Leighton's theorem and regular cube complexes

DANIEL J WOODHOUSE

Leighton's theorem and regular cube complexes

DANIEL J WOODHOUSE

Leighton's graph covering theorem states that two finite graphs with common universal cover have a common finite cover. We generalize this to a large family of nonpositively curved special cube complexes that form a natural generalization of regular graphs. This family includes both hyperbolic and nonhyperbolic CAT(0) cube complexes.

Leighton's graph covering theorem states that two finite graphs with isomorphic universal covers have isomorphic finite covers. First conjectured by Angluin [2] and proven by Leighton [16], whose background was in computer science and the study of networks, the topic has been picked up by topologists and group theorists interested in producing generalizations to graphs with extra structure, including colourings and line patterns; see Bass and Kulkarni [3], Neumann [18], Shepherd [21], and the author [24]. Although it is desirable to generalize such a theorem to higher dimensions, counterexamples are known even when the universal cover is the product of two trees. Standard arithmetic constructions were known to give irreducible lattices acting on the product of trees, and in the 90s nonresidually finite and even simple examples were given; see Burger and Mozes [6] and Wise [22].

A particularly exciting conjecture was made by Haglund in [11] that Leighton's graph covering theorem should generalize to special cube complexes. In the same paper Haglund proved the conjecture for the class of right-angled Fuchsian buildings (commonly referred to as "Bourdon buildings") and more generally for "type-preserving" lattices in the automorphism group of a building associated to a finite graph product of finite groups.

In this paper we will prove Haglund's conjecture for a large family of CAT(0) cube complexes which exhibit symmetry and homogeneity reminiscent of finite regular trees. Let L be a finite simplicial flag complex. An L–*cube-complex* X is a cube complex

such that every link is isomorphic to L. Given a flag complex, the Davis complex $D(L)$ of the associated right-angled Coxeter group is a CAT(0) L–cube-complex. In general, $D(L)$ is not the unique CAT(0) L–cube-complex, but in [15] Lazarovich shows that $D(L)$ is unique if and only if L is *superstar-transitive*. Recall that the *star* of simplex σ in L, denoted by $\mathrm{St}(\sigma)$, is the subcomplex given by the union of all simplices containing σ. We say that flag complex L is superstar-transitive if for any two simplices $\sigma, \sigma' \subseteq L$, any isomorphism $\mathrm{St}(\sigma) \to \mathrm{St}(\sigma')$ sending σ to σ' extends to an automorphism of L. Lazarovich also showed that in this case $\mathrm{Aut}(X)$ is virtually simple.

The principal set of examples of superstar transitive flag complexes presented by Lazarovich are *Kneser complexes*. Let Δ be a finite set. The Kneser complex $\mathfrak{K}_n(\Delta)$ is the simplicial flag complex defined with vertex set the n–element subsets of Δ, and edges joining disjoint n–element subsets. In the particular case that $|\Delta| = nd + 1$, the Kneser complex is superstar transitive and its automorphism group is precisely the natural action of the permutation group $\mathrm{Sym}(\Delta)$; see Section 1.2. We prove the following:

Theorem 0.1 *Let $n \geq 2$, $d \geq 1$ and Δ be a finite set of cardinality $nd + 1$. Let L be the Kneser complex $\mathfrak{K}_n(\Delta)$. Suppose that X_1 and X_2 are compact, L–cube-complexes such that all finite-index subgroups of the hyperplane subgroups are separable in $\pi_1 X_1$ and $\pi_1 X_2$, respectively. Then X_1 and X_2 have a common finite cover.*

If the hyperplane subgroups of a compact nonpositively curved cube complex are separable, then there is a finite cover such that the hyperplanes are 2–sided, embed, and do not self-osculate. If no interosculations could be added to this list, then the cube complex would be virtually special. Conversely, specialness implies separable hyperplane subgroups, and it is conjectured that the converse holds as well.

Note that in the case $d = 1$ that L is the set of $n + 1$ disconnected points, so the L–cube-complexes will be $(n+1)$–regular trees. If $n = 2$ and $d = 2$ then L is the famous Petersen graph. In the case when L has no induced squares (as in the case of the Petersen graph), the fundamental groups of X_1 and X_2 will be hyperbolic — see Moussong [17] — and as a consequence of Agol's proof of the virtual Haken conjecture [1], X_1 and X_2 are virtually special. Thus we have:

Corollary 0.2 *Let $n \geq 2$, $d = 1, 2$, $|\Delta| = nd + 1$ and $L = \mathfrak{K}_n(\Delta)$. If X_1 and X_2 are compact L–cube complexes then X_1 and X_2 have common finite covers.*

Proof In the case $d = 1$ the cube complexes are graphs, so it suffices to show that L is square free when $d = 2$. Let $\Delta = \{1, \ldots, 2n + 1\}$. Suppose that v_1, v_2, v_3 and v_4 are the vertices of an induced square in L. Then without loss of generality we can assume that $v_1 = \{1, \ldots, n\}$ and $v_2 = \{n + 1, \ldots, 2n\}$ since they are disjoint sets. Thus we can further assume that $v_3 = \{2, \ldots, n, 2n + 1\}$ since it must be an n–element set disjoint from v_2. Then we have a contradiction since v_4 must be an n–element subset disjoint from $v_1 \cup v_3 = \{1, \ldots, n, 2n + 1\}$, so $v_2 = v_4$. □

0.1 Strategy

The plan is to show (in Proposition 4.1) that each L–cube-complex has a finite cover X admitting a finite orbicovering $X \to X_L$, where X_L is the orbicomplex $W_L \backslash D(L)$. We seek to construct this orbicovering by identifying the link of the 0–cube in X_L with $\mathfrak{K}_n(\Delta)$ and finding a suitable map $\mathrm{lk}(x) \to \mathfrak{K}_n(\Delta)$ for each 0–cube x in X such that the orbicovering is defined. By associating a copy Δ_x of Δ with each 0–cube in X we identify $\mathrm{lk}(x)$ with $\mathfrak{K}_n(\Delta_x)$. The orbicovering is then locally defined by a choice of map $q_x \colon \Delta_x \to \Delta$; see Lemma 1.3.

In order for the set of q_x to define an orbicover we need to ensure that certain conditions are satisfied. If $e = (x, y)$ is a 1–cube, then we need to ensure that e will be mapped to the same half edge in X_L by the maps induced by q_x and q_y. Given a square in X, we also need to ensure that it will be mapped to a quarter-square in X_L.

In Section 3, we formulate the problem in the language of a Δ–*category*, which is a choice of bijection $\phi_e \colon \Delta_x \to \Delta_y$ for each edge $e = (x, y)$, satisfying certain conditions. Most of the action in this paper concerns being able to (virtually) construct a Δ–category. Once we have the Δ–category we obtain a holonomy

$$\Psi \colon \pi_1(X, x) \to \mathrm{Sym}(\Delta_x)$$

and the kernel of this holonomy will give a finite cover for which we can define suitable q_x; see Section 4.

0.2 Previous results and connections to QI–rigidity

A major motivation for proving Haglund's conjecture is the potential applications to Gromov's program of understanding groups up to quasi-isometry [10]. In [11], Haglund proved his conjecture for Bourdon buildings and his result can be combined with a result of Bourdon and Pajot [4] which says that each quasi-isometry of such a building

is finite distance from a unique automorphism. Thus we deduce that if G is a group quasi-isometric to the graph product W associated to such a Bourdon building B, then in fact it acts by isometries on B. By Agol's result [1], G will be virtually special, thus acting faithfully on B, and by Haglund G will be weakly commensurable with W. Thus W is quasi-isometrically rigid.

