Extending Continuous Maps: Polynomiality and Undecidability*

Martin Čadek
Department of Mathematics and Statistics
Masaryk University
Brno, Czech Republic
cadek@math.muni.cz

Marek Krčál
Department of Applied Mathematics
Charles University
Praha, Czech Republic
krcal@kam.mff.cuni.cz

Jiří Matoušek
Charles University in Prague and ETH Zurich
matousek@kam.mff.cuni.cz

Lukáš Vokřínek
Department of Mathematics and Statistics
Masaryk University
Brno, Czech Republic
koren@math.muni.cz

Uli Wagner
Institute of Science and Technology
Klosterneuburg, Austria
uli@ist.ac.at

ABSTRACT

We consider several basic problems of algebraic topology, with connections to combinatorial and geometric questions, from the point of view of computational complexity.

The *extension problem* asks, given topological spaces X, Y, a subspace $A \subseteq X$, and a (continuous) map $f: A \to Y$, whether f can be extended to a map $X \to Y$. For computational purposes, we assume that X and Y are represented as finite simplicial complexes, A is a subcomplex of X, and f is given as a simplicial map. In this generality the problem is undecidable, as follows from Novikov's result from the 1950s on uncomputability of the fundamental group $\pi_1(Y)$. We thus study the problem under the assumption that, for some $k \geq 2$, Y is $(k-1)$-connected; informally, this means that Y has "no holes up to dimension $k-1$" (a basic example of such a Y is the sphere S^k).

We prove that, on the one hand, this problem is still undecidable for $\dim X = 2k$. On the other hand, for every fixed $k \geq 2$, we obtain an algorithm that solves the extension problem in polynomial time assuming Y $(k-1)$-connected and $\dim X \leq 2k-1$. For $\dim X \leq 2k-2$, the algorithm also provides a classification of all extensions up to homotopy (continuous deformation). This relies on results of our SODA 2012 paper, and the main new ingredient is a machinery of *objects with polynomial-time homology*, which is a polynomial-time analog of objects with effective homology developed earlier by Sergeraert et al.

We also consider the computation of the higher homotopy groups $\pi_k(Y)$, $k \geq 2$, for a 1-connected Y. Their computability was established by Brown in 1957; we show that $\pi_k(Y)$ can be computed in polynomial time for every fixed $k \geq 2$. On the other hand, Anick proved in 1989 that computing $\pi_k(Y)$ is $\#P$-hard if k is a part of input, where Y is a cell complex with certain rather compact encoding. We strengthen his result to $\#P$-hardness for Y given as a simplicial complex.

*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. 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).

STOC'13, June 1-4, 2013, Palo Alto, California, USA.

Categories and Subject Descriptors

F.2.2 [**ANALYSIS OF ALGORITHMS AND PROBLEM COMPLEXITY**]: Nonnumerical Algorithms and Problems—*Geometrical problems and computations*

General Terms

Algorithms, Theory

Keywords

algebraic topology; homotopy theory; homotopy groups; extendability

1. INTRODUCTION

This extended abstract summarizes the results of the three papers [12, 4, 5]. Together these comprise well over 100 pages, and they deal with numerous moderately advanced topological concepts, as well as many algorithmic notions and results. Here we provide an overview of the background, results, and some of the main concepts and ideas. It is aimed at computer scientists with interest in topological questions.

1.1 The topological questions

We are going to talk about topological spaces or simply *spaces*. This is a very general notion but for algorithmic questions we are concerned only with spaces that can be represented as subspaces of finite-dimensional Euclidean spaces $\mathbb{R}^n$. Moreover, the considered spaces are locally nice and have a combinatorial description: each can be given as a *finite simplicial complex*; informally, this means that the space can be built from finitely many simplices of various dimensions by gluing some of them face-to-face. (Actually, we are using simplicial complexes mainly for public relations purposes, since we assume that most readers have heard about them. Our algorithms really work with the more flexible and more sophisticated notion of *simplicial sets*, to be introduced later, and some of the spaces encountered in the computations will even be infinite-dimensional.)

All maps up to homotopy. One of the central themes in algebraic topology is understanding the structure of all *continuous* maps $X \to Y$, for given spaces X and Y (all maps between topological spaces in this paper are assumed to be continuous). The set of all such maps is usually uncountable, but in algebraic topology, the maps are divided into equivalence classes according to *homotopy*: two maps $f, g\colon X \to Y$ are *homotopic* if one can be continuously deformed into the other.[1] The set of all homotopy classes of maps $X \to Y$ is denoted by $[X, Y]$, and in many cases of interest it exhibits interesting structure.

The simplest nontrivial example is $[S^1, S^1]$, self-maps of the unit circle. Up to homotopy, a map $S^1 \to S^1$ is uniquely described by its *winding number*, which is an integer (positive, negative, or zero) counting how many times the image goes around the target S^1. This is generalized by a famous result of Hopf from the 1930s, asserting that the homotopy class of a map $f\colon S^n \to S^n$, between two spheres of the same dimension, is in one-to-one correspondence with an integer parameter, the *degree* of f, which counts how many times the image "wraps around" the target. Another great discovery of Hopf, with ramifications in modern physics and elsewhere, was a map $\eta\colon S^3 \to S^2$, now called by his name, that is not homotopic to a constant map.

Homotopy groups. These are early results in the theory of *higher homotopy groups*, which belong among the most important invariants of a space. We recall that the kth homotopy group $\pi_k(Y)$ of a space Y is defined as the set of all homotopy classes of *pointed* maps $f\colon S^k \to Y$, i.e., maps f that send a distinguished point $s_0 \in S^k$ to a distinguished point $y_0 \in Y$ (and the homotopies F also satisfy $F(s_0, t) = y_0$ for all $t \in [0, 1]$).[2] The group operation in $\pi_k(Y)$ will be defined at the end of Section 2 below. The *fundamental group* $\pi_1(Y)$ need not be commutative, while for $k \geq 2$, the $\pi_k(Y)$ are commutative (Abelian).

One of the important challenges propelling the research in algebraic topology has been the computation of the *homotopy groups of spheres* $\pi_k(S^n)$. Many ingenious insights and amazing methods have been developed for this purpose, see, e.g., [19, 11], but still the $\pi_k(S^n)$ remain among the most puzzling objects of mathematics.[3]

The extension problem. This is another very basic topological question: given spaces X, Y, a subspace $A \subseteq X$, and a map $f\colon A \to Y$, can f be extended to a map $X \to Y$?

For example, the famous *Brouwer fixed-point theorem* can be re-stated as non-extendability of the identity map $S^n \to S^n$ to the ball D^{n+1}. A number of topological concepts and results, some of which may look quite advanced and esoteric to a newcomer, were motivated by attempts at a stepwise solution of the extension problem. In particular, *cohomology theory* was born in the study of $[X^k, S^k]$, where X^k is a k-dimensional space, investigating $[X^{k+1}, S^k]$ led to the discovery of *Steenrod squares*, and $[X^{k+2}, S^k]$ motivated *Adém's operation* and deeper study of the *Steenrod algebra*. We refer to [32] for a lucid exposition of early developments in this direction.

The equivariant setting. A $\mathbb{Z}/2$-*space* can be defined as a pair (X, ν), where $\nu\colon X \to X$ is a homeomorphism satisfying $\nu^2 = \mathrm{id}_X$ (which defines an action of the group $\mathbb{Z}/2$ on X—whence the name). A primary example of a $\mathbb{Z}/2$-space is a sphere S^k with the antipodal action $x \mapsto -x$. An *equivariant map* between $\mathbb{Z}/2$-spaces (X, ν) and (Y, ω) is a continuous map $f\colon X \to Y$ such that f commutes with the $\mathbb{Z}/2$-actions, i.e., $f\nu = \omega f$.

Our original motivation for working on the problems discussed in this paper was the computation of the $\mathbb{Z}/2$-*index* (or *genus*) $\mathrm{ind}(X)$ of a $\mathbb{Z}/2$-space X, i.e., the smallest k such that X can be equivariantly mapped into S^k.

This problem arises, among others, in the problem of *embeddability* of topological spaces, which is a classical and much studied area (see, e.g., the survey [29]). One of the basic questions here is, given a k-dimensional finite simplicial complex K, can it be (topologically) embedded in $\mathbb{R}^d$? The celebrated *Haefliger–Weber theorem* from the 1960s asserts that, in the *metastable range of dimensions*, i.e., for $k \leq \frac{2}{3}d - 1$, embeddability is equivalent to $\mathrm{ind}(K^2_\Delta) \leq d - 1$, where K^2_Δ is a certain $\mathbb{Z}/2$-space constructed from K (the *deleted product*). Thus, in this range, the embedding problem is, computationally, a special case of $\mathbb{Z}/2$-index computation; see [14] for a study of algorithmic aspects of the embedding problem, where the metastable range was left as one of the main open problems.

