Estimates for homological dimension of configuration spaces of graphs.

Jacek Świątkowski *

Abstract. We show that the homological dimension of a configuration space of a graph Γ is estimated from above by the number b of vertices in Γ whose valence is greater than 2. We show that this estimate is optimal for the n-point configuration space of Γ if $n \geq 2b$.

0. Introduction.

Let Γ be a finite graph and n a natural number. The *marked n-point configuration space of* Γ is a subspace $\widetilde{C_n\Gamma}$ in the n-th cartesian power of Γ defined by

$$\widetilde{C_n\Gamma} := \{(x_1, \ldots, x_n) \in \Gamma^n : x_i \neq x_j \text{ for } i \neq j\}.$$

Consider the natural free action of the symmetric group S_n on the space $\widetilde{C_n\Gamma}$ defined by $\sigma(x_1, \ldots, x_n) = (x_{\sigma(1)}, \ldots, x_{\sigma(n)})$ and put

$$C_n\Gamma := \widetilde{C_n\Gamma}/S_n.$$

The space $C_n\Gamma$ is called the (unmarked) *n-point configuration space of* Γ.

This paper reports a partial progress in the general task to understand the homology of configuration spaces of graphs, or even more general of compact polyhedra. The other recent result in that direction known to the author is [G].

We call a vertex v of Γ *branched*, if it is adjacent to at least three edges. We denote by $b = b(\Gamma)$ the number of branched vertices in Γ.

The main result of this paper is the following.

0.1. Theorem. Let Γ be a finite graph and n a natural number.
(1) There exists a cube complex $K_n\Gamma$ of dimension $\min(b(\Gamma), n)$ which embeds as a deformation retract into the configuration space $C_n\Gamma$.
(2) The fundamental group $\pi_1(C_n\Gamma)$ contains a subgroup isomorphic to the free abelian group $\mathbf{Z}^k$ with $k = \min(b(\Gamma), [n/2])$, where $[x]$ denotes the integer part of x.

For a topological space X define the *homological dimension* $\dim_h X$ of X by

$$\dim_h X := \max\{i : H_i(X, \Lambda) \neq 0 \text{ for some abelian group } \Lambda\}.$$

* The author was supported by the Polish State Commitee for Scietific Research (KBN) grant 2 P03A 023 14.

Note that if X is homotopically equivalent to a polyhedral complex K then $\dim_h X \leq \dim K$. On the other hand, if X is a $k(\pi, 1)$ space then $\dim_h X = homdim(G)$ for any subgroup G of the fundamental group $\pi_1 X$, where $homdim$ is the homological dimension of a group (see [CE], Ch. XIV.9, Application 1, p. 356). Configuration spaces of graphs are the $k(\pi, 1)$ spaces, as it was observed by M. Davis and T. Januszkiewicz. This fact follows also from the results of this paper (see Corollary 2.3.2). On the other hand, it is known that $homdim(\mathbf{Z}^k) = k$, and hence Theorem 0.1 implies the following.

0.2. Corollary.
(1) $\dim_h C_n\Gamma \leq \min(b(\Gamma), n)$.
(2) $\dim_h C_n\Gamma \geq \min(b(\Gamma), [n/2])$.

Note that if $n \geq 2b(\Gamma)$ then the estimates in above Corollary give the equality $\dim_h C_n\Gamma = b(\Gamma)$.

The main ingredient in proving the results of this paper is the construction of a cube complex $K_n\Gamma$ and its embedding into the configuration space $C_n\Gamma$ as a deformation retract. This is done in Section 1. The further study of the complex $K_n\Gamma$ exploits its natural geometry. In particular, it turns out that $K_n\Gamma$ is a nonpositively curved metric space in the comparison sense of Alexandrov and Toponogov, and hence a $k(\pi, 1)$ space. Nonpositive curvature allows to look for subgroups in the fundamental group $\pi_1(K_n\Gamma)$ in terms of locally convex subspaces in $K_n\Gamma$. Existence of a family of such subspaces homeomorphic to tori proves part (2) of Theorem 0.1. These geometrical aspects of the complex $K_n\Gamma$ are studied in Section 2.

