Removing degeneracy in LP-type problems revisited

Jiří Matoušek

Department of Applied Mathematics and
Institute of Theoretical Computer Science (ITI)
Charles University, Malostranské nám. 25
118 00 Praha 1, Czech Republic, and

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

Abstract

LP-type problems is a successful axiomatic framework for optimization problems capturing, e.g., linear programming and the smallest enclosing ball of a point set. In [Matoušek and Škovroň, *Theory of Computing* 3(2007) 159–177] it was proved that in order to remove degeneracies of an LP-type problem, we sometimes have to increase its combinatorial dimension by a multiplicative factor of at least $1+\varepsilon$ for a certain small positive constant ε. The proof went by checking the unsolvability of a system of linear inequalities, with several pages of calculations.

Here we prove by a short topological argument that the dimension sometimes has to increase at least twice. We also construct 2-dimensional LP-type problems with $-\infty$ for which removing degeneracies forces arbitrarily large dimension increase.

1 Introduction

Degeneracies, such as three points lying on a common line or four co-circular points, constitute a serious issue in geometric computations, and often it would be desirable to have a simple and general tool for removing them. In [MŠ07] we obtained a result indicating that the problem of removing degeneracies has no simple "abstract" solution.

We considered *LP-type problems*, a successful axiomatic framework for optimization problems due to Sharir and Welzl [SW92], capturing, e.g., linear programming and the smallest enclosing ball of a point set. Loosely speaking, for a fixed $\varepsilon > 0$ and for infinitely many values of D, we constructed an LP-type problem of combinatorial dimension D, such that removing degeneracies from it enforces increasing the combinatorial dimension to at least $(1+\varepsilon)D$ (the undefined terms will be explained below).

In the proof we first reduced the problem to showing that a certain poset (partially ordered set) cannot be partitioned into Boolean algebras in a certain way. This combinatorial statement was then proved by setting up a system of linear inequalities and showing its inconsistence by several pages of calculations and asymptotic estimates.

In the present paper we provide a different and quite short proof of that combinatorial statement, by a topological method. More precisely, we prove a stronger (and actually tight) lower bound for the partitioning problem, namely, increase from D to $1.5D$, but under stronger assumptions. As a consequence, the proof given here applies to LP-type problems but, unlike the proof in [MŠ07], it doesn't apply to the more general framework of *violator spaces* [GMRŠ07]. Thus, at present the method given here and the method of [MŠ07] can be regarded complementary, neither one completely superseding the other.

Next, we present a different LP-type problem, where removing degeneracy forces dimension increase from D to $2D$. Finally, if we allow for the so-called *LP-type problems with* $-\infty$, we prove the necessity of dimension increase from 2 to arbitrarily large.

2 LP-type problems

Here we formally introduce LP-type problems. We state only the definitions directly relevant to our results; for more background on LP-type problems we refer, e.g., to [MŠ07] or [GMRŠ07]. Then we state our results in the language of LP-type problems. Familiarity with LP-type problems is not really necessary for understanding the present paper, though, since in Section 3 we reformulate the results in a combinatorial language, speaking about a poset partitioning problem, and we prove them in that form.

LP-type problems and their dimension. Formally, an *LP-type problem* is a pair (H, w), where H is a finite set of *constraints* and $w\colon 2^H \to \mathbb{R}\cup\{-\infty\}$ is a function, obeying the two axioms below, that for every subset $G \subseteq H$

specifies its *value* $w(G) \in \mathbb{R} \cup \{-\infty\}$. Intuitively, $w(G)$ is the minimum value of a solution that satisfies all constraints in G. The value can be a real number or a special element $-\infty$, which is considered to be smaller than all real numbers. The axioms are:

Monotonicity:	For all $F \subseteq G \subseteq H$ we have $w(F) \leq w(G)$.
Locality:	For all $F \subseteq G \subseteq H$ and all $h \in H$ with $w(F) = w(G) = w(F \cup \{h\}) \neq -\infty$ we have $w(G \cup \{h\}) = w(G)$.