The $\mathbb{Z}/2$-index also appears as a fundamental quantity in combinatorial applications of topology. For example, the celebrated result of Lovász on Kneser's conjecture can nowadays be re-stated as $\chi(G) \geq \mathrm{ind}(B(G)) + 2$, where $\chi(G)$ is the chromatic number of a graph G, and $B(G)$ is a certain simplicial complex constructed from G (see, e.g., [13]).

1.2 Computational complexity

We are interested in the computational complexity of the problems from the previous section.

First of all, we should point out that, by classical uncomputability results in topology, most of these problems are *algorithmically unsolvable* if we place no restriction on the space Y. Indeed, given a finite simplicial complex Y, which may even be assumed to be 2-dimensional, it is undecidable whether $\pi_1(Y)$ is trivial (by a result of Adjan and of Rabin;

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

[2] Strictly speaking, one should really write $\pi_k(Y, y_0)$ but for a path-connected Y, the choice of y_0 does not matter.

[3] A dream project would be to show some kind of hardness for computing $\pi_k(S^n)$; at the moment we do not see any promising approach, though.

see, e.g., the survey [31]). The triviality of $\pi_1(Y)$ is equivalent to $[S^1, Y]$ having 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 $f\colon S^1 \to Y$ is homotopic to a constant map, and this homotopy triviality is equivalent to the extendability of f to the 2-dimensional ball D^2. Therefore, the extension problem is undecidable as well. (For undecidability results concerning numerous more loosely related topological problems we refer to [31, 18, 17] and references therein.)

In these undecidability results, all 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 $(k-1)$-*connected*, meaning that $\pi_i(Y)$ is trivial for all $i \le k-1$. 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 by Brown [2] from 1957. He showed that $[X, Y]$ is computable under the assumption that Y is 1-connected and all the higher homotopy groups $\pi_i(Y)$, $2 \le i \le \dim X$, are *finite* (this is a strong assumption, *not* satisfied by spheres, for example). Then he went on to show computability of $\pi_k(Y)$, $k \ge 2$, for all 1-connected 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 [26], by Smith [30], and by Sergeraert, Rubio, Dousson, and Romero (e.g., [27, 23, 21, 24]; also see [25] for an exposition). New algorithms for computing higher homotopy groups follow from these methods; see Real [20] for an algorithm based on Sergeraert et al. The problem of computing $[X, Y]$ and the extension problem were not addressed in those papers, but we rely on methods developed by Sergeraert et al. in implementing numerous operations in our algorithms.

In a previous paper [3], we gave an algorithm that computes $[X, Y]$ under the restriction that, for some $k \ge 2$, Y is $(k-1)$-connected and $\dim X \le 2k-2$. More precisely, it is known that $[X, Y]$ has a canonical structure of an Abelian group under these conditions, and the algorithm computes its isomorphism type. We conjectured that the algorithm can be made to run in polynomial time for every fixed k, and here we prove this conjecture (Corollary 1.4 below).

1.3 New results

Higher homotopy groups. For k fixed, $\pi_k(Y)$ is polynomial-time computable:

THEOREM 1.1. *For every fixed $k \ge 2$, there is a polynomial-time algorithm that, given a 1-connected finite simplicial complex Y, computes (the isomorphism type of) the kth homotopy group $\pi_k(Y)$.*

Here and in the sequel, the size of a simplicial complex is measured as the number of simplices. The isomorphism type of the Abelian group $\pi_k(Y)$ is represented as a direct sum of cyclic groups.

If k is a part of input, computing $\pi_k(Y)$ is #*P*-hard:

THEOREM 1.2. *The following problem is #P-hard: Given a finite 4-dimensional 1-connected simplicial complex Y and an integer k coded in unary, compute* $\operatorname{rank}(\pi_k(Y))$, *i.e., the number of direct summands of $\pi_k(Y)$ isomorphic to $\mathbb{Z}$.*

This looks very similar to a result of Anick [1], but there is an important difference: his input Y is a *cell complex*, rather than a simplicial complex, and if one converts his Y into a simplicial complex in a standard manner, the number of simplices is going to be *exponentially large* in the input size of the original cell complex, thus rendering the #P-hardness result meaningless for simplicial complexes. We devise another way of converting Anick's Y into a simplicial complex, which produces only polynomially many simplices.[4] The idea is simple, but there are several technical issues to be worked out; we refer to [5] for the derivation of Theorem 1.2 from Anick's result.

Computing a Postnikov system. The algorithm for computing $\pi_k(Y)$ in Theorem 1.1 is a by-product of an algorithm for computing the first k stages of a (standard) *Postnikov system* for Y. In this respect it is similar to the algorithm of Brown [2].

A Postnikov system of a space Y is, roughly speaking, a way of building Y (or rather, a space homotopy equivalent to Y) from "canonical pieces", called *Eilenberg–MacLane spaces*, whose homotopy structure is the simplest possible. A Postnikov system has countably many *stages* $P_0, P_1, \ldots$, where P_k reflects the homotopy properties of Y up to dimension k, and in particular, $\pi_i(P_k) \cong \pi_i(Y)$ for all $i \le k$, while $\pi_i(P_k) = 0$ for $i > k$. The isomorphisms of the homotopy groups for $i \le k$ are induced by maps $\varphi_i\colon Y \to P_k$, which are also part of the Postnikov system. Moreover, there is a mapping $\boldsymbol{k}_i$ defined on P_i, called the ith *Postnikov class*; together with the group $\pi_{i+1}(Y)$ it describes how P_{i+1} is obtained from P_i, and it is of fundamental importance for dealing with maps from a space X into Y. We will say more about Postnikov systems later on; now we state the result somewhat informally.

THEOREM 1.3 (INFORMAL). *For every fixed number $k \ge 2$, given a 1-connected space Y represented as a finite simplicial complex, a suitable representation of the first k stages of a Postnikov system for Y can be constructed in polynomial time, in such a way that each of the mappings φ_i, $i \le k$, and $\boldsymbol{k}_i$, $i \le k-1$, can be evaluated in polynomial time.*

Some of the tools and methods from the proof are sketched in Section 3, and for a full proof we refer to [4].

Computing the structure of all maps. In the algorithm for computing $[X, Y]$ from our previous paper [3] mentioned above, the stage P_{2k-2} of the Postnikov system of Y is used as an approximation to Y, since for every 1-connected Y and every X of dimension at most $2k-2$, there is an isomorphism $[X, Y] \cong [X, P_{2k-2}]$, induced by the composition with the mapping $\varphi_{2k-2}\colon Y \to P_{2k-2}$. At the same

[4] The simplicial complex is only *homotopy equivalent* to Y. We recall that spaces X and Y are homotopy equivalent if there are maps $f\colon X \to Y$ and $g\colon Y \to X$ such that the compositions $fg \sim \mathrm{id}_Y$ and $gf \sim \mathrm{id}_X$. Equivalently, homotopy equivalence of X and Y can be thought of as follows: there is another space Z that can be "continuously shrunk" to both X and Y (the technical term for continuous shrinking here is *deformation retraction*). As far as questions "up to homotopy" are concerned, homotopy equivalent X and Y are indistinguishable, and in particular, $\pi_k(X) = \pi_k(Y)$ for all $k \ge 0$.

time, the continuous maps $X \to P_{2k-2}$ are easier to handle than the maps $X \to Y$: each of them is homotopic to a simplicial, and thus combinatorially described, map, and it is possible to define (and implement) a binary operation on P_{2k-2} which induces the group structure on $[X, P_{2k-2}]$. This, in a nutshell, explains the usefulness of the Postnikov system for dealing with maps into Y.

As a direct consequence of the main result of [3] and of Theorem 1.3, we obtain the following.

Corollary 1.4 (based on [3]). *For every fixed $k \geq 2$ there is a* polynomial-time *algorithm that, given finite simplicial complexes X, Y, where $\dim(X) \leq 2k-2$ and Y is $(k-1)$-connected, computes the isomorphism type of $[X, Y]$ as an Abelian group.*

The extension problem. Here polynomial-time decidability can be pushed even by one dimension of X higher, compared to Corollary 1.4.