I would like to express my thanks to Tadek Januszkiewicz for his encouragement and inspiring discussions.

1. Complex $K_n\Gamma$ and its embedding in $C_n\Gamma$.

1.1. The natural combinatorial structure of a graph.

In the rest of this paper we allow that graphs contain loops or multiple edges. Given a finite graph Γ, its underlying space $|\Gamma|$ carries the unique combinatorial structure satisfying the property that no vertex separates locally $|\Gamma|$ into exactly two components. We call this structure the *natural combinatorial structure* on $|\Gamma|$. In what follows we assume that graphs are equipped with the natural combinatorial structure. This is clearly no loss of generality, since the configuration space depends in fact only on the underlying space of a graph.

Note that the vertices in a graph Γ (with the natural combinatorial structure) fall into two classes. Those which locally separate the underlying space $|\Gamma|$ into at least three components are called *branched*, while those which are locally non-separating are called *free*.

1.2. The cube complex $K_n\Gamma$.

Let $B = B_\Gamma$ and $E = E_\Gamma$ be respectively the set of all branched vertices and the set of all edges of a graph Γ (with respect to the natural combinatorial structure). Each edge in Γ (even if it is a loop) carries two distinct orientations. For an oriented edge s denote by $|s|$ the underlying unoriented edge, by $-s$ the same edge with the opposite orientation, and by v_s that of the vertices adjacent to $|s|$ which is determined by the orientation of s. Note that if s is an oriented loop in Γ then $s \neq -s$ as for all oriented edges, but $v_s = v_{-s}$.

Define an abstract graded poset $P_n\Gamma = (P_n^{(0)}\Gamma, \ldots, P_n^{(k)}\Gamma, \ldots)$, where $P_n^{(k)}\Gamma$ denotes the set of k-faces of $P_n\Gamma$, as follows. Let $P_n^{(k)}\Gamma$ be the set of all pairs (f, S) such that

(1) $f : E_\Gamma \cup B_\Gamma \to \mathbf{N} \cup \{0\}$ is a function;
(2) $S = \{s_1, \ldots, s_k\}$ is a set consisting of exactly k oriented edges of Γ;
(3) $v_{s_i} \in B_\Gamma$ for $i = 1, \ldots, k$ and $v_{s_i} \neq v_{s_j}$ for $i \neq j$;
(4) $f(b) \in \{0, 1\}$ for each $b \in B$ and $f(v_{s_i}) = 0$ for $i = 1, \ldots, k$;
(5) $\sum_{a \in B \cup E} f(a) = n - k$.

1.2.1. Remark. The above poset has the following intuitive interpretation. Its 0-faces correspond to all possible quantitative distributions of n distinct points among the branched vertices and the (remaining parts of) edges in the graph Γ. The value of the function f shows the number of points lying in the corresponding part of Γ (again, edges are considered without their branched endpoints). The 1-faces of the poset represent elementary "moves" between distributions. In each such "move" one point changes its position from a branched vertex to an edge, or in opposite direction. This moving point is represented by the unique oriented edge s in the set S. The orientation of s shows the direction in which the point moves along $|s|$ to approach the branched endpoint v_s. The distribution of the other points is represented by the function f. Finally, the k-faces with $k > 1$ represent systems of k independent "moves", where independent means that these moves can be performed in any order, without interfering with each other. These single "moves" are represented by the elements of the set S, while the distribution of the points not involved in moves is again represented by the function f.

Let (f_1, S) and $(f_2, S \cup \{s\})$ be two faces of $P_n\Gamma$, and assume that $s \notin S$. We say that $(f_1, S) \prec (f_2, S \cup \{s\})$ if one of the following two conditions holds:

(i) $f_1(a) = \begin{cases} f_2(a) + 1 & \text{if } a = |s|, \\ f_2(a) & \text{otherwise;} \end{cases}$

(ii) $f_1(a) = \begin{cases} f_2(a) + 1 & \text{if } a = v_s, \\ f_2(a) & \text{otherwise.} \end{cases}$

We extend the relation $\prec$ to the smallest partial order on the set of faces of $P_n\Gamma$, denote it also by $\prec$, and take $\prec$ as the relation of being a face in $P_n\Gamma$.

