}_0: X \rightarrow Y$ of f_0 and a homotopy $H: A \times [0, 1]$ between f_0 and some other map $f_1: A \rightarrow Y$, then H can be extended to a homotopy $\tilde{H}: X \times [0, 1] \rightarrow Y$ between $\tilde{f}_0$ and some extension $\tilde{f}_1: X \rightarrow Y$ of f_1 . The following is an immediate consequence.

Corollary 3.1 *For a cell complex X , subcomplex $A \subseteq X$, and a space Y , the extendability of a map $f: A \rightarrow Y$ to X depends only on the homotopy class of f in $[A, Y]$. Moreover, the map $f: A \rightarrow Y$ has an extension $\tilde{f}: X \rightarrow Y$ iff there exists a map $g: X \rightarrow Y$ such that the diagram*

$$\begin{array}{ccc} A & \xrightarrow{f} & Y \\ i \downarrow & \nearrow g & \\ X & & \end{array}$$

commutes up to homotopy, i.e., $gi \sim f$.

Cellular Maps and Cellular Approximation A map $f: X \rightarrow Y$ between cell complexes is called *cellular* if it maps skeletons to skeletons, i.e., $f(X^{(n)}) \subseteq Y^{(n)}$ for every n .

The *cellular approximation theorem* (see [9, Theorem 4.8]) states that every continuous map $f: X \rightarrow Y$ between cell complexes is homotopic to a cellular one; moreover, if the given map f is already cellular on some subcomplex $A \subseteq X$, then the homotopy can be taken to be *stationary* on A (i.e., the image of every point in A remains fixed throughout).

3.2 Simplicial Sets

For certain constructions it is advantageous to use a special type of cell complex with an additional structure that allows for a purely combinatorial description; the latter also facilitates representing and manipulating the objects in question, simplicial

sets, on a computer. We refer to [8] for a very friendly and thorough introduction to simplicial sets.

Intuitively, a simplicial set can be thought of as a kind of hybrid or compromise between a simplicial complex (more special) on the one hand and a cell complex (more general) on the other hand. As in the case of simplicial complexes, the building blocks (cells) of which a simplicial set is constructed are simplices (vertices, edges, triangles, tetrahedra, ...), and the boundary of each n -simplex Δ^n is attached to the lower dimensional skeleton by identifications that are linear on each proper face (subsimplex) of Δ^n ; thus, these identifications can be described combinatorially by maps between the vertex sets of the simplices.¹³ However, the attachments are more general than the one permitted for simplicial complexes; for example, one may have several 1-dimensional simplices connecting the same pair of vertices, a 1-simplex forming a loop, two edges of a 2-simplex identified to create a cone, or the boundary of a 2-simplex all contracted to a single vertex, forming an S^2 .

Moreover, one keeps track of certain additional information that might seem superfluous but turns out to be very useful for various constructions. For example, even if the identifications force some n -simplex to be collapsed to something lower dimensional (so that it could be discarded for the purposes of describing the space as a cell complex), it will still be formally kept on record as a *degenerate* n -simplex; for instance, the edges of the triangle with a boundary contracted to a point (the last example above) do not disappear—formally, each of them keeps a phantom-like existence of a degenerate 1-simplex.

Formally, a simplicial set X is given by a sequence $(X_0, X_1, X_2, \dots)$ of mutually disjoint sets, where the elements of X_n are called the n -simplices of X (we note that, unlike the case of simplicial complexes, a simplex in a simplicial set need not be determined by the set of its vertices; indeed, there can be many simplices with the same vertex set). The 0-simplices are also called *vertices*.

For every $n \geq 1$, there are $n + 1$ mappings $\partial_0, \dots, \partial_n: X_n \rightarrow X_{n-1}$ called *face operators*; the intuitive meaning is that for a simplex $\sigma \in X_n$, $\partial_i \sigma$ is the face of σ opposite to the i th vertex. Moreover, there are $n + 1$ mappings $s_0, \dots, s_n: X_n \rightarrow X_{n+1}$ called the *degeneracy operators*; the approximate meaning of $s_i \sigma$ is the degenerate simplex which is geometrically identical to σ , but with the i th vertex duplicated. A simplex is called *degenerate* if it lies in the image of some s_i ; otherwise, it is *non-degenerate*. We write X^{ndg} for the set of all nondegenerate simplices of X . A simplicial set is called *finite* if it has only finitely many nondegenerate simplices (if X is nonempty, there are always infinitely many degenerate simplices, at least one for every positive dimension).

¹³More precisely, the vertex set of each simplex is equipped with an ordering, and the identifications are required to be weakly order-preserving maps (not necessarily injective) between the vertex sets.

There are natural axioms that the ∂_i and the s_i have to satisfy, but we will not list them here, since we won't really use them. Moreover, the usual definition of simplicial sets uses the language of category theory and is very elegant and concise; see, e.g., [7, Sect. 4.2].

If A and X are simplicial sets such that $A_n \subseteq X_n$ for every n and the face and degeneracy operators of A are the restrictions of the corresponding operators of X , then we call A a *simplicial subset* of X .

Examples Here we sketch some basic examples of simplicial sets; again, we won't provide all details, referring to [8]. Let Δ^p denote the standard p -dimensional simplex regarded as a simplicial set. For $p = 0$, $(\Delta^0)_n$ consists of a single simplex, denoted by 0^n , for every $n = 0, 1, \dots$; 0^0 is the only nondegenerate simplex. The face and degeneracy operators are defined in the only possible way.

For $p = 1$, Δ^1 has two 0-simplices (vertices), say 0 and 1, and in general there are $n + 2$ simplices in $(\Delta^1)_n$; we can think of the i th one as containing i copies of the vertex 0 and $n + 1 - i$ copies of the vertex 1, $i = 0, 1, \dots, n + 1$. For p arbitrary, the n -simplices of Δ^p can be thought of as all nondecreasing $(n + 1)$ -term sequences with entries in $\{0, 1, \dots, p\}$; the ones with all terms distinct are nondegenerate.

In a similar fashion, every simplicial complex K can be converted into a simplicial set X in a canonical way; first, however, we need to fix a linear ordering of the vertices. The nondegenerate n -simplices of X are in one-to-one correspondence with the n -simplices of K , but many degenerate simplices show up as well.

Geometric Realization Like a simplicial complex, every simplicial set X defines a topological space $|X|$, the *geometric realization* of X , which is unique up to homeomorphism. More specifically, $|X|$ is a cell complex with one n -cell for every *nondegenerate* n -simplex of X , and these cells are glued together according to the identifications implied by the face and degeneracy operators (we omit the precise definition of the attachments, since we will not really use it and refer to the literature, e.g., to [8] or [7, Sect. 4.3]).

Simplicial Maps Simplicial sets serve as a combinatorial way of describing a topological space; in a similar way, simplicial maps provide a combinatorial description of continuous maps.

A *simplicial map* $f: X \rightarrow Y$ of simplicial sets X, Y consists of maps $f_n: X_n \rightarrow Y_n$, $n = 0, 1, \dots$, that commute with the face and degeneracy operators.

A simplicial map $f: X \rightarrow Y$ induces a continuous, in fact, a cellular map $|f|: |X| \rightarrow |Y|$ of the geometric realizations in a natural way (we again omit the precise definition). Often we will take the usual liberty of omitting $|\cdot|$ and not distinguishing between simplicial sets and maps and their geometric realizations.

Of course, not all continuous maps are induced by simplicial maps. However, simplicial maps can be used to *approximate* arbitrary continuous maps up to homotopy. The *simplicial approximation theorem* (which may be most familiar in the context of simplicial complexes) says that for an arbitrary continuous map $\varphi: |X| \rightarrow |Y|$ between the geometric realizations of simplicial sets, with X finite, there exist a *sufficiently fine subdivision* X' of X and a simplicial map $f: X' \rightarrow Y$ whose geometric realization is homotopic to φ ; see Sect. 3.4 for more details.

Encoding Finite Simplicial Sets A finite simplicial complex can be encoded in a straightforward way by listing the vertices of each simplex.

For simplicial sets, the situation is a bit more complicated, since the simplices are no longer uniquely determined by their vertices, but if X is finite, then we can encode X by the set X^{ndg} of its nondegenerate simplices (which we assume to be numbered from 1 to N , where N is the total number of nondegenerate simplices), plus a little bit of additional information.

The simple but crucial fact (see, e.g. [7, Theorem 4.2.3]) we need is that every simplex σ can be written uniquely as $\sigma = s\tau$, where τ is nondegenerate and s is a degeneracy, i.e., a composition $s = s_{i_k} \dots s_{i_1}$ of degeneracy operators where $k = \dim \sigma - \dim \tau$ (in particular, σ is nondegenerate itself if $\sigma = \tau$ and s is the identity). Thus, as mentioned above, degenerate simplices σ do not need to be encoded explicitly but can be represented by $s\tau$ when needed, where the degeneracy s can be encoded by the sequence $(i_k, \dots, i_1)$ of indices of its components.¹⁴ The extra information we need to encode X , in addition to the list of its nondegenerate simplices, is how these fit together. Specifically, for $\sigma \in X_n^{\text{ndg}}$ and $0 \leq i \leq n$, the i th face can be written uniquely as $\partial_i \sigma \in X_{n-1} = s\tau$ with τ nondegenerate, and for each σ , we record the $(n+1)$ -tuple of pairs (τ, s) .

Similarly, if $f: X \rightarrow Y$ is a simplicial map between finite simplicial sets, then given the encodings of X and Y , we can encode f by expressing, for each $\sigma \in X_n^{\text{ndg}}$, the image $f(\sigma) = s\tau$, with $\tau \in Y_m^{\text{ndg}}$ and recording the list of triples (σ, τ, s) .

For a finite simplicial set X , we define $\text{size}(X)$ as the number of nondegenerate simplices. If the dimension of X is bounded by some number d , then the number of bits in the encoding of X described above is bounded by $O(\text{size}(X) \log \text{size}(X))$, with the constant of proportionality depending only on d .

The notion of size will be a convenient tool that allows us to ensure that our reductions can be carried out in polynomial time, without analyzing the running time in complete detail, which we feel would be cumbersome and not very enlightening.

More specifically, our reductions will be composed of a sequence of various basic constructions of simplicial sets, which will be described in the next subsection.

For each of these basic constructions, it is straightforward to check¹⁵ that when we apply them to finite simplicial sets of *bounded dimension*, both the running time of the construction (the number of steps needed to compute the encoding of the output from the encoding of the input) as well as the size of the output simplicial set are polynomial in the size of the input. Thus, to ensure that the overall reduction is polynomial, it will be enough to take care that we combine only a polynomial number of such basic constructions, that the size of every intermediate simplicial set constructed during the reduction remains polynomial in the initial input, and that the dimension remains bounded.

¹⁴Moreover, this sequence is unique, by the simplicial set axioms that we have not specified, if one stipulates $i_k < \dots < i_1$.

¹⁵A notable exception are subdivisions, for which we provide more detail in Appendix A.

3.3 Basic Constructions

In this subsection, we review several basic constructions for cell complexes and simplicial sets. (One advantage of simplicial sets over simplicial complexes is that various operations on topological spaces, in particular Cartesian products and quotients, have natural counterparts for simplicial sets. This is where the degeneracy operators and degenerate simplices turn out to be necessary.) For more details, we refer to [7, 9].

Pointed and k -Reduced Simplicial Sets and Cell Complexes Several of the constructions are defined for pointed spaces. We recall that a *pointed space* (X, x_0) is a topological space X with a choice of a distinguished point $x_0 \in X$ (the *basepoint*). If X is a cell complex or a simplicial set, then we will always assume the basepoint to be a *vertex* (i.e., a 0-cell or 0-simplex, respectively). A *pointed map* $(X, x_0) \rightarrow (Y, y_0)$ of pointed spaces (cell complexes, simplicial sets) is a continuous (cellular, simplicial) map sending x_0 to y_0 . Homotopies of pointed maps are also meant to be pointed; i.e., they must keep the image of the basepoint fixed. The reader may recall that, for example, the homotopy groups $\pi_k(Y)$ are defined as homotopy classes of pointed maps. The set of pointed homotopy classes of pointed maps $X \rightarrow Y$ will be denoted by $[X, Y]_*$.

A simplicial set X is called *k -reduced*, $k \geq 0$, if it has a single vertex and no nondegenerate simplices in dimensions 1 through k . Similarly, a cell complex X is *k -reduced* if it has a single vertex and no cells of dimensions 1 up to k . It is then necessarily *k -connected*.

If (Y, y_0) is a 0-reduced cell complex (or simplicial set), then any cellular (or simplicial) map from a pointed complex (X, x_0) into Y is automatically pointed. Moreover, if Y is 1-reduced, then every homotopy is pointed, too, and thus $[X, Y]$ is canonically isomorphic to $[X, Y]_*$.

Products If X and Y are cell complexes, then their Cartesian product $X \times Y$ has a natural cell complex structure whose n -cells are products $e^p \times e^q$, where $p + q = n$ and e^p and e^q range over the p -cells of X and the q -cells of Y , respectively.

Furthermore, if X and Y are simplicial sets, then there is a formally very simple way to define their product $X \times Y$: one sets $(X \times Y)_n := X_n \times Y_n$ for every n , and the face and degeneracy operators work componentwise; e.g., $\partial_i(\sigma, \tau) := (\partial_i \sigma, \partial_i \tau)$. As one would expect from a good definition, the product of simplicial sets corresponds to the Cartesian product of their geometric realizations, i.e., $|X \times Y| \cong |X| \times |Y|$.¹⁶ The apparent simplicity of the definition hides some intricacies, though, as one can guess after observing that, for example, the product of two 1-simplices is not a simplex—so the above definition must imply some canonical way of triangulating the product.