Theorem 1.5. *Let $k \geq 2$ be fixed. Then there is a polynomial-time algorithm that, given finite simplicial complexes X and Y with Y $(k-1)$-connected and $\dim X \leq 2k-1$, a subcomplex $A \subseteq X$, and a simplicial map[5] $f\colon A \to Y$, decides whether f admits an extension to a map (*not *necessarily simplicial) $X \to Y$.*

This is a simple extension of Corollary 1.4, presented in [4]; for $\dim X \leq 2k-2$, the algorithm even finds a certain description of the set of *all* possible extensions up to homotopy.

Undecidability. Next, we present results showing that the assumption of Y $(k-1)$-connected and $\dim X \leq 2k-1$ in Theorem 1.5, which may look artificial at first sight, is sharp, in the sense that for $\dim X = 2k$ the extension problem becomes undecidable. Moreover, we show that either X or Y can be fixed once and for all, and undecidability still holds.

Theorem 1.6. *Let $k \geq 2$ be fixed.*

(a) 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 X and A, $A \subseteq X$, with $\dim X = 2k$, and a simplicial map $f\colon A \to Y$, decide whether there exists a continuous map $X \to Y$ extending f. For k even, Y_k can be taken as the sphere S^k.*

(b) 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 finite $(k-1)$-connected simplicial complex Y and a simplicial map $f\colon A \to Y$, decide whether there exists a continuous map $X \to Y$ extending f.

A proof is given in [5], and some of the ideas are outlined in Section 4.

While most of the previous undecidability results in topology rely on the word problem in groups and its relatives, our proof of Theorem 1.6 relies on 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

$$p_i(x_1, \ldots, x_n) = 0, \quad i = 1, 2, \ldots, m, \tag{1}$$

where $p_1, \ldots, p_m$ are n-variate polynomials with integer coefficients.

Remark on the equivariant setting. In a forthcoming paper, which is not discussed in this extended abstract, we extend some of the results presented above to the equivariant setting. In particular, we show that $\mathrm{ind}(X)$ can be computed if $\dim X \leq 2\,\mathrm{ind}(X) - 1$. On the other hand, we conjecture that, for every $k \geq 2$, the question "Does $\mathrm{ind}(X) = k$?" for a $(2k)$-dimensional $\mathbb{Z}/2$-space X is undecidable.

2. TOPOLOGICAL PRELIMINARIES

Simplicial sets. A simplicial set is a way of specifying a topological space in purely combinatorial terms. We refer to [7, 28] for a more comprehensive introduction.

Similar to a simplicial complex, a simplicial set is a space built of vertices, edges, triangles, and higher-dimensional simplices, but simplices are allowed to be glued to each other and to themselves in more general ways. 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.

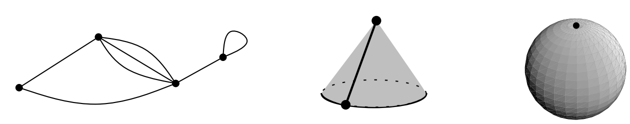

Another new feature of a simplicial set, in comparison with a simplicial complex, is the presence of *degenerate simplices*. For example, the edges of the triangle with a contracted boundary (in the last example above) do not disappear, but each of them becomes a degenerate 1-simplex.

A simplicial set X is a sequence $(X_0, X_1, \ldots)$ of mutually disjoint sets, where the elements of X_k are called the *k-simplices of X* (unlike for simplicial complexes, there can be many simplices with the same vertex set). In addition, for every $k \geq 1$, there are mappings $\partial_0, \ldots, \partial_k\colon X_k \to X_{k-1}$ called *face operators*; the intuitive meaning is that for a simplex $\sigma \in X_k$, $\partial_i \sigma$ is the face of σ opposite to the ith vertex. Moreover, for $k \geq 0$ there are mappings $s_0, \ldots, s_k\colon X_k \to X_{k+1}$ (opposite direction) called the *degeneracy operators*; the approximate meaning of $s_i\sigma$ is the degenerate simplex which is geometrically identical to σ, but with the ith vertex duplicated. A simplex is called *degenerate* if it lies in the image of some s_i; otherwise, it is *nondegenerate*. There are natural axioms that the ∂_i and the s_i have to satisfy, but we will not list them here, since we will not really use them.

There is a simple canonical way of converting a simplicial complex into a simplicial set—essentially we just add degenerate simplices in the only possible way.

Every simplicial set X specifies a topological space $|X|$, the *geometric realization* of X. It is obtained by assigning a geometric k-dimensional simplex to each nondegenerate k-simplex of X, and then gluing these simplices together

[5] A simplicial map $A \to Y$ maps each simplex of A onto a simplex of Y of the same or lower dimension. It is fully specified by a map of the vertex set of A into the vertex set of Y, and thus it provides a finite representation of a continuous map.

according to the face operators; we refer to the literature for the precise definition.

Product. The *product* $X \times Y$ of simplicial sets X and Y is the simplicial set whose k-simplices are ordered pairs (σ, τ), where $\sigma \in X_k$ and $\tau \in Y_k$. The face and degeneracy operators are applied to such pairs componentwise. On the level of geometric realizations, this product corresponds to the usual Cartesian product of spaces.

Simplicial maps. For simplicial sets X, Y, a *simplicial map* $f\colon X \to Y$ is a sequence $(f_k)_{k=0}^{\infty}$ of maps $f_k\colon X_k \to Y_k$ (every k-simplex is mapped to a k-simplex) that commute with the face and degeneracy operators, i.e., $\partial_i f_k = f_{k-1}\partial_i$ and $s_i f_k = f_{k+1} s_i$.

Every simplicial map $f\colon X \to Y$ defines a continuous map $\varphi\colon |X| \to |Y|$ of the geometric realizations. Of course, not all continuous maps are induced by simplicial maps, but there is a very important class of *Kan simplicial sets*, with the following crucial property: if Y is a Kan simplicial set and X is any simplicial set, then every continuous map $\varphi : |X| \to |Y|$ is homotopic to (the geometric realization of) some simplicial map $f\colon X \to Y$. This is essential in the algorithmic treatment of continuous maps. Here we omit the definition of a Kan simplicial set, since we will not directly use it. An example will be given a little later.

Chain complexes and homology. Besides the homotopy groups $\pi_k(X)$, the homology groups $H_k(X)$ and cohomology groups $H^k(X)$ belong to the most important topological invariants of a space X. Homology and cohomology are generally considered much easier to handle than homotopy, and in particular, computationally they are quite well understood and, for most purposes, they can be regarded as "easy."

Here we will use homology and cohomology as auxiliary devices, and we will only set up some basic notation, referring to standard textbooks for an introduction. For our purposes, in accordance with standard textbooks of topology, a *chain complex* C_* is a sequence $(C_k)_{k\in\mathbb{Z}}$ of *free* Abelian groups, together with a sequence $(d_k\colon C_k \to C_{k-1})_{k\in\mathbb{Z}}$ of group homomorphisms that satisfy the condition $d_{k-1}d_k = 0$. The C_k are the *chain groups*, their elements are called *k-chains*, and the d_k the *differentials*. We also recall that $Z_k = Z_k(C_*) := \ker d_k \subseteq C_k$ is the group of *cycles*, $B_k = B_k(C_*) := \operatorname{im} d_{k+1} \subseteq Z_k$ is the group of *boundaries*, and the quotient group $H_k(C_*) := Z_k/B_k$ is the kth *homology group* of C_*.

With a simplicial set X we associate the *normalized chain complex* $C_*(X)$, in which the kth chain group $C_k(X)$ consists of k-chains that are formal sums $c = \sum_{\sigma\in X_k^{\rm ndg}} \alpha_\sigma \cdot \sigma$, where $X_k^{\rm ndg}$ stands for the set of all k-dimensional nondegenerate simplices of X and the α_σ are integers, only finitely many of them nonzero. The differentials are defined in a standard way using the face operators: for k-chains of the form $1\cdot\sigma$, which constitute a basis of $C_k(X)$, we set $d_k(1\cdot\sigma) := \sum_{i=0}^{k}(-1)^i \cdot \partial_i\sigma$ (some of the $\partial_i\sigma$ may be degenerate simplices; then they are ignored in the sum), and this extends linearly to all k-chains. The homology of $C_*(X)$ coincides with the homology of the geometric realization of X defined in the usual way.

