Growth series for vertex-regular CAT(0) cube complexes

RICHARD SCOTT

We show that the known formula for the growth series of a right-angled Coxeter group holds more generally for any CAT(0) cube complex whose vertex links all have the same f–polynomial.

20F55, 20F67, 20F69; 05A15

1 Introduction

A *cube complex* is a regular cell complex X all of whose cells are cubes and such that the intersection of any two cells is a face of both. An *edge-path* in a cube complex X is a sequence $e_1, \dots, e_n$ of oriented edges such that the head of e_i coincides with the tail of e_{i+1} for $1 \le i \le n-1$. The number of edges in an edge-path is called the *length* of the path. Given two vertices x, x' in a cube complex X, we define the distance $d(x, x')$ to be the minimum length (possibly infinite) of an edge-path connecting x to x'. For each vertex x_0 in X, we let $G(X, x_0; t)$ denote the corresponding growth series for X. That is,

$$G(X, x_0; t) = \sum_{i=0}^{\infty} \sigma(i) t^i$$

where $\sigma(i)$ is the number of vertices in X whose distance to x_0 is i.

The link of any vertex in a cube complex X is a simplicial complex, and by a result of Gromov, X is nonpositively curved with respect to the standard piecewise Euclidean metric if and only if every link is a flag complex. If, in addition, X is simply-connected then it is CAT(0) (see, for example, Bridson and Haefliger [1]). The main result of this article is the following.

Theorem 1 *Let X be a connected n–dimensional CAT(0) cube complex with the property that the link of every vertex has the same number of i–simplices for each $i \in \{0, \dots, n-1\}$. Let $f(t)$ be the polynomial $f(t) = f_{-1} + f_0 t + f_1 t^2 + \cdots + f_{n-1} t^n$ where $f_{-1} = 1$ and f_i is the number of i–simplices in the link of a vertex (for $0 \le i < n$). Then the growth series of X is independent of the point x_0 and given by the formula*

$$\frac{1}{G(X, x_0; t)} = f\left(\frac{-t}{1+t}\right).$$

Published: 29 March 2007 DOI: 10.2140/agt.2007.7.285

In the case where X is the cube complex associated to a right-angled Coxeter group W this formula is well-known (it is a special case of the known formula for the growth series of a Coxeter group relative to the standard generators, see Steinberg [7, Theorem 1.25 and Corollary 1.29]). More generally, Theorem 1 applies to any group acting on a CAT(0) cube complex X whose action on the vertex set is simply-transitive. Such groups were considered by Noskov [4] who proved that the geodesic words corresponding to edge-paths in X form a regular language and that the corresponding growth series is a rational function. Not only does Theorem 1 give an explicit formula for that rational function, but it applies to even more general CAT(0) cube complexes. In particular, X need not admit a group action, and the vertex links need not even be isomorphic; the only requirement for the vertex links is that they all have the same f–polynomial.

The organization of the paper is as follows. In Section 2 we describe examples and applications of the formula. In Section 3 we summarize results of Sageev [5] concerning the geometry of CAT(0) cube complexes. In particular, we describe the notion of hyperplanes in a cube complex, and their manifestation in contracting disks as a collection of embedded arcs. By using Reidemeister-type moves on these contracting disks, we are then able to control the distance between certain minimal edge-paths. We use this in Section 4 to establish the distance from a fixed vertex x_0 to each vertex of an arbitrary cube in X. In particular, we show that any such cube has a unique closest vertex to x_0. In Section 5 we use these vertex distances to set up a recurrence relation for the number of k–cubes starting at distance l from x_0. We then derive formulas for growth series of k–cubes in X (one for each k), the formula in Theorem 1 being the $k=0$ case.

2 Examples and consequences

The simplest example of a CAT(0) cube complex to which the formula applies is a regular tree. In this case, the f–polynomial is of the form $f(t)=1+at$ where a is the degree of a vertex. The formula for the growth series simplifies to the usual one:

$$G(t)=\frac{1+t}{1-(a-1)t}=1+at+a(a-1)t^2+a(a-1)^2t^3+\cdots$$

Other examples are provided by CAT(0) cube complexes X that have vertex-transitive automorphism groups. Consider the special case where the automorphism group $\mathrm{Aut}(X)$ has a subgroup G that acts simply-transitively on the vertex set. In this case the group G can be identified with the vertex set of X, and if we let S denote the set of group elements that are adjacent to the vertex 1, then the Cayley graph of G with respect to S can be identified with the 1–skeleton of X. It follows that the growth series for the group G with respect to the word metric induced by S coincides with the growth series $G(X, x_0; t)$. We describe some examples of this situation.