This argument motivates the following problem:

Problem 0.3 *Let $L = \mathfrak{K}_n(\Delta)$, where $|\Delta| = nd + 1$. Is every quasi-isometry of $D(L)$ finite distance from an automorphism?*

A positive answer to Problem 0.3 in the hyperbolic case would immediately give quasi-isometric rigidity for the associated groups W_Γ by a similar argument to the case of Bourdon buildings. That is to say that any group quasi-isometric to W_Γ would be weakly commensurable with W_Γ. In the "higher rank" nonhyperbolic case one might look to Huang's results on the quasi-isometric rigidity of large families of right-angled Artin groups [13]. In this case following would need to be considered:

Problem 0.4 *Suppose that L is a Kneser complex as above, such that W_Γ is not hyperbolic. Are there groups acting geometrically on $D(\Gamma)$ that are not virtually special?*

1 Preliminaries

1.1 Right-angled Coxeter groups

We refer to Davis [8] for classical background on Coxeter groups and their geometry and to [7] for a recent survey of their large scale geometry.

Let L denote a finite simplicial flag complex. The right-angled Coxeter group W_L associated to L is given by the presentation:

$$W_L = \langle v \in L^{(0)} \mid v^2 = 1 \text{ and } [u, v] = 1 \text{ if } (u, v) \in L^{(1)} \rangle.$$

The Davis complex $D(L)$ is the CAT(0) cube complex obtained from the Cayley 2–complex constructed from the above presentation, after collapsing each v^2 bigon to a single edge, and inserting higher dimensional cubes wherever their 2–skeleta appear. The link of each vertex in $D(L)$ is isomorphic to L, which makes it an *L–cube-complex*. The following theorem tells us when $D(L)$ is the unique CAT(0) L–cube-complex:

Theorem 1.1 [15, Theorem 1.2] *The Davis complex $D(L)$ is the unique CAT(0) cube complex with each link isomorphic to L if and only if L is superstar-transitive.*

If we colour the edges in $D(L)$ according to the corresponding element of L, or alternatively the conjugacy class of the associated generator, we can identify W_L as the subgroup of $\mathrm{Aut}(D(L))$ that preserves the colours. Sometimes this subgroup is referred to as the *type-preserving automorphisms*. The quotient $X_L = W_L \backslash D(L)$ has the structure of an orbicomplex. Each face given by the intersection of k hyperplanes has the associated group $(\mathbb{Z}/2)^k$ with a factor corresponding to a hyperplane.

1.2 Kneser complexes

Let Δ be a finite set. The *Kneser complex* $\mathfrak{K}_n(\Delta)$ is the flag complex with underlying graph with vertex set given by n–elements subsets of Δ, and edges corresponding to disjoint n–element subsets. There is a natural action of $\mathrm{Sym}(\Delta)$ on $\mathfrak{K}_n(\Delta)$.

If $\mathfrak{K} := \mathfrak{K}_n(\Delta)$ is a Kneser complex, then we let $\mathfrak{s}_v = \mathfrak{s}(v) \subseteq \Delta$ denote the subset associated to $v \in \mathfrak{K}^{(0)}$.

Example 1.2 If $|\Delta| = 5$, then $P := \mathfrak{K}_2(\Delta)$ is the *Petersen graph*. It is a simple exercise to verify that P is triangle and square free; see Figure 1.

More generally, if $|\Delta| = nd + 1$, then $\mathfrak{K}_n(\Delta)$ is a $(d-1)$–dimensional flag simplicial complex with a superstar-transitive automorphism group; see [15]. We also note the following:

Lemma 1.3 [9, Corollary 7.8.2] *If $|\Delta| \neq 2n$, then $\mathrm{Aut}(\mathfrak{K}_n(\Delta))$ is equal to $\mathrm{Sym}(\Delta)$.*

Given a subset $\Sigma \subseteq \Delta$, the inclusion induces an embedding $\mathfrak{K}_n(\Sigma) \subseteq \mathfrak{K}_n(\Delta)$, where the vertex in $\mathfrak{K}_n(\Sigma)$ corresponding to $\mathfrak{s} \subseteq \Sigma \subseteq \Delta$ is sent to the corresponding vertex in $\mathfrak{K}_n(\Delta)$. Indeed, an automorphism $(\Delta, \Sigma) \to (\Delta, \Sigma)$ induces an automorphism of $\mathfrak{K}_n(\Delta)$ that restricts to an automorphism on $\mathfrak{K}_n(\Sigma)$. Conversely, by Lemma 1.3, provided $2n$ is not equal to $|\Delta|$ or $|\Sigma|$, an automorphism of $\mathfrak{K}_n(\Delta)$ that preserves $\mathfrak{K}_n(\Sigma)$ gives a self-bijection of Δ that preserves Σ.

Figure 1: The Petersen graph (courtesy of Tilman Piesk [19]).

Kneser complexes were presented by Lazarovich as a large and readily accessible set of superstar transitive graphs.

Theorem 1.4 [15, Corollary 5.5] *Let $n \geq 2$ and $d \geq 1$. Let $|\Delta| = nd + 1$ and $L := \mathfrak{K}_n(\Delta)$. Then $\mathrm{Aut}(D(L))$ is virtually simple.*

We note that $D(L)$ is Gromov hyperbolic if and only if L does not contain any induced squares [17]. Thus, it is an exercise to verify that $D(L)$ is hyperbolic only if $d \leq 2$.

Remark The most direct means that a result like Theorem 0.1 could be true is if the automorphism group of $D(L)$ were to act properly. In which case any other uniform lattice in $\mathrm{Aut}(D(L))$ would lie inside $\mathrm{Aut}(D(L))$ as a finite-index subgroup. Common covers of the corresponding quotient spaces could be constructed by taking the intersections of the associated lattices. In general the automorphism groups of universal covers will be far too large for this argument to work. Theorem 1.4 is the most extreme example of this: since W_L is residually finite (and indeed virtually special), it cannot lie inside a virtually simple group like $\mathrm{Aut}(D(L))$ as a finite-index subgroup.

2 Special cube complexes

We refer to [5; 12; 14; 20; 23] for more detailed background on nonpositive curvature, cube complexes, and specialness. We outline here the terminology that we will use.

An *n–cube* C is a metric space isometrically identified with $[-1, 1]^n$. A 0–cube is a singleton. A *subcube* $S \subseteq C$ of dimension m in an n–cube is the m–cube obtained by restricting $(m-n)$–many coordinates to 1 or -1. The i^{th} *midcube* $M \subset C$, for $1 \leq i \leq n$, is the $(n-1)$–cube obtained by restricting the i^{th} coordinate to 0.

The *reflection* of an n–cube over it's i^{th} midcube $M \subseteq C$ is the map $C \to C$ obtained by multiplying the i^{th} coordinate by -1. Note that all the reflections in a cube commute. The *antipodal* map $C \to C$ is obtained by reflection over all the midcubes in C.

The *link* $\text{lk}(x)$ of a 0–cube x in an n–cube C is the simplex σ given by the ϵ–sphere of x in the ℓ^1–metric (where $1 > \epsilon > 0$). Each subcube of C that contains x has a link at x that gives a corresponding face in σ. If x and y are 0–cubes in C, then x is mapped to y by the composition R of all the reflections over midcubes separating x and y. Thus R induces an isomorphism $\text{lk}(x) \to \text{lk}(y)$.