Let C_* and $\tilde{C}_*$ be two chain complexes. A *chain map* $f\colon C_* \to \tilde{C}_*$ is a sequence $(f_k)_{k\in\mathbb{Z}}$ of homomorphisms $f_k\colon C_k \to \tilde{C}_k$ satisfying $f_{k-1}d_k = \tilde{d}_k f_k$. A simplicial map $f\colon X \to Y$ of simplicial sets induces a chain map $f_*\colon C_*(X) \to C_*(Y)$ in the obvious way.

Cohomology. Given a chain complex C_* and an Abelian group π, we define its *kth cochain group with coefficients in* π as $C^k(C_*;\pi) := \operatorname{Hom}(C_k, \pi)$ (all homomorphisms of C_k into π) with pointwise addition. Its elements are called *k-cochains*. The *coboundary operator* $\delta_k\colon C^k(C_*;\pi) \to C^{k+1}(C_*;\pi)$ is given by $(\delta_k c^k)(c_{k+1}) := c^k(d_{k+1}c_{k+1})$ for every k-cochain c^k and every $(k+1)$-chain c_{k+1}.

If X is a simplicial set, we abbreviate $C^k(C_*(X);\pi)$ to $C^k(X;\pi)$. Here a k-cochain can be specified by its values on the standard basis of $C_k(X)$, and thus it can be viewed, in a more pedestrian way, as an arbitrary labeling of the nondegenerate k-simplices of X by elements of π. If X has infinitely many nondegenerate k-simplices, then a k-cochain is an infinite object, unlike a k-chain.

Given a chain complex C_*, $B^k = B^k(C_*;\pi) := \operatorname{im}\delta_{k-1}$ is the group of *k-coboundaries*, $Z^k = Z^k(C_*;\pi) := \ker\delta_k$ the group of *k-cocycles*, and $H^k = H^k(C^*;\pi) := Z^k/B^k$ is the *kth cohomology group with coefficients in* π.

Eilenberg–MacLane spaces, and $K(\mathbb{Z},1)$ in particular. As was mentioned in the introduction, Eilenberg–MacLane spaces are basic building blocks of Postnikov systems. For an Abelian group π and an integer $k \ge 1$, the *Eilenberg–MacLane space* $K(\pi,k)$ is defined as a topological space T with $\pi_k(T) \cong \pi$ and $\pi_i(T) = 0$ for all $i \ne k$. In the realm of "nice" spaces, namely, CW-complexes, this defines $K(\pi,k)$ uniquely up to homotopy equivalence.

We will work with a standard way of representing $K(\pi,k)$ as a Kan simplicial set, and we reserve the symbol $K(\pi,k)$ for this particular simplicial representation.

Crucially, the maps from a simplicial set X into $K(\pi,k)$ have a relatively simple description. Namely, all *simplicial* maps $X \to K(\pi,k)$ are in one-to-one correspondence with $Z^k(X;\pi)$, the k-cocycles on X, and homotopy classes of such maps correspond to cohomology classes on X: $[X, K(\pi,k)] \cong H^k(X;\pi)$.

Here we will not define $K(\pi,k)$ in general; we will consider only the particular case of $K(\mathbb{Z},1)$. It deserves a special attention for reasons given later, and it provides a good example of a Kan simplicial set.

The k-simplices of $K(\mathbb{Z},1)$ are represented by k-term sequences $\sigma = [a_1\,|\,a_2\,|\cdots|\,a_k]$ of integers (the "bar notation" is traditional). Nondegenerate simplices are exactly sequences with all terms nonzero. The face operators are given by $\partial_0\sigma = [a_2\,|\cdots|\,a_k]$, $\partial_k\sigma = [a_1\,|\cdots|\,a_{k-1}]$, and $\partial_i\sigma = [a_1\,|\cdots|\,a_{i-1}\,|\,a_i + a_{i+1}\,|\,a_{i+2}\,|\cdots|\,a_k]$, $1 \le i \le k-1$.

Topologically, $K(\mathbb{Z},1)$ is homotopy equivalent to S^1, and thus, in a sense, very simple. Yet, from an algorithmic point of view, handling it efficiently is one of the most demanding parts of our development.

Postnikov systems. We will not define a Postnikov system of a space Y precisely, since this would require too many auxiliary notions. In addition to the properties stated in the introduction, we mention that the kth stage P_k is defined as a *twisted product* of P_{k-1} with an Eilenberg–MacLane space:

$$K(\pi_k(Y),k) \times_{\boldsymbol{k}_{k-1}} P_{k-1},$$

where P_0 and P_1 are one-point spaces. (In [4], a slightly different but simplicially isomorphic representation is used, using *pullbacks* which we do not want to introduce here.) Twisted product is a simplicial analog of the topological notion of *fiber bundle* (this is a generalization of a *vector bundle*), which we will now outline to convey some intuition.

Let B, the *base space*, and F, the *fiber space*, be two spaces. The Cartesian product $F \times B$ can be thought of as a copy of F sitting above each point of B; for B the unit circle S^1 and F a segment this is indicated in the left picture:

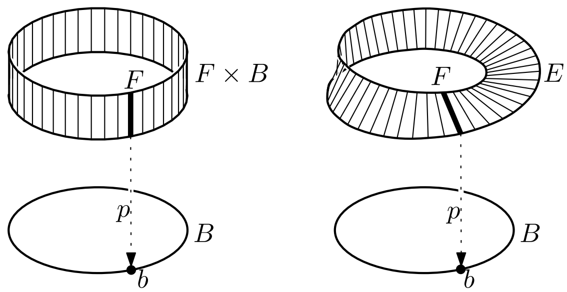

The product $F \times B$ is a *trivial* fiber bundle, while the right picture shows a nontrivial fiber bundle (a Möbius band in this case). Above every point $b \in B$, we still have a copy of F, and moreover, each such b has a small neighborhood U such that the union of all fibers sitting above U is homeomorphic to the product $F \times U$, a rectangle in the picture. However, globally, the union of the fibers above all of B forms a space E, the *total space* of the fiber bundle, that is in general different from $F \times B$.

In the Postnikov system setting, we can think of P_k as the total space of a fiber bundle with base space P_{k-1} and fiber $K(\pi_k(Y), k)$. The way the bundle is "twisted" is determined by the Postnikov class $\boldsymbol{k}_{k-1}$, which can be regarded as a simplicial map $P_{k-1} \to K(\pi_k(Y), k+1)$, or alternatively, using the correspondence mentioned earlier, as a cocycle in $Z^{k+1}(P_{k-1}, \pi_k(Y))$. Technically, the twisting is achieved by forming the ordinary simplicial product $K(\pi_k(Y), k) \times P_{k-1}$, and then modifying the 0th face operator in it using $\boldsymbol{k}_{k-1}$—we refer to [4, 15] for the details.

Lifting a map by one stage. The Postnikov class $\boldsymbol{k}_{k-1}$ is also crucial if we consider a map $f\colon X \to P_{k-1}$ from some space to a stage of the Postnikov system, and we ask when there is a map $\overline{f}\colon X \to P_k$ to the next stage that *lifts* f, i.e., satisfies $f = p_k\overline{f}$, where $p_k\colon P_k \to P_{k-1}$ is the projection. The answer is that $\overline{f}$ exists iff the composition $\boldsymbol{k}_{k-1}f$ is homotopic to a constant map. Since this composition goes into $K(\pi_k(Y), k+1)$, it can also be interpreted as a cocycle in $Z^{k+1}(X; \pi_k(Y))$, and the liftability condition then says that this cocycle must represent 0 in cohomology, which is algorithmically testable. This result can be regarded as a version of the classical *obstruction theory*, and it plays a crucial role in the algorithms of [3].

Mapping cylinder and mapping cone. Let $f\colon X \to Y$ be a map of topological spaces. Then the *mapping cylinder* $\mathrm{Cyl}(f)$ is obtained by gluing the product ("cylinder") $X \times [0,1]$ to Y via the identification of $(x,0)$ with $f(x) \in Y$, for all $x \in X$. The *mapping cone* of f is obtained by contracting the "top copy" $X \times \{1\}$ in the mapping cylinder to a point.

The group operation in higher homotopy groups. We recall how one adds homotopy classes of maps in $\pi_k(Y)$. For this, it is convenient to regard S^k as the cube I^k, $I = [0,1]$, with the boundary identified to a single point, which is the basepoint s_0 of S^k. Then a pointed map $S^k \to Y$ can be represented as a map $I^k \to Y$ that sends all of the boundary to the basepoint of Y. Given two such maps $f, g\colon I^k \to Y$, we form a new map h of this kind, representing the sum $[f] + [g]$ in $\pi_k(Y)$, by re-interpreting I^k as a stack of two cubes sharing a facet, and using f on the bottom cube and g on the top one.