Observe that for each element $F \in P_n^{(k)}\Gamma$ the subposet $\{A : A \prec F\}$ in $P_n\Gamma$ is isomorphic to the face poset of the k-dimensional cube. It follows that $P_n\Gamma$ is a face poset of a uniquely determined cube complex, which we denote by $K_n\Gamma$.

1.2.2. Remark. The cube complex $K_n\Gamma$ defined above does not necessarily satisfy the property that its cells are determined by the sets of their vertices. In particular, $K_n\Gamma$ can contain multiple edges. However, the closure of each cell in $K_n\Gamma$ is isomorphic to the cube of the corresponding dimension, so that in particular no 1-cell of $K_n\Gamma$ is a loop. Moreover, each k-cell with $k > 1$ is uniquely determined by its boundary in $K_n\Gamma$.

1.2.3. Lemma. $\dim K_n\Gamma = \min(b(\Gamma), n)$.

Proof: It is clear that $\dim K_n\Gamma = \max\{k : P_n^{(k)}\Gamma \neq \emptyset\}$. It follows then from condition (5) above that $\dim K_n\Gamma \leq n$, and from condition (3) that $\dim K_n\Gamma \leq b(\Gamma)$. Hence we get $\dim K_n\Gamma \leq \min(b(\Gamma), n)$. The opposite inequality follows from the obvious existence of a pair (f, S) as required with $card(S) = \min(b(\Gamma), n)$.

1.3. Embedding of $K_n\Gamma$ into $C_n\Gamma$.

In this subsection we construct an embedding $i : K_n\Gamma \to C_n\Gamma$ of (the underlying space of) the cube complex $K_n\Gamma$ into the configuration space $C_n\Gamma$. This construction explains more precisely than Remark 1.2.1 the ideas behind the definitions of the poset $P_n\Gamma$ and the complex $K_n\Gamma$.

Consider a segment AB, and let s be the oriented segment underlied by AB, with $v_s = B$ and $v_{-s} = A$. Given $n \in \mathbf{N} \cup \{0\}$ and real numbers $t_s, t_{-s} \in [0, 1]$, let $D_{AB}(n, (t_s)_{|s|=AB})$ be the set $\{A_1, \ldots, A_n\}$ of n points of AB satisfying the following properties:

(1) $A \leq A_1 < A_2 < \ldots < A_n \leq B$ with respect to the natural order on AB;
(2) $|A_1A_2| = |A_2A_3| = \ldots = |A_{n-1}A_n|$, where $|XY|$ denotes the the length of the segment XY;
(3) if $n \geq 2$ then $|AA_1| = t_{-s} \cdot |A_1A_2|$ and $|A_nB| = t_s \cdot |A_1A_2|$;
(4) if $n = 1$ and $(t_s, t_{-s}) \neq (0, 0)$ then $t_s \cdot |AA_1| = t_{-s} \cdot |A_1B|$;
(5) if $n = 1$ and $(t_s, t_{-s}) = (0, 0)$ then the set $D_{AB}(n, (t_s)_{|s|=AB})$ is not defined.

Note that if $n = 0$ then $D_{AB}(n, (t_s)_{|s|=AB}) = \emptyset$.

If e is an edge (of a graph) which is not a loop, define the set $D_e(n, (t_s)_{|s|=e})$ in an analogous way as for the segment AB. If e is a loop, identify e with the segment AB in which the endpoints A and B are identified. Define the set $D_e(n, (t_s)_{|s|=e})$ as the pull-back of the set $D_{AB}(n, (t_s)_{|s|=AB})$ by this identification. $D_e(n, (t_s)_{|s|=e})$ is then a well defined set of n elements of e, except when $n \geq 2$ and $(t_s, t_{-s}) = (0, 0)$, since then A_1 and A_n coincide.