2.1 Right-angled Coxeter groups

Let Γ be a graph with vertex set V and edge set E. The *right-angled Coxeter group* with defining graph Γ is the group W given by the presentation

$$W = \langle V \mid v^2 = 1 \text{ for all } v \in V \text{ and } uv = vu \text{ for all } \{u, v\} \in E \rangle.$$

There is a natural CAT(0) cube complex X (called the *Davis complex*) on which W acts. We give a rough description here, and refer the reader to Davis [3] for details. The Cayley 2–complex of the presentation for W is a square complex with the property that the link of every vertex can be naturally identified with the graph Γ. The cube complex X is obtained by attaching higher dimensional cubes in such a way that every clique (complete subgraph) in every link gets "filled in". Thus all vertex links in X are isomorphic to this "flag completion" of Γ, so the f–polynomial is $k_\Gamma(t) = 1 + k_1 t + k_2 t^2 + \cdots$ where k_i denotes the number of i–cliques in Γ. The formula for the growth series of X (hence for W) in this case is a (well-known) special case of our theorem:

$$\frac{1}{G(t)} = k_\Gamma\left(\frac{-t}{1+t}\right).$$

Example 2 Let Γ be the graph with $V = \{a, b, c, d\}$ shown in Figure 1 (on the left). Then W is the group $(\mathbb{Z}_2)^3 *_{\mathbb{Z}_2} (\mathbb{Z}_2)^2$, and the Cayley 2–complex of the presentation is shown in Figure 1 (on the right). By filling in all of the 3–cubes, we obtain the Davis complex. The f–polynomial for the link is $f(t) = 1 + 4t + 4t^2 + t^3$, so the growth series for X (and hence W) is

$$G(t) = \frac{(1+t)^3}{1-t-t^2} = 1 + 4t + 8t^2 + 15t^3 + 23t^4 \cdots .$$

2.2 Right-angled mock reflection groups

More generally, suppose Γ is a graph as above and for each vertex $v \in V$, one specifies an involution j_v defined on vertices adjacent to v. Let J denote the collection $\{j_v\}$ of these "local involutions". For any pair of adjacent vertices v_0, v_1, one can then define a sequence $v_0, v_1, \ldots$ inductively by the formula $v_{k+1} = j_{v_k}(v_{k-1})$. We call

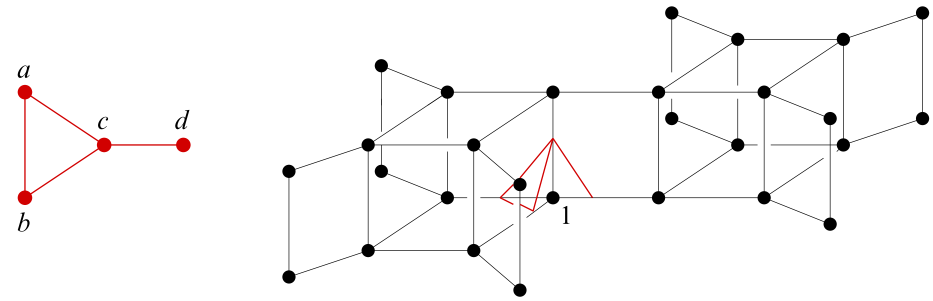

Figure 1: A graph Γ and the corresponding Davis complex

such a sequence a *trajectory*, and assume that all trajectories are 4–periodic (that is, $v_k = v_{k+4}$ for all k). We define the group $W(J)$ by the presentation

$$W(J) = \langle V \mid v^2 = 1 \text{ for } v \in V \text{ and } v_0 v_1 v_2 v_3 = 1 \text{ for all trajectories } v_0, v_1, \ldots \text{ in } \Gamma \rangle,$$

noting that if all of the involutions in J are trivial, this reduces to the (ordinary) right-angled Coxeter group W described above.

With some additional assumptions on the local involutions, one can mimic the Davis complex construction to get a CAT(0) cube complex $X(J)$ with an action of $W(J)$ that is simply transitive on the vertex set. We refer the reader to [6] for details. For such groups $W(J)$ (the ones that act on CAT(0) cube complexes) we call the graph with local involutions a *mock reflection system*, and we call the group $W(J)$ a *mock reflection group*. The link of every vertex in $X(J)$, as for the ordinary Davis complex, is again obtained by filling in all cliques in the graph Γ. Thus, by Theorem 1, the growth series for $W(J)$ (with respect to the generators V) depends only on the underlying graph Γ, not on the choice of local involutions J. This is not an obvious fact, considering that the complexes $X(J)$ definitely do depend on J.

Example 3 Let Γ be the same graph as in Example 2. For the vertices b, c, and d, we define the corresponding local involutions to be the identity, and for the vertex a, we define j_a to be the involution that swaps b and c. If an involution at a vertex v interchanges two adjacent vertices u and w, then we indicate this in the diagram for Γ by connecting the edges vu and vw by an arc at the vertex v (Figure 2). This collection of local involutions determines a mock reflection system, and the corresponding mock reflection group is

$$W(J) = \langle a, b, c, d \mid a^2 = b^2 = c^2 = d^2 = 1, abac = bcbc = cdcd = 1 \rangle.$$

The cube complex $X(J)$ which is shown in Figure 2 is clearly not isomorphic to the Davis complex in the previous example, but by Theorem 1 it does have the same growth series.