¹⁶To be more precise, the above equality holds literally, with the product topology on the right-hand side, only under suitable assumptions on X and Y , e.g., if both X and Y have only countably many simplices. In the general case, one has to interpret the product $|X| \times |Y|$ differently, in the category of so-called *k -spaces*, and the same subtlety arises for products of cell complexes; see, e.g., the discussion in the respective appendices in [7, 9]. For the spaces we will encounter, however, this issue will not arise, and the product will be the same as the usual product of topological spaces.

Remark 3.2 A pair $(s\sigma, t\tau)$ of degenerate simplices in the factors may yield a nondegenerate simplex in the product, if the degeneracies s and t are composed of different degeneracy operators s_i . However, $\dim(X \times Y) = \dim X + \dim Y$, so the product contains no nondegenerate simplices of dimension larger than $\dim X + \dim Y$, and hence $\text{size}(X \times Y)$ is at most $\text{size}(X) \times \text{size}(Y)$ times some factor that depends only on the dimension¹⁷ $\dim(X \times Y)$.

Moreover, if the dimensions are bounded, the product can be constructed in polynomial time.

Quotients and Attachments If X , Y , and A are cell complexes with $A \subseteq X$ and if $f: A \rightarrow Y$ is a cellular map, then the space $X \cup_f Y$ obtained by attaching X to Y along f is also a cell complex in a natural way (see, e.g., [7, Sect. 2.3]). In particular, X/A is a cell complex, with cells corresponding to the cells of X not contained in A , plus one additional 0-cell (corresponding to the image of A under the quotient map).

Similarly, if X is a simplicial set and if $\sim$ is an equivalence relation on each X_n that is compatible with the face and degeneracy operators, then the quotient $X/\sim$ is also a simplicial set. In particular, this includes *simplicial attachments* $X \cup_f Y$ of simplicial sets along a simplicial map $f: A \rightarrow Y$ defined on a simplicial subset $A \subseteq X$, and quotients X/A by simplicial subsets. These constructions are compatible with geometric realizations, e.g., $|X \cup_f Y| \cong |X| \cup_{|f|} |Y|$.

Moreover, the size of $X \cup_f Y$ is at most the size of X plus the size of Y , and in bounded dimension, the attachment can be constructed in polynomial time.

Wedge Sum (or Wedge Product) If $X_1, \dots, X_m$ are pointed spaces, then their *wedge sum* $X_1 \vee \dots \vee X_m$ is simply the disjoint union of the X_i with the basepoints identified (this is a very special type of attachment). If the X_i are cell complexes or simplicial sets, then so is their wedge sum.

Later we will need the following bijection:

$$[X_1 \vee X_2 \vee \dots \vee X_m, Y]_* \xrightarrow{\cong} [X_1, Y]_* \times [X_2, Y]_* \times \dots \times [X_m, Y]_* \quad (5)$$

where the components of this map are given by the restrictions to the respective X_i .

Mapping Cylinder and Mapping Cone For a map $f: X \rightarrow Y$, the *mapping cylinder* of f is the space $\text{Cyl}(f)$ defined as the quotient of $(X \times [0, 1]) \sqcup Y$ under the identifications $(x, 0) \sim f(x)$ for each $x \in X$. The *mapping cone* $\text{Cone}(f)$ is defined as

¹⁷This follows from the fact about realizations mentioned above. Another way of seeing this is that if $\dim \sigma = p$, $\dim \tau = q$, and $\dim(s\sigma) = \dim(t\tau) = n > p + q$, then s and t involve $n - p$ and $n - q$ degeneracy operators s_i with $i \leq n$, respectively, so there must be a repetition since $n - p + n - q > n$. Without further reflection, this immediately implies that $\text{size}(X \times Y) \leq \text{size}(X) \cdot \text{size}(Y) \cdot (\dim X)! (\dim Y)!$.

In fact, the factor is only singly exponential in the dimensions. For instance, for a product $\Delta^p \times \Delta^q$ of two standard simplices, the vertices of $\Delta^p \times \Delta^q$ correspond to the grid points in $\{0, \dots, p\} \times \{0, \dots, q\}$, and the nondegenerate k -simplices correspond to subsets of size $k + 1$ of the grid that are weakly monotone in both coordinates (weakly monotone paths of length k). Thus, the number of nondegenerate simplices of full dimension $p + q$ equals $\binom{p+q}{p}$, and the number of all nondegenerate simplices is at most 4^{p+q} , say. Thus, $\text{size}(X \times Y) \leq \text{size}(X) \cdot \text{size}(Y) \times 4^{\dim X + \dim Y}$, say.

the quotient $\text{Cyl}(f)/(X \times \{1\})$ of $\text{Cyl}(f)$ with the subspace $X \times \{1\}$ collapsed into a point.

By the discussion concerning attachments, if X and Y are cell complexes and f is cellular, then $\text{Cyl}(f)$ and $\text{Cone}(f)$ are cell complexes as well. Moreover, if f is a simplicial map between simplicial sets, then by taking the analogous simplicial attachments and quotients, we obtain simplicial sets, denoted by $\text{Cyl}(f)$ and $\text{Cone}(f)$ as well, and called the *simplicial mapping cylinder* and *simplicial mapping cone*, respectively. The simplicial constructions are compatible with geometric realizations, for example, $|\text{Cyl}(f)| \cong \text{Cyl}(|f|)$.

We will use the mapping cylinder in our construction to replace an arbitrary map $f: X \rightarrow Y$ by an inclusion $X \hookrightarrow \text{Cyl}(f)$, which has the same homotopy properties as f . A more precise statement is given in the following lemma (see, e.g., [9, Corollary 0.21]).

Lemma 3.3 *Let $f: X \rightarrow Y$ be a continuous map between topological spaces. We consider $X \cong X \times \{1\}$ and Y as subspaces of $\text{Cyl}(f)$ and denote the corresponding inclusion maps¹⁸ by $i_X: X \hookrightarrow \text{Cyl}(f)$ and $i_Y: Y \hookrightarrow \text{Cyl}(f)$.*

- (a) *Y is a strong deformation retract¹⁹ of $\text{Cyl}(f)$.*
- (b) *X (considered as a subspace via i_X) is a strong deformation retract of $\text{Cyl}(f)$ iff f is a homotopy equivalence.*
- (c) *$i_X \sim i_Y f$ are homotopic as maps $X \rightarrow \text{Cyl}(f)$.*
- (d) *If $f: X \rightarrow Y$ is a homotopy equivalence and if $g: Y \rightarrow X$ is a homotopy inverse for f , then $i_X g \sim i_Y$ as well.*

Reduced Mapping Cone and Mapping Cylinder If X and Y and f are pointed, with basepoints x_0 and y_0 , it will be technically convenient, particularly in Sect. 6, to consider the spaces $\widehat{\text{Cyl}}(f)$ and $\widehat{\text{Cone}}(f)$, called the *reduced mapping cylinder* and the *reduced mapping cone*, respectively, that are obtained from $\text{Cyl}(f)$ and $\text{Cone}(f)$ by

¹⁸More precisely, the inclusion maps are given as the composition of the respective inclusions $X \cong X \times \{1\} \subseteq X \times [0, 1] \sqcup Y \subseteq X \times [0, 1] \sqcup Y$ with the quotient map $X \times [0, 1] \sqcup Y \rightarrow \text{Cyl}(f)$.

¹⁹We recall that a deformation retraction of a space X onto a subspace A is a map $H: X \times [0, 1] \rightarrow X$ such that $H(x, 0) = x$ and $H(x, 1) \in A$ for all $x \in X$ and $H(a, 1) = a$ for all $a \in A$. Thus, a deformation retraction witnesses that the inclusion map $i_A: A \hookrightarrow X$ is a homotopy equivalence with a homotopy inverse $r = H(\cdot, 1): X \rightarrow A$ that is a retraction, i.e., that restricts to the identity on A .

A deformation retraction is called *strong* if it keeps A fixed pointwise throughout, i.e., if $H(a, t) = a$ for all $a \in A$ and $t \in [0, 1]$ (some authors include this directly in the definition of a deformation retraction).

collapsing the segment $x_0 \times [0, 1]$ (whose lower end is identified with y_0) to a single point. We will apply this construction only to cellular or simplicial mapping cylinders and cones, in which case contracting the subcomplex $x_0 \times [0, 1]$ is a homotopy equivalence.

Moreover, if f is a homotopy equivalence, then we may assume that its homotopy inverse g is pointed as well and that the homotopies $fg \simeq \text{id}_Y$ and $gf \simeq \text{id}_X$ keep the basepoints fixed (see [9, Corollary 0.19]). It follows that Lemma 3.3 remains true if we take $C = \widetilde{\text{Cyl}}(f)$ as the *reduced* mapping cylinder (the inclusions are given as those into $\text{Cyl}(f)$, followed by the quotient map $\text{Cyl}(f) \rightarrow \widetilde{\text{Cyl}}(f)$, which does not make any identifications within X or within Y).

By the remarks concerning the size of simplicial products and attachments, the size of the (reduced or unreduced) simplicial mapping cylinder or cone is at most the size of X plus the size of Y , times a factor depending only on $\dim X$.

3.4 Subdivisions and Simplicial Approximation

For simplicial complexes, there is the well-known notion of *barycentric subdivision* (see, e.g., [18, §15]). An analogous notion of subdivision, called *normal subdivision*, can also be defined for simplicial sets. Informally speaking, the normal subdivision $\text{Sd}(X)$ of a simplicial set X is defined by barycentrically subdividing each simplex of X and then gluing these subdivided simplices together according to the identifications implied by the face and degeneracy operators of X . We refer to [7, Sect. 4.6] for the precise formal definition and just state the facts that we will need in what follows.

For the standard simplex Δ^p , the nondegenerate k -simplices of $\text{Sd}(\Delta^p)$ correspond to chains of proper inclusions of nondegenerate simplices (faces) of Δ^p . It follows that $\text{Sd}(\Delta^p)$ has $(p+1)!$ nondegenerate p -simplices and, in general, at most $2^{p+1}(p+1)!$ nondegenerate simplices of any dimension. Consequently, for any simplicial set X , the size of $\text{Sd}(X)$ is at most the size of X times a factor that depends only on $\dim X$ and which can be bounded from above by $2^{\dim X+1}(\dim X+1)!$. Moreover, if the dimension is bounded, $\text{Sd}(X)$ can be constructed in time polynomial in $\text{size}(X)$.

If $f: X \rightarrow Y$ is a simplicial map, then subdivision also induces a map $\text{Sd}(f): \text{Sd}(X) \rightarrow \text{Sd}(Y)$, and this is compatible with compositions, i.e., $\text{Sd}(fg) = \text{Sd}(f)\text{Sd}(g)$.

For each simplicial set X , there is a simplicial map $\text{lastv}_X: \text{Sd}(X) \rightarrow X$, called the *last vertex map*,²⁰ which is a *homotopy equivalence* that is compatible with simplicial maps $f: X \rightarrow Y$, i.e., $f \text{lastv}_X = \text{lastv}_Y \text{Sd}(f)$.^{21,22}

²⁰On the standard simplex Δ^n , seen as a simplicial set, this map is defined by sending a chain $(\sigma_0, \dots, \sigma_k)$ (a k -simplex of $\text{Sd}(\Delta^n)$) to the simplex $[v_0, \dots, v_k]$, where v_i is the last vertex of the simplex σ_i (recall that the vertices in each simplex are ordered).

²¹In the language of category theory, Sd is a functor and lastv is a natural transformation between Sd and the identity functor on simplicial sets.

²²In fact, it is true that X and $\text{Sd}(X)$ are not only homotopy equivalent but homeomorphic (as one might expect given the terminology “subdivision”). However, for simplicial sets this is a decidedly nontrivial result; see [7, Corollary 4.6.5]. The difficulty is related to the fact that there is no way of defining this homeomorphism for all simplicial sets in such a way that it becomes compatible with simplicial maps. For our purposes, the natural homotopy equivalence lastv_X will be sufficient and more convenient.

There is also a simplicial approximation theorem for simplicial sets, which uses iterated normal subdivisions. Specifically, the t -fold iterated normal subdivision of a simplicial set is defined inductively as $\text{Sd}^t(X) := \text{Sd}(\text{Sd}^{t-1}(X))$, where $\text{Sd}^0(X) := X$.

Theorem 3.4 [7, Theorem 4.6.25] *Let X and Y be simplicial sets such that X has only finitely many nondegenerate simplices, and let $f: |X| \rightarrow |Y|$ be a continuous map. Then there exist a finite integer t (which depends on f) and a simplicial map $g: \text{Sd}^t(X) \rightarrow Y$ such that $|g|$ is homotopic to the composition $f|_{\text{lastv}_X^t}|$ of f with the iterated last vertex map $\text{lastv}_X^t: \text{Sd}^t(X) \rightarrow Y$.*

To convert arbitrary simplicial sets into homotopy equivalent (in fact, homeomorphic) simplicial complexes, another subdivision-like operation is needed (see, e.g., [11]). Given a simplicial set Z , one can define a simplicial complex $B_*(Z)$ inductively, by introducing a new vertex v_σ for every nondegenerate simplex σ , and then replacing σ by the cone with apex v_σ over $B_*(\partial\sigma)$. If the simplicial set Z has a certain regularity property—which is satisfied, for instance, if $Z = \text{Sd}(X)$ —then $B_*(Z)$ and Z are homotopy equivalent (in fact, homeomorphic).²³ We summarize the properties that we need in the following proposition (for completeness, we provide a proof in Appendix A).