By a *cube complex* X we will mean a topological space that decomposes into cubes $\mathcal{C}(X)$, such that every subcube of a cube in $\mathcal{C}(X)$ is a cube in $\mathcal{C}(X)$, and such that the intersection of any two cubes $C, C' \in \mathcal{C}(X)$ give subcubes of C and C', or the intersection is empty. The link $\text{lk}(x)$ of a 0–cube x in X is the complex given by the union of all the links of all the cubes containing x, with inclusion of simplices induced by inclusion of subcubes. Alternatively, it can also be thought of as the ϵ neighbourhood of x inside X itself. A cube complex X is *nonpositively curved* if the link of each vertex is a simplicial flag complex. Each n–simplex σ in $\text{lk}(x)$ corresponds to a unique $(n+1)$–cube $C(\sigma)$ in X containing x. Conversely, each $(n+1)$–cube C that contains x corresponds to a simplex $\sigma(C)$ in $\text{lk}(x)$.

Unless otherwise noted, our 1–cubes will be *directed* in the sense that $e = (x, y)$ comes with an initial and terminal 0–cube, denoted by $\iota e = x$ and $\tau e = y$. The *reversed 1–cube* with the opposite direction will be denoted by $\bar{e} = (y, x)$. Let X be a compact nonpositively curved cube complex. A hyperplane Λ in X is an equivalence class of directed 1–cubes generated by the relation $e \sim e'$ if they are opposite faces of a square in X or $\bar{e} = e'$. Associated to the equivalence class is the *realization* of Λ. This is a nonpositively curved cube complex, which we will also denote by Λ, constructed from the midcubes dual to the edges in the equivalence class that immerses by a local isometry $\Lambda \looparrowright X$. Note that this immersion is only a cellular map when both Λ and X have been cubically subdivided. The *hyperplane subgroup* associated to Λ is the image of $\pi_1(\Lambda)$ in $\pi_1(X)$ under the injective homomorphism given by the immersion.

A hyperplane is *embedded* if no two edges in the equivalence class form the corner of a square (that is to say a 2–cube) in X. Equivalently a hyperplane is embedded if the immersion of the realization is an embedding. The *carrier* $N(\Lambda) \subseteq X$ of a hyperplane is the subcomplex obtained by taking all cubes that contain an edge in the associated equivalence class. We say that a hyperplane is *fully clean and 2–sided* if $N(\Lambda) \cong \Lambda \times [-1, 1]$. That is to say that we can extend the embedding of the realization to an embedding $N(\Lambda) = \Lambda \times [-1, 1] \hookrightarrow X$. If the hyperplane subgroups of $\pi_1(X)$ are separable then there is a finite cover of X such that the hyperplanes are fully clean and 2–sided. Indeed, fully clean follows from [12, Lemma 9.14], and with hyperplanes embedded a standard cut-and-paste argument applied to a 1–sided hyperplane yields a degree 2 cover with a two sided hyperplane; see also the proof of [12, Proposition 3.10]. *Thus, we will now assume going forward that all hyperplanes satisfy this condition.* In terms of the definition of specialness, this is equivalent to the hyperplanes being 2–sided, embedded, and without self-osculations. Such a cube complex may fail to be special since interosculations do not contradict this assumption (see Figure 2 for an illustration of the hyperplane pathologies). In terms of the assumptions of Theorem 0.1, if the finite-index subgroups of hyperplane subgroups are separable in $\pi_1 X$, then this remains true of the hyperplane subgroups in a finite cover.

A 0–cube x is *incident* to Λ if it is contained in $N(\Lambda)$. An edge $e = (x, y)$ is *parallel* to Λ if it is contained in $N(\Lambda)$ without being dual to Λ. Under the assumption that the hyperplane Λ is fully clean and 2–sided, the immersion of the realization extends to an embedding $\Lambda \times [-1, 1] \hookrightarrow X$, where the realization is the 0 fiber. The edges parallel to Λ are contained in the -1 and 1 fibers. We will refer to the subcomplexes of X given by the ± 1 fibers as the *sides of the carrier*.

2.1 The adjacency map

Let $e = (x, y)$ be an edge in X dual to Λ and let v be the vertex in $\text{lk}(x)$ corresponding to e, and u be the vertex in $\text{lk}(y)$ corresponding to e. The *star* $\text{Star}(\sigma)$ of a simplex σ in a simplicial complex is the subcomplex spanned by the union of all simplices containing σ. We note that in [15] the star of a simplex is defined by Lazarovich to be the combinatorial 1–neighbourhood. The two notions only coincide in the case when the simplex is the singleton. This alternative notion, which we denote by $\text{St}(\sigma)$ in the introduction, applies to the definition of superstar transitive, but will not be otherwise relevant to the content of this paper.
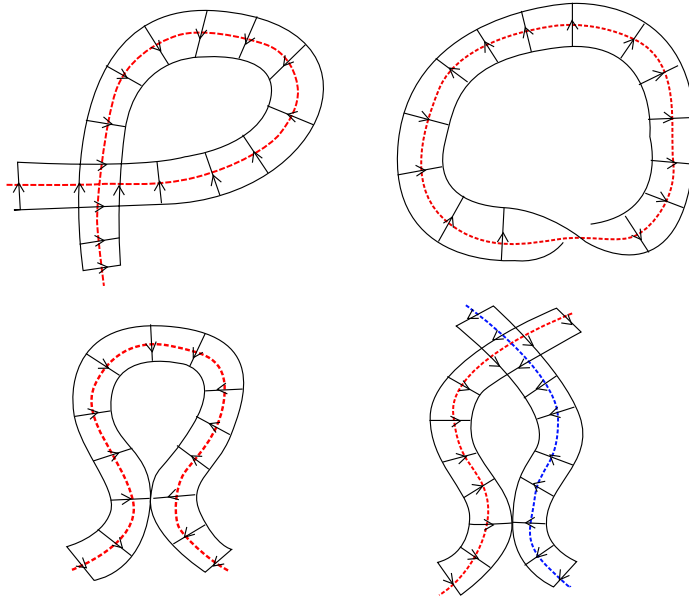
Figure 2: An illustration of the standard hyperplane pathologies. The dotted line depicts the topological realization of the hyperplane. The edges in the equivalence classes are given arrows indicating the direction. The top left depicts a self intersection. The top right depicts a 1–sided hyperplane, and the edges with the arrows reversed also belong to the equivalence class. The bottom left depicts a direct self-osculation. The bottom right depicts an interosculation.

The *adjacency map* for e is the natural isomorphism

$$\mathrm{ad}_e : \mathrm{Star}(v) \to \mathrm{Star}(u)$$

such that if $v \in \sigma$ then $\mathrm{ad}_e(\sigma)$ is the unique simplex such that $C(\mathrm{ad}_e(\sigma)) = C(\sigma)$. (This is referred to as the *transfer map* in [15].)

More generally, let x and y be 0–cubes in X that belong to some n–cube. Let C be the minimal such n–cube in X containing x and y. Let $\sigma_x \subseteq \mathrm{lk}(x)$ and $\sigma_y \subseteq \mathrm{lk}(y)$ be the simplices corresponding to C. Then we have a natural adjacency map for C given by the natural isomorphism

$$\mathrm{ad}_C : \mathrm{Star}(\sigma_x) \to \mathrm{Star}(\sigma_y)$$

such that if σ is a simplex in $\mathrm{lk}(x)$ containing σ_x, then ad_C on σ is induced by the composition of reflections in $C(\sigma)$ over the midcubes separating x and y. Note that $C(\mathrm{ad}_C(\sigma)) = C(\sigma)$.