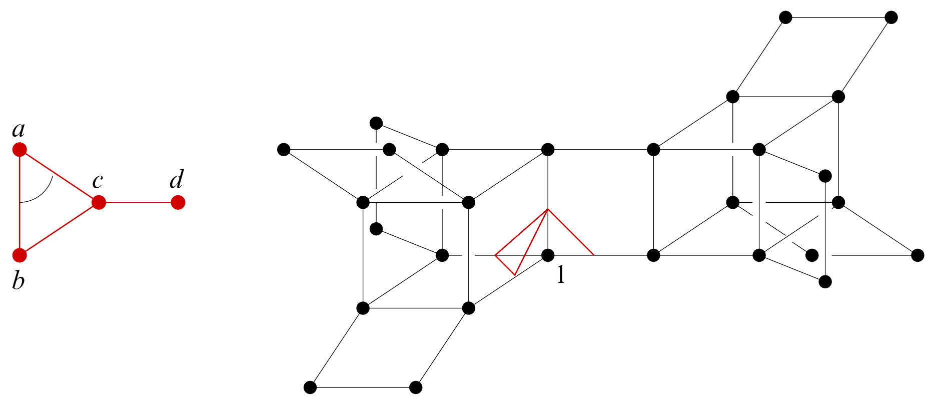

Figure 2: A mock reflection system and the corresponding complex $X(J)$

2.3 Right-angled Artin groups

By removing the involution relations from the presentations for right-angled Coxeter groups, one obtains the class of right-angled Artin groups. That is, given a graph Γ with vertex set V and edge set E, the corresponding *right-angled Artin group* is the group A given by

$$A = \langle V \mid uv = vu \text{ for all } \{u, v\} \in E \rangle.$$

The group A also acts on a CAT(0) cube complex Y (namely, the universal cover of the "Salvetti complex" $\widetilde{\Sigma}$ in Charney and Davis [2].) In this case the action on Y is free, and the quotient Y/A is a finite $K(A,1)$–space. The group A acts simply-transitively on the vertices of Y, and if we let S be the subset $V \cup V^{-1}$, then the 1–skeleton of Y coincides with the Cayley graph of A with respect to S. Thus, again, the growth series for Y coincides with the growth series for the Artin group A with respect to its standard generating set.

The link of a vertex in Y is well-understood (see, for example, [2]). In particular, if $\widehat{\Gamma}$ denotes the simplicial complex obtained by filling in all of the cliques in Γ, and L denotes the link of a vertex in Y, then each i–simplex in $\widehat{\Gamma}$ corresponds to 2^{i+1} simplices in L of the same dimension. It follows that the f–polynomial for L is

$f(t) = k_\Gamma(2t) = 1 + 2k_1 t + 4k_2 t^2 + \cdots$ and, by Theorem 1, that the growth series for Y (and hence A) is determined by

$$\frac{1}{G(t)} = k_\Gamma\left(\frac{-2t}{1+t}\right).$$

Remark There is also a notion of a right-angled *mock* Artin group. Again one starts with a graph Γ with local involutions J, and removes the involution relations from the presentation for the mock reflection group $W(J)$. There is a corresponding complex $Y(J)$ in this case, and the link of a vertex coincides with the link of a vertex in the (ordinary) Artin group associated to the underlying graph Γ. (See [6] for the details.) In particular, the growth series for a mock Artin group $A(J)$ does not depend on the involutions J and coincides with the growth series for the corresponding Artin group for Γ

3 Hyperplanes, contracting disks and pictures

Given an n–dimensional cube Q and an edge $e \subset Q$, let $Q(e)$ denote the $(n-1)$–dimensional subcube obtained by intersecting Q with the hyperplane orthogonal to e passing through the midpoint of e. Following Sageev [5], we call $Q(e)$ a *dual block* in Q. The dual blocks in Q determines an equivalence relation on edges of Q by $e \sim e' \Leftrightarrow Q(e) = Q(e')$. More generally, given a cube complex X, we consider the equivalence relation on edges generated by this relation on each cell. That is, e and e' are equivalent if and only if there exists a sequence of edges $e = e_0, e_1, \dots, e_n = e'$ and a sequence of cubes $Q_1, \dots, Q_n$ in X such that for $0 \le i < n$, $e_i \sim e_{i+1}$ in Q_i. Given an equivalence class ϵ of edges, we then define its *dual hyperplane* to be the union of dual blocks $H(\epsilon) = \bigcup Q(e)$ where the union is taken over all $e \in \epsilon$ and all cells Q in X (we adopt the obvious convention that $Q(e) = \varnothing$ if e is not an edge of Q). We let $\mathcal{H}$ denote the collection of all hyperplanes in X.