Let $F = (f, S)$ be a cell of the complex $K_n\Gamma$. Consider the set $[0, 1]^S$ of all functions $t : S \to [0, 1]$, and view it as a cube of dimension equal to the cardinality of the set S. Let $\tau : F \to [0, 1]^S$ be the linear isomorphism of cubes such that for each vertex $p = (\psi, \emptyset)$ of F the function $\tau(p) \in [0, 1]^S$ is given by

$$\tau(p)(s) = 1 - \psi(v_s) \text{ for all } s \in S.$$

For each point $x \in F$ the image $\tau(x) : S \to [0, 1]$ is a function. Extend $\tau(x)$ to the function $\tau_0(x)$ defined on the set of all oriented edges in Γ by

$$\tau_0(x)(s) = \begin{cases} \tau(x)(s) & \text{if } s \in S \\ 1 & \text{if } s \notin S. \end{cases}$$

View the elements of $C_n\Gamma$ as subsets of Γ consisting of n elements, and define the mapping $i_F : F \to C_n\Gamma$ by

$$i_F(x) = \{b \in B_\Gamma : f(b) = 1\} \cup \bigcup_{e \in E_\Gamma} D_e(\tilde{f}(e), (\tau_0(x)(s))_{|s|=e}),$$

where $\tilde{f}(e) := f(e) + card\{s \in S : |s| = e\}$. A straightforward verification shows that the mapping i_F is well defined, continuous and injective. Moreover, if F_1 is a face of F_2 in $K_n\Gamma$ then $i_{F_1} = i_{F_2}|_{F_1}$. This means that the family $i_F : F \in P_n\Gamma$ of mappings defines a continuous mapping $i : K_n\Gamma \to C_n\Gamma$, which is easily verified to be an embedding. We omit further details.

1.4. Deformation retraction of $C_n\Gamma$ onto $K_n\Gamma$.

In this subsection we construct a retraction of the configuration space $C_n\Gamma$ onto the image $i(K_n\Gamma)$ of the complex $K_n\Gamma$ by the embedding i. We also show that this retraction is a deformation retraction. Together with Lemma 1.2.3 this completes the proof of part (1) of Theorem 0.1.

Recall that given a topological space X and its subspace Y, a map $r : X \to X$ is a *retraction of* X *onto* Y if $Y = \operatorname{im} r$ and $r|_Y = \mathrm{id}_Y$. It is called a deformation retraction if moreover r is homotopic to the identity map id_X by a homotopy $H : [0,1] \times X \to X$ such that $H(t,y) = y$ for each $t \in [0,1]$ and each $y \in Y$.

Fix the length metric d on Γ for which each edge of Γ has length 1. Let $C \in C_n\Gamma$ be an n-point configuration in Γ (which we view as a subset of Γ). Points of C (together with the vertices of Γ) subdivide Γ into a collection of segments. For each oriented edge s in Γ for which v_s is a branched vertex of Γ let d_s^C be the length of that segment of above subdivision which is contained in the edge $|s|$ and adjacent to v_s from the side determined by the orientation of s (note that if s is a loop then there might be two segments contained in $|s|$ and adjacent to v_s). For an edge e of Γ let e^0 denote the closed edge e with its branched vertices deleted. Put

$$n_e^C := card(C \cap e^0)$$

to be the number of points of C contained in e^0. The edge e is then subdivided into $n_e^C + 1$ segments, the average length of which is equal to $1/(n_e^C + 1)$. For each oriented edge s of Γ consider the ratio δ_s^C of the length d_s^C defined above and the average length $1/(n_{|s|^C} + 1)$, i.e.

$$\delta_s^C := d_s^C \cdot (n_{|s|}^C + 1).$$

Define the coefficient $t_s^C \in [0,1]$ by

$$t_s^C := \begin{cases} 1 & \text{if } v_s \text{ is a free vertex of } \Gamma \text{ or } v_s \in C \\ \min(1, \delta_s^C / \min\{\delta_{s'}^C : s' \neq s, v_{s'} = v_s\}) & \text{otherwise.} \end{cases}$$

Under the notation of the previous subsection put

$$r(C) := (C \cap B_\Gamma) \cup \bigcup_{e \in E_\Gamma} D_e(n_e^C, (t_s^C)_{|s|=e}).$$

A straightforward verification shows that the map $r : C_n\Gamma \to C_n\Gamma$ defined as above is continuous, its image coincides with the image of the embedding i, and for each configuration $C \in i(K_n\Gamma)$ we have $r(C) = C$. Hence r is a retraction as required.