3. POLYNOMIAL-TIME HOMOLOGY AND POSTNIKOV SYSTEMS

Locally effective objects. In many areas where computer scientists seek efficient algorithms, both the input objects and intermediate results in the algorithms are finite, and they can be explicitly represented in the computer memory. In contrast, in the algorithms considered here, we need to deal with infinite objects. For example, even if the input is a finite simplicial complex, its Postnikov system is made of Eilenberg–MacLane spaces, such as $K(\mathbb{Z}, 1)$, represented as Kan simplicial sets, and these are necessarily infinite.

For algorithmic purposes, we thus represent a simplicial set X by a collection of several algorithms, which allow us to access certain information about X (this is also called a *black box* or *oracle* representation of X). Concretely, let X be a simplicial set, and suppose that some encoding for the simplices of X by strings (finite sequences over some fixed alphabet, say $\{0,1\}$) has been fixed. We say that X is a *locally effective simplicial set* if an algorithm is available that, given (an encoding of) a k-simplex σ of X and $i \in \{0, 1, \ldots, k\}$, computes the simplex $\partial_i\sigma$, and similarly for the degeneracy operators s_i.

A *locally effective chain complex* is defined similarly: the addition and subtraction of chains and the differentials are computable maps. We note that if X is a locally effective simplicial set, then its chain complex $C_*(X)$ is locally effective automatically.

Computing global information. The locally effective representation of a simplicial set X is typically insufficient for computing global information about X, such as the kth homology group $H_k(X)$.

Sergeraert and his co-authors have developed a way of augmenting a locally effective simplicial set X with homological information, which is captured in the notion of a *simplicial set with effective homology.* These simplicial sets allow for computing homology groups, but they are also equipped with additional information, which makes them stable under a large repertoire of operations: if we apply some of the "classical" operations, such as product, classifying space, loop space, etc. to them, the result is again a simplicial set with effective homology.

Equipping X with effective homology means associating with the chain complex $C_*(X)$, whose chain groups may have infinite ranks, another, typically much smaller locally effective chain complex EC_*, which we can think of as a finitary approximation of $C_*(X)$.

We assume that EC_* is *(globally) effective* in the following sense: there is an algorithm that, given k, outputs a distinguished basis Bas_k of the chain group EC_k (and in particular, EC_k has a finite rank). Together with the computability of the differential in EC_*, this allows us to compute the homology groups of EC_*, for example (using a Smith normal form algorithm, as is explained in standard textbooks such as [16]).

Reductions and strong equivalences. Now we discuss the way of associating a "small" chain complex EC_* with a "big" chain complex C_*.

If $f, g\colon C_* \to \tilde{C}_*$ are two chain maps, then a *chain homotopy* of f and g is a sequence $(h_k)_{k\in\mathbb{Z}}$ of homomorphisms, where $h_k\colon C_k \to \tilde{C}_{k+1}$ (raising the dimension by one), such that $g_k - f_k = \tilde{d}_{k+1}h_k + h_{k-1}d_k$. Chain maps and chain homotopies can be regarded as algebraic counterparts of continuous maps of spaces and their homotopies, respectively. In particular, two chain-homotopic chain maps induce the same map in homology.

Let C_* and $\tilde{C}_*$ be chain complexes. A *reduction* ρ from C_* to $\tilde{C}_*$ consists of three maps f, g, h such that $f\colon C_* \to \tilde{C}_*$ and $g\colon \tilde{C}_* \to C_*$ are chain maps; the composition $fg\colon \tilde{C}_* \to \tilde{C}_*$ is equal to the identity $\mathrm{id}_{\tilde{C}_*}$, while the composition $gf\colon C_* \to C_*$ is chain-homotopic to id_{C_*}, with $h\colon C_* \to C_*$ providing the chain homotopy; and $fh = 0$, $hg = 0$, and $hh = 0$. We write $C_* \Rrightarrow \tilde{C}_*$ if there is a reduction from C_* to $\tilde{C}_*$.

A *strong equivalence* of chain complexes $C_*, \tilde{C}_*$, in symbols $C_* \Lleftarrow\!\Rrightarrow \tilde{C}_*$, means that there exists another chain complex $\tilde{\tilde{C}}_*$ and reductions $C_* \Lleftarrow \tilde{\tilde{C}}_* \Rrightarrow \tilde{C}_*$. Strong equivalence is transitive; this requires a nontrivial argument (see [4, 25]).

Effective homology. A locally effective simplicial set is equipped with *effective homology* if there is a globally effective chain complex EC_* and a strong equivalence $C_*(X) \Lleftarrow\!\Rrightarrow EC_*$, such that the intermediate chain complex in the strong equivalence is locally effective and all of the maps involved in the reductions are computable.

Using the strong equivalence with EC_*, we can compute the homology groups of X (represented as abstract Abelian groups, i.e., as direct sums of cyclic groups). Given an element of $H_k(X)$, we can also compute a cycle representing this homology class, and conversely. Moreover, given a cycle c representing 0 in $H_k(X)$, we can compute a $(k+1)$-chain for which c is the boundary. With these computational primitives, we can solve most homological problems for $C_*(X)$ of interest.

Polynomiality and parameterization. In order to define a polynomial-time counterpart of effective homology, the first natural, but not quite sufficient, idea is to require all the algorithms (black boxes) involved in the definition of a simplicial sets with effective homology, i.e., the ones witnessing local effectivity, as well as those for evaluating the maps in the reductions, to run in polynomial time.

This is fine as long as we deal with an individual simplicial set, such as $K(\mathbb{Z},1)$. But if, in some subroutine of our Postnikov system algorithm, we want to deal with $K(\mathbb{Z}^r,1)$, for example, where r is a parameter that depends on the input simplicial complex Y, then we certainly do not want the black boxes for $K(\mathbb{Z}^r,1)$ to run in time n^r, where n is the size of the input to the black box. Instead, the running time should also be bounded polynomially in r.

As another example, the kth stage $P_k = P_k(Y)$ of a Postnikov system for an input simplicial complex Y depends on Y, and again we want the size of Y to enter at most polynomially in the running time of the black boxes for $P_k(Y)$. In order to deal with this issue, we do not define polynomial-time homology for individual simplicial sets, but rather for *families* of simplicial sets, typically infinite ones.

A *parameter set* is a set $\mathcal{I}$ on which an injective mapping $\mathsf{enc}\colon \mathcal{I} \to \{0,1\}^*$ is defined, specifying an encoding of each element of $\mathcal{I}$ by a string. An example of a parameter set in our algorithms is the set of all 1-connected finite simplicial complexes, and another is the set of all (isomorphism types of) finitely generated Abelian groups.

We define a *parameterized simplicial set* as a mapping X that, for some parameter set $\mathcal{I}$, assigns to each $I \in \mathcal{I}$ a simplicial set $X(I)$. We often write such a parameterized simplicial set as $(X(I) : I \in \mathcal{I})$. We also assume that an encoding of simplices by strings has been fixed for each of the simplicial sets $X(I)$.

Such an X is a *locally polynomial-time simplicial set* if, for each k, there is an algorithm that, given $I \in \mathcal{I}$, a k-dimensional simplex $\sigma \in X(I)_k$, and $i \in \{0, 1, \ldots, k\}$, computes $\partial_i\sigma$ in time polynomial in $\mathsf{size}(I) + \mathsf{size}(\sigma)$ (where the polynomial may depend on k), and there is a similar algorithm for evaluating the degeneracy operators $s_i\sigma$. Here $\mathsf{size}(I)$ is the length of the string encoding I, and similarly for $\mathsf{size}(\sigma)$.

Similarly, for parameterized simplicial sets $(X(I) : I \in \mathcal{I})$ and $(Y(I) : I \in \mathcal{I})$, a *polynomial-time simplicial map* $X \to Y$ is a collection $(f_I)_{I\in\mathcal{I}}$, where f_I is a simplicial map $X(I) \to Y(I)$ and for each $k \geq 0$, there is an algorithm that, given $I \in \mathcal{I}$ and $\sigma \in X(I)_k$, computes $f_I(\sigma)$ in time polynomial in $\mathsf{size}(I) + \mathsf{size}(\sigma)$.