Proposition 3.5 *If X is a simplicial set, then the twofold subdivision $B_*(\text{Sd}(X))$ is a simplicial complex. Moreover, there is a simplicial map $\gamma_X: B_*(\text{Sd}(X)) \rightarrow X$, which is a homotopy equivalence. For a simplicial subset $A \subseteq X$, $B_*(\text{Sd}(A))$ is a subcomplex of $B_*(\text{Sd}(X))$ and $\gamma_X|_A = \gamma_A$.²⁴*

If X is finite and of bounded dimension, there are algorithms that construct the simplicial complex $B_(\text{Sd}(X))$ and evaluate the map γ_X , both in polynomial time.*

4 Homotopy Groups

We review some further facts about homotopy groups that we will need. For more details see, e.g., [9, Sect. 4.1].

4.1 Basic Facts

So far, we used the definition of the n th homotopy group $\pi_n(X, x_0)$ of a pointed space (X, x_0) as the set of homotopy classes of pointed maps $(S^n, p_0) \rightarrow (X, x_0)$, where $p_0 \in S^n$ is an arbitrarily chosen basepoint, and the homotopies are required to keep the basepoint fixed. Equivalently, the elements of $\pi_n(X, x_0)$ can be viewed as

²³As an illustration that this fails for general simplicial sets, consider the case where $Z = \Sigma^p$ is the simplicial set model of the d -sphere with only two nondegenerate simplices, one in dimension 0 and one in dimension d . In this case, $B_*(\Sigma^d)$ is a 1-dimensional simplex.

²⁴In fact, the construction $B_*\text{Sd}$ is functorial and γ is a natural transformation (like the construction Sd and the map lastv), but we will never use this stronger fact.

homotopy classes $[f]$ of maps $f: (D^n, \partial D^n) \rightarrow (X, x_0)$ sending all of ∂D^n to x_0 , modulo homotopies that keep the image of ∂D^n fixed (as before, we will often drop the basepoint from the notation).²⁵

In what follows, we will also need the Abelian group operation in $\pi_n(X, x_0)$, $n \geq 2$, which can be defined as follows: Suppose $f_1, \dots, f_m$ are maps $(D^n, \partial D^n) \rightarrow (X, x_0)$. Suppose we have a cellular decomposition of D^n as a cell complex $D^n/_m$ with n -cells $e_1^n, \dots, e_m^n$ (in Sect. 6.1 we will provide a concrete geometric construction of $D^n/_m$). Then we can define a map f from $D^n \cong D^n/_m$ to (X, x_0) representing the homotopy class $[f_1] + \dots + [f_m]$ by sending the $(n-1)$ -skeleton of $D^n/_m$ to x_0 , and by defining the restriction of f to each open cell e_i^n to be f_i .

A very important special case of homotopy groups is that of spheres. We will use the following well-known facts:

- The sphere S^n is $(n-1)$ -connected.
- For all $n \geq 1$, $\pi_n(S^n)$ is isomorphic to $\mathbb{Z}$ and generated by the homotopy class ι of the identity. For each map $\varphi: S^n \rightarrow S^n$, there is a unique integer $a \in \mathbb{Z}$ such that $[\varphi] = a\iota$; it is called the *degree* of φ and denoted by $\deg \varphi$. The degree is obviously invariant under homotopy.
- We have $\pi_3(S^2) \cong \mathbb{Z}$. The group is generated by the famous *Hopf map*²⁶ $\eta: S^3 \rightarrow S^2$.

We will also need the following simple fact.

Lemma 4.1 *Let $g: S^n \rightarrow S^n$ be a map of degree $b \in \mathbb{Z}$. Then, for any map $f: S^n \rightarrow X$, we have $[fg] = b \cdot [f] \in \pi_n(X)$.*

Proof Consider the n -dimensional unit cube $I^n \cong D^n$, where $I = [0, 1]$ is the unit interval. We identify S^n with the quotient $I^n / \partial I^n$. From the map of sets $(I^n, \partial I^n) \rightarrow (I^n, \partial I^n)$ given by $(s_1, \dots, s_{n-1}, s_n) \mapsto (s_1, \dots, s_{n-1}, bs_n \bmod 1)$ we obtain $g_0: S^n \rightarrow S^n$ by passing to quotients. By the definition of the addition of homotopy classes, on the one hand, $[g_0]$ is the b -fold sum of the identity, and hence a particular example of a map of degree b . On the other hand, fg_0 is a representative of $b \cdot [f]$, the b -fold sum of $[f]$. Since $[fg]$ depends only on the homotopy class $[g]$, which is uniquely determined by the degree of g , the lemma follows. $\square$

Let X be a cell complex and $A \subseteq X$ a subcomplex. Then the homotopy groups of the spaces A , X , and X/A in a certain range are connected by an exact sequence.

Theorem 4.2 *Let $A \subseteq X$ be cell complexes. Let $p, q \geq 0$ be integers, $q \leq p+1$. If A is p -connected and X is q -connected, then X/A is also q -connected and there is an exact sequence*

$$\pi_{p+q}(A) \rightarrow \dots \rightarrow \pi_i(A) \rightarrow \pi_i(X) \rightarrow \pi_i(X/A) \rightarrow \pi_{i-1}(A) \rightarrow \dots \rightarrow \pi_{q+1}(X/A).$$

²⁵The claimed equivalence is obtained by identifying S^n with the quotient $D^n / \partial D^n$ of the n -disk by its boundary and p_0 with the image of ∂D^n under the quotient map.

²⁶See [9, Example 4.45] for the definition (which is not difficult, but which we will not need).

Here the maps $\pi_i(A) \rightarrow \pi_i(X)$ and $\pi_i(X) \rightarrow \pi_i(X/A)$ are induced by the inclusion and the projection, respectively, and the exactness means that the kernel of each homomorphism equals the image of the preceding one.

Proof One can define homotopy groups of any pair (X, A) , $A \subseteq X$, and these homotopy groups fit into the following exact sequence:

$$\cdots \rightarrow \pi_i(A) \rightarrow \pi_i(X) \rightarrow \pi_i(X, A) \rightarrow \pi_{i-1}(A) \rightarrow \cdots, \quad i \geq 1;$$

see [9], Sect. 4.1. From the exactness and the connectivity assumptions it is easy to show that $\pi_i(X, A) = 0$ for $i \leq q$. Then, according to [9, Proposition 4.28], the map $\pi_i(X, A) \rightarrow \pi_i(X/A)$ induced by the quotient map $X \rightarrow X/A$ is an isomorphism for $i \leq p + q$. Substituting $\pi_i(X/A)$ in this range into the exact sequence above, we get the exact sequence from the statement of the theorem. $\square$

In the proof of Theorem 1.1 we will need a description of the n th homotopy group of a cell complex Y obtained from T by attaching $(n + 1)$ -cells e_ℓ , $1 \leq \ell \leq m$, by attaching maps $\varphi_\ell: S^n \rightarrow T$.

Proposition 4.3 *Let $n \geq 2$ be an integer. Suppose that T is a 1-connected cell complex and Y is a cell complex obtained from T as described above. Then*

$$\pi_n(Y) \cong \pi_n(T) / \langle [\varphi_1], [\varphi_2], \dots, [\varphi_m] \rangle,$$

where $\langle [\varphi_1], [\varphi_2], \dots, [\varphi_m] \rangle$ is the subgroup of $\pi_n(T)$ generated by the homotopy classes of φ_ℓ , $1 \leq \ell \leq m$.

Proof It is sufficient to prove the statement for a single cell attached; then we can proceed by induction. Consider the reduced mapping cylinder $\widetilde{\text{Cyl}}\varphi$, together with the inclusions of S^n and T into it and the projection onto $\widetilde{\text{Cyl}}\varphi/S^n = \widetilde{\text{Cone}}\varphi = Y$. The situation is summarized in the diagram

$$\begin{array}{ccccc} S^n & \longrightarrow & \widetilde{\text{Cyl}}\varphi & \longrightarrow & \widetilde{\text{Cyl}}\varphi/S^n = Y \\ & \searrow \varphi & \uparrow \sim & & \\ & & T & & \end{array}$$

which commutes up to homotopy. Applying Theorem 4.2 for $A = S^n$, $X = \widetilde{\text{Cyl}}\varphi$, $p = n - 1$, and $q = 1$, we obtain the exact sequence

$$\pi_n(S^n) \rightarrow \pi_n(\widetilde{\text{Cyl}}\varphi) \rightarrow \pi_n(Y) \rightarrow \pi_{n-1}(S^n) = 0.$$

If we replace the inclusion $S^n \hookrightarrow \widetilde{\text{Cyl}}\varphi$ by the map $\varphi: S^n \rightarrow T$, we get

$$\pi_n(S^n) \xrightarrow{\varphi_*} \pi_n(T) \rightarrow \pi_n(Y) \rightarrow 0.$$

Hence $\pi_n(Y) = \pi_n(T) / \langle \varphi_*(\iota) \rangle$, where ι is the homotopy class of the identity on S^n , and thus $\varphi_*(\iota)$ is the homotopy class of φ . $\square$

4.2 Whitehead Products and Wedge Sums of Spheres

Whitehead Products There is another type of operation on elements of homotopy groups that we will need. Consider two spheres S^k and S^ℓ with their standard structures as cell complexes (one vertex and one cell of the top dimension). Then the product $S^k \times S^\ell$ is also a cell complex, with one vertex, one respective cell e^k and e^ℓ in dimensions k and ℓ , and one cell $e^k \times e^\ell$ in dimension $k + \ell$. In particular, the $(k + \ell - 1)$ -skeleton of the product is a wedge $S^k \vee S^\ell$, to which the $(k + \ell)$ -cell is attached via a map $\varphi: S^{k+\ell-1} \cong \partial(D^{k+\ell}) \rightarrow S^k \vee S^\ell$.

Now, if $f: S^k \rightarrow X$ and $g: S^\ell \rightarrow X$ are (pointed) maps, we can combine them into a map $f \vee g: S^k \vee S^\ell \rightarrow X$. If we compose this with the attachment map φ discussed before, we get a map $[f, g]: S^{k+\ell-1} \rightarrow X$, called the *Whitehead product* of f and g . The homotopy class of this product clearly depends only on the homotopy classes of the factors, so we get a well-defined product $\pi_k(X) \times \pi_\ell(X) \rightarrow \pi_{k+\ell-1}(X)$, again denoted by $[\cdot, \cdot]$. As a quite trivial but nonetheless useful example, if $X = S^k \times S^\ell$, then the attachment map φ itself equals the Whitehead product $[\iota_{S^k}, \iota_{S^\ell}]$ of the two inclusions $\iota_{S^k}: S^k \hookrightarrow S^k \vee S^\ell$ and $\iota_{S^\ell}: S^\ell \hookrightarrow S^k \vee S^\ell$.

In our proofs we will use the fact that the Whitehead product is natural, graded commutative, and bilinear, i.e.,

$$\begin{aligned} f_*[\alpha, \beta] &= [f_*\alpha, f_*\beta], \\ [\alpha, \beta] &= (-1)^{k\ell}[\beta, \alpha], \\ [\alpha + \gamma, \beta] &= [\alpha, \beta] + [\gamma, \beta], \\ [\alpha, \beta + \delta] &= [\alpha, \beta] + [\alpha, \delta] \end{aligned}$$

where $\alpha, \gamma \in \pi_k(X)$, $\beta, \delta \in \pi_\ell(X)$, and $f: X \rightarrow Y$. For the proof see [34], Chap. X, 7.2, Corollary 7.12 and Corollary 8.13.

In the proof of Theorem 1.1 we will need some facts about the homotopy groups of spheres and their wedge sums.

Theorem 4.4 [9, Corollary 4B.2], [34, XI, Theorem 2.5] *There is a homomorphism (called the Hopf invariant) $H: \pi_{2k-1}(S^k) \rightarrow \mathbb{Z}$ such that for d even $H([\iota, \iota]) = \pm 2$.*

Let us note that for k odd the Whitehead product $[\iota, \iota] \in \pi_{2k-1}(S^k)$ is of order 2, i.e., $2[\iota, \iota] = 0$. Whitehead products play a crucial role in Hilton's theorem, which converts the computation of homotopy groups of a wedge of spheres to the computations of homotopy groups of spheres. We do not need this theorem in its full generality as it was proved in [10], and so we restrict ourselves to a special case.

Let $k \geq 2$ and $r, s \geq 1$ be integers. Let

$$T = S_1^k \vee \dots \vee S_r^k \vee S_1^{2k-1} \vee \dots \vee S_s^{2k-1} \quad (6)$$

be the wedge sums of r copies of S^k and s copies of S^{2k-1} . Denote by v_i and μ_q the homotopy classes of the inclusions $S^k \hookrightarrow T$ and $S^{2k-1} \hookrightarrow T$ onto the i th copy of S^k and the q th copy of S^{2k-1} , respectively. Then the homotopy groups $\pi_k(T)$ and $\pi_{2k-1}(T)$ can be described by the following special case of Hilton's theorem.

Theorem 4.5 [10, Theorem A] *With the notation as above, there are isomorphisms*

$$\begin{aligned}\pi_k(T) &\cong \bigoplus_{1 \leq i \leq r} \pi_k(S_i^k), \\ \pi_{2k-1}(T) &\cong \bigoplus_{1 \leq i \leq r} \pi_{2k-1}(S_i^k) \oplus \bigoplus_{1 \leq q \leq s} \pi_{2k-1}(S_q^{2k-1}) \oplus \bigoplus_{1 \leq i < j \leq r} \pi_{2k-1}(S_{ij}^{2k-1}).\end{aligned}$$

An element $\beta \in \pi_{2k-1}(S_i^k)$ corresponds to the composition $v_i \beta \in \pi_{2k-1}(T)$, an element $\beta \in \pi_{2k-1}(S_q^{2d-1})$ to the composition $\mu_q \beta \in \pi_{2k-1}(T)$, and an element $\beta \in \pi_{2k-1}(S_{ij}^{2k-1})$ to the composition $[v_i, v_j] \beta \in \pi_{2k-1}(T)$.