Figure 3 shows a cube complex and the hyperplane dual to the equivalence class consisting of vertical edges. (The collection $\mathcal{H}$ in this case consists of this hyperplane together with three other hyperplanes that are not shown.)

Now suppose X is a CAT(0) cube complex, and suppose $\gamma = e_1, \dots, e_n$ is an edge-path that starts and ends at the same vertex x_0 (that is, γ is an *edge-loop*). Since X is simply-connected, there exists a 2–disk D and a map $f\colon D \to X$ that restricts to a map $f\colon \partial D \to \gamma$. Given such a contracting map f, let D_f denote the union of the preimages $f^{-1}(H)$ as H runs over all hyperplanes in $\mathcal{H}$.

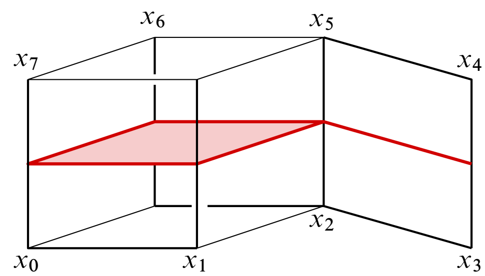

Figure 3: A hyperplane in a cube complex

Proposition 4 (Sageev [5, Theorem 4.4]) *Let X be a CAT(0) cube complex and let γ be an edge-loop. Then there exists a contracting map $f\colon D \to X$ for γ satisfying*

(1) *$f(D)$ is contained in the 2–skeleton $X^{(2)}$*

(2) *The subset $D_f \subset D$ is the union of a collection $\mathcal{A}$ of embedded arcs with endpoints on ∂D and such that any two arcs intersect at most once.*

(3) *Any point on the boundary of D is an endpoint of at most one arc in $\mathcal{A}$, and any point in the interior of D is contained in at most two arcs of $\mathcal{A}$.*

An example of such a contracting map is shown in Figure 4. Here the edge-loop is the one passing through the vertices $x_0, x_1, \ldots, x_7$, and the map $f\colon D \to X$ maps the disk homeomorphically onto the front three faces of the 3–cube and the adjoining 2–cube. The four hyperplanes in X meet the image of this disk in the four arcs indicated (the arc corresponding to the shaded hyperplane is in bold).

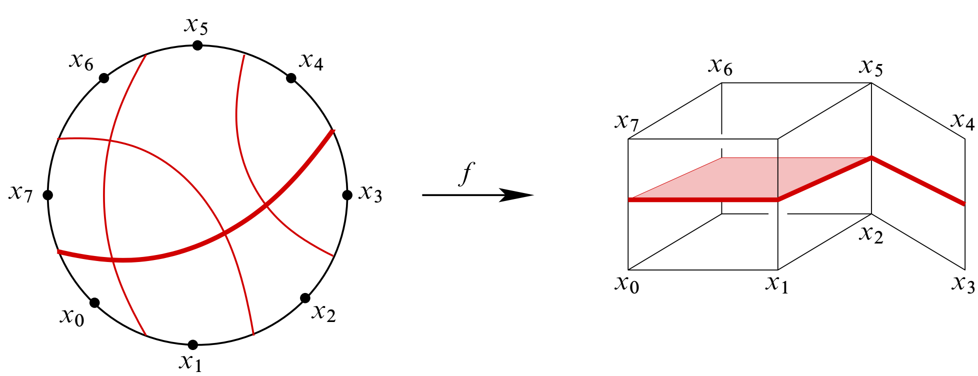

Figure 4: A contracting map $f\colon D \to X$ for an edge-loop

To prove the proposition, one first uses transversality results to make the hyperplane preimages D_f a collection of immersed closed curves and arcs meeting in general position. Using the fact that links of vertices in X are all flag simplicial complexes, one then uses certain Reidemeister-type moves to simplify this collection of curves. The complete argument can be found in Sageev [5].

To simplify our exposition, we shall call the pair (D, D_f) a *picture* if it satisfies all of the conditions in Proposition 4. We will need to be able to modify our pictures using one of the Reidemeister moves mentioned above: the so-called "triangle move". In a picture, the disk D gets broken up into a collection of contractible regions, each of which is bounded by a finite number of sub-arcs. We call such a region a *triangle* if it is bounded by precisely 3 sub-arcs.

