SIMPLICES AT INFINITY IN CAT(0) CUBE COMPLEXES

MARK F. HAGEN

Abstract. These are notes for lectures given in Nice in March 2016, on simplicial boundaries of CAT(0) cube complexes. They cover the definition, basic properties, some applications to cubulated groups, and some open problems.

Contents

Motivation and outline

"Question A". Here is an arcane-looking question that I hope to show is interesting:

Question (Question A')**.** Suppose that the proper CAT(0) cube complex $\mathcal{X}$ admits a geometric action by some group G. Consider the smallest collection $\mathfrak{F}$ of convex subcomplexes so that:

(1) $\mathcal{N}(H) \in \mathfrak{F}$ for each hyperplane H, where $\mathcal{N}(H)$ denotes the carrier;
(2) if $F, F' \in \mathfrak{F}$, then $\mathfrak{g}_F(F') \in \mathfrak{F}$, where $\mathfrak{g}_F(F')$ is the convex subcomplex of F spanned by the vertices of F which are closest-point projections on F of vertices of F'.

Does there exist $N \in \mathbb{N}$ so that each $x \in \mathcal{X}$ lies in at most N elements of $\mathfrak{F}$?

These notes are a circuitous route to this question, via the *simplicial boundary* $\partial_\triangle\mathcal{X}$ of $\mathcal{X}$, which is a combinatorial analogue of the Tits boundary, introduced for quite different reasons.

Date: February 26, 2016.

A positive answer to Question A' would have implications for the structure of $\partial_{\triangle}\mathcal{X}$ when $\mathcal{X}$ admits a proper, cocompact group action. Specifically, we would be able to use $\partial_{\triangle}\mathcal{X}$ to study quasi-isometry invariants like thickness and divergence without introducing extra hypotheses. This application of Question A' and this viewpoint on $\partial_{\triangle}\mathcal{X}$ are not the main focus of these notes (but see the problems at the end).

Instead, it turns out that when Question A' has a positive answer, there is an alternative description of $\partial_{\triangle}\mathcal{X}$ in terms of a collection of quasi-trees associated to the elements of $\mathfrak{F}$. This is intimately related to the "hierarchical structure" of such cube complexes discussed in [BHS14, BHS15], and this new way of looking at $\partial_{\triangle}\mathcal{X}$ is a special case of what can be done for general "hierarchically hyperbolic spaces" [DHS15]. One consequence of this is that the simplicial boundary can be retopologized in a fairly gentle way (i.e. the identity from $\partial_{\triangle}\mathcal{X}$ to the newly-topologized thing is an embedding on simplices), yielding a compactification of $\mathcal{X}$ that encodes more data than the visual boundary or the Roller compactification. One can then do various cool things. As an example, we'll consider G–stationary measures on this compactification to give a new proof of the Caprace-Sageev rank-rigidity theorem [CS11] (in many cases).

Question A' is reasonable because it has a positive answer for every known example. For instance, if $\overline{\mathcal{X}}$ is a compact nonpositively-curved cube complex admitting a local isometry to the Salvetti complex S_Γ of the right angled Artin group A_Γ, for some graph Γ, then Question A' has a positive answer for the universal cover of $\overline{\mathcal{X}}$. Hence, here's a punchier question:

Question (Question A)**.** Let $\mathcal{X}$ be a proper CAT(0) cube complex admitting a proper, cocompact action by some group G. Does there exist a finite graph Γ and a cubical embedding $\mathcal{X} \to \widetilde{S}_\Gamma$ whose image is a convex subcomplex?

(It's called "Question A" because it is Question A in [BHS15].) Here, $\widetilde{S}_\Gamma$ is the universal cover of the Salvetti complex S_Γ of the right-angled Artin group A_Γ. Observe that we are *not* asking for any kind of relationship between G and A_Γ, or any kind of equivariance of the map. This question appears to be less nutty than it perhaps looks.

Outline. In Part 1 defines the simplicial boundary $\partial_{\triangle}\mathcal{X}$ of $\mathcal{X}$, describes some of its basic properties, and discusses how it compares to other boundaries. This part of the notes also covers *visibility*, which is about associating simplices in $\partial_{\triangle}\mathcal{X}$ to geodesic rays in $\mathcal{X}$. Whether visibility is guaranteed by geometric group actions is unknown and related to Question A.