We say that (H, w) is an *LP-type problem without* $-\infty$ if $w(G) \neq -\infty$ for all $G \subseteq H$. (We note that for sets with value $-\infty$ we permit an "exception to locality" in the definition, and so in LP-type problems without $-\infty$ locality must hold without exceptions.)

Let $G \subseteq H$. A *basis of* G is any inclusion-minimal subset $B \subseteq G$ with $w(B) = w(G)$. A set $B \subseteq H$ is a *basis* if it is a basis of some $G \subseteq H$, and the *combinatorial dimension* $\dim(H, w)$ of an LP-type problem (H, w) is the maximum cardinality of a basis.

An example of a useful LP-type problem (one without $-\infty$) is the problem of computing the smallest disk containing a given planar point set. Here H is a finite point set in $\mathbb{R}^2$ and $w(G)$ is the radius of the smallest circular disk that encloses all points of G. The combinatorial dimension is typically 3, but for some particular choices of H it can also be 2, 1, or 0.

Nondegenerate refinement. A reasonable notion of degeneracy in LP-type problems is an "overdetermined" solution, that is, a set with several distinct bases. Following [Mat95], we call an LP-type problem (H, w) *nondegenerate* if $w(B_1) \neq w(B_2)$ for any two distinct bases B_1, B_2. Consequently, in a nondegenerate LP-type problem, every $G \subseteq H$ has exactly one basis.

For removing degeneracies, we want to break the ties $w(B_1) = w(B_2)$ by slightly modifying the values of w while retaining all strict inequalities among the original values. We say that an LP-type problem (H, w') is a *refinement* of an LP-type problem (H, w) on the same set of constraints if for all $F, G \subseteq H$ with $w(F) < w(G)$ we have $w'(F) < w'(G)$. To *remove degeneracies* of an LP-type problem means to find a nondegenerate refinement of it.

In order to produce a nondegenerate refinement, we sometimes need to introduce new bases, and sometimes even larger ones than those in the original LP-type problem. The following LP-type problem $(H_{\mathrm{sq}}, w_{\mathrm{sq}})$, which we call the *square example*, was introduced in [Mat95]: We have $H_{\mathrm{sq}} =$

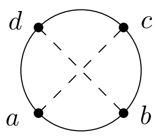

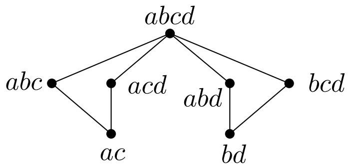

Figure 1: The square example (left) and its top poset (right).

$\{a, b, c, d\}$, where $a, b, c, d \in \mathbb{R}^2$ are the vertices of a square (Fig. 1 left), and for $G \subseteq H_{\rm sq}$, $w_{\rm sq}(G)$ is the radius of the smallest enclosing disk of G. The set $H_{\rm sq}$ has two bases $B_1 = \{a, c\}$ and $B_2 = \{b, d\}$, and $\dim(H_{\rm sq}, w_{\rm sq}) = 2$. However, it is not hard to check that any nondegenerate refinement has combinatorial dimension at least 3.

Top posets. Let (H, w) be an LP-type problem. The *top poset* of (H, w) is the set system

$$\mathrm{Topp}(H, w) := \{G \subseteq H : w(G) = w(H)\}$$

ordered by inclusion. For example, Fig. 1 right shows the Hasse diagram of the top poset for the square example. As will be stated precisely in the next section, a necessary condition for the existence of a nondegenerate refinement of (H, w) of combinatorial dimension at most d is the existence of a certain partition of the top poset.