To prove that r is a deformation retraction we need to construct a homotopy H between r and $\mathrm{id}_{C_n\Gamma}$ such that if $C \in i(K_n\Gamma)$ then $H(t,C) = C$ for each $t \in [0,1]$.

Note that, by definition of r, for each $C \in C_n\Gamma$ we have $C \cap B_\Gamma = r(C) \cap B_\Gamma$. Moreover, for each $e \in E_\Gamma$ we have $card(C \cap e^0) = card(r(C) \cap e^0)$. For each $e \in E_\Gamma$ let $C_e(t) : t \in [0,1]$ be the unique continuous 1-parameter family of configurations in e^0 which connects $C \cap e^0$ with $r(C) \cap e^0$ and for which positions of all points (measured in length parameter in e^0) depend linearly on t. Put

$$H(t,C) = (C \cap B_\Gamma) \cup \bigcup_{e \in E_\Gamma} C_e(t).$$

Then H is a homotopy between r and $\mathrm{id}_{C_n\Gamma}$ as required, and thus r is a deformation retraction.

2. Geometry of the complex $K_n\Gamma$ and its applications.

In this section we view the cube complex $K_n\Gamma$ as a metric space, and derive some geometric properties of it that have useful topological consequences. In particular, we prove part (2) of Theorem 0.1 and the fact that $K_n\Gamma$ is a $k(\pi,1)$ space. We make use of a well developed theory of metric spaces with nonpositive curvature, and our main reference for this subject is the book of M. Bridson and A. Haefliger [BH].

2.1. Nonpositively curved metric spaces.

Let (X,d) be a metric space. A *geodesic segment* in X is an isometric embedding $\sigma : [a,b] \to X$, i.e. a mapping such that for all $x,y \in [a,b]$ we have

$$d(\sigma(x),\sigma(y)) = |x-y|.$$

We say that a geodesic segment as above *connects* the endpoints $\sigma(a)$ and $\sigma(b)$. A metric space X is called *geodesic*, if its any two points can be connected by at least one geodesic segment.

A *geodesic triangle* in X is a collection of three geodesic segments pairwise connecting three points $p_1, p_2, p_3 \in X$. Given a geodesic triangle T in X, a *comparison triangle* for T is a geodesic triangle T' in the euclidean plane with the same side lengths as T. Note that the comparison triangle always exists and is unique up to an isometry of the plane.

A geodesic triangle T is said to *satisfy the CAT(0) inequality*, if the distance of any two points on its sides is not bigger than the distance of the corresponding points on the sides of the comparison triangle T'. A geodesic space X is a *CAT(0) space*, if each geodesic triangle in X satisfies the CAT(0) inequality. A metric space X is a *nonpositively curved metric space* if it is a CAT(0) space locally.

We briefly recall from [BH] some topological properties of nonpositively curved metric spaces.

2.1.1. Fact ([BH], Chapter II, Corollary 1.5, p. 161). Any CAT(0) space is contractible.

2.1.2. Fact ([BH], Chapter II, Theorem 4.1 (2), p. 194). If X is a complete connected nonpositively curved metric space, then its universal covering $\widetilde{X}$ (with the induced length metric) is a CAT(0) space.

These two results imply the following.