We will say that some elements $x_1, \dots, x_r$ of an Abelian group are *integrally independent* if the only valid relation $a_1 x_1 + \dots + a_r x_r = 0$ with coefficients $a_i \in \mathbb{Z}$ is that with all a_i zero. The following statement is an immediate consequence of Theorems 4.5 and 4.4.

Corollary 4.6 *If $k \geq 2$ is odd, then the elements μ_q , $1 \leq q \leq s$, and $[v_i, v_j]$, $1 \leq i < j \leq r$ are integrally independent in $\pi_{2k-1}(T)$.*

If $k \geq 2$ is even, then the elements μ_q , $1 \leq q \leq s$ and $[v_i, v_j]$, $1 \leq i \leq j \leq r$ are integrally independent in $\pi_{2k-1}(T)$.

Proof The reason is that every element in the list comes from a different direct summand and is of infinite order. $\square$

In the case $k = 2$ and $s = 0$ we can say even more.

Corollary 4.7 [9, Example 4.52] *The homotopy group $\pi_3(\bigvee_{i=1}^r S_i^2)$ is a free Abelian group generated by the Whitehead products $[v_i, v_j]$, $1 \leq i < j \leq r$, and homotopy classes $v_i \eta$, where $\eta: S^3 \rightarrow S^2$ is the Hopf map.*

5 The Constructions for Theorem 1.1 Presented as Cell Complexes

Here we present the essence of the proof of Theorem 1.1. Namely, for every system of quadratic Diophantine equations of the form (Q-SYM) (for k even) or (Q-SKEW) (for k odd), we construct cell complexes A , X , Y , and a continuous map $f: A \rightarrow Y$, where Y is $(k-1)$ -connected and $\dim X = 2k$, such that f is extendable to X iff the Diophantine system has a solution. Moreover, one of (X, A) and Y can be assumed to be fixed, as in Theorem 1.1(a) and (b). We will also see the role of Whitehead products and Hilton's theorem in the proof.

What remains for the next section is to convert X , A , Y into finite simplicial complexes and f into a simplicial map, so that the solvability of the extension problem remains unchanged. Moreover, the construction has to be algorithmic.

While discussing the cellular constructions of X , A , Y , it is also natural to describe the cell complex used by Anick in the proof of his #P-hardness result. Indeed, his construction uses tools very similar to those employed in our constructions.

The Generalized Extension Problem In order to simplify the presentation, it is convenient to remove the assumption in the extension problem that A is a subspace of X , or in other words, that the map $A \rightarrow X$ is an inclusion. Instead, we consider three spaces A, W, Y and (arbitrary) maps $g: A \rightarrow W$ and $f: A \rightarrow Y$, and we ask if there is a map $h: W \rightarrow Y$ making the following diagram commutative up to homotopy:

$$\begin{array}{ccc} A & \xrightarrow{f} & Y \\ g \downarrow & \nearrow h & \\ W & & \end{array} \quad (\text{GEP})$$

For a generalized extension problem as above, we obtain an equivalent extension problem by setting $X = \text{Cyl}(g)$ (where A is considered as a subspace of the cylinder in the usual way). This is easy to see, but we nonetheless briefly describe a proof of this fact that only uses the homotopy extension property for pairs of cell complexes and the properties of the mapping cylinder summarized in Lemma 3.3. This means that we can replace the mapping cylinder $\text{Cyl}(g)$ by any other cell complex that has these properties, and the same proof will still apply. This will be useful for our simplicial constructions later on, for which it will be convenient to work with what are called *generalized mapping cylinders* (see Sect. 6.2).

Let i_A and i_W be the inclusions of A and W into X . On the one hand, given a solution $\tilde{f}: X \rightarrow Y$ of the extension problem, i.e., $\tilde{f}i_A = f$, we can define $h := \tilde{f}i_W$ as the restriction of $\tilde{f}$ to W . Then $hg = \tilde{f}i_Wg \sim \tilde{f}i_A = f$, so h is a solution to the generalized extension problem.

On the other hand, given a solution h for the generalized extension problem (GEP), let $r: X \rightarrow W$ be the retraction from X to W and define $\tilde{f} := hr$. Then $\tilde{f}i_A = hri_A \sim hg \sim f$, so $\tilde{f}$ is an extension of a map homotopic to f , and since extendability depends only on the homotopy class of a map (Corollary 3.1), f can be extended as well.

Thus, we are free to consider the generalized extension problem with $\dim W \leq 2k$ and $\dim A \leq 2k - 1$.

5.1 Fixed Target

We describe an instance of the generalized extension problem for part (a) of Theorem 1.1, where the target Y is fixed. The system of equations will be encoded into cell complexes A, W and the maps $g: A \rightarrow W, f: A \rightarrow Y$.

Fixed Target with k Even Here $Y = S^k$,

$$A = S_1^{2k-1} \vee \cdots \vee S_s^{2k-1}, \quad W = S_1^k \vee \cdots \vee S_r^k. \quad (8)$$

Then the diagram (GEP) becomes

$$\begin{array}{ccc} A = S_1^{2k-1} \vee \dots \vee S_s^{2k-1} & \xrightarrow{f} & S^k \\ g \downarrow & \nearrow h & \\ W = S_1^k \vee \dots \vee S_r^k & & \end{array}$$

According to (5) from the discussion of wedge sums in Section 3.3, the homotopy class of $f: A \rightarrow S^k$ is specified completely by the homotopy classes of its restrictions to the spheres forming A . We will use $f_*\mu_q = [f]\mu_q \in \pi_{2k-1}(S^k)$ to denote these, where μ_q is the homotopy class of the inclusion of the q th sphere S_q^{2k-1} into A . Our particular choice is

$$f_*\mu_q = b_q[\iota, \iota]. \quad (9)$$

Similarly, the homotopy class of $g: A \rightarrow W$ is given by its restrictions as

$$g_*\mu_q = \sum_{1 \leq i < j \leq r} a_{ij}^{(q)}[v_i, v_j], \quad (10)$$

where v_i is the homotopy class of the inclusion of the i th sphere S_i^k into W . Finally, let $h: W \rightarrow S^k$ be an arbitrary map and write $h_*v_i = x_i\iota \in \pi_k(S^k)$ for some integers $x_i \in \mathbb{Z}$. According to (5) again, the diagram (GEP) commutes up to homotopy iff $(hg)_*\mu_q = f_*\mu_q$, i.e., iff

$$h_*\left(\sum_{1 \leq i < j \leq r} a_{ij}^{(q)}[v_i, v_j]\right) = b_q[\iota, \iota] \quad (11)$$

for all $1 \leq q \leq s$. Using the naturality and bilinearity of the Whitehead product, the left-hand side equals

$$\sum_{1 \leq i < j \leq r} a_{ij}^{(q)}[h_*v_i, h_*v_j] = \sum_{1 \leq i < j \leq r} a_{ij}^{(q)}[x_i\iota, x_j\iota] = \sum_{1 \leq i < j \leq r} a_{ij}^{(q)}x_ix_j[\iota, \iota].$$

According to Corollary 4.6 the homotopy class $[\iota, \iota]$ is of infinite order, and so the system of Eqs. (11) is equivalent to (Q-SYM). We get the following proposition.

Proposition 5.1 *Let the maps $f: A \rightarrow S^k$ and $g: A \rightarrow W$ be as in (9) and (10) above. Then f can be extended to $X = \widetilde{\text{Cyl}}(g)$ if and only if the system (Q-SYM) has a solution.*

Fixed Target with k Odd The element $[\iota, \iota] \in \pi_{2k-1}(S^k)$ has order 2, so we cannot use $Y = S^k$. However, leaving A , W and $g: A \rightarrow W$ as before, we can take $Y = S^k \vee S^k$ and specify f by

$$f_*\mu_q = b_q[\iota_1, \iota_2], \quad (12)$$

where ι_1 and ι_2 are inclusions of S^k onto the first and the second summand in $S^k \vee S^k$, respectively. Using Hilton's theorem (Theorem 4.5) for $\pi_k(S^k \vee S^k)$, the homotopy class of a general map $h: W \rightarrow Y$ satisfies

$$h_* v_i = x_i \iota_1 + y_i \iota_2, \quad x_i, y_i \in \mathbb{Z}. \quad (13)$$

Using the fact that $[\iota_1, \iota_2] = -[\iota_2, \iota_1]$, it is easy to show that the commutativity of the diagram (GEP) is equivalent to the system of s equations in $\pi_{2k-1}(S^k \vee S^k)$,

$$\begin{aligned} & \left(\sum_{i < j} a_{ij}^{(q)} (x_i y_j - x_j y_i) \right) [\iota_1, \iota_2] + \left(\sum_{i, j} a_{ij}^{(q)} x_i x_j \right) [\iota_1, \iota_1] \\ & + \left(\sum_{i, j} a_{ij}^{(q)} y_i y_j \right) [\iota_2, \iota_2] = b_q [\iota_1, \iota_2]. \end{aligned}$$

By Corollary 4.6 of Hilton's theorem the element $[\iota_1, \iota_2] \in \pi_{2k-1}(S^d \vee S^d)$ is of infinite order, while $[\iota_1, \iota_1]$ and $[\iota_2, \iota_2]$ are of order 2. Multiplying all the equations in (Q-SKEW) by 2, we get an equivalent system, in which all the $a_{ij}^{(q)}$ are even. For this system, the above equation is exactly the one from (Q-SKEW). We get the following.

Proposition 5.2 *Let the maps $f: A \rightarrow S^k \vee S^k$ and $g: A \rightarrow W$ be as in (12) and (10) above. Then f can be extended to $X = \widehat{\text{Cyl}}(g)$ if and only if the system (Q-SKEW) has a solution.*

5.2 Fixed Source

The idea for the first step of the proof of Theorem 1.1(b) is similar; however, the constructions involve attaching cells, and also the usage of Hilton's theorem is more substantial.

Fixed Source with k Even We put

$$A = S^{2k-1}, \quad W = S^k$$

where the homotopy class of $g: A \rightarrow W$ is $[g] = [\iota, \iota]$. A given system of equations (Q-SYM) will be encoded in the target space Y and in the homotopy class of $f: A \rightarrow Y$. The target space Y is a cell complex obtained from the wedge of spheres T defined in (6) by attaching $(2k)$ -cells e_{ij} , $1 \leq i < j \leq r$ and e_{ii} , $1 \leq i \leq r$, i.e.,

$$Y = \underbrace{(S_1^k \vee \dots \vee S_r^k \vee S_1^{2k-1} \vee \dots \vee S_s^{2k-1})}_T \cup \bigcup_{1 \leq i < j \leq r} e_{ij} \cup \bigcup_{1 \leq i \leq r} e_{ii}. \quad (14)$$

The attaching maps for the cells are the maps $S^{2k-1} \rightarrow T$ whose homotopy classes are, respectively,

$$\varphi_{ij} = [v_i, v_j] - \sum_{1 \leq q \leq s} a_{ij}^{(q)} \mu_q, \quad \varphi_{ii} = [v_i, v_i].$$

Denote the images of the homotopy classes $\mu_q \in \pi_{2k-1}(T)$, $1 \leq q \leq s$ and $v_i \in \pi_k(T)$, $1 \leq i \leq r$, in Y by μ'_q and v'_i , respectively. Further, take a map $f: A \rightarrow Y$ of the homotopy class

$$[f] = 2b_1\mu'_1 + 2b_2\mu'_2 + \cdots + 2b_s\mu'_s.$$

Since $\pi_k(Y) \cong \pi_k(T) \cong \pi_k(S_1^k) \oplus \cdots \oplus \pi_k(S_r^k)$ by Theorem 4.5, a general map $h: W \rightarrow Y$ has a homotopy class

$$[h] = x_1v'_1 + x_2v'_2 + \cdots + x_rv'_r$$

with arbitrary integer coefficients x_i . To show that the commutativity of the diagram (GEP) (up to homotopy) is equivalent to the satisfaction of the system (Q-SYM), we will need the following lemma.

Lemma 5.3 *Let Y be the cell complex as above. Then the classes $\mu'_q \in \pi_{2k-1}(Y)$, $1 \leq q \leq s$, are integrally independent and*

$$\begin{aligned} [v'_i, v'_j] &= \sum_{1 \leq q \leq s} a_{ij}^{(q)} \mu'_q, \quad 1 \leq i < j \leq r, \\ [v'_i, v'_i] &= 0, \quad 1 \leq i \leq r. \end{aligned}$$

Proof The statement is a consequence of Proposition 4.3 and Corollary 4.6. $\square$

Using this lemma and the bilinearity and graded commutativity of the Whitehead product, we compute $[hg] \in \pi_{2k-1}(Y)$ as

$$\begin{aligned} h_*[g] &= h_*[t, \iota] = [h_*t, h_*\iota] \\ &= \left[\sum_{1 \leq i \leq r} x_i v'_i, \sum_{1 \leq j \leq r} x_j v'_j \right] = \sum_{1 \leq i, j \leq r} x_i x_j [v'_i, v'_j] \\ &= 2 \sum_{1 \leq i < j \leq r} x_i x_j [v'_i, v'_j] + \sum_{1 \leq i \leq r} x_i^2 [v'_i, v'_i] \\ &= 2 \sum_{1 \leq i < j \leq r} x_i x_j \left(\sum_{q=1}^s a_{ij}^{(q)} \mu'_q \right) \\ &= 2 \sum_{1 \leq q \leq s} \left(\sum_{1 \leq i < j \leq r} a_{ij}^{(q)} x_i x_j \right) \mu'_q. \end{aligned}$$

Comparing with $[f]$ and using the fact that μ'_q are integrally independent, we obtain the system (Q-SYM).