Proposition 5 *Let $f\colon D \to X$ define a picture for the closed edge-loop γ, and suppose this picture has a triangle region with bounding arcs α_1, α_2, and α_3. Then there exists another contracting map $g\colon D \to X$ for γ such that (D, D_g) is a picture identical to (D, D_f) except that the arc α_1 is on the other side of the intersection point $\alpha_2 \cap \alpha_3$.*

Proof The triangle in (D, D_f) corresponds to a vertex v in the 2–skeleton $X^{(2)}$ where 3 squares meet like the corner of a 3–cube. The flag condition on the link of this vertex in X ensures that there is, in fact, a 3–cube in X that is attached to these three squares. Replacing these three squares with the opposite three squares in this cube yields a homotopic contracting map g with the desired picture. For example, the contracting map in Figure 4 can be modified so that the disk maps onto the *back* three faces of the 3–cube. The resulting picture is shown on the right in Figure 5. □

4 Cube positions in a CAT(0) cube complex

In this section we now fix a vertex x_0 in X. We shall say that a vertex x is *at level l* if $d(x, x_0) = l$. An immediate consequence of the existence of pictures for edge-loops in X is that any edge-loop has even length (twice the number of embedded arcs). This means that the vertices of any 1–cube in X must be at different levels. In fact they must be at levels l and $l+1$ for some $l \geq 0$. In general we have the following.

Lemma 6 *Let Q be a k–dimensional cube in X and let l be the minimum level attained by vertices of Q. Then for each $j \in \{0, \dots k\}$, there are precisely $\binom{k}{j}$ vertices of Q at level $l+j$ (Figure 6). In particular, Q has a unique (closest) vertex at level l and a unique (farthest) vertex at level $l+k$.*

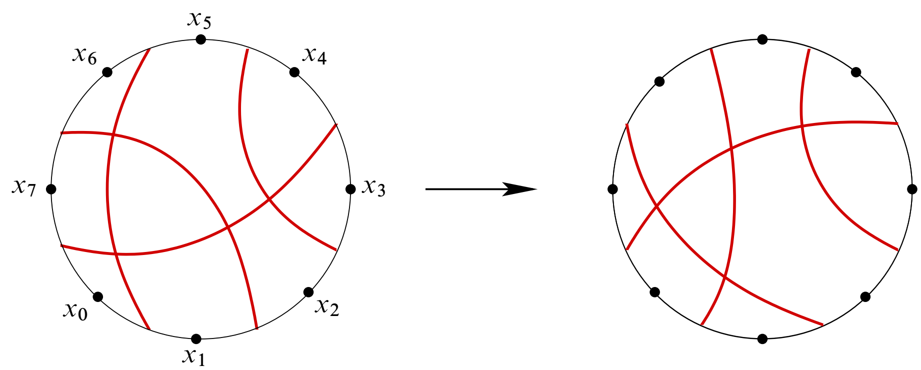

Figure 5: A triangle move on a picture

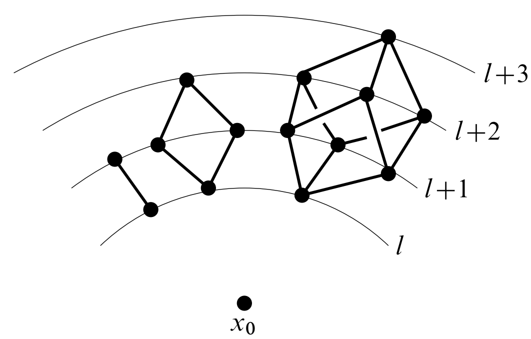

Figure 6: Level sets in a CAT(0) cube complex

Proof First we show that there exists a unique vertex of Q at level l. For suppose x and x' are two different vertices of Q at level l. Let α be a minimal edge-path connecting x_0 to x, let α' be a minimal edge-path connecting x' to x_0, and let β be an edge-path in Q connecting x to x' that has minimal length (in Q). Then the composition of α, β, and α' is an edge-loop γ, so by Proposition 4, there exists a contracting map $f\colon D \to X$ such that (D, D_f) is a picture for γ. We can assume this picture is minimal in the sense that it has the minimum number of interior crossings among all pictures for γ. We claim that such a picture has the following properties:

(1) If two arcs each have an endpoint on α (respectively, α'), then they do not intersect.

(2) If two arcs each have an endpoint on β, then they do not intersect.

(3) No arc has both endpoints on α, α', or β.

(4) If p is the last arc endpoint along α and q is the first arc endpoint along β (that is, p and q are the closest arc endpoints to x) then p and q do *not* belong to the same arc. Similarly for the two arc endpoints that are closest to x'.

Assuming the claim, it is not hard to see that no such picture can exist. For consider the arc A with endpoint closest to x along α and the arc B with endpoint closest to x along β (Figure 7). By (4), A and B are different. By (3) A cannot have both endpoints on α and B cannot have both endpoints along β. It follows that B must have second endpoint along α' (if it were along α, A and B would be forced to cross and each would have an endpoint on α, contradicting (1)). By symmetry, the arc B' that has endpoint closest to x' along β must have second endpoint along α. But then B and B' must intersect, violating (2).