Part 2 is about "rank-one" phenomena from the boundary viewpoint. This is closely related to another object associated to $\mathcal{X}$, the *contact graph* $\mathcal{C}\mathcal{X}$, which is the intersection graph of the hyperplane-carriers. This graph is always a quasi-tree, on which $\mathrm{Aut}(\mathcal{X})$ acts [Hag14b]. The simplicial boundary was actually introduced as a bookkeeping device in an attempt to answer the question of when $\mathcal{C}\mathcal{X}$ is bounded (as it is when, for instance, $\mathcal{X}$ splits as a nontrivial product).

Geodesic rays in $\mathcal{X}$ project to subgraphs in $\mathcal{C}\mathcal{X}$ in a natural way, so one way to answer the question of boundedness of $\mathcal{C}\mathcal{X}$ is to look for a ray γ in $\mathcal{X}$ whose shadow in $\mathcal{C}\mathcal{X}$ is unbounded. This turns out to hold for any γ for which the obvious obstructions are absent: γ must be "rank-one" in the appropriate sense, and it must not lie uniformly close to any hyperplane. "Rank-one" turns out to mean the same thing as "γ represents an isolated point in $\partial_{\triangle}\mathcal{X}$". We'll state some results, which make use of lemmas of Caprace-Sageev [CS11], that relate $\partial_{\triangle}\mathcal{X}$, isometries of $\mathcal{X}$, their actions on $\mathcal{C}\mathcal{X}$, and product decompositions of $\mathcal{X}$.

We then re-examine the projection of γ to $\mathcal{C}\mathcal{X}$ from a more sophisticated point of view. Roughly, if γ fails to project quasi-geodesically to $\mathcal{C}\mathcal{X}$, then a more refined version of the above analysis provides a collection of hyperplanes on which γ has large projection, and we can consider (roughly) the projection of γ to the contact graphs of these hyperplanes, etc. To make all of this work out technically, we need to impose conditions on $\mathcal{X}$, which would be guaranteed to hold if the answer to Question A were "yes". More precisely, we can often associate to $\mathcal{X}$ a collection of quasi-trees, the *factored contact graphs*, which completely encode the geometry of $\mathcal{X}$ (up to quasi-isometry). Part 3 is about an alternative definition of the boundary in terms of the boundaries of the factored contact graphs. This yields the "compact" version of the simplicial boundary, and we close by discussing some uses of this object.

Notes on exercises. This text contains various exercises and problems designed to stimulate discussion and/or build familiarity with the notions being used, as well as a few open problems. The non-open problems are either very straightforward, or solutions to them are explained in [Hag13, Hag14b] or the forthcoming[1] paper [DHS15]. The open or tricky problems are starred.

Acknowledgments. I am very grateful to Indira Chatterji for organizing the meeting which provided the opportunity to write these notes, and for her interest in the subject. Also, many thanks to Mike Carr, Ruth Charney, Dave Futer, Dan Guralnik, Jingyin Huang, Alessandra Iozzi, Brian Rushton, Michah Sageev, Yulan Qing, and Dani Wise for interesting discussions, observations, and questions about the objects studied here. Part of the notes are based on joint work with subsets of {Jason Behrstock, Matt Durham, Alessandro Sisto}, with whom it's been extremely eye-opening and enjoyable to collaborate.

1. Part I: the simplicial boundary

We start with the definition and fundamental properties of the simplicial boundary of a CAT(0) cube complex, along with examples and pictures one should have in mind.

Throughout, $\mathcal{X}$ is a CAT(0) cube complex, $\mathcal{H}$ is the set of hyperplanes in $\mathcal{X}$. The *crossing* relation on $\mathcal{H}$ is: $H, H' \in \mathcal{H}$ *cross*, written $H \perp H'$, if $H \cap H' \neq \emptyset$. If $H, H' \in \mathcal{H}$ are not separated by some third hyperplane, they *contact*, written $H \perp\!\!\!\!\smile H'$. Note that crossing hyperplanes contact.

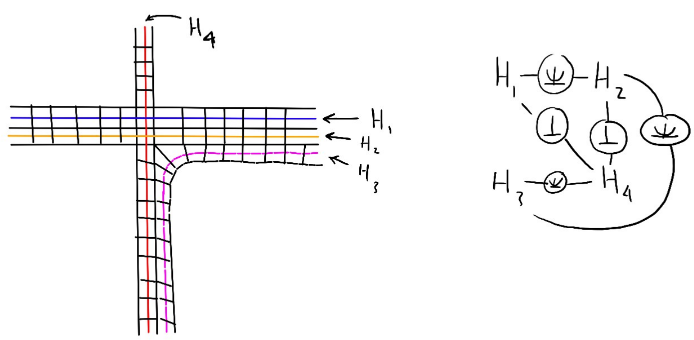

Figure 1. Some crossing and osculating hyperplanes. Hyperplanes contact if they either cross or osculate.