Furthermore, suppose that x, y and z are 0–cubes in C such that C is the minimal cube containing x and z. If C_1 and C_2 are the minimal subcubes in C containing x, y and y, z respectively, then $\text{ad}_C = \text{ad}_{C_1} \circ \text{ad}_{C_2}$, where each ad_{C_i} is suitably restricted.

3 Constructing Δ–categories

This section will be devoted to constructing a Δ–category on a compact L–cube-complex X such that all finite-index subgroups of the hyperplane subgroups are separable in $\pi_1 X$, where L is the Kneser graph as specified in the statement of Theorem 0.1. We will assume, as stated in Section 2, that we have passed to a finite-index cover such that the hyperplanes are fully clean and 2–sided.

3.1 A note on notation

In what follows we will be constructing a category over a cube complex. We will be doing this by assigning objects to 0–cubes and assigning morphisms to each 1–cube. For example we might denote the morphism associated to e by ϕ_e. In this case, given an edge path $\gamma = (e_1, \ldots, e_n)$, we will let ϕ_γ denote the composition $\phi_{e_n} \circ \cdots \circ \phi_{e_1}$. If all the edges are parallel to a given hyperplane Λ, then we will call γ a *parallel path*.

3.2 Our objects

Let $n \geq 2$, $d \geq 1$ and Δ be a finite set with $|\Delta| = nd + 1$. Let X be a compact, nonpositively curved cube complex with 2–sided hyperplanes such that $\text{lk}(x)$ is isomorphic to $\mathfrak{K}_n(\Delta)$. To each 0–cube x in X let Δ_x be a copy of Δ, and identify $\text{lk}(x)$ with the associated Kneser complex $\mathfrak{K}_n(\Delta_x)$. Let Λ be a hyperplane incident to x. Let e be the 1–cube dual to Λ with $\tau e = x$. Let $v = \sigma(e)$ be the vertex in $\text{lk}(x)$ corresponding to e. The identification of $\text{lk}(x)$ with $\mathfrak{K}(\Delta_x)$ allows us to define $\Lambda_x := \mathfrak{s}(v) \subseteq \Delta_x$. We will also let $\mathfrak{s}(e) := \mathfrak{s}(v)$ when it is clear which 0–cube link we are working with. This is well defined since Λ is fully clean, so the 1–cube e is the only 1–cube dual to Λ incident to x.

3.3 Δ–categories on X

Definition 3.1 A Δ–*category on X* is a collection of bijections

$$\phi_e \colon \Delta_x \to \Delta_y,$$

one for each 1–cube $e = (x, y)$ in X, such that the following conditions are satisfied:

Figure 3: The square. The hyperplane Λ^1 is depicted as the vertical dotted line with the arrows on e_1 and e_1' giving the direction. The hyperplane Λ^2 is the horizontal dotted line with the arrows on e_2 and e_2' giving the direction.

(1) **Invertibility** If e is a directed one cube then $\phi_{\bar{e}} = \phi_e^{-1}$.

(2) Let $e_1 = (x, y)$, $e_2 = (y, z)$, $e_1' = (y', z)$, $e_2' = (x, y')$ be the edges bounding a square S, and Λ^i the hyperplane dual to e_i and e_i' (see Figure 3). Then:

    (a) **Commutativity** $\phi_{e_2} \circ \phi_{e_1} = \phi_{e_1'} \circ \phi_{e_2'}$.

    (b) **Parallel transport** $\phi_{e_1}(\Lambda_x^2) = \Lambda_y^2$ and $\phi_{e_2'}(\Lambda_x^1) = \Lambda_{y'}^1$.

Remark The parallel transport condition applied to all squares containing e allows us to deduce that $\phi_{e_1}(\Lambda_x^1) = \Lambda_y^1$.

Let $\{\phi_e\}$ be a Δ–category on X, and $f : \widehat{X} \to X$ a cover. By identifying each link $\mathrm{lk}(\hat{x})$ in $\widehat{X}$ with $\mathfrak{K}_n(\widehat{\Delta}_{\hat{x}})$, where $\widehat{\Delta}_{\hat{x}}$ is the copy of Δ assigned to $\hat{x}$, by Lemma 1.3 the induced isomorphism between the links

$$f_{\hat{x}} : \mathrm{lk}(\hat{x}) \to \mathrm{lk}(f(\hat{x}))$$

induces an isomorphism

$$\mathfrak{f}_{\hat{x}} : \widehat{\Delta}_{\hat{x}} \to \Delta_{f(\hat{x})}.$$

Thus we can lift the Δ–category $\{\phi_e\}$ on X to a unique Δ–category on $\widehat{X}$ such that the following diagram commutes, for each 1–cube $\hat{e} = (\hat{x}, \hat{y})$ in $\widehat{X}$ mapping to $e = (x, y)$ in X:

$$
\begin{array}{ccc}
\widehat{\Delta}_{\hat{x}} & \xrightarrow{\hat{\phi}_{\hat{e}}} & \widehat{\Delta}_{\hat{y}} \\
{\scriptstyle \mathfrak{f}_{\hat{x}}} \downarrow & & \downarrow {\scriptstyle \mathfrak{f}_{\hat{y}}} \\
\Delta_x & \xrightarrow{\phi_e} & \Delta_y
\end{array}
$$

It is straightforward to verify that $\{\hat{\phi}_{\hat{e}}\}$ satisfies the invertibility and commutativity conditions, since $\mathfrak{f}_{\hat{x}}$ is invertible, and since the squares in X lift to squares in $\widehat{X}$. Parallel

transport holds for $\{\hat{\phi}_{\hat{e}}\}$ by tracing the correspondence of n–element subsets of $\hat{\Delta}_{\hat{x}}$ to vertices in $\mathrm{lk}(\hat{x})$, which then correspond to hyperplanes incident to $\hat{x}$. Consider a square in $\hat{X}$ covering the square in Figure 3, labelled with the vertices $\hat{x}$, $\hat{y}$, $\hat{y}'$ and $\hat{z}$, and bounded by edges $\hat{e}_1$, $\hat{e}_2$, $\hat{e}_1'$ and $\hat{e}_2'$. Then for the hyperplane $\hat{\Lambda}^2$ covering Λ^2, we deduce that

$$\hat{\phi}_{\hat{e}_1}(\hat{\Lambda}_{\hat{x}}^2) = \hat{\phi}_{\hat{e}_1} \circ \mathfrak{f}_{\hat{x}}^{-1}(\Lambda_x^2) = \mathfrak{f}_{\hat{y}}^{-1} \circ \phi_{e_1}(\Lambda_x^2) = \mathfrak{f}_{\hat{y}}^{-1}(\Lambda_y^2) = \hat{\Lambda}_{\hat{y}}^2.$$

The second equality follows from commutativity of the above square, and the third from parallel transport for $\{\phi_e\}$ in X. The corresponding conclusion follows similarly for $\hat{\Lambda}^1$.

3.4 Constructing a Δ–category

We will construct our Δ–category in two stages. In the first stage we will define functions ϕ_e^* that will be defined on subsets of the domain Δ_x. We note that in this section we will be composing functions whose domain and ranges will be subsets of larger sets. In this case the composition will be given by restricting to the intersection of the corresponding domains and ranges.