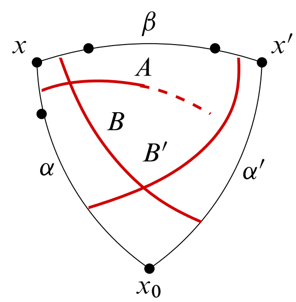

Figure 7

To establish (1) of the claim, it suffices to show that no two arcs with consecutive endpoints along α can intersect. For suppose A and B are two such arcs that intersect in the point p. By applying repeated triangle moves one can first ensure that all arcs crossing A and B to the left of the intersection point are parallel (Figure 8). Then one can use repeated triangle moves to move all of these arcs to the right of p. We now have a triangular region in the picture with one edge along α. The intersection point p is dual to a square R in $X^{(2)}$ having two consecutive edges along α. Replacing these edges with the opposite two edges of R would result in fewer interior crossings (Figure 9), contradicting our choice of a picture that minimizes these crossings.

The proof of (2) is identical to the proof of (1) except that one needs to observe that the square R is in fact a *face of* Q (thus replacing the two edges along β with the opposite edges of the square, still gives an edge-path *in* Q). But this is clear since the intersection of Q and R must be a face of both, hence it must be the entire square R.

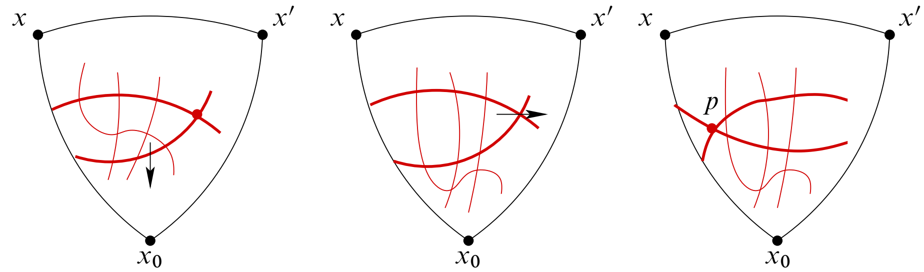

Figure 8

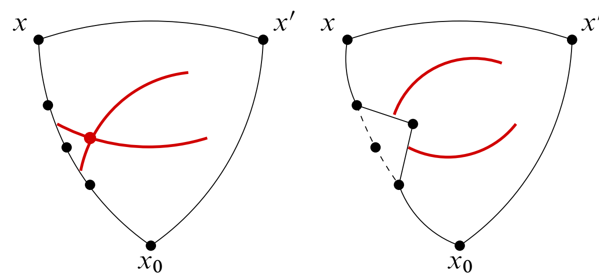

Figure 9

To prove (3), suppose without loss of generality that there exists an arc with both endpoints along α. No other arc can cross this one since it would either have to cross it twice (which is not allowed) or it would result in two crossing arcs that both meet α (violating (1)). It follows that we can choose such an arc A so that no other arcs touch α between the endpoints of A. In this case, the endpoints of A must map (under f) to the midpoint of the same 1–cell in X. If we let y denote the endpoint of this 1–cell that is not enclosed by A, then we see (in the left-hand picture in Figure 10) that the edge-path α can be shortened, contradicting its minimality.

For (4), we simply note that the existence of such an arc connecting p to q would imply that the last edge in the edge-path α coincides with the first edge of β. In particular, if y is the endpoint of this edge that is opposite x, then y would be a point in Q closer to x_0 than x (the right-hand picture in Figure 10).

Finally, to see that there are $\binom{k}{j}$ vertices of Q at level $l+j$, let x be the closest vertex (at level l). Index the vertices of Q using subsets of $\{1,2\ldots,k\}$ so that $v_\varnothing = x$ and v_I is adjacent to v_J if and only if the symmetric difference of I and J has one element. It suffices to show then that the vertex v_I is at level $l+|I|$ for each subset $I\subset\{1,\ldots,k\}$. We proceed by induction on $|I|$, the case $|I|=0$ being trivial.

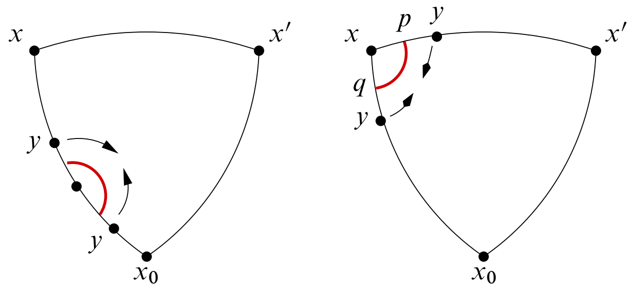

Figure 10

Consider the vertex v_I where $|I| > 0$. For every proper subset $J \subset I$, the vertex v_J is at level $l + |J|$ (by induction). The level of v_I is at most $l + |I|$ since removing an element from I gives an adjacent vertex at level $l + |I| - 1$. Suppose the level of v_I were less than $l + |I|$, say $l + j$. Pick a subset $J \subset I$ with $|J| = j$ and consider the face $Q_J \subset Q$ consisting of vertices v_K with $J \subset K \subset I$. Then both v_I and v_J would attain the minimum level for vertices of Q_J, contradicting the uniqueness of a closest vertex. Hence v_I must be at level $l + |I|$. □