Joins. In [MŠ07] we have defined the operation of *join* of two LP-type problems (H_1, w) and (H_2, w_2), resulting in another LP-type problem $(H, w) = (H_1, w_1) * (H_2, w_2)$, where H is the disjoint union of H_1 and H_2 (in particular, if $H_1 = H_2$, we must take two "distinct copies" of H_1), and for $G \subseteq H$ we define $w(G) = w_1(G_1) + w_2(G_2)$, where G_i is the part of G contained in H_i, $i = 1, 2$. We have $\dim(H, w) = \dim(H_1, w_1) + \dim(H_2, w_2)$.

The results. In [MŠ07], we proved the following:

Theorem 2.1 ([MŠ07]) *For every m, the m-fold join of $(H_{\rm sq}, w_{\rm sq})$ is an LP-type problem without $-\infty$ of combinatorial dimension $2m$ such that any nondegenerate refinement has combinatorial dimension at least $(2 + \varepsilon)m$, where ε is a positive constant (independent of m).*

Here we give an alternative short proof of Theorem 2.1, which moreover

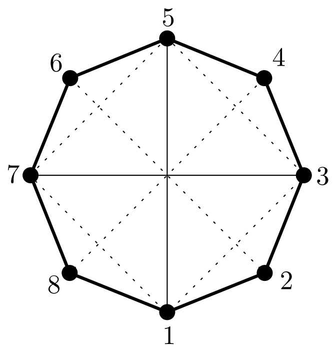

Figure 2: The octagon example.

yields the best possible value of ε, namely, $\varepsilon = 1$ (since there indeed is a nondegenerate refinement of combinatorial dimension $3m$).

Next, we define an LP-type problem $(H_{\mathrm{oct}}, w_{\mathrm{oct}})$, the *octagon exapmle*, of combinatorial dimension 2; see Fig. 2. There are 8 constraints marked by full circles, and weights of subsets $G \subseteq H_{\mathrm{oct}}$ are defined as follows:

- $\emptyset$ has weight -1 and every one-point set has weight 0.
- The bold edges (2-element sets) have weight 4 (the maximum weight), the light solid edges weight 3, the dashed edges weight 2, and all other 2-point sets weight 1.
- The weight of any set of G more than 2 elements is given as the maximum weight of a 2-element subset of G.

This problem was found by a computer search (using a program developed by P. Škovroň). It is clear that the weight function thus defined satisfies monotonicity, and also that the dimension is 2. Checking locality is more demanding, but it can still be done by hand (first one can check that if the locality axiom is violated by some F, G, and h, then it is also violated by F, G, h with $|F| \leq 2$ and $|G| \leq 3$, and then one can distinguish a small number of cases, using the symmetry of the problem).[1]

[1] Interestingly, it can be shown that if we want a 2-dimensional LP-type problem where the maximum-weight bases form exactly the edges of the 8-cycle, there is no way of setting the weights of the remaining pairs so that they are invariant under all symetries of the 8-cycle.

Theorem 2.2 *For every m, the m-fold join of $(H_{\rm oct}, w_{\rm oct})$ is an LP-type problem without $-\infty$ of combinatorial dimension $2m$ such that any nondegenerate refinement has combinatorial dimension at least $4m$.*

Finally, we turn to LP-type problems with $-\infty$. We set $H_t = \{a_i, b_i : i = 1, 2, \ldots, t\}$, and for $G \subseteq H_t$ we define $w_t(G) = 1$ if G contains at least one of the pairs $\{a_i, b_i\}$, $i = 1, 2, \ldots, t$, and $w_t(G) = -\infty$ otherwise. Then (H_t, w_t) is clearly a 2-dimensional LP-type problem with $-\infty$.

Theorem 2.3 *For every positive integer t, any nondegenerate refinement of the 2-dimensional LP-type problem (H_t, w_t) has combinatorial dimension larger than t.*

It seems very likely that this theorem can also be proved by the method of [MŠ07]; the calculations should even be easier than those in [MŠ07], and similar to the situation in Theorem 2.1, they would also yield nonexistence of nondegenerate refinement of bounded dimension in the realm of violator spaces.