Lemma 3.2 *There exists a unique family of functions*

$$\{\phi_e^* \colon (\Delta_x - \Lambda_x) \to (\Delta_y - \Lambda_y) \mid \Lambda \text{ is dual to } e = (x, y) \in X^{(1)}\}$$

such that:

(1) $\phi_{\bar{e}}^* = (\phi_e^*)^{-1}$.

(2) *Let $e_1 = (x, y)$, $e_2 = (y, z)$, $e_1' = (y', z)$ and $e_2' = (x, y')$ be the edges bounding a square S, and Λ^i the hyperplane dual to e_i and e_i' (see Figure 3). Then*

(a) *after suitably restricting domains,*

$$\phi_{e_2}^* \circ \phi_{e_1}^* = \phi_{e_2'}^* \circ \phi_{e_1'}^* \colon (\Delta_x - \Lambda_x^1 - \Lambda_x^2) \to (\Delta_z - \Lambda_z^1 - \Lambda_z^2),$$

(b) $\phi_{e_1}^*(\Lambda_x^2) = \Lambda_y^2$ *and* $\phi_{e_2'}^*(\Lambda_x^1) = \Lambda_{y'}^1$.

Proof Let $e = (x, y)$ be a directed 1–cube in X dual to Λ. Let v be the vertex in $\mathrm{lk}(x)$ corresponding to $\bar{e}$, and u be the vertex in $\mathrm{lk}(y)$ corresponding to e. Then $\mathrm{Star}(v)$ decomposes as the simplicial join $v * \mathfrak{K}_n(\Delta_x - \Lambda_x)$ and similarly $\mathrm{Star}(u)$ decomposes as $u * \mathfrak{K}_n(\Delta_y - \Lambda_y)$. Thus the adjacency map ad_e restricts to an isomorphism

$$\mathfrak{K}_n(\Delta_x - \Lambda_x) \to \mathfrak{K}_n(\Delta_y - \Lambda_y).$$

Since $|\Delta_x - \Lambda_x| = |\Delta_y - \Lambda_y| = n(d-1) + 1$, by Lemma 1.3 this isomorphism is induced by the bijection

$$\phi_e^* : (\Delta_x - \Lambda_x) \to (\Delta_y - \Lambda_y).$$

(This requires checking that $n(d-1) + 1 \neq 2n$ for $n \geq 2$ and $d \geq 1$.)

In the case that $d = 1$ there are no squares in X, so conditions (2)(a)–(b) are satisfied automatically. So we assume $d \geq 2$. Suppose that $e_1 = (x, y)$, $e_2 = (y, z)$, $e_1' = (y', z)$ and $e_2' = (x, y')$ are the edges bounding a square S, and Λ^i is the hyperplane dual to e_i and e_i' (see Figure 3). We now check that conditions (2)(a)–(b) are satisfied.

Verifying (2)(b) follows from observing that $\Lambda_x^2 \subseteq \Delta_x - \Lambda_x^1$ corresponds to a vertex $u \in \mathrm{lk}(x) = \mathfrak{K}_n(\Delta_x)$ and $\Lambda_y^2 \subseteq \Delta_y - \Lambda_y^1$ corresponds to a vertex v in $\mathrm{lk}(y) = \mathfrak{K}_n(\Delta_y)$ such that $\mathrm{ad}_{e_1}(u) = v$. (Stare at Figure 3.) Thus $\phi_{e_1}^*(\Lambda_x^2) = \Lambda_y^2$ and similarly $\phi_{e_2'}^*(\Lambda_x^1) = \Lambda_{y'}^1$.

We now consider (2)(a). Observe that (2)(b) implies

$$\phi_{e_2}^* \circ \phi_{e_1}^* ((\Delta_x - \Lambda_x^2) - \Lambda_x^1) = \phi_{e_2}^*((\Delta_y - \Lambda_y^1) - \Lambda_y^2) = \Delta_z - \Lambda_z^1 - \Lambda_z^2.$$

Combined with the similar set of equalities for $\phi_{e_2'}^* \circ \phi_{e_1'}^*$ this verifies (2)(b) when $d = 2$ since there is only one possible map between singletons.

In the case that $d > 2$, let $\sigma_x \subseteq \mathrm{lk}(x)$, $\sigma_y \subseteq \mathrm{lk}(y)$, $\sigma_{y'} \subseteq \mathrm{lk}(y')$ and $\sigma_z \subseteq \mathrm{lk}(z)$ denote the 1–simplices corresponding to the square S. We know that

$$\mathrm{ad}_S = \mathrm{ad}_{e_2} \circ \mathrm{ad}_{e_1} = \mathrm{ad}_{e_1'} \circ \mathrm{ad}_{e_2'} : \mathrm{Star}(\sigma_x) \to \mathrm{Star}(\sigma_z).$$

We also have the decomposition

$$\mathrm{Star}(\sigma_x) = \sigma_x * \mathfrak{K}_n(\Delta_x - \Lambda_x^1 - \Lambda_x^2)$$

and similar decompositions for the stars of σ_y, $\sigma_{y'}$ and σ_z. The adjacency map ad_S therefore restricts to an isomorphism

$$\mathfrak{K}_n(\Delta_x - \Lambda_x^1 - \Lambda_x^2) \to \mathfrak{K}_n(\Delta_z - \Lambda_z^1 - \Lambda_z^2)$$

which, by Lemma 1.3, is induced by an isomorphism

$$(\Delta_x - \Lambda_x^1 - \Lambda_x^2) \to (\Delta_z - \Lambda_z^1 - \Lambda_z^2)$$

that must coincide with the compositions $\phi_{e_2}^* \circ \phi_{e_1}^*$ and $\phi_{e_1'}^* \circ \phi_{e_2'}^*$, as their restrictions induce the same isomorphism. Thus $\phi_{e_2}^* \circ \phi_{e_1}^* = \phi_{e_1'}^* \circ \phi_{e_2'}^*$.

Finally, we show uniqueness of the family $\{\phi_e^*\}$. Note that $|\Delta_x - \Lambda_x| = (n-1)d + 1$. Therefore, if $n = 1$ then uniqueness is trivial as the maps are between singletons. Otherwise, for $n > 1$ each element of $\Delta_x - \Lambda_x$ is given by the intersection of the n–element subsets Λ_y that contain the given element and are disjoint from Λ_x. Thus, condition (2)(b) applied to each these Λ_y allows us to deduce that ϕ_e^* is uniquely determined on the given element of $\Delta_x - \Lambda_x$. $\qquad\square$

We will refer to the maps $\{\phi_e^*\}$ as the *pre–Δ–category*. Note that if $f : \hat{X} \to X$ is a cover, then we can lift the pre–Δ–category to $\hat{X}$ and check that the conditions are satisfied, in the same way we checked for the Δ–category. Alternatively, since such pre–Δ–categories are unique, we could instead verify that the following square commutes:

$$
\begin{array}{ccc}
(\hat{\Delta}_{\hat{x}} - \hat{\Lambda}_{\hat{x}}) & \xrightarrow{\;\hat{\phi}_{\hat{e}}^*\;} & (\hat{\Delta}_{\hat{y}} - \hat{\Lambda}_{\hat{y}}) \\[2pt]
{\scriptstyle f_{\hat{x}}}\downarrow & & \downarrow{\scriptstyle f_{\hat{y}}} \\[2pt]
(\Delta_x - \Lambda_x) & \xrightarrow{\;\phi_e^*\;} & (\Delta_y - \Lambda_y)
\end{array}
$$

This would follow from the parallel transport conditions and the correspondence between hyperplanes and the corresponding subsets of $\hat{\Delta}_{\hat{x}}$, in a similar fashion to the argument given for lifting Δ–categories.