Convention 1.1. Throughout these lectures, we make only one standing finiteness assumption about $\mathcal{X}$: **any set of pairwise-crossing hyperplanes is finite.** This is weaker than requiring $\mathcal{X}$ to be finite-dimensional or locally finite. (In applications, we often impose one of the latter two conditions, or the even stronger condition that 0–cubes of $\mathcal{X}$ have uniformly bounded degree.)

Remark 1.2 (Convexity)**.** What really makes CAT(0) cube complexes so organized and so nice as a setting for doing geometry is their hyperplanes, and the combinatorics of how the hyperplanes interact. One of the main ways that this niceness manifests is in the very rich notion of convexity that the hyperplane structure provides. Whether one adopts the median

[1] Very soon, as of February 26, 2016

graph viewpoint, or uses disk diagrams over cube complexes, convexity usually lurks near the crux of whatever one is doing.

Let $\mathcal{X}$ be a CAT(0) cube complex. A subcomplex $\mathcal{Y}$ is *full* if $\mathcal{Y}$ contains every cube of $\mathcal{X}$ whose 0–skeleton appears in $\mathcal{Y}$. A full subcomplex $\mathcal{Y}$ is *convex* if $\mathcal{Y}^{(1)}$ is metrically convex in $\mathcal{X}^{(1)}$. There are numerous equivalent formulations.

First, if $\mathcal{Y} \subseteq \mathcal{X}$ is convex, then $\mathcal{Y}$ inherits a CAT(0) cubical structure from $\mathcal{X}$, and the hyperplanes of $\mathcal{Y}$ have the form $H \cap \mathcal{Y}$, where H is a hyperplane of $\mathcal{X}$. Convexity of $\mathcal{Y}$ means that any two hyperplanes of $\mathcal{X}$ that intersect $\mathcal{Y}$ intersect one another in $\mathcal{X}$ if and only if their intersections with $\mathcal{Y}$ intersect one another.

Second, any subspace $\mathcal{Y}$ of $\mathcal{X}$ has a convex hull, which is just the intersection of all convex subcomplexes containing $\mathcal{Y}$. This has an attractive alternative characterization, in terms of halfspaces. The archetypal convex subcomplex is the *combinatorial hyperplane*. A combinatorial hyperplane is the image of $H \times \{\pm\frac{1}{2}\}$ under the map $H \times [-\frac{1}{2}, \frac{1}{2}] \cong \mathcal{N}(H) \hookrightarrow \mathcal{X}$, where $\mathcal{N}(H)$ denotes the *carrier* of H (i.e. the union of all closed cubes intersecting H). A *combinatorial halfspace* is defined as follows: let $\overrightarrow{H}$ be a halfspace associated to H (i.e. a component of $\mathcal{X} - H$). There are two associated combinatorial halfspaces, namely $\overrightarrow{H} \cup \mathcal{N}(H)$ and the closure of its complement; these are respectively bounded by the two different combinatorial hyperplanes associated to H. A combinatorial halfspace is convex, and the convex hull of $\mathcal{Y}$ is the intersection of all combinatorial halfspaces containing $\mathcal{Y}$.

1.1. **The definition.** The simplicial boundary of $\mathcal{X}$ is, like the Tits boundary and the visual boundary, designed to encode the different directions in which one can move off to infinity inside $\mathcal{X}$. Instead of thinking about a space of geodesic rays, or infinite sequences, we'll use the hyperplanes to define "directions". First, consider a (graph-metric) geodesic ray $\gamma \subset \mathcal{X}^{(1)}$, which is characterized by the property that each hyperplane intersects γ in at most one point. The set $\mathcal{H}(\gamma)$ of hyperplanes H with $H \cap \gamma \neq \emptyset$ has a few interesting properties:

(1) $|\mathcal{H}(\gamma)| = \infty$;
(2) $\mathcal{H}(\gamma)$ contains no *facing triple*: if $H, H', H'' \in \mathcal{H}(\gamma)$ are pairwise-disjoint, then one must separate the other two;
(3) $\mathcal{H}(\gamma)$ is *closed under separation*: if $H, H' \in \mathcal{H}(\gamma)$ are separated by some hyperplane V, then $V \in \mathcal{H}(\gamma)$;
(4) $\mathcal{H}(\gamma)$ is *unidirectional*: if $H \in \mathcal{H}(\gamma)$, then at most one of the halfspaces associated to H contains infinitely many elements of $\mathcal{H}(\gamma)$.