A natural question is, can we get an analogue of Theorem 2.3 for LP-type problems *without* $-\infty$ (whose structure is considerably more restricted)? Or more generally, is there a function f such that every LP-type problem without $-\infty$ of combinatorial dimension d has a nondegenerate refinement of combinatorial dimension at most $f(d)$? So far this remains open and Theorem 2.2 is the strongest result in this direction.

3 Poset partitioning problems

Cube partitions. Let H be a fixed ground set. For sets $B \subseteq C \subseteq H$, we define the *cube* $[B, C]$ as the set system $\{G : B \subseteq G \subseteq C\}$ ordered by inclusion. It is isomorphic to the poset $(2^{C \setminus B}, \subseteq)$, i.e., to a Boolean algebra with $|C \setminus B|$ atoms. The *bottom dimension* of the cube $[B, C]$ is defined to be $|B|$.

Let $\mathcal{P} \subseteq 2^H$ be a set system. A *cube partition* of $\mathcal{P}$ is a partition of $\mathcal{P}$ into disjoint cubes, and the *bottom dimension* of such a cube partition is the maximum of the bottom dimensions of the cubes involved in the partition.

As shown in [MŠ07], if an LP-type problem (H, w) has a nondegenerate refinement (H, w') of combinatorial dimension at most d, then the top poset $\mathrm{Topp}(H, w)$ has a cube partition of bottom dimension at most d. Each of the cubes in the partition has the form $w'^{-1}(x)$ for some $x \in \mathbb{R}$; that is, it consists of all $G \subseteq H$ with some given w'-value.

An additional property of these cube partitions, which was not used in [MŠ07] but which we will need here, is *acyclicity*. For two cubes $[B, C]$ and $[B', C']$, let us define $[B, C] \preceq [B', C']$ if $B \subseteq C'$ (equivalently, if $G \subseteq G'$ for some $G \in [B, C]$ and $G' \in [B', C']$). A cube partition of some $\mathcal{P}$ is called *acyclic* if the relation $\preceq$ on its cubes is acyclic. The cube partition of $\mathrm{Topp}(H, w)$ induced by a nondegenerate refinement (H, w') is obviously acyclic (since the w'-values are linearly ordered).

We also need to see how the join of LP-type problem is reflected in the top posets. First we introduce the following general definition: Let V, W be disjoint sets and let $\mathcal{P} \subseteq 2^V$ and $\mathcal{Q} \subseteq 2^W$ be set systems. The *join* of $\mathcal{P}$ and $\mathcal{Q}$ is $\mathcal{P} * \mathcal{Q} = \{P \cup Q : P \in \mathcal{P}, Q \in \mathcal{Q}\}$. If we want to consider the join $\mathcal{P} * \mathcal{Q}$ for $\mathcal{P} = \mathcal{Q}$ (or in other situations where the ground sets are not disjoint), we first form isomorphic copies of $\mathcal{P}$ and $\mathcal{Q}$ with disjoint ground sets, and then we use the previous definition. Now if (H_1, w_1) and (H_2, w_2) are LP-type problems, we have $\mathrm{Topp}((H_1, w_1) * (H_2, w_2)) = \mathrm{Topp}(H_1, w_1) * \mathrm{Topp}(H_2, w_2)$.

Poset versions of Theorems 2.1–2.3. Let $\mathcal{S}_m$ denote the top poset of the join of m copies of the square example as in Theorem 2.1. Explicitly, we can describe $\mathcal{S}_m$ as follows: Let us write the $4m$-element set H of constraints as $\{a_i, b_i, c_i, d_i : i = 1, 2, \ldots, m\}$; then a set $G \subseteq H$ belongs to $\mathcal{S}_m$ if, for every $i = 1, 2, \ldots, m$, it contains at least one of the pairs $\{a_i, c_i\}$ and $\{b_i, d_i\}$.