3.5 The hyperplane parallel holonomy

As a consequence of Lemma 3.2 we deduce that if an edge $e = (x, y)$ is parallel to Λ then we have a bijection

$$\psi_e : \Lambda_x \to \Lambda_y$$

obtained by restricting ϕ_e^* as given by Lemma 3.2. Indeed, if Λ' is the hyperplane dual to e, then $\Lambda_x \subseteq \Delta_x - \Lambda_x'$. We note that this is a category, with the n–element set Λ_x associated to each vertex x that Λ is incident to, and there is a morphism ψ_e associated to each edge e parallel to Λ. In fact, since Λ is 2–sided, there is a category corresponding to each side.

Thus if we fix a choice of side of Λ and a 0–cube p in Λ as a basepoint, we obtain a *parallel holonomy*

$$\Psi_p : \pi_1(\Lambda, p) \to \mathrm{Sym}(\Lambda_x).$$

If e' is the edge dual to Λ with midpoint p such that $\tau e' = x$ lies on the given side, this holonomy is given by identifying Λ with the side of the hyperplane carrier containing

the basepoint x, and letting the equivalence class of a parallel path $[\gamma] = [e_1, \dots, e_n]$ based at x map to

$$\Psi_p([\gamma]) = \psi_\gamma,$$

where ψ_γ denotes the composition $\psi_{e_n} \circ \cdots \circ \psi_{e_1}$. Conditions (1) and (2)(a) in Lemma 3.2 ensure that this does not depend on the choice of representative.

We note that the triviality of the holonomy does not depend on the choice of basepoint p (but may depend on the side of the carrier that is chosen). Indeed, given another 1–cube e'' dual to Λ, with $\tau e'' = y$ on the same side of Λ, with midpoint p', we can check the following diagram commutes:

$$
\begin{array}{ccc}
\pi_1(\Lambda, p) & \xrightarrow{\ \Psi_p\ } & \mathrm{Sym}(\Lambda_x) \\
\downarrow & & \downarrow \\
\pi_1(\Lambda, p') & \xrightarrow{\ \Psi_{p'}\ } & \mathrm{Sym}(\Lambda_y)
\end{array}
$$

We have chosen some path γ connecting x to y in $\tau(\Lambda)$. The left vertical map is given by conjugating closed loops by $[\gamma]$, in the standard fashion, and the right vertical map is given by conjugating by ψ_γ.

The kernel of Ψ_p is a finite-index normal subgroup of $\pi_1(\Lambda)$, and by the assumptions of Theorem 0.1 will be separable in $\pi_1 X$.

Lemma 3.3 *There exists a finite cover $\hat{X} \to X$ such that the parallel holonomies in $\hat{X}$ are trivial.*

Proof Let Ψ be a parallel holonomy for some hyperplane Λ, and some choice of side and basepoint. The kernel of Ψ is a finite-index normal subgroup of $\pi_1(\Lambda)$, and therefore, by the assumption of Theorem 0.1, will be separable in $\pi_1 X$. Let $\{\mathrm{id}, g_1, \dots, g_\ell\}$ be a minimal set of representatives for the left cosets of $\mathrm{ker}(\Psi)$ in $\pi_1(\Lambda)$. As $g_i \notin \mathrm{ker}(\Psi)$, by separability there exists a finite-index subgroup $N_i \leqslant \pi_1(X)$ such that $\mathrm{ker}(\Psi) \subseteq N_i$ and $g_i \notin N_i$. Thus $\mathrm{ker}(\Psi) = \bigcap_{i=1}^{\ell} N_i \cap \pi_1(\Lambda)$, since we know $\mathrm{ker}(\Psi) \subseteq \bigcap_{i=1}^{\ell} N_i$ and that if $g_i h \in \bigcap_{i=1}^{\ell} N_i \cap \pi_1(\Lambda)$, where $h \in \mathrm{ker}(\Psi)$, then $g_i \in \bigcap_{i=1}^{\ell} N_i$. The normal core, $\mathrm{Core}(\bigcap_{i=1}^{\ell} N_i)$, is a finite-index normal subgroup of $\pi_1 X$ such that $\pi_1(\Lambda) \cap \mathrm{Core}(\bigcap_{i=1}^{\ell} N_i)$ is contained in $\mathrm{ker}(\Psi)$.

By repeating this for each side of each hyperplane, and intersecting all the resulting normal cores, we obtain a finite-index normal subgroup $N \leqslant \pi_1(X)$ such that for each hyperplane Λ, the intersection $N \cap \pi_1(\Lambda)$ is contained in the kernel of the parallel holonomies on either side of Λ. Then the desired finite cover $f : \hat{X} \to X$ is given by N.

Let $\{\hat{\phi}_{\hat{e}}^*\}$ denote the lift of the pre–Δ–category on X to $\hat{X}$. Then the following diagram commutes, where hyperplane $\hat{\Lambda}$ covers Λ, and the bottom arrow is the isomorphism induced by conjugation by $\mathfrak{f}_{\hat{x}}$:

$$
\begin{array}{ccc}
\pi_1(\hat{\Lambda}) & \xrightarrow{\ f_*\ } & \pi_1(\Lambda) \\
{\scriptstyle \hat{\Psi}_{\hat{p}}} \downarrow & & \downarrow {\scriptstyle \Psi_p} \\
\mathrm{Sym}(\hat{\Delta}_{\hat{x}}) & \longrightarrow & \mathrm{Sym}(\Delta_x)
\end{array}
$$

Indeed, if we take a combinatorial path $[\hat{\gamma}]$ given by the edge sequence $\hat{e}_1, \ldots, \hat{e}_n$ that traversed the vertices $\hat{x} = \hat{x}_0, \hat{x}_1, \ldots, \hat{x}_{n-1}$, and let $f([\hat{\gamma}]) = [\gamma]$ with $f(\hat{e}_i) = e_i$ and $f(\hat{x}_i) = x_i$, we deduce that

$$
\begin{aligned}
\Psi_p \circ f_*([\hat{\gamma}]) = \Psi_p([\gamma]) = \psi_\gamma &= \psi_{e_n} \circ \cdots \circ \psi_{e_1} \\
&= \mathfrak{f}_{\hat{x}_0} \circ \hat{\psi}_{\hat{e}_n} \circ \mathfrak{f}_{\hat{x}_{n-1}}^{-1} \circ \mathfrak{f}_{\hat{x}_{n-1}} \circ \cdots \circ \mathfrak{f}_{\hat{x}_1} \circ \hat{\psi}_{\hat{e}_1} \circ \mathfrak{f}_{\hat{x}_0}^{-1} \\
&= \mathfrak{f}_{\hat{x}_0} \circ \hat{\psi}_{\hat{e}_n} \circ \cdots \circ \hat{\psi}_{\hat{e}_1} \circ \mathfrak{f}_{\hat{x}_0}^{-1} \\
&= \mathfrak{f}_{\hat{x}_0} \circ \hat{\Psi}_{\hat{p}} \circ \mathfrak{f}_{\hat{x}_0}^{-1}.
\end{aligned}
$$

Thus the square commutes and the parallel holonomies in $\hat{X}$ are trivial. $\qquad\square$

3.6 Extending the maps ϕ_e^*

By Lemma 3.3, we now assume that we have passed to a suitable finite cover such that X has trivial parallel holonomies in its pre–Δ–category. Given an edge e dual to Λ, it remains to extend ϕ_e^*, and this means making a choice of bijection $\Lambda_x \to \Lambda_y$. We can certainly make such choices so that the inversion condition (2)(a) is satisfied, and condition (2)(b) holds as it holds for ϕ_e^*. It therefore remains to ensure we can make our choices so that the commutativity condition (2)(a) is satisfied.