2.1.3. Corollary. Any complete connected nonpositively curved metric space is a $k(\pi, 1)$ space.

2.2. Nonpositively curved cube complexes.

A *cube complex* is a space obtained from a disjoint union of cubes (of any dimensions) by a family of face identifications. We assume furthermore that each of the initial cubes embeds (injectively) into the whole complex, or equivalently that no two faces of a cubical cell of the complex are identified. The complex $K_n\Gamma$ defined in Section 1 is a cube complex in this sense.

Given a cube complex K, equip each of its cells with the euclidean metric in which the cell is isometric to the unit cube (of the corresponding dimension). Extend this partial metric on K to the unique length metric, and recall that if K is connected and finite dimensional, then it is a complete geodesic space with respect to this metric ([BH], Chapter I, remark 7.33, p. 112).

Let v be a vertex of a cube complex K. The poset of all cells of K containing v and distinct from v is then isomorphic with the face poset of a uniquely determined simplicial complex K_v, called the *link* of K at v. A k-cell of K containing v corresponds by this isomorphism to a $(k-1)$-simplex of K_v.

A simplicial complex L is a *flag complex*, if every finite set of its vertices that is pairwise joined by edges spans a simplex of L.

2.2.1. Gromov's Lemma ([BH], Chapter II, Theorem 5.20, p. 212). A finite dimensional cube complex is a nonpositively curved metric space (with respect to the canonical metric mentioned above) if and only if the link at each of its vertices is a flag complex.

2.3. Nonpositive curvature of $K_n\Gamma$.

The crucial geometric property of the complex $K_n\Gamma$ is the following.

2.3.1. Proposition. The cube complex $K_n\Gamma$ is nonpositively curved.

Proof: In view of Gromov's Lemma it is sufficient to show that the link of $K_n\Gamma$ at each vertex is a flag complex. Let $x \in P_n^{(0)}\Gamma$ be a vertex of $K_n\Gamma$. Then $x = (\varphi, \emptyset)$, where the function loop $\varphi : E_\Gamma \cup B_\Gamma \to \mathbf{N} \cup \{0\}$ satisfies the following conditions:
(1) $\varphi(b) \in \{0, 1\}$ for each $b \in B$;
(2) $\sum_{a \in B \cup E} \varphi(a) = n$.
We will describe the link of the complex $K_n\Gamma$ at the vertex x.

We start with describing the set of cells of $K_n\Gamma$ which contain x. These are clearly the cells F for which $x \prec F$. Such cells are in one-to-one correspondence with the sets S of oriented edges in Γ which satisfy the following conditions:
(1) for each $s \in S$ the vertex v_s is branched;
(2) $v_{s_1} \neq v_{s_2}$ for any two distinct elements s_1, s_2 of S;

(3) if $s \in S$ then $\varphi(|s|) + \varphi(v_s) \geq 1$;
(4) if $s \in S$ and $-s \in S$ then $\varphi(|s|) + \varphi(v_s) + \varphi(v_{-s}) \geq 2$.
Namely, a set S as above corresponds to the cell (f, S), where the function $f : E \cup B \to \mathbf{N} \cup \{0\}$ is given by

$$f(a) = \begin{cases} \varphi(a) & \text{if } a \in B_\Gamma \setminus \{v_s : s \in S\} \\ 0 & \text{if } a \in \{v_s : s \in S\}, \\ \varphi(a) + \sum_{s \in S: |s| = a} (\varphi(v_s) - 1) & \text{if } a \in E_\Gamma. \end{cases}$$

From what was said above it is not difficult to realize that the subposet of cells of $K_n\Gamma$ which contain x (and which are distinct from x) coincides with the poset P_φ of the nonempty sets S as above with the order given by inclusion. Poset P_φ has the property that if $S \in P_\varphi$ and $\emptyset \neq R \subset S$ then $R \in P_\varphi$, and therefore it is the face poset of a simplicial complex K_φ the k-simplices of which correspond to sets $S \in P_\varphi$ with cardinality $k + 1$. But P_φ is also the face poset of the link of $K_n\Gamma$ at x, and thus this link is isomorphic to the simplicial complex K_φ.