Analogously we define *locally polynomial-time chain complex*, *polynomial-time chain map*, a *polynomial-time reduction* $C_* \overset{\mathrm{P}}{\Rrightarrow} \tilde{C}_*$ between parameterized chain complexes, and a *polynomial-time strong equivalence* $C_* \overset{\mathrm{P}}{\Lleftarrow\!\Rrightarrow} \tilde{C}_*$ (in which we require all the chain complexes involved to be locally polynomial-time). A *globally polynomial-time chain complex* EC_* is a locally polynomial-time chain complex such that, for every fixed k, there is an algorithm that outputs a distinguished basis $\mathrm{Bas}(I)_k$ of $EC(I)_k$ in time polynomial in $\mathsf{size}(I)$.

We say that a locally polynomial-time simplicial set X is equipped with *polynomial-time homology* if there is a globally polynomial-time chain complex EC_* and a polynomial-time strong equivalence $C_*(X) \overset{\mathrm{P}}{\Lleftarrow\!\Rrightarrow} EC_*$. For such an X, we can perform in polynomial time all of the homological computations listed above in connection with effective homology (for this, we need to employ a polynomial-time algorithm for the Smith normal form, which is nontrivial but known—see [10, 33]).

About half of our paper [4] is devoted to showing that various operations with simplicial sets preserve polynomial-time homology: if $X^{(1)}, \ldots, X^{(t)}$ are simplicial sets equipped with polynomial-time homology and Φ is a "reasonable" way of constructing a new simplicial set from t old ones, then the simplicial set $\Phi(X^{(1)}, \ldots, X^{(t)})$ can also be equipped with polynomial-time homology. This is done by adapting known methods, developed earlier for effective homology and based on much older work by algebraic topologists. In most cases the adaptation is straightforward, but there are cases where polynomiality requires extra tricks, analysis, or assumptions. We omit all of this in this extended abstract.

Polynomial-time homology for $K(\mathbb{Z},1)$ from a discrete vector field. Let us regard the Eilenberg–MacLane space $K(\pi,k)$, with k fixed, as a simplicial set parameterized with the Abelian group π. Here π is given in the form $\mathbb{Z}^r \oplus (\mathbb{Z}/m_1) \oplus (\mathbb{Z}/m_2) \oplus \cdots \oplus (\mathbb{Z}/m_s)$, as a direct sum of cyclic groups, and we define its encoding size as

$r + \sum_{i=1}^{s} \lfloor 1 + \log_2 m_i \rfloor$. In the Postnikov system algorithm, we need to equip $K(\pi, k)$ with polynomial-time homology.

There are general operations on simplicial sets that allow us to obtain polynomial-time homology for $K(\pi, k+1)$ from that for $K(\pi, k)$; that for $K(\pi_1 \times \pi_2, 1)$ from that for $K(\pi_1, 1)$ and $K(\pi_2, 1)$, and that for $K(\mathbb{Z}/m, 1)$ from that for $K(\mathbb{Z}, 1)$. This leaves us with $K(\mathbb{Z}, 1)$ as the base case, and here we cannot use the classical way of obtaining effective homology, going back to Eilenberg and Mac Lane, since it is not polynomial. We thus had to develop a new way in [12].

Our method uses a suitable *discrete vector field* on $K(\mathbb{Z}, 1)$. Discrete vector fields originate in discrete Morse theory, developed by Forman [6], and Romero and Sergeraert in the preprint [22] discovered that a suitable vector field on a simplicial set X can be used to equip X with effective homology.

We refer to [12, 22] for the general definition of discrete vector fields and deriving effective or polynomial-time homology from it; here we formulate concretely what a discrete vector field on $K(\mathbb{Z}, 1)$ looks like and what it must satisfy in order to yield polynomial-time homology.

We recall that the nondegenerate simplices σ of $K(\mathbb{Z}, 1)$ are sequences $[a_1|\cdots|a_k]$ of nonzero integers. We want to partition these simplices into three classes $\mathcal{S}$, $\mathcal{T}$, and $\mathcal{C}$ (the *source simplices*, *target simplices*, and *critical simplices*), and construct a bijection $V : \mathcal{S} \to \mathcal{T}$ (a *discrete vector field*), such that for every $\sigma \in \mathcal{S}$, we have $\sigma = \partial_i V(\sigma)$ for exactly one i.

Given such a V, let us consider a simplex $\tilde{\sigma} \in \mathcal{S}$ of some dimension k, and let us say that a simplex τ (of dimension k or $k+1$) is *reachable* from $\tilde{\sigma}$ if it can be reached from $\tilde{\sigma}$ by finitely many moves, where the allowed moves are passing from a current simplex $\sigma \in \mathcal{S}$ to the simplex $\tau = V(\sigma) \in \mathcal{T}$, and passing from a current simplex $\tau \in \mathcal{T}$ to a simplex $\sigma = \partial_i \tau \in \mathcal{S} \cup \mathcal{C}$ such that $\tau \neq V(\sigma)$, where $i \in \{0, 1, \ldots, k+1\}$.

With these definitions, we moreover require the following:

(i) For every k, $\mathcal{C}$ contains only finitely many k-dimensional simplices.

(ii) Starting with any $\tilde{\sigma}$, we can never make an infinite sequence of allowed moves; that is, we can reach only finitely many simplices, and we also cannot get into a cycle.

(iii) For every k-dimensional simplex $\tilde{\sigma}$, the sum of $\mathsf{size}(\sigma)$ over all σ reachable from $\tilde{\sigma}$ is bounded by a polynomial (depending on k) in $\mathsf{size}(\tilde{\sigma})$. Here the size of a simplex $\sigma = [a_1|\cdots|a_k]$ is the total number of bits in the binary encoding of $a_1, \ldots, a_k$.

If all these conditions are met, and $V(\sigma)$ can be computed from σ in polynomial time, then we can use V to equip $K(\mathbb{Z}, 1)$ with polynomial-time homology.

To illustrate these definitions, let us present a classical vector field V_{EML} due to Eilenberg and Mac Lane, which satisfies (i) and (ii) (and yields effective homology for $K(\mathbb{Z}, 1)$) but not (iii). There are only two critical simplices, the 0-dimensional $[\,]$ (the empty sequence) and the 1-dimensional $[1]$. The remaining simplices are source or target depending on whether $a_1 \neq 1$ or $a_1 = 1$, respectively.

For $\sigma = [a_1|\cdots|a_k] \in \mathcal{S}$, $a_1 \neq 1$, we have $V_{\mathrm{EML}}(\sigma) = [1|a_1-1|a_2|\cdots|a_k]$ for $a_1 > 1$, and $V_{\mathrm{EML}}(\sigma) = [1|a_1|a_2|\cdots|a_k]$ for $a_1 < 0$. It can be checked that, for any starting $\tilde{\sigma}$, the sequence of moves is determined uniquely (there is no branching). It is easy to see that, for a positive integer a, the sequence of moves starting from $[a]$ is $[a] \to [1|a-1] \to [a-1] \to [1|a-2] \to [a-2] \to \ldots$; there are about a moves, and this is *exponential* in the number of bits of a. Thus, condition (iii) above indeed fails.

Next, we define another vector field V that does satisfy (iii). Actually, by an auxiliary step, we can restrict ourselves only to simplices $[a_1|\cdots|a_k]$ with all the a_i positive, and define the vector field only on them.

The only critical simplices are again $[]$ and $[1]$. There are two types of source simplices. The first type are simplices of the form $\sigma = \left[2^{i_1} \,|\, 2^{i_2} \,|\, \cdots \,|\, 2^{i_q} \,|\, b \,|\, a_{q+2} \,|\, \cdots \,|\, a_k\right]$, where $2^{i_1} \leq 2^{i_2} \leq \cdots \leq 2^{i_q} < b$ and b is not a power of two. In this case we set

$$V(\sigma) := \left[2^{i_1} \,|\, 2^{i_2} \,|\, \cdots \,|\, 2^{i_q} \,|\, \mathsf{lpow}(b) \,|\, \mathsf{ltrim}(b) \,|\, a_{q+2} \,|\, \cdots \,|\, a_k\right],$$

where $\mathsf{lpow}(b)$ is the largest power of 2 not exceeding b, and $\mathsf{ltrim}(b) := b - \mathsf{lpow}(b)$. That is, $V(\sigma)$ is obtained by splitting b into two components, $\mathsf{lpow}(b)$ and $\mathsf{ltrim}(b)$; informally, we can think of this as "chipping off" the leading bit of b.

The second type of source simplices are $\sigma = \left[2^{i_1} \,|\, \cdots \,|\, 2^{i_k}\right]$ with $2^{i_1} \leq 2^{i_2} \leq \cdots \leq 2^{i_{k-1}} < 2^{i_k}$ and $i_k \geq 1$ (this last condition is important only for $k = 1$). In this case we set $V(\sigma) := \left[2^{i_1} \,|\, \cdots \,|\, 2^{i_{k-1}} \,|\, 2^{i_k - 1} \,|\, 2^{i_k - 1}\right]$; i.e., we split the last component of σ into two equal halves.