Rather than define a boundary in terms of geodesic rays, we work with sets of hyperplanes modelled on $\mathcal{H}(\gamma)$:

Definition 1.3 (Boundary set, almost-containment, equivalence, minimality)**.** A *boundary set* is an infinite set of hyperplanes that is unidirectional, closed under separation, and contains no facing triple. Given boundary sets $\mathcal{H}, \mathcal{H}'$, we say $\mathcal{H}'$ *almost contains* $\mathcal{H}$, written $\mathcal{H} \preceq \mathcal{H}'$, if $|\mathcal{H} - \mathcal{H}' \cap \mathcal{H}| < \infty$. If $\mathcal{H} \preceq \mathcal{H}'$ and $\mathcal{H}' \preceq \mathcal{H}$, then $\mathcal{H}, \mathcal{H}'$ are *equivalent*. The boundary set $\mathcal{H}$ is *minimal* if $\mathcal{H}'$ and $\mathcal{H}$ are equivalent whenever $\mathcal{H}' \preceq \mathcal{H}$.

Example 1.4. Here are some examples of boundary sets:

(1) If γ is a geodesic ray in $\mathcal{X}$, then $\mathcal{H}(\gamma)$ is a boundary set.
(2) If $\mathcal{X}$ is a tree, then there is a bijection between the geodesic rays in $\mathcal{X}$ and the boundary sets, given by $\gamma \to \mathcal{H}(\gamma)$, which descends to a bijection between $\partial\mathcal{X}$ (Gromov boundary) and the set of equivalence classes of boundary sets.
(3) If $\mathcal{X} = \alpha \times \beta$, where α, β are combinatorial rays, then every boundary set is equivalent to $\mathcal{H}(\alpha), \mathcal{H}(\beta)$, or $\mathcal{H}(\alpha) \cup \mathcal{H}(\beta)$.

(4) More generally, if $\mathcal{X} = \mathcal{A} \times \mathcal{B}$, where $\mathcal{A}, \mathcal{B}$ are CAT(0) cube complexes, then every boundary set in $\mathcal{X}$ is equivalent to a boundary set in $\mathcal{A}$, a boundary set in $\mathcal{B}$, or a union $\mathcal{H} \cup \mathcal{V}$, where $\mathcal{H}$,$\mathcal{V}$ are boundary sets of $\mathcal{A}, \mathcal{B}$ respectively.

(5) This is a very important example: start with the obvious tiling of $[0, \infty)^2$ by 2–cubes, let $f : [0, \infty) \to [0, \infty)$ be some unbounded nondecreasing function with $f(0) = 0$, and let $\mathcal{X}$ be the union of all closed 2–cubes intersecting the set $\{(x, y) \in [0, \infty)^2 \mid y \leq f(x)\}$. This is a *staircase* and is a major source of CAT(0) cube complex pathology. See Figure 2.

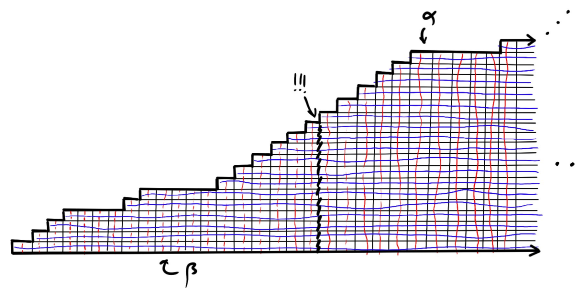

FIGURE 2. A staircase. The set of all hyperplanes is a boundary set, in fact equal to $\mathcal{H}(\alpha)$. The set of vertical hyperplanes is a boundary set, $\mathcal{H}(\beta)$. The set of horizontal hyperplanes is also a boundary set, but is not equal to $\mathcal{H}(\gamma)$ for any geodesic ray γ – where would such a ray start?

(6) Next, we have the *ziggurat*. Two types of cross-section – hyperplane! – are coarse rays that look like "staircases with plateaux", while the other type is a Euclidean quadrant. In Figure 3, you can see a red boundary set, a blue boundary set, a pink boundary set. Any union of two or more of these boundary sets is again a boundary set.