Lemma 7 *Let x be a vertex at level l, and let S be any set of vertices at level $l-1$ that are adjacent to x. Then there exists a unique cube of dimension $|S|$ that contains all of the vertices $S \cup \{x\}$. (By the previous lemma, this cube starts at level $l - |S|$.)*

Proof If S is empty or consists of a single vertex, the statement is trivial. Suppose S is a two-element set $\{y, y'\}$. Let α be a minimal edge-path from x_0 to y, and let α' be a minimal edge-path from y' to x_0. Let e be the oriented edge from y to x, and let e' be the oriented edge from x to y'. Then composing the paths α, e, e', and α', we obtain an edge-loop γ, and we let (D, D_f) be a picture for γ. As in the previous proof, we assume that α, and α' are chosen so that the number of interior crossings in this picture is minimized. Let A denote the arc that meets the midpoint of the edge e and let A' denote the arc that meets the midpoint of e' (Figure 11). Since $y \neq y'$, we know $e \neq e'$, so the arcs A and A' cannot coincide. By an argument similar to the previous proof we know that the arc A must connect to the boundary segment α', and the arc A' must connect to α, hence A and A' must cross. By using triangle moves, we can assume this crossing is "at the top" (that is, the crossing point of A and A' is the closest crossing point to e along A and the closest crossing point to e' along A'). This new crossing point then has a dual square R having e and e' as consecutive edges. This R is the desired 2–cube; it is unique since the intersection of two cubes must be a (single) face of each.

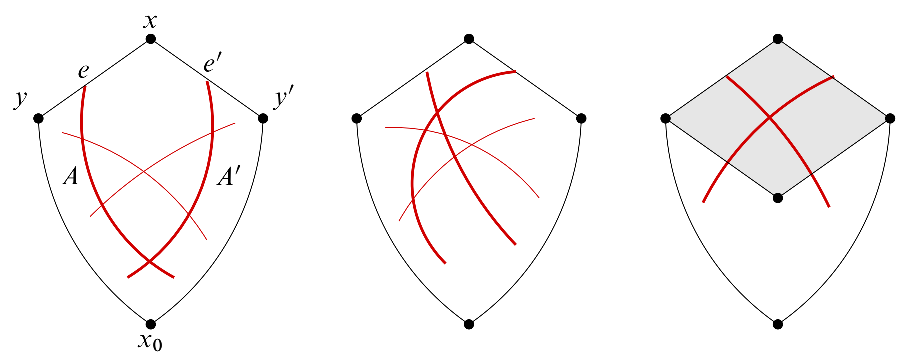

Figure 11

Now suppose S has more than two elements. These elements correspond to vertices in the link of x. By the previous paragraph, any two-element subset determines a square and hence an edge connecting the corresponding vertices in the link of x. Since the link of x is a flag complex, the vertices in the link that correspond to S must span a simplex; hence, there exists a cube of dimension $|S|$ containing $S \cup \{x\}$. Uniqueness follows (again) from the fact that the intersection of two cubes must be a face of both. □

5 Generating functions

Let X and x_0 be as above, and let $s_{k,l}$ denote the number of k cells in X that start at level l. In particular, the numbers $s_{k,0}$ for $k = 0, \dots, n$ are the coefficients of the f–polynomial for the link of the vertex x_0, and $s_{0,l}$ is the number of vertices in X at level l.

Lemma 8 *We have the following identities for the $s_{k,l}$:*

(1) $s_{k,l} = 0$ *if* $k < 0$ *or* $l < 0$.

(2) $\sum_{k=0}^{l} (-1)^k s_{k,l-k}$ *is* 1 *if* $l = 0$ *and* 0 *if* $l > 0$.

(3) $s_{k,0} s_{0,l} = \sum_{j=0}^{k} \binom{k}{j} s_{k,l-j}$ *if* $k > 0$.

Proof The first statement is obvious. For the second, note that by Lemma 6, the sum in question can be interpreted as

$$\sum (-1)^{\dim Q}$$

where the sum is taken over all cubes Q that end at level l. By Lemma 7, any such cube is contained in a unique maximal cube (ending at level l). If we restrict the sum to one of these maximal cubes Q_0 we get a sum of the form

$$\sum_{I \subset J} (-1)^{|I|}$$

where J is the set of vertices of Q_0 at level $l-1$. Since this sum is zero for each maximal cube (it's the binomial expansion of $(1-1)^{|J|}$), the result follows. The third identity corresponds to two different ways of counting the number of $(k-1)$–simplices in the links of all of the vertices at level l. Since $s_{0,l}$ is the number of vertices at level l and $s_{k,0}$ is the number of $(k-1)$–simplices in the link of each vertex, the left hand side $s_{k,0}s_{0,l}$ certainly gives this number. On the other hand, since $s_{k,l-j}$ is the number of k–cells starting at level $l-j$, and (by Lemma 6) each such k–cell contributes a $(k-1)$–simplex to the links of its $\binom{k}{j}$ vertices at level l, the right-hand sum also yields this number. $\square$

For $k = 0, \dots, n$ we let $g_k(t)$ be the generating function

$$g_k(t) = \sum_{l=0}^{\infty} s_{k,l} t^l.$$

That is, $g_k(t)$ is the growth series for k–cells in X. In particular, $g_0(t)$ is the growth series $G(X, x_0; t)$.