Similarly, for the top poset $\mathcal{O}_m$ of the join of m copies of the octagon example, we write $H = \{v_{ij} : j = 1, 2, \ldots, 8, i = 1, 2, \ldots, m\}$, and $G \subseteq H$ belongs to $\mathcal{O}_m$ if it contains, for every $i = 1, 2, \ldots, m$, at least one of the edges of the 8-cycle on $\{v_{i1}, \ldots, v_{i8}\}$.

Finally, for the top poset $\mathcal{M}_t$ of the LP-type problem from Theorem 2.3, with $H_t = \{a_i, b_i : i = 1, 2, \ldots, t\}$, a set $G \subseteq H_t$ belongs to $\mathcal{M}_m$ if it contains at least one of the pairs $\{a_i, b_i\}$.

Theorem 3.1

(i) (A combinatorial version of Theorem 2.1) *The poset $\mathcal{S}_m$ has no acyclic cube partition of bottom dimension smaller than $3m$.*

(ii) (A combinatorial version of Theorem 2.2) *The poset $\mathcal{O}_m$ has no acyclic cube partition of bottom dimension smaller than $4m$.*

(iii) (A combinatorial version of Theorem 2.3) *The poset $\mathcal{M}_t$ has no acyclic cube partition of bottom dimension t or smaller.*

The case $m = 2$ in (i) was proved by Škovroň [Ško06] by case analysis. In [MŠ07], we have shown that $\mathcal{S}_m$ has no cube partition of bottom dimension smaller than $(2 + \varepsilon)m$, without requiring the partition to be acyclic. The acyclicity requirement may indeed make a difference: While $\mathcal{S}_2$ does have a cube partition of bottom dimension 5 (shown in [MŠ07]), it has no acyclic one.

4 Topological conditions for acyclic cube partitions

Here we formulate a necessary condition for the existence of an acyclic cube partition of a given bottom dimension for a given poset in terms of topological connectivity of an associated simplicial complex.

We will recall most of the necessary topological definitions and facts, some of them here and others in the course of the arguments. However, our presentation is not fully self-contained in the sense that we don't elaborate on some claims that are standard exercises in elementary algebraic topology, but which may be harder to establish without any topological background. For such background we refer, e.g., to the reference work [Bjö95] written for combinatorialists, or to standard textbooks like [Hat01] or [Mun84].

Topological preliminaries. Let V be a finite ground set. A simplicial complex with vertex set V is a hereditary family $\mathsf{K} \subseteq 2^V$ (that is, $\sigma \in \mathsf{K}$ and $\sigma' \subset \sigma$ implies $\sigma' \in \mathsf{K}$). The sets of K are called *simplices* or *faces* of K, and the *dimension* of a face $\sigma \in \mathsf{K}$ is $|\sigma| - 1$.

The simplicial complex K is assigned a topological space $\|\mathsf{K}\|$, the *geometric realization* of K, in a canonical way (uniquely up to homeomorphism), in which each k-dimensional face $\sigma \in \mathsf{K}$ corresponds to a k-dimensional geometric simplex contained in the space $\|\mathsf{K}\|$.

The *join* $\mathsf{K} * \mathsf{L}$ of simplicial complexes K and L is a particular case of our definition of the join of set systems in the previous section.

A topological space X is *k-connected* if every continuous map $f\colon S^j \to X$ can be extended to a continuous map $\bar{f} : B^{j+1} \to X$ for all $j \leq k$, where B^{j+1} denotes the unit ball in $\mathbb{R}^{j+1}$ and S^j denotes its boundary sphere. A simplicial complex K is said to be k-connected if its geometric realization $\|\mathsf{K}\|$ is k-connected. We will use the fact that S^k is not k-connected, which is a simple consequence of the famous Borsuk–Ulam antipodality theorem.