The proof of the required properties is not simple and currently it needs a careful case analysis; we refer to [12].

Notes on the Postnikov system algorithm. Since we haven't really defined a Postnikov system, we will not present the algorithm in any detail either; we refer to [4, Section 4] instead.

The algorithm is similar to the algorithm of Brown [2] (but, unlike Brown's, it can handle all Y, not only those with all homotopy groups finite). It works by induction on k, assuming that the input simplicial set Y comes with polynomial-time homology (this is trivially satisfied for a finite simplicial complex, but in Theorems 1.1, 1.3, and 1.5, as well as in Corollary 1.4, Y can actually be any simplicial set with polynomial-time homology), P_{k-1} has been constructed with polynomial-time homology, and $\varphi_{k-1} : Y \to P_{k-1}$ is a polynomial-time simplicial map that induces isomorphism of homotopy groups up to dimension $k-1$. For starting the construction, we use that P_0 and P_1 consist of a single point.

In the inductive step, we first want to compute $\pi_k(Y)$. For this, we use the *Hurewicz isomorphism*, which in its simplest form asserts that, for a 1-connected Y, the first nonzero homotopy group of Y occurs in the same dimension as the first nonzero homology group and these two groups are isomorphic. We thus need a construction that "kills" the first $k-1$ homotopy groups of Y and leaves the kth one intact. The mapping cone M of φ_{k-1} is such a construction[6]; we construct it with polynomial-time homology and compute the isomorphism type of $\pi_k(Y)$ as $H_{k+1}(M)$ (there is a shift by one in dimension).

To get the stage P_k, we need to form the twisted product $K(\pi_k(Y), k) \times_{\boldsymbol{k}_{k-1}} P_{k-1}$. As was noted earlier, we can equip $K(\pi_k(Y), k)$, with polynomial-time homology, and it remains to define (and implement) $\boldsymbol{k}_{k-1}$ and φ_k, which we omit here.

[6]Mapping cone for a map of spaces was defined at the end of Section 2; in the algorithm, we work with an appropriate simplicial version.

The polynomiality of the algorithm follows easily, since we start with Y, which has polynomial-time homology by assumption, and we make some number of steps, depending only on k and thus constant, and in each of them we produce $K(\pi_i(Y), i)$ with polynomial-time homology, $i \le k$, or take simplicial sets with polynomial-time homology and produce a new one, by one of the operations from our repertoire. Thus, we end up with P_k with polynomial-time homology.

4. UNDECIDABILITY

We sketch the main steps of the proof of the undecidability result, Theorem 1.6. We first focus on part (a), i.e., fixed target, and on the case of *even* k—this is technically the simplest case in the theorem. We show that an algorithm solving the extension problem could be used to decide the solvability of the quadratic system

$$\sum_{1\le i<j\le r} a_{ij}^{(q)} x_i x_j = b_q, \quad q = 1, 2, \ldots, s. \qquad (2)$$

in unknowns $x_1, \ldots, x_r \in \mathbb{Z}$. The latter is undecidable, as follows from the undecidability of Hilbert's tenth problem (1) by an easy reduction.

Given the integer vectors $\mathbf{a} = (a_{ij}^{(q)})_{ijq}$ and $\mathbf{b} = (b_q)_q$, we first describe spaces $X, Y, A \subseteq X$ and a map $f\colon A \to Y$ such that f is extendable to X iff (2) is solvable. Then we still need to work further to convert these spaces into finite simplicial complexes and f into a simplicial map.

As was mentioned in the theorem, in part (a) and k even we take $Y = S^k$. The space $X = X_{\mathbf{a}}$ depends on the vector $\mathbf{a}$. To construct it, we start with r disjoint copies $S_1^k, \ldots, S_r^k$ of the k-dimensional sphere, and we glue all of them together at a single point, forming the *wedge* $W := \bigvee_{i=1}^r S_i^k$; let $\iota_i\colon S_i^k \to W$, $i = 1, 2, \ldots, r$, be the inclusion maps.

Then we take a space A as the wedge of s disjoint copies $S_1^{2k-1}, \ldots, S_s^{2k-1}$ of S^{2k-1}. Let $\varphi_q\colon S_q^{2k-1} \to W$ be the map $\sum_{1\le i<j\le r} a_{ij}^{(q)} [\iota_i, \iota_j]$. The addition and multiplication by integers in this formula is the group operation on pointed maps in the homotopy group $\pi_{2k-1}(W)$, and $[\iota_i, \iota_j]$ denotes a particular map $S^{2k-1} \to S_i^k \vee S_j^k \subseteq W$, the *Whitehead product* of the inclusions ι_i and ι_j.

The Whitehead product in general is a binary operation that, for a space Z, assigns to maps $\alpha\colon S^k \to Z$ and $\beta\colon S^\ell \to Z$ a map $[\alpha, \beta]\colon S^{k+\ell-1} \to Z$. It respects homotopy and thus it induces a binary operation $\pi_k(Z) \times \pi_\ell(Z) \to \pi_{k+\ell-1}(Z)$ on homotopy groups, also denoted by $[\cdot, \cdot]$; see, e.g., [8, Example 4.51]. For our particular case of $[\iota_i, \iota_j]$, if we consider the $2k$-dimensional torus $T = S_i^k \times S_j^k$, then T can be obtained by attaching a $2k$-dimensional ball D^{2k} by its boundary sphere S^{2k-1} to the wedge $S_i^k \vee S_j^k$, and $[\iota_i, \iota_j]$ is the appropriate attaching map.

Finally, $X = X_{\mathbf{a}}$ is the mapping cylinder[7] of the map $\varphi\colon A \to W$ that equals φ_q on S_q^{2k-1}, $q = 1, 2, \ldots, s$.

A mapping cylinder is always homotopy equivalent to the target space (range) of the mapping; in our case, $X \simeq W$. The homotopy equivalence is given by the inclusion $i\colon W \to X$ and by the map $h\colon X \to W$ that is the identity on W and sends a point $(a, t) \in A \times [0, 1)$ to $\varphi(a)$.

[7] Actually, for technical convenience, we take the *reduced* mapping cylinder, which is obtained from the ordinary mapping cylinder by collapsing the segment $\{x_0\} \times [0, 1]$ over the basepoint to a single point.

Thus, homotopy classes of maps $\bar{f}\colon X \to S^k$ are in one-to-one correspondence with homotopy classes of maps $\tilde{f}\colon W \to S^k$ (the correspondence is given by $\bar{f} \mapsto \tilde{f}h$), and the latter have a simple description: each class is specified by a vector $\mathbf{x} = (x_1, \ldots, x_r) \in \mathbb{Z}^r$, where x_i is the degree of the restriction of $\tilde{f}$ to the ith sphere S_i^k in the wedge W.

In a similar vein, each vector $\mathbf{b} = (b_1, \ldots, b_s) \in \mathbb{Z}^s$ specifies, up to homotopy, a map $f = f_{\mathbf{b}}\colon A \to S^k$, namely, such that the restriction of f to the qth sphere S_q^{2k-1} is (homotopic to) $b_q[\iota, \iota]$, where $[\iota, \iota]$ is the Whitehead square of the identity $\iota\colon S^k \to S^k$. It is known that $[\iota, \iota]$ is an element of infinite order in $\pi_{2k-1}(S^k)$, and no two of the $f_{\mathbf{b}}$ are homotopic.

Given an instance of the quadratic system (2), we let $X = X_{\mathbf{a}}$ and $f = f_{\mathbf{b}}\colon A \to S^k$ be as above. It remains to see when a map $\bar{f}\colon X \to S^k$ specified by $\mathbf{x}$ is a solution (up to homotopy) of the extension problem given by f. In other words, we ask for the homotopy class of the restriction of $\bar{f}$ to A. By the above, this restriction can be written as $\tilde{f}h|_A = \tilde{f}\varphi$, where $\tilde{f}$ is a map $W \to S^k$ of degree x_i on S_i^k, $i = 1, 2, \ldots, r$. Using the definition of φ and bilinearity and naturality properties of the Whitehead product, one can calculate that $\tilde{f}\varphi$ is homotopic to f iff the quadratic system (2) holds.