Fixed Source with Odd k As in the fixed target case, we resolve the problem of $[\iota, \iota]$ being of order 2 by replacing it with $[\iota_1, \iota_2]$. In this case, it means that we set

$A = S^{2k-1}$, $W = S^k \vee S^k$. The target space Y remains the same as for k even. We take f to be any map with

$$[f] = b_1\mu'_1 + b_2\mu'_2 + \cdots + b_s\mu'_s$$

and g has the advertised homotopy class $[g] = [\iota_1, \iota_2]$, where ι_1 and ι_2 are the homotopy classes of the inclusions of the two copies of S^k into $W = S^k \vee S^k$. The homotopy class of a map $h: W \rightarrow Y$ is again determined by its restrictions along ι_1, ι_2 , namely

$$h_*\iota_1 = x_1v'_1 + x_2v'_2 + \cdots + x_rv'_r, \quad h_*\iota_2 = y_1v'_1 + y_2v'_2 + \cdots + y_rv'_r,$$

where the x_i and y_i can be arbitrary integers. The composition $[hg] \in \pi_{2k-1}(Y)$ equals

$$\begin{aligned} h_*[g] &= \left[\sum_{1 \leq i \leq r} x_i v'_i, \sum_{1 \leq j \leq r} y_j v'_j \right] = \sum_{1 \leq i, j \leq r} x_i y_j [v'_i, v'_j] \\ &= \sum_{1 \leq i < j \leq r} (x_i y_j - x_j y_i) [v'_i, v'_j] + \sum_{1 \leq i \leq r} x_i y_i [v'_i, v'_i] \\ &= \sum_{1 \leq i < j \leq r} (x_i y_j - x_j y_i) \left(\sum_{1 \leq q \leq s} a_{ij}^{(q)} \mu'_q \right) \\ &= \sum_{1 \leq q \leq s} \left(\sum_{1 \leq i < j \leq r} a_{ij}^{(q)} (x_i y_j - x_j y_i) \right) \mu'_q \end{aligned}$$

(again using Lemma 5.3). Since the μ'_q are integrally independent, the comparison with $[f]$ leads to the system (Q-SKEW).

Summarizing our findings, for k both even and odd we get the following.

Proposition 5.4 *For each $k \geq 2$ let the maps $f: A \rightarrow Y$ and $g: A \rightarrow W$ be as above. Then f can be extended to $X = \widetilde{\text{Cyl}}(g)$ if and only if there is a solution to the system (Q-SYM) when k is even, or (Q-SKEW) when k is odd.*

5.3 Anick's 4-Dimensional Cell Complexes

Here we introduce complexes constructed by Anick [2, p. 42] for his hardness result. These are compact 4-dimensional cell complexes which arise from the wedge $W = S^2_1 \vee \cdots \vee S^2_r$ of r copies of S^2 by attaching s 4-cells. According to Corollary 4.7, the homotopy class of a general attaching map must be an integral combination of the homotopy classes $v_i \eta$ and the Whitehead products $[v_i, v_j]$, where η is the homotopy class of the Hopf map $S^3 \rightarrow S^2$ and v_i is the homotopy class of the inclusion $S^2 \rightarrow W$ on the i th copy of S^2 .

Together with [9, Proposition 0.18] this implies that, up to homotopy equivalence, a completely general way of attaching 4-cells to W is described by integers $a_{ij}^{(q)}$ for

$1 \leq i \leq j \leq r$ and $q = 1, 2, \dots, s$. Specifically, the q th 4-cell is attached via a map $S_q^3 \rightarrow W$ representing the homotopy class in $\pi_3(W)$ defined by

$$\varphi_q = \sum_{1 \leq i \leq r} a_{ii}^{(q)} \iota_i \eta + \sum_{1 \leq i < j \leq r} a_{ij}^{(q)} [\iota_i, \iota_j].$$

Therefore, the *homotopy type* of the resulting *Anick complex*, i.e., its class of homotopy equivalence, which we denote by Y_a^4 , is completely determined by the vector $\mathbf{a} = (a_{ij}^{(q)})_{1 \leq i \leq j \leq r}^{1 \leq q \leq s}$ of integer coefficients.

In Anick's #P-hardness result, the input complex Y_a^4 (whose higher homotopy groups are to be computed) is encoded very concisely by the vector $\mathbf{a}$ of integers, represented in binary.

Theorem 5.5 (Anick [2]) *It is #P-hard to compute the rank of $\pi_n(Y_a^4)$ for a given integer $n \geq 2$ (encoded in unary) and a given integer vector $\mathbf{a}$ (represented in binary).*

In Sect. 6, we will show that, given $\mathbf{a}$, we can construct, in polynomial time, a finite 4-dimensional simplicial complex homotopy equivalent to Y_a^4 . Together with Anick's result, this will imply Theorem 1.2.

6 Simplicial Constructions

In this section, we prove that the constructions of cell complexes and cellular maps from the last section can be converted into homotopy equivalent finite simplicial complexes and simplicial maps. The constructions of such simplicial sets and maps involve only simplicial products, attachments (in particular, mapping cylinders), quotients, and subdivisions, which are all algorithmic. It is crucial, that in the case of Anick's complexes, the algorithms are polynomial time.

We remark that the conversion of the cell complexes and maps for the extension problem (those from Propositions 5.1, 5.2, and 5.4) runs in time polynomial in the size of the binary encoding of the coefficients a_{ij}^q and the unary encoding of the coefficients b_q . The running time could be made polynomial with respect to the binary encoding of all the coefficients including b_q ; however, it would complicate the presentation slightly, and it is unnecessary for the undecidability results.

Let us denote by Σ^p a “model” of the sphere S^p as a simplicial set with only two nondegenerate simplices, one in dimension 0 and the other in dimension p .

6.1 Constructing the Sum of Several Maps $S^p \rightarrow Y$

In this short section we describe how, given simplicial maps $f_1, \dots, f_m: \Sigma^p \rightarrow Y$, we can construct a simplicial representative of the sum $[f_1] + \dots + [f_m] \in \pi_p(Y)$. To this end, we have to change the domain to a simplicial set with a larger number of simplices.

We define the simplicial set $D_{/m}^p$, which is a union of m copies of Δ^p , where the i th copy is glued by its ∂_1 -face to the ∂_0 -face of the $(i+1)$ st copy. The union of the

remaining faces (the ∂_1 -face of the first copy, the ∂_0 -face of the last copy, and all the ∂_i -faces with $i > 1$) is denoted by $\partial D^p_{/m}$. Here is a picture of $D^2_{/m}$:

with the double arrows denoting the boundary $\partial D^2_{/m}$. Another point of view is, that $D^1_{/m}$ is a chain of m copies of the 1-simplex Δ^1 and each $D^p_{/m}$ is a cone over $D^{p-1}_{/m}$.

There is a simplicial map $D^p_{/m} \rightarrow \Delta^p$ that sends the first copy of Δ^p in $D^p_{/m}$ onto Δ^p by the identity, while the rest is sent to the degeneracy of the ∂_0 -face of Δ^p . It induces a simplicial map

$$q: D^p_{/m}/\partial D^p_{/m} \rightarrow \Sigma^p,$$

which is a homotopy equivalence (it is easy to see this, e.g., from homology). There is another simplicial map that collapses the whole $(p-1)$ -skeleton of $D^p_{/m}$. The map factors through $D^p_{/m}/\partial D^p_{/m}$ as

$$D^p_{/m} \longrightarrow D^p_{/m}/\partial D^p_{/m} \xrightarrow{\delta} \Sigma^p \vee \dots \vee \Sigma^p.$$

We specify a simplicial map $f: \Sigma^p \vee \dots \vee \Sigma^p$ by mapping the i th copy of Σ^p to Y by f_i .

The maps q , δ , and f fit into a diagram

$$\begin{array}{ccc} D^p_{/m}/\partial D^p_{/m} & \xrightarrow{\delta} & \Sigma^p \vee \dots \vee \Sigma^p \xrightarrow{f} Y \\ q \downarrow \sim & & \nearrow \\ \Sigma^p & \text{---} & \end{array}$$

Since q has a continuous homotopy inverse, there is a unique homotopy class of maps $\Sigma^p \rightarrow Y$ extending $[f\delta]$ up to homotopy, namely the homotopy class of $[f_1] + \dots + [f_m]$.

By the naturality of the subdivision, we also have maps

$$\begin{aligned} q: \text{Sd}^t(D^p_{/m}/\partial D^p_{/m}) &\rightarrow \Sigma^p, \\ \delta: \text{Sd}^t(D^p_{/m}/\partial D^p_{/m}) &\rightarrow \text{Sd}^t(\Sigma^p) \vee \dots \vee \text{Sd}^t(\Sigma^p) \end{aligned}$$

(the map q is the composition of the subdivision of the original q with the iterated last vertex map $\text{Sd}^t(\Sigma^p) \rightarrow \Sigma^p$) that will serve to add representatives $f_1, \dots, f_m: \text{Sd}^t(\Sigma^p) \rightarrow Y$.

6.2 Generalized Mapping Cylinders

In the above approach, in order to construct simplicial maps $\Sigma^p \rightarrow Y$, we replaced the domain Σ^p by a homotopy equivalent simplicial set. This will be very useful for the proof of part (a) of Theorem 1.1. For part (b), the domain has to be left unchanged, and thus a different construction must be used.²⁷ It is roughly “dual” to the previous one: It replaces the target Y by a homotopy equivalent simplicial set.

Thus, instead of subdividing the sphere Σ^p , we will replace the target space Y by a “generalized mapping cylinder.” This solution also works for domains other than Σ^p . Thus, for a map $f: X \rightarrow Y$, we will be interested in diagrams, commutative up to homotopy, of the following form:

$$\begin{array}{ccc} & & Y \\ & \nearrow f & \downarrow i_Y \\ X & \xrightarrow{\quad} & M \end{array} \quad \begin{array}{c} \sim \\ i_X \end{array}$$

Definition 6.1 Let M be a pointed simplicial set with two simplicial subsets $X, Y \subseteq M$ containing the basepoint of M . Let $i_X: X \rightarrow M$ and $i_Y: Y \rightarrow M$ be the corresponding inclusion maps, and let $f: |X| \rightarrow |Y|$ be a pointed continuous map. We say that M is a *generalized mapping cylinder* for f , with *upper rim* X and *lower rim* Y , if i_Y is a homotopy equivalence and $i_X \sim i_Y f$. We denote this situation by $M: X \xrightarrow{f} \rightsquigarrow Y$.

We remark that the above definition depends only on the homotopy class of f ; we may thus say that M is a generalized mapping cylinder for the homotopy class $[f]$.

By Lemma 3.3 and the remark following it, the reduced mapping cylinder $\widetilde{\text{Cyl}}(f)$ of $f: X \rightarrow Y$ is an example of such a generalized mapping cylinder with upper rim X and lower rim Y . Moreover, if f is a homotopy equivalence with a homotopy inverse g , one can easily see from the definition that a generalized mapping cylinder M for f is also a generalized mapping cylinder M^{op} for g with upper rim Y and lower rim X (i.e., the roles of upper and lower rim are interchanged).²⁸

The important property of generalized mapping cylinders, which we are going to use heavily, is that they may be used for attaching cells.

Proposition 6.2 Let $M: \Sigma_1^p \vee \cdots \vee \Sigma_m^p \xrightarrow{f} \rightsquigarrow Y$ be a generalized mapping cylinder for a pointed map f , whose restriction to the i th summand is $f_i: |\Sigma_i^p| \rightarrow |Y|$. Then the

²⁷There is a further issue with the subdivision—it is not polynomial. The construction of a representative of a multiple $b[f]$ of a map $f: \Sigma^p \rightarrow Y$ requires b simplices, and this number is exponential in the number of bits of b .

²⁸When f is injective, one can use Y as a generalized mapping cylinder for f .

composition $Y \xrightarrow{i_Y} M \xrightarrow{\text{proj}} M/(\Sigma_1^p \vee \dots \vee \Sigma_m^p)$ extends to a homotopy equivalence

$$Y \cup (e_1^{p+1} \cup \dots \cup e_m^{p+1}) \xrightarrow{\sim} M/(\Sigma_1^p \vee \dots \vee \Sigma_m^p)$$

where the cell e_i^{p+1} on the left is attached to Y along the map f_i .

Proof Set $X = \Sigma_1^p \vee \dots \vee \Sigma_m^p$. Then the space from the statement, obtained from Y by attaching cells, is the mapping cone of f . By [3, Theorem I.14.19] the mapping cone of $f: X \rightarrow Y$ is homotopy equivalent to that of $i_Y f: X \rightarrow M$ (since i_Y is a homotopy equivalence). Further, by [3, Theorem I.14.18] it is also homotopy equivalent to the mapping cone of $i_X: X \rightarrow M$ (since $i_Y f \sim i_X$). By [3, Theorem VII.1.6] the mapping cone of i_X is homotopy equivalent to M/X .

All the involved maps respect Y (which is naturally a subspace of all the mapping cones and also maps by $\text{proj } i_Y$ to the quotient M/X), proving that the resulting homotopy equivalence is indeed an extension of the composition $\text{proj } i_Y$. $\square$

Generalized mapping cylinders can be composed in an obvious way.

Lemma 6.3 *Let $f: X \rightarrow Y$ and $g: Y \rightarrow Z$ be pointed continuous maps, and let M and N be generalized mapping cylinders for f and g , respectively. Let $NM := N \cup_Y M$ be the simplicial set obtained by identifying the lower rim of M with the upper rim of N . Then NM is a generalized mapping cylinder for gf .*

Proof Consider the diagram

$$\begin{array}{ccccc}
 X & & & & \\
 \downarrow f & \searrow i_X & & & \\
 & & M & \xrightarrow{i_M} & N \cup_Y M \\
 & \nearrow i_Y & & & \\
 Y & & & & \\
 \downarrow g & \searrow j_Y & & & \\
 & & N & \xrightarrow{i_N} & N \cup_Y M \\
 & \nearrow j_Z & & & \\
 Z & & & &
 \end{array}$$

where $i_X, i_Y, j_Y, i_Z, i_M, i_N$ are inclusions, both triangles commute up to homotopy, and the square commutes strictly. Consequently, the triangle formed by the spaces X, Z and $N \cup_Y M$ commutes up to homotopy, too. To show that $N \cup_Y M$ is a generalized mapping cylinder for gf , it suffices to prove that the inclusion i_N is a homotopy equivalence.