For each hyperplane Λ let $e = (x, y)$ be a choice of edge dual to Λ. We make a choice of map

$$
\phi_e^\circ \colon \Lambda_x \to \Lambda_y
$$

that extends ϕ_e^* to ϕ_e.

Suppose that e' is some other edge dual to Λ such that $\tau e'$ lies on the same side of Λ as τe. Then let $\gamma = (e_1, \ldots, e_p)$ be an edge path parallel to Λ that connects τe to $\tau e'$. We also let $\gamma' = (e_1', \ldots, e_q')$ be an edge path parallel to Λ that connects ιe to $\iota e'$. Then we define

$$
\phi_{e'}^\circ = \psi_{e_p} \circ \cdots \circ \psi_{e_1} \circ \phi_e^\circ \circ \psi_{e_1'}^{-1} \circ \cdots \circ \psi_{e_q'}^{-1},
$$

where ψ_{e_i} and $\psi_{e_i'}$ are the parallel holonomies on either side of Λ. Since the parallel holonomies are trivial, $\phi_{e'}^{\circ}$ will not depend on the choice of paths γ and γ'. We let $\phi_{\bar{e}}^{\circ} = (\phi_e^{\circ})^{-1}$ and recover that $\phi_{\bar{e}'}^{\circ} = (\phi_{e'}^{\circ})^{-1}$.

It remains to check that $\{\phi_e\}$, as defined, satisfy our commutativity relations. Let $e_1 = (x, y)$, $e_2 = (y, z)$, $e_1' = (y', z)$ and $e_2' = (x, y')$ be edges bounding a square, and let Λ^i be the hyperplane dual to e_i and e_i' (see Figure 3). Then we consider the separate cases

$$\phi_{e_2} \circ \phi_{e_1} = \begin{cases} \phi_{e_2}^* \circ \phi_{e_1}^* : (\Delta_x - \Lambda_x^1 - \Lambda_x^2) \to (\Delta_z - \Lambda_z^1 - \Lambda_z^2), \\ \phi_{e_2}^\circ \circ \phi_{e_1}^* : \Lambda_x^2 \to \Lambda_z^2, \\ \phi_{e_2}^* \circ \phi_{e_1}^\circ : \Lambda_x^1 \to \Lambda_z^1, \\ \phi_{e_2}^\circ \circ \phi_{e_1}^\circ : \varnothing \to \varnothing. \end{cases}$$

It follows from Lemma 3.2 that $\phi_{e_2}^* \circ \phi_{e_1}^* = \phi_{e_1'}^* \circ \phi_{e_2'}^*$. By considering the parallel holonomies with respect to Λ^2 we can see that

$$\phi_{e_2}^\circ \circ \phi_{e_1}^* = \phi_{e_2}^\circ \circ \psi_{e_1} = \psi_{e_1'} \circ \phi_{e_2'}^\circ = \phi_{e_1'}^* \circ \phi_{e_2'}^\circ.$$

A similar sequence of equalities gives that $\phi_{e_2}^* \circ \phi_{e_1}^\circ = \phi_{e_1'}^\circ \circ \phi_{e_2'}^*$. Altogether this allows us to conclude that $\phi_{e_2} \circ \phi_{e_1} = \phi_{e_1'} \circ \phi_{e_2'}$, and that $\{\phi_e\}$ is a Δ–category, and that we have proven the following:

Proposition 3.4 *Let $n \geq 2$, $d \geq 1$ and Δ be a finite set of cardinality $nd + 1$. Let L be the Kneser complex $\mathfrak{K}_n(\Delta)$. Suppose that X is an L–cube-complex such that hyperplane subgroups have separable finite-index subgroups. Then there exists a finite cover $\hat{X} \to X$, such that there is a Δ–category over X.*

4 The holonomy

Given a Δ–category $\{\phi_e\}$ for X we obtain a holonomy map

$$\Phi_x : \pi_1(X, x) \to \mathrm{Sym}(\Delta_x),$$

where the homotopy class $[\gamma] = [e_1, \dots, e_n]$ of the edge path based at x has image

$$\Phi_x([\gamma]) = \phi_\gamma.$$

The invertibility and commutativity conditions guarantee that this does not depend on the choice of representative of the homotopy class. Note that if Φ_x is trivial, then the

holonomy is trivial with respect to any basepoint since the following diagram commutes:

$$
\begin{array}{ccc}
\pi_1(X, x) & \xrightarrow{\;\Phi_x\;} & \mathrm{Sym}(\Delta_x) \\
\downarrow & & \downarrow \\
\pi_1(X, y) & \xrightarrow{\;\Phi_y\;} & \mathrm{Sym}(\Delta_y)
\end{array}
$$

If γ is an edge path connecting x to y, then the vertical left arrow is the isomorphism given by conjugating a homotopy class of based loops by $[\gamma]$, and the vertical right arrow is the isomorphism given by conjugating by ϕ_γ.

The kernel of Φ_x is a finite-index normal subgroup of $\pi_1 X$ and corresponds to a finite-sheeted, regular cover $f : \widehat{X} \to X$. Lift the Δ–category on X to a Δ–category $\{\hat{\phi}_{\hat{e}}\}$ on $\widehat{X}$. We can check that the following diagram commutes:

$$
\begin{array}{ccc}
\pi_1(\widehat{X}, \hat{x}) & \xrightarrow{\;\widehat{\Phi}_{\hat{x}}\;} & \mathrm{Sym}(\Delta_{\hat{x}}) \\
f_* \downarrow & & \downarrow \\
\pi_1(X, x) & \xrightarrow{\;\Phi_x\;} & \mathrm{Sym}(\Delta_x)
\end{array}
$$

The 0–cube $\hat{x}$ is chosen so that $f(\hat{x}) = x$, and the right vertical arrow is the isomorphism given by conjugation by $\mathfrak{f}_{\hat{x}}$. Thus we conclude that the holonomy $\widehat{\Phi}_x$ on $\widehat{X}$ is trivial. If the holonomy on X obtained from a Δ–category is trivial, then we say that the Δ–category itself is *flat*.

4.1 Constructing the orbicover

Proposition 4.1 *Let $L = \mathfrak{K}_n(\Delta)$ where $|\Delta| = nd + 1$. Let X be a compact L–cube-complex that has a flat Δ–category on X. Then there is an orbicomplex cover $X \to X_L$, where $X_L = W_L \backslash D(L)$.*

Proof Let $\{\phi_e\}$ be the flat Δ–category on X. For a basepoint x, fix an identification $q_x \colon \Delta_x \to \Delta$. For any other 0–cube y in X, let $q_y = q_x \circ \phi_\gamma$ where γ is an edge path connecting y to x. Note that q_y does not depend on the choice of γ since the Δ–category is flat.