The conditions (1)-(4) satisfied by the sets $S \in P_\varphi$ have the property that a set S of oriented edges in Γ satisfies them iff each subset of S with cardinality 1 or 2 satisfies them. This implies that the simplicial complex K_φ is a flag complex, and the lemma follows.

The above proposition, together with Corollary 2.1.3 and part (1) of Theorem 0.1, imply the following.

2.3.2. Corollary. Configuration spaces of graphs are $k(\pi, 1)$ spaces.

2.4. Locally isometric maps.

A map $f : Y \to X$ between metric spaces is *locally isometric* if each point $y \in Y$ has a neighbourhood $U \subset Y$ such that the restricted map $f|_U : U \to X$ is an isometry. The importance of this notion in our context is due to the following.

2.4.1. Fact ([BH], Chapter II, Proposition 4.14 (1), p. 201). Let X and Y be complete connected metric spaces, and suppose that X is nonpositively curved and Y is geodesic. Suppose also that $f : Y \to X$ is a localy isometric mapping. Then for every $y_0 \in Y$ the induced homomorphism of fundamental groups

$$f_* : \pi_1(Y, y_0) \to \pi_1(X, f(y_0))$$

is injective.

Suppose now that X and Y are cube complexes with the canonical piecewise Euclidean length metrics. Let $f : Y \to X$ be a nondegenerate combinatorial map. For each vertex v of Y the map f induces in a unique way a nondegenerate simplicial map $f_v : Y_v \to X_{f(v)}$ of the links at corresponding vertices.

2.4.2. Fact ([DJS], Proposition 1.7.1, p. 514). Let X and Y be finite dimensional cube complexes. A nondegenerate combinatorial map $f : Y \to X$ is a local isometry if and only if for each vertex v of Y
(1) the map $f_v : Y_v \to X_{f(v)}$ is an embedding;

(2) the image $f_v(Y_v)$ is a full subcomplex in $X_{f(v)}$, i.e. it contains each simplex of $X_{f(v)}$ all vertices of which are in $f_v(Y_v)$.

2.5. Essential tori in $K_n\Gamma$ and the proof of Theorem 0.1 (2).

In this subsection we construct locally isometric maps of k-dimensional tori into the complex $K_n\Gamma$, for $k \le \min(b(\Gamma), [n/2])$. Since, due to Fact 2.4.1, such maps are π_1-injective, their existence proves part (2) of Theorem 0.1.

Let Δ be a tree consisting of three edges meeting at a common vertex q. Then the configuration space C_2Q is homotopy equivalent to the circle S^1 and the complex K_2Q is isomorphic to the boundary of a hexagon. Let $(K_2Q)^k = K_2Q \times \ldots \times K_2Q$ be the cartesian product of k copies of the complex K_2Q. Then $(K_2Q)^k$ is a cube complex whose underlying space is homeomorhic to the k-dimensional torus. For each $k \le \min(b(\Gamma), [n/2])$ we will construct a locally convex mapping of the torus $(K_2Q)^k$ into the complex $K_n\Gamma$.

Fix the subset $\mathcal{B} = \{b_1, \ldots, b_k\}$ of the set B_Γ and a function $\psi : E_\Gamma \cup B_\Gamma \to \mathbf{N} \cup \{0\}$ such that $\psi(a) \in \{0,1\}$ for each $a \in B_\Gamma$, $\psi(b_i) = 0$ for $i = 1, \ldots, k$ and $\sum_{a \in E \cup B} \psi(a) = n - 2k$. For each $1 \le i \le k$ choose a combinatorial immersion $\delta_i : Q \to \Gamma$ for which $\delta_i(q) = b_i$. Recalling that faces of the torus $(K_2Q)^k$ correspond to tuples $((g_1, R_1), \ldots, (g_k, R_k))$ of faces of the complex K_2Q, put

$$k_{\mathcal{B},\psi}((g_1, R_1), \ldots, (g_k, R_k)) = (\psi + \sum_{i=1}^{k} (\delta_i)_* g_i, \sum_{i=1}^{k} \delta_i(R_i)),$$

where the functions $(\delta_i)_* g_i : E \cup B \to \mathbf{N} \cup \{0\}$ are defined by

$$(\delta_i)_* g_i(a) = \sum_{a' : \delta_i(a') = a} g_i(a').$$

It is easy to check that $k_{\mathcal{B},\psi}$ is a well defined nondegenerate combinatorial map of the torus $(K_2Q)^k$ into the complex $K_n\Gamma$.