Now we have the appropriate X, A, Y, f, but we still need to make everything finite and simplicial. There is a way of doing this using the *simplicial approximation theorem*, which can be made algorithmic with some effort (although it does not supply any bound on the number of simplices). Actually, we can do even better, since an algorithm from the proof of Theorem 1.2 on #P-hardness can compute a finite simplicial complex homotopy equivalent to X in polynomial time (although here we don't really care about polynomiality). This finishes the sketch of the proof of part (a) with k even.

The just given argument for even k hinges on the fact that $[\iota, \iota]$ has infinite order in $\pi_{2k-1}(S^k)$; for k odd, though, $[\iota, \iota]$ has order 2 (and actually, all of the homotopy groups $\pi_n(S^k)$, $n > k$, are finite). For odd k, we thus take $Y = S^k \vee S^k$, and the role of $[\iota, \iota]$ in the previous proof is played by $[\iota_1, \iota_2]$, where ι_1, ι_2 are the inclusions of the two copies of S^k into Y.

Otherwise, the proof is similar to the case of k; however, instead of (2), it leads to a skew-symmetric bilinear system $\sum_{i<j} a_{ij}^{(q)}(x_i y_j - x_j y_j)$, $q = 1, 2, \ldots, s$, with unknowns $x_1, \ldots, x_r, y_1, \ldots, y_r$. Proving undecidability of that system is somewhat more demanding.

For part (b) of Theorem 1.6, fixed source, we briefly mention only the case of even k. Here X is the mapping cylinder of the Whitehead product $[\iota, \iota]\colon A = S^{2k-1} \to W = S^k$. To construct $Y = Y_{\mathbf{a}}$, we take the wedge $T := S_1^k \vee \cdots \vee S_r^k \vee S_1^{2k-1} \vee \cdots \vee S_s^{2k-1}$ and we attach to it disjoint copies D_{ij}^{2k}, $1 \le i \le j \le r$, of the ball D^{2k} along their boundaries, according to suitable attaching maps $\varphi_{ij}\colon S^{2k-1} \to T$—these depend on $\mathbf{a}$. For this construction, one thus needs to understand the homotopy groups $\pi_{2k-1}(T)$, which is described by a special case of a theorem of Hilton [9]. The rest of the analysis resembles that for the fixed source case, and we refer to [5] for more details.

Acknowledgment. We would like to thank Francis Sergeraert for many useful discussions and extensive advice; although he chose not to be listed as a co-author of this survey, his contribution is significant, not speaking of his leading role in the development of effective homology, which constitutes the foundation of our algorithmic methods. We also thank Marek Filakovský for useful discussions, and J. Maurice Rojas for kind replies to our questions concerning variants of Hilbert's tenth problem.

5. REFERENCES

[1] D. J. Anick. The computation of rational homotopy groups is $\#\wp$-hard. Computers in geometry and topology, Proc. Conf., Chicago/Ill. 1986, Lect. Notes Pure Appl. Math. 114, 1–56, 1989.

[2] E. H. Brown (jun.). Finite computability of Postnikov complexes. *Ann. Math. (2)*, 65:1–20, 1957.

[3] M. Čadek, M. Krčál, J. Matoušek, F. Sergeraert, L. Vokřínek, and U. Wagner. Computing all maps into a sphere. Preprint, arXiv:1105.6257, 2011. Extended abstract in *Proc. ACM–SIAM Symposium on Discrete Algorithms* (SODA 2012).

[4] M. Čadek, M. Krčál, J. Matoušek, L. Vokřínek, and U. Wagner. Polynomial-time computation of homotopy groups and Postnikov systems in fixed dimension. Preprint, arXiv:1211.3093, 2012.

[5] M. Čadek, M. Krčál, J. Matoušek, L. Vokřínek, and U. Wagner. Extendability of continuous maps is undecidable. Preprint, arXiv:1302.2370, 2013.

[6] R. Forman. Morse theory for cell complexes. *Adv. Math.*, 134(1):90–145, 1998.

[7] G. Friedman. An elementary illustrated introduction to simplicial sets. Rocky Mountain J. Math., to appear. Preprint arXiv:math/0809.4221v3, 2011.

[8] A. Hatcher. *Algebraic Topology.* Cambridge University Press, Cambridge, 2001. Electronic version available at `http://math.cornell.edu/hatcher#AT1`.

[9] J. P. Hilton. On the homotopy groups of union of spheres. *J. London Math. Soc.*, 3:154–172, 1955.

[10] R. Kannan and A. Bachem. Polynomial algorithms for computing the Smith and Hermite normal forms of an integer matrix. *SIAM J. Computing*, 8:499–507, 1981.

[11] S. O. Kochman. *Stable homotopy groups of spheres. A computer-assisted approach.* Lecture Notes in Mathematics 1423. Springer-Verlag, Berlin etc., 1990.

[12] M. Krčál, J. Matoušek, and F. Sergeraert. Polynomial-time homology for simplicial Eilenberg–MacLane spaces. Preprint, arXiv:1201.6222, 2011.

[13] J. Matoušek. *Using the Borsuk-Ulam theorem (revised 2nd printing).* Universitext. Springer-Verlag, Berlin, 2007.

[14] J. Matoušek, M. Tancer, and U. Wagner. Hardness of embedding simplicial complexes in $\mathbb{R}^d$. *J. Eur. Math. Soc.*, 13(2):259–295, 2011.

[15] J. P. May. *Simplicial objects in algebraic topology.* Chicago Lectures in Mathematics. University of Chicago Press, Chicago, IL, 1992. Reprint of the 1967 original; the page numbers do not quite agree with the 1967 edition.

[16] J. R. Munkres. *Elements of Algebraic Topology.* Addison-Wesley, Reading, MA, 1984.

[17] A. Nabutovsky and S. Weinberger. Algorithmic unsolvability of the triviality problem for multidimensional knots. *Comment. Math. Helv.*, 71(3):426–434, 1996.

[18] A. Nabutovsky and S. Weinberger. Algorithmic aspects of homeomorphism problems. In *Tel Aviv Topology Conference: Rothenberg Festschrift (1998)*, volume 231 of *Contemp. Math.*, pages 245–250. Amer. Math. Soc., Providence, RI, 1999.

[19] D. C. Ravenel. *Complex Cobordism and Stable Homotopy Groups of Spheres (2nd ed.).* Amer. Math. Soc., 2004.

[20] P. Real. An algorithm computing homotopy groups. *Mathematics and Computers in Simulation*, 42:461—465, 1996.

[21] A. Romero, J. Rubio, and F. Sergeraert. Computing spectral sequences. *J. Symb. Comput.*, 41(10):1059–1079, 2006.

[22] A. Romero and F. Sergeraert. Discrete vector fields and fundamental algebraic topology. Preprint arXiv:1005.5685, an updated version at `http://www-fourier.ujf-grenoble.fr/~sergerar/Papers/`, 2011.

[23] J. Rubio and F. Sergeraert. Constructive algebraic topology. *Bull. Sci. Math.*, 126(5):389–412, 2002.

[24] J. Rubio and F. Sergeraert. Algebraic models for homotopy types. *Homology, Homotopy and Applications*, 17:139–160, 2005.

[25] J. Rubio and F. Sergeraert. Constructive homological algebra and applications. Preprint, arXiv:1208.3816, 2012. Written in 2006 for a MAP Summer School at the University of Genova.

[26] R. Schön. Effective algebraic topology. *Mem. Am. Math. Soc.*, 451:63 p., 1991.

[27] F. Sergeraert. The computability problem in algebraic topology. *Adv. Math.*, 104(1):1–29, 1994.

[28] F. Sergeraert. Introduction to combinatorial homotopy theory. Available at `http://www-fourier.ujf-grenoble.fr/~sergerar/Papers/`, 2008.

[29] A. B. Skopenkov. Embedding and knotting of manifolds in Euclidean spaces. In *Surveys in contemporary mathematics*, volume 347 of *London Math. Soc. Lecture Note Ser.*, pages 248–342. Cambridge Univ. Press, Cambridge, 2008.

[30] J. R. Smith. m-Structures determine integral homotopy type. Preprint, arXiv:math/9809151v1, 1998.

[31] R. I. Soare. Computability theory and differential geometry. *Bull. Symbolic Logic*, 10(4):457–486, 2004.

[32] N. E. Steenrod. Cohomology operations, and obstructions to extending continuous functions. *Advances in Math.*, 8:371–416, 1972.

[33] A. Storjohann. Near optimal algorithms for computing Smith normal forms of integer matrices. In *International Symposium on Symbolic and Algebraic Computation*, pages 267–274, 1996.