It is well known (see [9, Theorem 4.5]) that to every inclusion $i_Y: Y \rightarrow M$ which is a homotopy equivalence there is a deformation retraction $r: M \rightarrow Y$. Then the map $h: N \cup_Y M \rightarrow N$ defined as $j_Y r$ on M and as the identity on N is a homotopy inverse to i_N . $\square$

We will also need simplicial maps representing (the homotopy classes of) a constant number of specific maps, such as the Whitehead product $[\iota, \iota]: S^{2d-1} \rightarrow S^d$ of

the identity on S^d with itself, the Whitehead product $[l_1, l_2]: S^{2d-1} \rightarrow S_1^d \vee S_2^d$ of the inclusions $S_i^d \rightarrow S_1^d \vee S_2^d$, and the Hopf map $\eta: S^3 \rightarrow S^2$. In each case, it is possible to construct these explicitly, but we will use the following general lemma.

Lemma 6.4 *Let X and Y be finite simplicial sets and let $f: |X| \rightarrow |Y|$ be an arbitrary but fixed pointed continuous map. Then there exists a generalized mapping cylinder $X \xrightarrow{f} Y$. It is of dimension $\max\{\dim X + 1, \dim Y\}$.*

The point here is that, in contrast with Theorem 3.4, we can prescribe the exact triangulations of the upper and lower rims, which will make it easy to compose the resulting generalized mapping cylinders.

Proof By the simplicial approximation theorem for simplicial sets (Theorem 3.4), there exist an iterated barycentric subdivision $X' = \text{Sd}^t(X)$ of X and a simplicial map $g: X' \rightarrow Y$ homotopic to $f|_{\ell}$, where $\ell: X' \rightarrow X$ is the natural homotopy equivalence (the iterated last vertex map).

Let $M := \text{Cyl}(\ell)$ and $N := \text{Cyl}(g)$ be the corresponding reduced simplicial mapping cylinders. Since ℓ is a homotopy equivalence, we can also view M as a generalized mapping cylinder M^{op} for a homotopy inverse $h: |X| \rightarrow |X'|$ of $|\ell|$, with upper rim X and lower rim X' . Thus, NM^{op} is a generalized mapping cylinder for $|g|h \sim f|_{\ell}h \sim f$. $\square$

The following proposition plays a crucial role in our simplicial constructions.

Proposition 6.5 *Let Y be a finite simplicial set and let $f_1, \dots, f_m: \text{Sd}^t \Sigma^p \rightarrow Y$ be given simplicial maps. Then there is an algorithm that, given an integer vector $\mathbf{c} =$*

$(c_1, \dots, c_m)$, constructs a generalized mapping cylinder $\Sigma^p \xrightarrow{f} Y$ for the homotopy class

$$[f] = c_1[f_1] + \dots + c_m[f_m] \in \pi_p(Y) = [\Sigma^p, Y]_*,$$

in time polynomial in the (binary) encoding size of $\mathbf{c}$.

The proof will be given in a series of lemmas. Before going into the proof, we will generalize this proposition slightly. A homotopy class of a pointed map $f: \Sigma_1^p \vee \dots \vee \Sigma_s^p \rightarrow Y$ is determined uniquely by its restrictions $\Sigma_k^p \rightarrow Y$. When each restriction is expressed as an integral combination of the f_i , we may use Proposition 6.5 together with Lemma 6.7 to provide a generalized mapping cylinder for f .

Lemma 6.6 *There is an algorithm that, given an integer of the form $c = \pm 2^d$, $d \in \mathbb{N}$, constructs a generalized mapping cylinder $N_c: \Sigma^p \xrightarrow{\sim} \Sigma^p$ for the map of degree c . Moreover, if p is fixed, then $\text{size}(N_c)$ is linear in $d + 1$ and the running time of the algorithm is polynomial in $d + 1$, which is the encoding size of c .*

Proof Consider maps $g_{-1}, g_2: |\Sigma^p| \rightarrow |\Sigma^p|$ of degrees -1 and 2 , respectively, and choose a generalized mapping cylinder N_i for each g_i according to Lemma 6.4. For p constant, these are fixed simplicial sets.

Thus, by Lemma 6.3,

$$N_{2^d} := \underbrace{N_2 \cdots N_2}_{d \text{ factors}}$$

is a generalized mapping cylinder for $(g_2)^d$, a map of degree 2^d . By further composing this with N_{-1} , if necessary, we obtain a generalized mapping cylinder N_c for a map of degree c .

Moreover, if p is fixed, then we can precompute the generalized mapping cylinders N_{-1}, N_2 , which leads to $\text{size}(N_c)$ and running time as requested. $\square$

Lemma 6.7 *Suppose that $X_1, \dots, X_m$ and Y are pointed simplicial sets and that $M_i: X_i \rightsquigarrow Y$ are generalized mapping cylinders for pointed maps $f_i: |X_i| \rightarrow |Y|$, $1 \leq i \leq m$. Then there is an algorithm that constructs a generalized mapping cylinder*

$M: X_1 \vee \cdots \vee X_m \xrightarrow{f} Y$ *for the map $f: |X_1 \vee \cdots \vee X_m| \rightarrow |Y|$ with restrictions $f|_{|X_i|} = f_i$.*

Moreover, if p is fixed, then $\text{size}(N)$ is linear in $\sum_i \text{size}(M_i)$ and the running time is polynomial.

Proof The wedge sum $M' = M_1 \vee \cdots \vee M_m$ is a generalized mapping cylinder

$$M': X_1 \vee \cdots \vee X_m \xrightarrow{f_1 \vee \cdots \vee f_m} Y \vee \cdots \vee Y.$$

We attach to M' the mapping cylinder of the folding map $\nabla: Y \vee \cdots \vee Y \rightarrow Y$ to obtain the required generalized mapping cylinder M . $\square$

Lemma 6.8 *Let $M_1, \dots, M_m: \Sigma^p \rightsquigarrow Y$ be generalized mapping cylinders for $[f_1], \dots, [f_m] \in \pi_p(Y)$. Then there is an algorithm that constructs a generalized mapping cylinder $M: \Sigma^p \rightsquigarrow Y$ for the homotopy class $[f_1] + \cdots + [f_m] \in \pi_p(Y)$, in polynomial time if p is fixed.*

Proof Let us consider the following chain of maps:

$$\Sigma^p \xleftarrow{q} D^p_{/m} / \partial D^p_{/m} \xrightarrow{\delta} \Sigma_1^p \vee \cdots \vee \Sigma_m^p \xrightarrow{f} Y,$$

where f restricts to f_i on the i th summand. The first two maps are simplicial, and thus their mapping cylinders provide generalized mapping cylinders for any homotopy inverse $\bar{q}$ of q and for δ , respectively. A generalized mapping cylinder for f was constructed in Lemma 6.7. Composing these cylinders gives the result, since $f\delta\bar{q} \sim f_1 + \cdots + f_m$; see Sect. 6.1. $\square$

Proof of Proposition 6.5 Let $M_i: \Sigma^p \xrightarrow{f_i} Y$, $i = 1, \dots, m$, be generalized mapping cylinders.

Using the binary expansion of an integer c , Lemma 6.8, and Lemma 6.6, we can construct the generalized mapping cylinder N_c for every map $\Sigma^p \rightarrow \Sigma^p$ of degree c in time polynomial (at most quadratic) in the bit length of c . The composition $M_i N_{c_i}$ is a generalized mapping cylinder for $c_i[f_i]$. Lemma 6.8 then constructs a generalized mapping cylinder for the sum $c_1[f_1] + \cdots + c_m[f_m]$. $\square$

6.3 Proofs of the Main Results

In Sect. 5 we described the relevant spaces as cell complexes. It remains to construct them as finite simplicial complexes and the map $f: A \rightarrow Y$ as a simplicial map.

Proof of Theorem 1.1(a) We will give details only for k even. Using the notation from Sect. 5.1, we triangulate the target sphere $Y = S^k$ in an arbitrary manner and fix simplicial maps

$$w_+, w_-: \text{Sd}^t(\Sigma^{2k-1}) \rightarrow S^k$$

that represent the homotopy classes of the Whitehead square and its negative,

$$[w_{\pm}] = \pm[l, \iota] \in \pi_{2k-1}(S^k)$$

(by the simplicial approximation theorem, a sufficiently fine subdivision $\text{Sd}^t(\Sigma^{2k-1})$ and the required simplicial maps exist, and they can be hardwired into the algorithm). Let now $\mathbf{b}$ be the vector of the right-hand sides of an arbitrary system of the form (Q-SYM) and let $1 \leq q \leq s$. By adding $|b_q|$ times the map $w_{\pm}$, we obtain a simplicial map

$$A'_q := \text{Sd}^t(D^{2k-1}/\partial D^{2k-1}) \xrightarrow{f'_q} S^k$$

that represents $b_q[l, \iota]$.²⁹ Finally, we take $A' = A'_1 \vee \cdots \vee A'_s$ and specify $f': A' \rightarrow S^k$ by its restrictions to the A'_q , namely, the maps f'_q .

We recall that the space X is constructed as the mapping cylinder of a map $g: A \rightarrow W$ that was expressed in terms of the Whitehead products $[v_i, v_j]$ and the coefficients of the system (Q-SYM). In the simplicial setup it will be more convenient to use generalized mapping cylinders for this purpose. As explained during the discussion of the generalized extension problem in the beginning of Sect. 5, the extension problems are equivalent. Using a fixed representative $w'_{\pm}: \text{Sd}^t(\Sigma^{2k-1}) \rightarrow \Sigma^k \vee \Sigma^k$, we may construct the generalized mapping cylinder $X': A' \rightsquigarrow W$, with the inclusion denoted by $i': A' \rightarrow X'$, of the composition

$$\underbrace{A'_1 \vee \cdots \vee A'_s}_{A'} \xrightarrow{q \vee \cdots \vee q} \underbrace{\Sigma_1^{2k-1} \vee \cdots \vee \Sigma_s^{2k-1}}_A \xrightarrow{g} \underbrace{\Sigma_1^k \vee \cdots \vee \Sigma_r^k}_W.$$

²⁹When $b_q = 0$, we take $A'_q = \text{Sd}^t(\Sigma^{2k-1})$ and $f'_q: \text{Sd}^t(\Sigma^{2k-1}) \rightarrow S^k$ the constant map onto the base-point.

Thus, we have constructed an extension problem, given by i' and f' , and by Proposition 5.1, its solvability is equivalent to the solvability of the system (Q-SYM) that we started with.

Finally, we replace the simplicial sets A' and X' by the simplicial complexes $B_*(\text{Sd}(A'))$ and $B_*(\text{Sd}(X'))$ (see Proposition 3.5). The map f' is replaced by the composition $f'\gamma_{A'}$ in the diagram

$$\begin{array}{ccccc} B_*(\text{Sd}(A')) & \xrightarrow{\gamma_{A'}} & A' & \xrightarrow{f'} & S^k \\ \downarrow \cap & & \downarrow \cap & & \\ B_*(\text{Sd}(X')) & \xrightarrow{\gamma_{X'}} & X' & & \end{array}$$

where the maps $\gamma_{A'}$ and $\gamma_{X'}$ were also defined in Proposition 3.5.

Since both $\gamma_{X'}$ and $\gamma_{A'}$ are homotopy equivalences, the extendability of $|f'|$ is equivalent to that of $|f'\gamma_{A'}|$ by Corollary 3.1. $\square$

Proof of Theorem 1.1(b) Again, we work out the case k even. Let A, X, Y, f , and g be as in Proposition 5.1 and fix an arbitrary pair of simplicial sets (X', A') whose geometric realization is homotopy equivalent to (X, A) . Using generalized mapping cylinders for this purpose, we may assume that $A' = \Sigma^{2k-1}$. We fix some simplicial maps

$$w'_\pm: \text{Sd}^t(\Sigma^{2k-1}) \rightarrow \Sigma^k \vee \Sigma^k$$

representing the Whitehead product and its negative. According to Proposition 6.2, the cell complex Y of (14) is homotopy equivalent to the quotient M/S , where M

is an arbitrary generalized mapping cylinder $M: S \xrightarrow{\varphi} T$ for the map $\varphi: |S| \rightarrow |T|$ between the geometric realizations of the simplicial sets

$$S = \bigvee_{i < j} \Sigma_{ij}^{2k-1} \vee \bigvee_i \Sigma_{ii}^{2k-1}, \quad T = \bigvee_i \Sigma_i^k \vee \bigvee_q \Sigma_q^{2k-1},$$

whose restrictions to the spheres of S are the attaching maps φ_{ij} and φ_{ii} for the cells of Y ; see Sect. 5. The generalized mapping cylinder M is constructed by Proposition 6.5.

Since the image of $f: A \rightarrow Y$ lies in T , the replacement of Y by M/S results in replacing f by the composition

$$\tilde{f}: \Sigma^{2k-1} \xrightarrow{f} T \xrightarrow{i_T} M \xrightarrow{\text{proj}} M/S.$$

(the homotopy equivalence $Y \simeq M/S$ restricts to $\text{proj} i_T$ on T by Proposition 6.2).