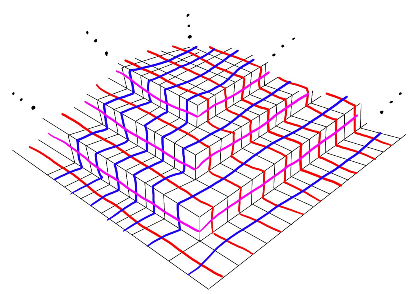

FIGURE 3. A ziggurat and its hyperplanes.

(7) *Cliffs.* The cliffs are to the ziggurat as the staircase is to the quadrant, in a way. See Figure 4. There is a red boundary set, whose hyperplanes are compact, and a blue boundary set whose hyperplanes are staircases. There is also a pink boundary set, whose hyperplanes are products of rays with increasingly long intervals. The union of all three sets is a boundary set, and the same is true for the union of blue and pink and the union of the red and pink sets. However, the union of the red and blue sets is not closed under separation! We will return to this example (which was pointed out by Dan Guralnik and Alessandra Iozzi) later in the lecture.

Two facts about the structure of boundary sets are needed to define the simplicial boundary:

Proposition 1.5 (Finding minimal boundary sets)**.** *Every boundary set contains a minimal boundary set. If $\mathcal{X}$ is locally finite, then every infinite subset of $\mathcal{H}$ contains a boundary set.*

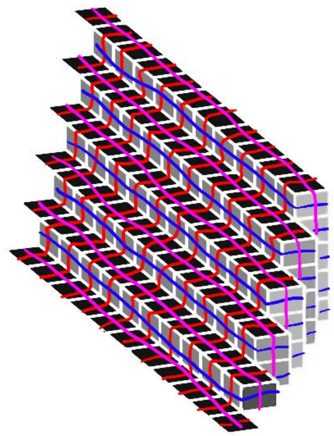

FIGURE 4. Cliffs.

Proof. See exercises. □

Proposition 1.6 (Canonical decomposition of boundary sets)**.** *Let $\mathcal{H}$ be a boundary set. Then there exist pairwise-disjoint minimal boundary sets $\mathcal{H}_0, \ldots, \mathcal{H}_k$ so that:*

(1) $\mathcal{H}$ is equivalent to $\sqcup_{i=0}^{k}\mathcal{H}_i$;
(2) for $0 \le i < j$, the set $\mathcal{H}_j$ dominates *$\mathcal{H}_i$, i.e. for all $H \in \mathcal{H}_j$, H crosses all but finitely many elements of $\mathcal{H}_i$.*

Moreover, this decomposition is essentially unique: if $\mathcal{H}$ is equivalent to some other union $\sqcup_{i=0}^{k'}\mathcal{H}'_i$ of minimal boundary sets, then $k = k'$ and, up to relabeling, $\mathcal{H}_i$ and $\mathcal{H}'_i$ are equivalent.

Sketch. If $\mathcal{H}$ is minimal, we're done. Otherwise, find a minimal subset $\mathcal{H}_0$ using Proposition 1.5. Check that either $\mathcal{H} - \mathcal{H}_0$ is finite, or contains an infinite subset closed under separation; then apply Proposition 1.5 to find a minimal boundary set $\mathcal{H}_1 \subseteq \mathcal{H} - \mathcal{H}_0$. Figure 5 offers the hint on how to do this.

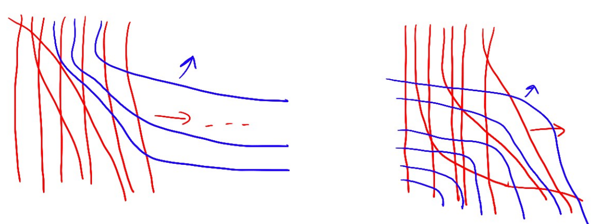

FIGURE 5. The inductive step in the proof of Proposition 1.6. The red minimal boundary set is given; any further infinite subset of $\mathcal{H}$ behaves like one of the blue ones.

Continue in this way. To prove uniqueness, use the following fact: if $\mathcal{M}, \mathcal{M}'$ are minimal boundary sets, then either they are equivalent or their intersection is finite... □

The *dimension* of a boundary set (equivalence class of boundary sets) $\mathcal{H}$ ($[\mathcal{H}]$) is the number k in the decomposition of $\mathcal{H}$ from Proposition 1.6.

Definition 1.7 (Boundary simplex, simplicial boundary)**.** The *simplicial boundary* $\partial_{\triangle}\mathcal{X}$ of $\mathcal{X}$ is the simplicial complex with a k–simplex for each k–dimensional equivalence class of boundary sets. Such a *boundary simplex*, a simplex corresponding to a boundary set $\mathcal{H}$, is a face of the simplex corresponding to $\mathcal{H}'$ if $\mathcal{H} \preceq \mathcal{H}'$.