Lemma 9 *These generating functions satisfy the following identities:*

(1) $(1+t)^k g_k(t) = s_{k,0} g_0(t)$ *for all* $k \geq 0$.

(2) $\sum_{k=0}^{\infty} (-t)^k g_k(t) = 1$.

Proof These are just generating function versions of the identities (2) and (3) in Lemma 8. For the first identity, the case $k = 0$ is trivial, and for $k \geq 1$ we have

$$s_{k,0} g_0(t) = \sum_{l=0}^{\infty} s_{0,l} s_{k,0} t^l = \sum_{l=0}^{\infty} \sum_{j=0}^{k} \binom{k}{j} s_{k,l-j} t^l$$

by (3) in Lemma 8. Interchanging the sums and noting that $s_{k,l-j} = 0$ for $l < j$ gives

$$\sum_{j=0}^{k} \sum_{l=j}^{\infty} s_{k,l-j} t^{l-j} \binom{k}{j} t^j = \sum_{j=0}^{k} g_k(t) \binom{k}{j} t^j = g_k(t)(1+t)^k,$$

as desired.

For the second identity, we have

$$\sum_{k=0}^{\infty}(-t)^k g_k(t) = \sum_{k=0}^{\infty}(-t)^k \sum_{i=0}^{\infty} s_{k,i}t^i.$$

Substituting $l-k$ for i, this becomes

$$\sum_{k=0}^{\infty}\sum_{l=k}^{\infty}(-1)^k s_{k,l-k}t^l.$$

Interchanging the sums then gives

$$\sum_{l=0}^{\infty}\left(\sum_{k=0}^{l}(-1)^k s_{k,l-k}\right)t^l$$

which, by (2) in Lemma 8, reduces to 1. □

The formula for the growth series $G(X, x_0; t)$ given in the introduction is the special case $k=0$ of the following theorem.

Theorem 10 *The generating functions* $g_k(t)$ *are given by*

$$\frac{1}{g_k(t)} = \frac{(1+t)^k}{f_{k-1}} f\left(\frac{-t}{1+t}\right).$$

Proof Since $f(t) = \sum s_{k,0}t^k$ where the sum is taken over all $k \geq 0$, we have

$$\begin{aligned} f\left(\frac{-t}{1+t}\right) &= \sum_{k=0}^{\infty} s_{k,0}\left(\frac{-t}{1+t}\right)^k \\ &= \sum_{k=0}^{\infty}\left(\frac{g_k(t)(1+t)^k}{g_0(t)}\right)\left(\frac{-t}{1+t}\right)^k \\ &= \frac{1}{g_0(t)}\sum_{k=0}^{\infty}(-t)^k g_k(t) \\ &= \frac{1}{g_0(t)} \end{aligned}$$

where the second line follows from (1) of Lemma 9, and the last line follows from (2). This gives the desired formula in the case $k=0$. The general formula then follows again from (1) in Lemma 9. □

References

[1] **M R Bridson**, **A Haefliger**, *Metric spaces of non-positive curvature*, Grundlehren der Mathematischen Wissenschaften 319, Springer, Berlin (1999) MR1744486

[2] **R Charney**, **M W Davis**, *Finite $K(\pi,1)$s for Artin groups*, from: "Prospects in topology (Princeton, NJ, 1994)", Ann. of Math. Stud. 138, Princeton Univ. Press, Princeton, NJ (1995) 110–124 MR1368655

[3] **M W Davis**, *Nonpositive curvature and reflection groups*, from: "Handbook of geometric topology", (R Daverman, R Sher, editors), North-Holland, Amsterdam (2002) 373–422 MR1886674

[4] **G A Noskov**, *Growth of certain non-positively curved cube groups*, European J. Combin. 21 (2000) 659–666 MR1771984

[5] **M Sageev**, *Ends of group pairs and non-positively curved cube complexes*, Proc. London Math. Soc. (3) 71 (1995) 585–617 MR1347406

[6] **R Scott**, *Right-angled mock reflection and mock Artin groups*, Trans. Amer. Math. Soc. (to appear)

[7] **R Steinberg**, *Endomorphisms of linear algebraic groups*, Memoirs of the American Mathematical Society 80, American Mathematical Society, Providence, R.I. (1968) MR0230728

Department of Mathematics and Computer Science, Santa Clara University
Santa Clara CA 95053, USA

rscott@math.scu.edu

Received: 21 November 2006