We will need a few things about homology and cohomology, but it will be enough to talk about *Betti numbers*: The kth Betti number $\beta_k(X)$ is the rank of the k-dimensional reduced homology group of X (with coefficients in $\mathbb{Z}$). By a theorem of Hurwitz, a k-connected simplicial complex K has $\beta_i(\mathsf{K}) = 0$ for all $i \leq k$. The Künneth formula for the homology of the join implies that $\beta_k(\mathsf{K} * \mathsf{L}) = \sum_{i+j=k-1} \beta_i(\mathsf{K})\beta_j(\mathsf{L})$, and consequently, if the join $\mathsf{K} * \mathsf{K} * \cdots * \mathsf{K}$ of m copies of K is $m(k+1) - 1$-connected, then $\beta_i(\mathsf{K}) = 0$ for all $i \leq k$.

We also recall that topological spaces X and Y are *homotopy equivalent* if there exist continuous maps $f\colon X \to Y$ and $g : Y \to X$ such that the composed map $f \circ g$ is homotopic to the identity map of Y, and $g \circ f$ is homotopic to the identity map of X. Here two continuous maps $h_0, h_1\colon S \to T$ of topological spaces are *homotopic* if there exists a continuous $H\colon S \times [0,1] \to T$ such that $h_0(x) = H(x, 0)$ and $h_1(x) = H(x, 1)$ for all $x \in S$. Homotopy equivalence preserves k-connectivity.

Proposition 4.1 (Necessary conditions for acyclic cube partitions)

(i) *Let H be an n-point set and let $\mathcal{P} \subseteq 2^H$ be an up-set in the poset $(2^H, \subseteq)$; i.e., $S \in \mathcal{P}$ and $S' \supseteq S$ implies $S' \in \mathcal{P}$. If $\mathcal{P}$ has an acyclic cube partition of bottom dimension d, then the simplicial complex $\mathcal{P}^c := \{H \setminus S : S \in \mathcal{P}\}$ on vertex set H is $(n-d-2)$-connected.*

(ii) *Let H and $\mathcal{P}$ be as in (i), and let $\mathcal{P}^{*m}$ be the join of m copies of $\mathcal{P}$. If $\mathcal{P}^{*m}$ has an acyclic cube parititon of bottom dimension at most $mk-1$, k integral, then $\beta_i(\mathcal{P}^c) = 0$ for all $i \leq n-k-1$.*

For 2-dimensional LP-type problems, nondegenerate refinements can also be related to the independence complex of a graph, a notion investigated in a number of papers (e.g., [Koz99], [Mes01], [Mes03], [EH06]). We recall that the *independence complex* $\mathsf{I}(G)$ of a graph $G = (V, E)$ has vertex set V and its simplices are all independent sets in G.

Corollary 4.2 *Let H be an n-point set and let $\mathcal{P} \subseteq 2^H$ be an up-set such that every inclusion-minimal set in $\mathcal{P}$ has size 2. Let G be the graph with vertex set H and edge set $\{e \in \mathcal{P} : |e| = 2\}$. If $\mathcal{P}^{*m}$ has acyclic cube partition of bottom dimension $mk-1$, k integral, then $\beta_i(\mathsf{I}(G)) = 0$ for all $i \geq k-2$.*

Now we start working towards the proof of Proposition 4.1.

($\geq d$)-collapsibility. Wegner [Weg75] introduced the notion of the *d-collapsibility* of a simplicial complex, in connection with investigating intersection patterns of convex sets in $\mathbb{R}^d$. We use a variation of his definition, where the parameter d plays a different role:

Definition 4.3 *Let d be a given natural number, let K be a simplicial complex, and let $\sigma \in \mathsf{K}$ be a face contained in a unique inclusion-maximal face $\tau \in \mathsf{K}$. Let us also assume that $|\tau| \geq d$. The* elementary ($\geq d$)-collapse *of σ transforms K into the simplicial complex $\mathsf{K} \setminus \{\sigma' \in \mathsf{K} : \sigma \subseteq \sigma' \subseteq \tau\}$. We say that K is ($\geq d$)-collapsible if it can be transformed into the simplicial complex $\emptyset$ by a sequence of elementary ($\geq d$)-collapses.*

The next lemma is the first main step in the proof of Proposition 4.1.