It remains to replace $\tilde{f}$ by a simplicial map. But since f is a combination of the Whitehead products and the remaining maps i_T and proj are simplicial, we may achieve this by replacing M/S further by the generalized mapping cylinder

$Y': \Sigma^{2k-1} \xrightarrow{\tilde{f}} M/S$ as in Sect. 6.2. We denote the inclusion $\Sigma^{2k-1} \rightarrow Y'$ by f' . From the definition of the generalized mapping cylinder, $i_{M/S}\tilde{f} \sim f'$ and $i_{M/S}$ is a homotopy equivalence, and therefore the extendability of $\tilde{f}$ is equivalent to that of f' . This finishes the construction of a simplicial replacement of the extension problem.

To make everything into simplicial complexes, we apply $B_*\text{Sd}$ to all the involved simplicial sets $A' = \Sigma^{2k-1}$, X' , Y' and the simplicial map f' . $\square$

Proof of Theorem 1.2 Let us fix some simplicial representatives

$$h: \text{Sd}^t(\Sigma^3) \rightarrow \Sigma^2, \quad w: \text{Sd}^t(\Sigma^3) \rightarrow \Sigma^2 \vee \Sigma^2, \quad m_2, m_{-1}: \text{Sd}^t(\Sigma^2) \rightarrow \Sigma^2$$

for the Hopf map η , the Whitehead product $[\iota_1, \iota_2]$, and the maps of degree 2 and -1 , respectively. We may then build the generalized mapping cylinder

$$M: \Sigma_1^3 \vee \dots \vee \Sigma_s^3 \xrightarrow{\varphi_a} \Sigma_1^2 \vee \dots \vee \Sigma_r^2,$$

for the map φ_a whose restriction to the k th summand Σ_k^3 is given by

$$\varphi_q = \sum_{1 \leq i \leq r} a_{ii}^{(q)} \iota_i \eta + \sum_{1 \leq i < j \leq r} a_{ij}^{(q)} [\iota_i, \iota_j].$$

We construct Anick's simplicial complex Y_a^4 as

$$Y_a^4 = B_*\text{Sd}(M/(\Sigma_1^3 \vee \dots \vee \Sigma_s^3)).$$

By Proposition 6.2, it is homotopy equivalent to the cellular complex obtained from the wedge $\Sigma_1^2 \vee \dots \vee \Sigma_r^2$ by attaching 4-cells along the maps with homotopy classes φ_q .

To get the statement of Theorem 1.2, it is now sufficient to realize that the algorithmic construction above can be carried out in the time polynomial in the binary encoding of the vector $a_{ij}^{(q)}$, $1 \leq i \leq j \leq r$, $1 \leq q \leq s$. $\square$

Acknowledgements We are grateful to Maurice Rojas for providing us useful information on Diophantine equations.

This research was supported by the ERC Advanced Grant No. 267165. The research of M.Č. was supported by the project CZ.1.07/2.3.00/20.0003 of the Operational Programme Education for Competitiveness of the Ministry of Education, Youth and Sports of the Czech Republic. The research by M.K. and J.M. was supported by the Center of Excellence, Inst. for Theor. Comp. Sci., Prague (project P202/12/G061 of GA ČR). The research of L.V. was supported by the Center of Excellence, Eduard Čech Institute (project P201/12/G028 of GA ČR). The research by U.W. was supported by the Swiss National Science Foundation (grants SNSF-200020-138230 and SNSF-PP00P2-138948).

Appendix A: Subdivisions of Simplicial Sets

Here we outline a proof of Proposition 3.5.

Since a simplicial set also encodes an ordering of the vertices within each simplex, there is another “barycentric subdivision” $\text{Sd}_*(X)$ associated with any simplicial set X , obtained by reversing the order of the vertices in every simplex of $\text{Sd}(X)$ (in $\text{Sd}(X)$, the inclusion chains of simplices are ordered according to ascending dimension, and in $\text{Sd}_*(X)$ according to descending dimension). Thus, for example, $\text{Sd}(\Delta^1)$ can be described pictorially as $\bullet \longrightarrow \bullet \longleftarrow \bullet$ while $\text{Sd}_*(\Delta^1)$ is $\bullet \longleftarrow \bullet \longrightarrow \bullet$. The barycentric subdivision $\text{Sd}_*(X)$ is related to the original simplicial set X via an initial vertex map $\text{Sd}_*(X) \rightarrow X$, which is a homotopy equivalence.

Moreover, there is a universal way of associating a simplicial complex with any simplicial set X : It has the same vertex set as X and a collection of vertices spans a (unique) simplex if and only if there *exists* a simplex in the original simplicial set X with this vertex set. An alternative, equivalent definition of B_*X is that it is the simplicial complex associated in this way with $\text{Sd}_*(X)$.

Proof of Proposition 3.5 The face operators ∂_i can be iterated to obtain more general face operators. Since each ∂_i leaves out the i -th vertex of a simplex, by iterating we obtain face operators that leave out a set of vertices. When this set is $I \subseteq \{0, \dots, n\}$, we write the corresponding operator as ∂_I . It is easy to observe that we can express ∂_I as

$$\partial_I = \partial_{i_1} \cdots \partial_{i_k},$$

where $i_1 < \dots < i_k$ is the ordering of the elements of $I = \{i_1, \dots, i_k\}$. We call k the *codimension* of ∂_I . Similarly, we can iterate the s_i and obtain general degeneracy operators s_I .

Since we are interested in the computational side of the story, we will describe the simplicial complex $B_*(\text{Sd}(X))$ explicitly. Its vertices are the chains

$$\sigma = \sigma_0 \xrightarrow{f_1} \sigma_1 \xrightarrow{f_2} \cdots \xrightarrow{f_k} \sigma_k,$$

where $\sigma_0, \dots, \sigma_k$ are simplices of X with σ_k nondegenerate, and each f_i is a face operator of codimension at least 1, for which $\sigma_{i-1} = f_i \sigma_i$.

We say that a chain

$$\tau = \tau_0 \xrightarrow{g_1} \tau_1 \xrightarrow{g_2} \cdots \xrightarrow{g_\ell} \tau_\ell$$

(still with τ_ℓ nondegenerate and all face operators g_i of codimension at least 1) is a subchain of σ , which we write as $\sigma > \tau$, if there exists an injective monotone map (a subsequence) $\varphi: \{0, \dots, \ell\} \rightarrow \{0, \dots, k\}$ with $\ell < k$ and a commutative diagram

$$\begin{array}{ccccccc} \sigma_{\varphi(0)} & \xrightarrow{\quad} & \sigma_{\varphi(1)} & \xrightarrow{\quad} & \cdots & \xrightarrow{\quad} & \sigma_{\varphi(\ell)} \\ \downarrow & & \downarrow & & & & \downarrow \\ \tau_0 & \xrightarrow{g_1} & \tau_1 & \xrightarrow{g_2} & \cdots & \xrightarrow{g_\ell} & \tau_\ell \end{array}$$

(the top maps are the appropriate compositions of the f_i), where every arrow $\sigma_{\varphi(i)} \twoheadrightarrow \tau_i$ is an iterated degeneracy operator p_i for which $p_i \tau_i = \sigma_{\varphi(i)}$. The commutativity means that the respective compositions of operators are equal. For each φ , there exists at most one subchain τ , but for a given τ , the choice of φ is not unique. The composition $\sigma_0 \rightarrow \sigma_{\varphi(0)} \rightarrow \tau_0$ gives a canonical operator $\sigma_0 \rightarrow \tau_0$. It is not too hard to show³⁰ that this operator depends only on σ and τ and not on the choice of φ .

The n -simplices of $B_*(\text{Sd}(X))$ are then formed by the subsets $\{\sigma^0, \dots, \sigma^n\}$ for all decreasing sequences $\sigma^0 > \dots > \sigma^n$ of chains; we order the vertices in each simplex according to the subchain relation.

The simplicial map $\gamma: B_*(\text{Sd}(X)) \rightarrow X$ is defined on vertices by sending σ to the last vertex of σ_0 . For a simplex specified by $\sigma^0 > \dots > \sigma^n$, we have a canonical chain

$$(\sigma^0)_0 \longrightarrow (\sigma^1)_0 \longrightarrow \dots \longrightarrow (\sigma^n)_0$$

of operators and we use these to map the last vertices $\text{lastv}((\sigma^i)_0)$ of the faces $(\sigma^i)_0$ to $(\sigma^n)_0$. In this way, we obtain an (ordered) collection of vertices of $(\sigma^n)_0$. The value of γ on the sequence $\sigma^0 > \dots > \sigma^n$ is then the simplex of $(\sigma^n)_0$ spanned by these vertices (it might be degenerate, e.g., when some of the $(\sigma^{i-1})_0 \rightarrow (\sigma^i)_0$ preserve the last vertex).

According to [7, Proposition 4.6.3] and [11, Corollary 4.3], the horizontal map and the vertical map in the triangle

$$\begin{array}{ccc} \text{Sd}_*(\text{Sd}(X)) & \xrightarrow{\pi} & B_*(\text{Sd}(X)) \\ \downarrow & \swarrow \gamma & \\ X & & \end{array}$$

are homotopy equivalences (the vertical map is the composition of the initial vertex map with the last vertex map). Since the diagram commutes,³¹ the map γ must be a homotopy equivalence, too. $\square$

Appendix B: Extending Maps into $(k - 1)$ -Reduced Simplicial Sets

Here we prove the claim made after Theorem 1.1 regarding the construction of the target space Y as a $(k - 1)$ -reduced simplicial set. It is usual in effective algebraic

³⁰The important ingredients are that every simplex can be expressed uniquely as a degeneracy of a nondegenerate simplex, and that every operator can be written uniquely as a degeneracy of a face.

³¹This is not too hard to show, but we do not want to dwell into the exact definition of $\text{Sd}_*(\text{Sd}(X))$. The main point is that the preimages under π of the simplex $\sigma^0 > \dots > \sigma^n$ are given by the choices of the subsequences φ . The commutativity is then implied by the independence of the operators $(\sigma^{i-1})_0 \rightarrow (\sigma^i)_0$ on these choices.

topology that certain computations with simplicial sets only work when at least some of the inputs are 0-reduced or 1-reduced. A typical example is the computation of the homology groups of a loop space ΩX of a simplicial set X . First, we remark that it is impossible to compute these homology groups for general X , as otherwise we would obtain an algorithmic computation of $\pi_1(X)$, which is known to be impossible by [19]. On the positive side, there is a very old method for the computation of these homology groups which, however, works only for 1-reduced X (using the so-called cobar construction and homological perturbation theory; see [27, Corollary 147] or [1]).³²

It is thus natural to ask if the undecidability of the extension problem of Theorem 1.1 might only be caused by Y being $(k - 1)$ -connected but not $(k - 1)$ -reduced. In this section, we will prove a version of Theorem 1.1 with $(k - 1)$ -reduced Y . The simplicial set Y' appearing in the proof of part (a) of Theorem 1.1 might be chosen to be either Σ^k or $\Sigma^k \vee \Sigma^k$ (depending on the dimension), both of which are $(k - 1)$ -reduced and no further work is needed.

To finish the proof of part (b), we need to replace Y' by some $(k - 1)$ -reduced simplicial set without changing the extendability. To this end, we introduce a very useful notion of an n -equivalence. Let $0 \leq n \leq \infty$. A continuous map $Y \rightarrow Z$ is said to be an n -equivalence if it induces an isomorphism on all homotopy groups up to dimension $n - 1$ and a surjection in dimension n (the ∞ -equivalences are usually called weak homotopy equivalences). We will need the following basic property of n -equivalences.

Proposition B.1 *Let X be a cell complex of dimension n , $f: A \rightarrow Y$ a continuous map defined on a subcomplex $A \subseteq X$, and $h: Y \rightarrow Z$ an n -equivalence. Then f is extendable to X if and only if hf is extendable to X .*

Proof If f has an extension g , then hf extends to hg . The other direction is [32, Sect. 7.6, Theorem 22]. $\square$

With the previous proposition in mind, we construct a replacement of Y by brute force, i.e., by going through all “candidate replacements” and checking if they give equivalent extension problems. In detail, we fix (X', A') as in Sect. 6.3 and make a list of all pairs (Z', h') , where Z' is a finite $(k - 1)$ -reduced simplicial set and $h': Y' \rightarrow Z'$ a simplicial map. In each step we test whether h' is a $(2k)$ -equivalence. If that is the case, then the problem of extending the composition $h'f': A' \rightarrow Z'$ to X' is equivalent to that of $f': A' \rightarrow Y'$, which we proved to be undecidable.

By the Hurewicz theorem, h' is a $(2k)$ -equivalence if and only if $\text{Cone}(h')$ has zero homology groups up to dimension $2k$. This can be tested easily using a Smith normal form algorithm. The pair (Z', h') with the above properties exists by the following theorem, finishing the proof of Theorem 1.1 with $(k - 1)$ -reduced target.

³²On the other hand, it is possible to compute these homology groups for any 1-connected X . This is another application of the “brute force” version of Theorem B.2 that is explained just prior to its statement.

Theorem B.2 *Let $n \geq k \geq 2$ and let Y be a $(k-1)$ -connected simplicial set whose homology groups $H_i(Y)$, $i \leq n$, are finitely generated. Then there exist a finite $(k-1)$ -reduced simplicial set Z and an n -equivalence $\psi: Y \rightarrow Z$.*

We believe that the theorem should be known in some form, but we were not able to find it in the literature. For its proof, we will need a couple of more advanced notions. Accordingly, the proof will be more sketchy.

Let Δ^n be the standard n -simplex (regarded as a simplicial set). For $0 \leq k \leq n$, the k th n -horn is the simplicial subset $\Lambda_k^n \subseteq \Delta^n$ spanned by all the proper faces of Δ^n with the exception of the k -th face.

A simplicial set Z is said to be a *Kan complex* if every simplicial map $f: \Lambda_k^n \rightarrow Z$ can be extended to a simplicial map $\Delta^n \rightarrow Z$. The map f is called a *horn* in Z and we say that this horn *can be filled* if the extension exists.