2.5.1. Lemma. The map $k_{\mathcal{B},\psi}$ is locally isometric.

Proof: The proof is based on Fact 2.4.2. Let $((g_1, R_1), \ldots (g_k, R_k))$ be a tuple representing a cell F of the complex $(K_2Q)^k$. Note that each set R_i is either empty or contains exactly one element. The dimension of the cell F is then equal to the number of the sets R_i which are nonempty. Sets $\delta_i(R_i)$ are pairwise disjoint, since the oriented edges contained in them point at distinc vertices b_i of $\mathcal{B}$. It follows that if $j \ge 1$ then distinct j-cells of the torus $(K_2Q)^k$ are mapped by $k_{\mathcal{B},\psi}$ onto distinct j-cells of $K_n\Gamma$. In particular, the induced maps $(k_{\mathcal{B},\psi})_v$ for links are embeddings and thus satisfy condition (1) of Fact 2.4.2.

To check condition (2) of Fact 2.4.2, let $v = ((h_1, \emptyset), \ldots, (h_k, \emptyset))$ be a vertex of $(K_2Q)^k$ and let $\mathcal{E}$ be a set of edges of $(K_2Q)^k$ which contain v. It is sufficient to prove that if all the edges in the image set $k_{\mathcal{B},\psi}(\mathcal{E})$ are contained in a single cell of $K_n\Gamma$ then all the edges in $\mathcal{E}$ are contained in a single cell of $(K_2Q)^k$.

The edges in $(K_2Q)^k$ correspond to tuples $((g_i, R_1), \ldots, (g_k, R_k))$ with exactly one R_i nonempty. If the images of the edges from the set $\mathcal{E}$ by the map $k_{\mathcal{B},\psi}$ are all contained in a single cell of $K_n\Gamma$ then the corresponding nonempty sets R_i for those edges contain oriented edges s with distinct vertices v_s. This implies that for distinct edges in $\mathcal{E}$ the indices i of the corresponding nonempty sets R_i are distinct. We may assume then that $\mathcal{E} = \{\epsilon_1, \ldots, \epsilon_m\}$ and $\epsilon_i = ((g_{1i}, R_{1i}), \ldots, (g_{ki}, R_{ki}))$, where $R_{ii} \neq \emptyset$ for $i = 1, \ldots, m$.

Since all the edges ϵ_i are adjacent to the vertex v, the cell $((g'_1, R'_1), \ldots, (g'_k, R'_k))$ with

$$(g'_i, R'_i) = \begin{cases} (g_{ii}, R_{ii}) & \text{if } 1 \leq i \leq m \\ (h_i, \emptyset) & \text{if } i > m \end{cases}$$

is well defined and contains all the edges from the set $\mathcal{E}$.

In view of Fact 2.4.2, this finishes the proof.

Part (2) of Theorem 0.1 follows directly from Fact 2.4.1 and the above lemma.

References

[BH] M. Bridson, A. Haefliger, *Metric Spaces of Non-Positive Curvature,* Springer, 1999.

[CE] H. Cartan, S. Eilenberg, *Homological Algebra*, Princeton University Press, New Jersey, 1956.

[DJS] M. Davis, T. Januszkiewicz, R. Scott, *Nonpositive curvature of blow-ups,* Selecta Mathematica, New ser. 4 (1998), 491–547.

[G] S. R. Gal, *Euler characteristic of a configuration space of a complex,* this volume of Colloquium Mathematicum.

Instytut Matematyczny
Uniwersytet Wrocławski
pl. Grunwaldzki 2/4
51-642 Wrocław

Instytut Matematyczny PAN
ul. Śniadeckich 8
00-950 Warszawa