Lemma 4.4 *If K is ($\geq d$)-collapsible, then it is $(d-2)$-connected.*

Proof. It suffices to check that if K' arises from K by an elementary ($\geq d$)-collapse and it is $(d-2)$-connected, then K is $(d-2)$-connected as well.

Let σ and τ be as in Definition 4.3. First, if $\sigma = \tau$, then K is obtained from K' by adding a new face of dimension at least $d-1$, and it is well known that such an operation preserves $(d-2)$-connectivity.

Second, if σ is a proper face of τ, then the elementary ($\geq d$)-collapse of σ is an elementary collapse in the sense of Whitehead (see, e.g., [Bjö95]), and hence K and K' are homotopy equivalent. The lemma is proved. □

The next lemma is the second main part of the proof of Proposition 4.1

Lemma 4.5 *Let H be an n-point set and let $\mathcal{P} \subseteq 2^H$ be an up-set. Then the following are equivalent:*

(i) *$\mathcal{P}$ has an acyclic cube partition of bottom dimension at most d.*

(ii) *The simplicial complex $\mathsf{K} = \mathcal{P}^c$ is ($\geq n-d$)-collapsible.*

Proof. First let us assume (i), and let $[B_1, C_1], [B_2, C_2], \ldots, [B_k, C_k]$ be a linear ordering of the cubes in some acyclic cube partition of $\mathcal{P}$ that extends the (acyclic) relation $\preceq$. In particular, for $i < j$ we have $[B_j, C_j] \not\preceq [B_i, C_i]$, and hence no set of $[B_i, C_i]$ contains any set of $[B_j, C_j]$. It follows that for every i, $\mathcal{P}_i := \bigcup_{j \geq i} [B_j, C_j]$ is an up-set, and hence $\mathsf{K}_i := \{H \setminus G : G \in \mathcal{P}_i\}$ is a simplicial complex.

We now claim that K_{i+1} is obtained from K_i by an elementary $(\geq n-d)$-collapse of the face $\sigma_i := H \setminus C_i$. The faces of K_i containing σ_i are the complements of all sets in $\mathcal{P}_i$ contained in C_i, and latter are precisely the sets in the cube $[B_i, C_i]$. In particular, $\tau_i := H \setminus B_i$ is the unique inclusion-maximal face of K_i containing σ_i. Moreover, since the considered cube partition has bottom dimension at most d, we have $|\tau_i| = n - |B_i| \geq n - d$. Thus K_{i+1} indeed arises from K_i by an elementary $(\geq n-d)$-collapse of σ_i, and (ii) follows.

The implication (ii) $\Rightarrow$ (i), which we actually don't need for the purposes of this paper, is checked very similarly. $\square$

Proof of Proposition 4.1. Part (i) is an immediate consequence of Lemmas 4.4 and 4.5. Part (ii) then follows from (i) and the fact about the connectivity of joins mentioned in the topological preliminaries. $\square$

Proof of Corollary 4.2. We recall that if K is a simplicial complex with vertex set V, the *(combinatorial) Alexander dual* of K is the simplicial complex $\mathsf{K}^* = \{V \setminus S : S \notin \mathsf{K}\}$. The combinatorial Alexander duality, see, e.g., [Bjö95], implies $\beta_i(\mathsf{K}^*) = \beta_{n-3-i}(\mathsf{K})$ for all $i = -1, 0, 1, \ldots, n-2$ (with the convention that $\beta_{-1}(\mathsf{K}) = 0$ for any nonempty K but $\beta_{-1}(\emptyset) = 1$). It remains to observe that, in the situation of the corollary, the Alexander dual of the simplicial complex $\mathcal{P}^c$ is exactly the independence complex of G. $\square$

5 Proof of Theorem 3.1

Part (i) of the theorem follows from Proposition 4.1(ii) with $n = 4$ and $k = 3$, since if $\mathcal{P} = \mathcal{S}_1$ is the top poset of the square example, the simplicial complex $\mathcal{P}^c$ consists of two disjoint edges, and thus $\beta_0(\mathcal{P}^c) \neq 0$.