A usual method for constructing Kan complexes is the *successive filling of horns*, which is described, e.g., in [7, Proofs of Propositions 4.5.5 and 4.5.6].

Given a simplicial set Y and a horn $f: \Lambda_k^n \rightarrow Y$ in Y , we can form a larger simplicial set $Y \cup_f \Delta^n$ by attaching Δ^n to Y along f ; this larger simplicial set (continuously) deformation retracts to Y , and, by construction, the horn $f: \Lambda_k^n \rightarrow Y$ can be filled in the larger simplicial set. We refer to this operation as a *single horn filling*.

If we simultaneously attach fillings for all unfillable horns in Y , we obtain a simplicial set $\mathcal{K}(Y)$ that contains Y and such that all horns in Y can be filled in $\mathcal{K}(Y)$. Iterating this procedure,³³ we obtain a sequence $Y \subseteq \mathcal{K}(Y) \subseteq \mathcal{K}^2(Y) \subseteq \dots$, where, by construction, every horn in $\mathcal{K}^n(Y)$, can be filled in $\mathcal{K}^{n+1}(Y)$. Let $\mathcal{K}^\infty(Y)$ be the union of the simplicial sets $\mathcal{K}^n(Y)$, $n \in \mathbb{N}$.

Proposition B.3

- The simplicial set $\mathcal{K}^\infty(Y)$ is a Kan complex.
- The inclusion $Y \rightarrow \mathcal{K}^\infty(Y)$ is an ∞ -equivalence.
- If Y is $(d-1)$ -reduced, then so is $\mathcal{K}(Y)$ and consequently also $\mathcal{K}^\infty(Y)$.
- Let $L \subseteq \mathcal{K}^\infty(Y)$ be a finite simplicial subset. Then L lies in a simplicial subset obtained from Y by a finite sequence of single horn fillings.

Proof Any horn in $\mathcal{K}^\infty(Y)$ lies in some $\mathcal{K}^n(Y)$ and is thus fillable in $\mathcal{K}^{n+1}(Y) \subseteq \mathcal{K}^\infty(Y)$. The proof of the second point is similar, using a deformation retraction of $\mathcal{K}^n(Y)$ onto Y . The third point is clear. For the last point, L lies in some $\mathcal{K}^n(Y)$. By induction, $L \cap \mathcal{K}^{n-1}(Y)$ uses only a finite number of single horn fillings. A finite number of single horn fillings is required to cover L . $\square$

Thus, there exists a homotopy equivalence $Y \rightarrow \mathcal{K}^\infty(Y)$ of any simplicial set with a Kan complex. For the proof of Theorem B.2, we will need every Kan complex to contain a minimal Kan complex as a deformation retract [16, Theorem 9.5] and that a minimal $(k-1)$ -connected Kan complex is in fact $(k-1)$ -reduced. Thus, composing the inclusion $Y \rightarrow \mathcal{K}^\infty(Y)$ with the deformation retraction yields an ∞ -equivalence $\iota: Y \rightarrow Z$ of Y with a $(k-1)$ -reduced Kan complex Z .

³³Formally, we define $\mathcal{K}^0(Y) := Y$ and $\mathcal{K}^n(Y) := \mathcal{K}(\mathcal{K}^{n-1}(Y))$ for $n \geq 1$.

Proof of Theorem B.2 Let $Z^0 \subseteq Z$ denote the image of ι and $\iota_0: Y \rightarrow Z^0$ the restriction of ι . By the following proposition, there exists a simplicial set Z' , containing Z^0 as a subset, and a simplicial $(n+1)$ -equivalence $\psi: Z' \rightarrow Z$. Since this map is the identity on Z^0 , there is a canonical map $h': Y \rightarrow Z'$ making the diagram

$$\begin{array}{ccc}
 & & Z' \\
 & \nearrow h & \downarrow \psi \\
 Y' & \xrightarrow{\iota_0} & Z^0 \\
 & \searrow \iota & \downarrow \psi \\
 & & Z
 \end{array}$$

commutative (namely the composition of $\iota_0: Y \rightarrow Z^0$ with the inclusion $Z^0 \hookrightarrow Z'$). Since ι is an ∞ -equivalence and ψ induces an isomorphism on homotopy groups up to dimension n , the same is true for h' and in particular, it is an n -equivalence. $\square$

Proposition B.4 *Let $n \geq k \geq 2$ and let Z be a $(k-1)$ -connected Kan complex whose homology groups $H_i(Z)$, $i \leq n$, are finitely generated. Then there exist a finite $(k-1)$ -reduced simplicial set Z' and an n -equivalence $\psi: Z' \rightarrow Z$.*

If Z^0 is an arbitrary finite $(k-1)$ -reduced subset of Z , then Z' can be chosen to contain Z^0 as a subset ψ to be the identity on Z^0 .

To prove Proposition B.4, we follow the argument in [9, Proposition 4C.1] but make the attachment maps simplicial by using the idea of filling horns described above.

Proof of Proposition B.4 We proceed by induction on n . For $n \leq k-1$, we can take $Z' = Z^0$.

Assume that we have constructed a finite simplicial set Z^{n-1} and a map $\psi^{n-1}: Z^{n-1} \rightarrow Z$ that is an $(n-1)$ -equivalence.

Let $\hat{Z}$ be the simplicial set obtained from the simplicial mapping cylinder $\text{Cyl}(\psi^{n-1})$ of $\psi^{n-1}: Z^{n-1} \rightarrow Z$ by collapsing $Z^0 \times \Delta^1$ onto the base Z^0 . Since ψ^{n-1} is the identity on Z^0 , the usual deformation retraction of $\text{Cyl}(\psi^{n-1})$ onto Z induces a deformation retraction of $\hat{Z}$ onto Z . We enlarge the pair $(\hat{Z}, Z^{n-1})$ to a Kan pair (K, L) by filling horns. Since Z^{n-1} is $(k-1)$ -reduced, so is L .

By the assumption on ψ^{n-1} and by the Hurewicz theorem, we have $H_i(K, L) = 0$ for $i \leq n-1$. Consider the exact sequence of homology groups for the pair (K, L) :

$$\cdots \rightarrow H_n(K) \rightarrow H_n(K, L) \rightarrow H_{n-1}(L) \rightarrow H_{n-1}(K) \rightarrow H_{n-1}(K, L) = 0.$$

Pick generators γ_j of $H_n(K, L)$. Since $H_n(K) \cong H_n(Z)$ and $H_{n-1}(L) \cong H_{n-1}(Z^{n-1})$ are finitely generated, a finite number of generators suffices. By the relative Hurewicz theorem, $H_n(K, L) \cong \pi_n(K, L, *)$ (simplicial homotopy groups, since we are working with a Kan pair). Thus, we can choose n -simplices g_j of K representing the γ_j , whose faces lie in L .

Let K' be the simplicial subset of K spanned by L and the simplices g_j . Then the natural homomorphism $H_n(K', L) \rightarrow H_n(K, L)$ is surjective by the choice of the simplices g_j . Since these relative homology groups are zero in lower dimensions, it follows from the long exact sequence of the triple (K, K', L) that $H_i(K, K') = 0$ for $i \leq n$. In effect, the inclusion $K' \rightarrow K$ is an n -equivalence (by the relative Hurewicz theorem again). Composing with an arbitrary deformation retraction of K onto Z we obtain an n -equivalence $K' \rightarrow Z$ satisfying all the required properties except that K' is infinite.

Thus, it remains to replace K' by a finite simplicial set. Since there are only finitely many simplices g_j , there is a simplicial set L' with $Z^{n-1} \subseteq L' \subseteq L$, obtained from Z^{n-1} by filling finitely many horns, and such that the boundaries of all the simplices g_j lie in L' . We take Z^n to be the finite simplicial set spanned by L' and the simplices g_j . Since $|L'|$ is a deformation retract of $|L|$, we get that $|Z^n|$ is a deformation retract of $|K'|$. Thus, Z^n has all the required properties. $\square$

References

1. Adams, J.F.: On the cobar construction. *Proc. Natl. Acad. Sci. USA* **42**(7), 409 (1956)
2. Anick, D.J.: The computation of rational homotopy groups is $\#P$ -hard. In: *Computers in Geometry and Topology*, Proc. Conf., Chicago, IL, 1986. *Lect. Notes Pure Appl. Math.*, vol. 114, pp. 1–56 (1989)
3. Bredon, G.: *Topology and Geometry*. Graduate Texts in Mathematics, vol. 139. Springer, Berlin (1993)
4. Brown, E.H. Jr.: Finite computability of Postnikov complexes. *Ann. Math.* (2) **65**, 1–20 (1957)
5. Čadek, M., Krčál, M., Matoušek, J., Sergeraert, F., Vokřínek, L., Wagner, U.: Computing all maps into a sphere. [arXiv:1105.6257](https://arxiv.org/abs/1105.6257) (2011). Extended abstract in *Proc. ACM–SIAM Symposium on Discrete Algorithms* (SODA 2012)
6. Čadek, M., Krčál, M., Matoušek, J., Vokřínek, L., Wagner, U.: Polynomial-time computation of homotopy groups and Postnikov systems in fixed dimension. [arXiv:1211.3093](https://arxiv.org/abs/1211.3093) (2012)
7. Fritsch, R., Piccinini, R.A.: *Cellular Structures in Topology*. Cambridge Studies in Advanced Mathematics, vol. 19. Cambridge University Press, Cambridge (1990)
8. Friedman, G.: An elementary illustrated introduction to simplicial sets. *Rocky Mt. J. Math.* **42**(2), 353–423 (2012)
9. Hatcher, A.: *Algebraic Topology*. Cambridge University Press, Cambridge (2001). Electronic version available at <http://math.cornell.edu/hatcher#AT1>
10. Hilton, J.P.: On the homotopy groups of union of spheres. *J. Lond. Math. Soc.* **3**, 154–172 (1955)
11. Jardine, J.F.: Simplicial approximation. *Theory Appl. Categ.* **12**(2), 34–72 (2004)
12. Krčál, M., Matoušek, J., Sergeraert, F.: Polynomial-time homology for simplicial Eilenberg–MacLane spaces. [arXiv:1201.6222](https://arxiv.org/abs/1201.6222) (2011). *J. Found. Comput. Math.* (to appear)
13. Kochman, S.O.: *Stable Homotopy Groups of Spheres. A Computer-Assisted Approach*. Lecture Notes in Mathematics, vol. 1423. Springer, Berlin (1990)
14. Matiyasevich, Yu.: The Diophantine-ness of enumerable sets. *Dokl. Akad. Nauk SSSR* **191**, 279–282 (1970)
15. Matiyasevich, Y.V.: *Hilbert’s Tenth Problem*. Foundations of Computing Series. MIT Press, Cambridge (1993). Translated from the 1993 Russian original by the author, with a foreword by Martin Davis
16. May, J.P.: *Simplicial Objects in Algebraic Topology*. Chicago Lectures in Mathematics. University of Chicago Press, Chicago (1992). Reprint of the 1967 original; the page numbers do not quite agree with the 1967 edition
17. Mazur, B.: Questions of decidability and undecidability in number theory. *J. Symb. Log.* **59**(2), 353–371 (1994)
18. Munkres, J.R.: *Elements of Algebraic Topology*. Addison-Wesley, Reading (1984)

19. Novikov, P.S.: Ob algoritmičeskoj nerazrešimosti problemy toždestva slov v teorii grupp (On the algorithmic unsolvability of the word problem in group theory). *Tr. Mat. Inst. Steklova* **44**, 1–143 (1955)
20. Nabutovsky, A., Weinberger, S.: Algorithmic unsolvability of the triviality problem for multidimensional knots. *Comment. Math. Helv.* **71**(3), 426–434 (1996)
21. Nabutovsky, A., Weinberger, S.: Algorithmic aspects of homeomorphism problems. In: *Tel Aviv Topology Conference: Rothenberg Festschrift* (1998). *Contemp. Math.*, vol. 231, pp. 245–250. Amer. Math. Soc., Providence (1999)
22. Ravenel, D.C.: *Complex Cobordism and Stable Homotopy Groups of Spheres*, 2nd edn. Am. Math. Soc., Providence (2004)
23. Real, P.: An algorithm computing homotopy groups. *Math. Comput. Simul.* **42**, 461–465 (1996)
24. Romero, A., Rubio, J., Sergeraert, F.: Computing spectral sequences. *J. Symb. Comput.* **41**(10), 1059–1079 (2006)
25. Rubio, J., Sergeraert, F.: Constructive algebraic topology. *Bull. Sci. Math.* **126**(5), 389–412 (2002)
26. Rubio, J., Sergeraert, F.: Algebraic models for homotopy types. *Homol. Homotopy Appl.* **17**, 139–160 (2005)
27. Rubio, J., Sergeraert, F.: Constructive homological algebra and applications. [arXiv:1208.3816](https://arxiv.org/abs/1208.3816) (2012). Written in 2006 for a MAP Summer School at the University of Genova
28. Schön, R.: Effective algebraic topology. *Mem. Am. Math. Soc.* **451**, 63 (1991)
29. Sergeraert, F.: The computability problem in algebraic topology. *Adv. Math.* **104**(1), 1–29 (1994)
30. Smith, J.R.: m-Structures determine integral homotopy type. [arXiv:math/9809151v1](https://arxiv.org/abs/math/9809151v1) (1998)
31. Soare, R.I.: Computability theory and differential geometry. *Bull. Symb. Log.* **10**(4), 457–486 (2004)
32. Spanier, E.H.: *Algebraic Topology*. McGraw-Hill, New York (1966)
33. Steenrod, N.E.: Cohomology operations, and obstructions to extending continuous functions. *Adv. Math.* **8**, 371–416 (1972)
34. Whitehead, G.W.: *Elements of Homotopy Theory*. Graduate Texts in Mathematics, vol. 61. Springer, New York (1978)