We will prove the claim by producing an orbicomplex cover $X \to X_L$. First we map all 0–cubes in X to the unique 0–cube in X_L. We can extend X to the 1–skeleton of X by mapping each 1–cube $e = (x, y)$ dual to Λ to the half 1–cube corresponding to $q(\Lambda_x)$. This makes sense since we know that $q_x(\Lambda_x) = q_y \circ \phi_e(\Lambda_x) = q_y(\Lambda_y)$ by the remark following Definition 3.1, so e and $\bar{e}$ are mapped to the same half edge.

Now we want to extend $X^{(1)} \to X_L$ to the 2–skeleton. Let $e_1 = (x, y)$, $e_2 = (y, z)$, $e'_1 = (y', z)$ and $e'_2 = (x, y')$ be the directed 1–cubes bounding a square S in X such that e_i and e'_i are dual to the hyperplane Λ^i (as in Figure 3). We want to show that e_i and e'_i map to the same half edge, and the e_1 and e'_2 map to half edges that bound a quarter-square in X_L. The first fact follows from the parallel transport property since $\phi_{e_1}(\Lambda^2_x) = \Lambda^2_y$ so $q_x(\Lambda^2_x) = q_y \circ \phi_{e_1}(\Lambda^2_x) = q_y(\Lambda^2_y)$. The second follows from the fact that $\Lambda^1_x \cap \Lambda^2_x = \varnothing$ since e_1 and e'_2 bound the corner of a square, so $q_x(\Lambda^1_x) \cap q_x(\Lambda^2_x) = \varnothing$.

It is immediate that we can extend $X^{(2)} \to X_L$ to the entire skeleton since the higher dimension cubes are entirely determined by the 1–skeleton. In this particular case, we have an orbicovering since the induced maps on the vertex links are isomorphisms. Thus we can lift this orbicovering to an isomorphism $\widetilde{X} \to D(L)$ such that the deck transformation group $\pi_1(X)$ is a subgroup of W_L. □

Proof of Theorem 0.1 Let X_1 and X_2 be our L–cube-complexes. Finite-index subgroups of the hyperplane subgroups are separable, so by Proposition 3.4 there is a finite cover $X'_i \to X_i$ such that there is a Δ–category over X'_i. By considering the holonomy given by the Δ–category, we can pass to a further finite cover $\widehat{X}_i \to X'_i$ such that the induced Δ–category is flat. By Proposition 4.1, there are finite orbicovers $f_i \colon \widehat{X}_i \to X_L$. The common cover is then obtained by taking the intersection of the corresponding deck transformation groups inside of W_L. □

References

[1] **I Agol**, *The virtual Haken conjecture*, Doc. Math. 18 (2013) 1045–1087 MR Zbl

[2] **D Angluin**, *Local and global properties in networks of processors (extended abstract)*, from "Proceedings of the twelfth annual ACM symposium on theory of computing", Association for Computing Machinery, New York, NY (1980) 82–93

[3] **H Bass**, **R Kulkarni**, *Uniform tree lattices*, J. Amer. Math. Soc. 3 (1990) 843–902 MR Zbl

[4] **M Bourdon**, **H Pajot**, *Rigidity of quasi-isometries for some hyperbolic buildings*, Comment. Math. Helv. 75 (2000) 701–736 MR Zbl

[5] **M R Bridson**, **A Haefliger**, *Metric spaces of non-positive curvature*, Grundl. Math. Wissen. 319, Springer (1999) MR Zbl

[6] **M Burger**, **S Mozes**, *Finitely presented simple groups and products of trees*, C. R. Acad. Sci. Paris Sér. I Math. 324 (1997) 747–752 MR Zbl

[7] **P Dani**, *The large-scale geometry of right-angled Coxeter groups*, from "Handbook of group actions, V" (L Ji, A Papadopoulos, S-T Yau, editors), Adv. Lect. Math. 48, International, Somerville, MA (2020) 107–141 MR Zbl

[8] **M W Davis**, *The geometry and topology of Coxeter groups*, London Mathematical Society Monographs Series 32, Princeton Univ. Press (2008) MR Zbl

[9] **C Godsil**, **G Royle**, *Algebraic graph theory*, Graduate Texts in Mathematics 207, Springer (2001) MR Zbl

[10] **M Gromov**, *Hyperbolic groups*, from "Essays in group theory" (S M Gersten, editor), Math. Sci. Res. Inst. Publ. 8, Springer (1987) 75–263 MR Zbl

[11] **F Haglund**, *Commensurability and separability of quasiconvex subgroups*, Algebr. Geom. Topol. 6 (2006) 949–1024 MR Zbl

[12] **F Haglund**, **D T Wise**, *Special cube complexes*, Geom. Funct. Anal. 17 (2008) 1551–1620 MR Zbl

[13] **J Huang**, *Commensurability of groups quasi-isometric to RAAGs*, Invent. Math. 213 (2018) 1179–1247 MR Zbl

[14] **R Kropholler**, *Special cube complexes*, from "Geometric and cohomological group theory" (P H Kropholler, I J Leary, C Martínez-Pérez, B E A Nucinkis, editors), London Math. Soc. Lecture Note Ser. 444, Cambridge Univ. Press (2018) 46–66 MR Zbl

[15] **N Lazarovich**, *On regular* CAT(0) *cube complexes and the simplicity of automorphism groups of rank-one* CAT(0) *cube complexes*, Comment. Math. Helv. 93 (2018) 33–54 MR Zbl

[16] **F T Leighton**, *Finite common coverings of graphs*, J. Combin. Theory Ser. B 33 (1982) 231–238 MR Zbl

[17] **G Moussong**, *Hyperbolic Coxeter groups*, PhD thesis, Ohio State University (1988) MR Available at `https://www.proquest.com/docview/303686972`

[18] **W D Neumann**, *On Leighton's graph covering theorem*, Groups Geom. Dyn. 4 (2010) 863–872 MR Zbl

[19] **T Piesk**, public domain image Available at `https://en.wikipedia.org/wiki/Petersen_graph#/media/File:Kneser_graph_KG(5,2).svg`

[20] **M Sageev**, CAT(0) *cube complexes and groups*, from "Geometric group theory" (M Bestvina, M Sageev, K Vogtmann, editors), IAS/Park City Math. Ser. 21, Amer. Math. Soc., Providence, RI (2014) 7–54 MR Zbl

[21] **S Shepherd**, *Two generalisations of Leighton's theorem*, Groups Geom. Dyn. 16 (2022) 743–778 MR Zbl

[22] **D T Wise**, *Non-positively curved squared complexes*: *Aperiodic tilings and non-residually finite groups*, PhD thesis, Princeton University (1996) MR Available at `https://www.proquest.com/docview/304259249`

[23] **D T Wise**, *From riches to raags: 3–manifolds, right-angled Artin groups, and cubical geometry*, CBMS Regional Conference Series in Mathematics 117, Amer. Math. Soc., Providence, RI (2012) MR Zbl

[24] **D J Woodhouse**, *Revisiting Leighton's theorem with the Haar measure*, Math. Proc. Cambridge Philos. Soc. 170 (2021) 615–623 MR Zbl

University of Oxford
Oxford, United Kingdom

mr.daniel.woodhouse@gmail.com

www.djwoodhouse.com

ALGEBRAIC & GEOMETRIC TOPOLOGY

msp.org/agt

ALGEBRAIC & GEOMETRIC TOPOLOGY