In part (ii), we can use Proposition 4.1(ii) again, with $n = 8$ and $k = 4$, and directly calculate that $\beta_3(\mathcal{O}_1^c) \neq 0$. Alternatively, we can use a result of Kozlov [Koz99], stating that the independence complex of a cycle of length $3q-1$ is homotopy equivalent to an S^{q-1}. So in our case the independence complex is an S^2 with a nonzero β_2, and Corollary 4.2 also shows that no acyclic cube partition of $\mathcal{O}_m$ of bottom dimension smaller than $4m$ exists.

In part (iii) we can again apply Corollary 4.2 conveniently. The inlcusion-minimal sets of $\mathcal{M}_t$ are t disjoint pairs, and so the graph G in the corollary is a matching of t edges. The independence complex is thus the boundary of the t-dimensional crosspolytope, which is homeomorphic to an S^{t-1}, and so $\beta_{t-1}(\mathsf{I}(G)) = 1$. Corollary 4.2, with $m = 1$, yields that $\mathcal{M}_t$ has no acyclic cube partition of bottom dimension t or smaller. (This also gives an alternative proof of part (i), since $\mathcal{S}_1$ is isomorphic to $\mathcal{M}_2$.)

Alternatively, it is not hard to show directly that the simplicial complex $\mathcal{M}_t^c$ is homotopy equivalent to an S^{t-2} and use Proposition 4.1(i). □

References

[Bjö95] A. Björner. Topological methods. In R. L. Graham, M. Grötschel, and L. Lovász, editors, *Handbook of Combinatorics*, chapter 34, pages 1819–1872. North-Holland, Amsterdam, 1995.

[EH06] R. Ehrenborg and G. Hetyei. The topology of the independence complex. *Eur. J. Comb.*, 27(6):906–923, 2006.

[GMRŠ07] B. Gärtner, J. Matoušek, L. Rüst, and P. Škovroň. Violator spaces: Structure and algorithms. In *Proc. 14th Annual European Symposium on Algorithms (ESA)*, pages 387–398, 2007. Full version in arXiv.org e-Print archive cs.DM/0606087.

[Hat01] A. Hatcher. *Algebraic Topology*. Cambridge University Press, Cambridge, 2001. Electronic version available at `http://math.cornell.edu/~hatcher#AT1`.

[Koz99] D. Kozlov. Complexes of directed trees. *J. Comb. Theory, Ser. A*, 88(1):112–122, 1999.

[Mat95] J. Matoušek. On geometric optimization with few violated constraints. *Discrete Comput. Geom.*, 14:365–384, 1995.

[Mes01] R. Meshulam. The clique complex and hypergraph matching. *Combinatorica*, 21(1):89–94, 2001.

[Mes03] R. Meshulam. Domination numbers and homology. *J. Comb. Theory, Ser. A*, 102(2):321–330, 2003.

[MŠ07] J. Matoušek and P. Škovroň. Removing degeneracy may require a large dimension increase. *Theory of Computing*, 3:159–177, 2007.

[Mun84] J. R. Munkres. *Elements of Algebraic Topology*. Addison-Wesley, Reading, MA, 1984.

[Ško06] P. Škovroň. Removing degeneracies in LP-type problems may need to increase dimension. In Jana Šafránková and Jiří Pavlů, editors, *Proc. 15th Week of Doctoral Students (WDS)*, pages I:196–207. Matfyzpress, Prague, 2006.

[SW92] M. Sharir and E. Welzl. A combinatorial bound for linear programming and related problems. In *Proc. 9th Sympos. Theoret. Aspects Comput. Sci.*, volume 577 of *Lecture Notes Comput. Sci.*, pages 569–579. Springer-Verlag, 1992.

[Weg75] G. Wegner. d-collapsing and nerves of families of convex sets. *Arch. Math.*, 26:317–321, 1975.