Remark 1.8 (Warning)**.** The example of cliffs shows that, while every maximal simplex of $\partial_\triangle \mathcal{X}$ is actually a boundary simplex corresponding to some equivalence class of boundary sets, a simplex of $\partial_\triangle \mathcal{X}$ can have a proper face which does not correspond to an equivalence class of boundary sets. See Exercise (2). However, $\partial_\triangle \mathcal{X}$ is the union of genuine boundary simplices, and in practice one gets away with just thinking about these simplices. Under the assumption of *full visibility* discussed below, this weirdness disappears.

Example 1.9 (First examples)**.** Here are some examples to keep in mind.

(1) $\partial_\triangle \mathcal{X}$ is a discrete set, equal to the set of ends of $\mathcal{X}$, when $\mathcal{X}$ is a tree.
(2) Suppose that $\mathcal{X}$ is hyperbolic. Then $\partial_\triangle \mathcal{X}$ is a discrete set; see Exercise (5).
(3) If $\mathcal{X} = \mathcal{A}\times\mathcal{B}$, then $\partial_\triangle \mathcal{X} = \partial_\triangle \mathcal{A}\star\partial_\triangle \mathcal{B}$, where $\star$ denotes the simplicial join. This uses the fact that product decompositions of $\mathcal{X}$ correspond to partitions of the set of hyperplanes of $\mathcal{X}$ into disjoint sets $\mathcal{V},\mathcal{H}$ so that each element of $\mathcal{V}$ crosses each element of $\mathcal{H}$.
(4) The simplicial boundary of a staircase is a 1–simplex; this is also the simplicial boundary we'd get if we filled in the rest of the quadrant.
(5) Let $\mathcal{X}$ be the universal cover of the Salvetti complex of the Croke-Kleiner RAAG: $\langle a,b,c,d \mid [a,b],[b,c],[c,d]\rangle$. Then $\partial_\triangle \mathcal{X}$ has many isolated 0–simplices coming from rank-one elements (more on this later), but the interactions between the "standard flats" give rise to the subcomplex in Figure 6:

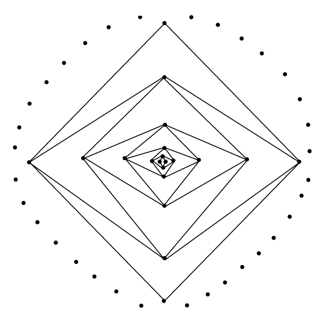

FIGURE 6. Eye of Sauron.

Theorem 1.10 (Basic properties)**.** *$\partial_\triangle \mathcal{X}$ has the following properties:*

(1) It is a flag complex.
(2) Every simplex of $\partial_\triangle \mathcal{X}$ is contained in a finite-dimensional maximal simplex.
(3) Let $\mathcal{Y} \subset \mathcal{X}$ be a convex subcomplex. Then there is a natural simplicial embedding $\partial_\triangle \mathcal{Y} \to \partial_\triangle \mathcal{X}$ whose image is a full subcomplex. This holds in particular when $\mathcal{Y}$ is the carrier of a hyperplane.

Proof. Assertion (1) follows from the definition. Assertion (2) follows from Proposition 1.6 and the fact that there is no infinite family of pairwise-crossing hyperplanes.

It remains to prove assertion (3). By convexity, the hyperplanes of $\mathcal{Y}$ have the form $H\cap\mathcal{Y}$, where H is a hyperplane of $\mathcal{X}$ intersecting $\mathcal{Y}$. Moreover, $H\cap H'\cap\mathcal{Y}\neq\emptyset$ if and only if $H\cap H'\neq\emptyset$, whenever H, H' are hyperplanes both intersecting $\mathcal{Y}$ (again by convexity). It follows that each k–dimensional boundary set of $\mathcal{Y}$ has the form $\{H\cap\mathcal{Y} \mid H\in\mathcal{H}\}$, where $\mathcal{H}$ is a k–dimensional boundary set in $\mathcal{X}$, and the claim follows. □

1.2. **Visibility.** The definition of a boundary set was motivating by looking at the set of hyperplanes crossing a combinatorial geodesic ray. How far away can the definition wander from this example? To what extent is $\partial_\triangle \mathcal{X}$ really a "space of directions" in $\mathcal{X}$?

Definition 1.11 (Visible simplex, visible pair, full visibility)**.** The simplex $h \subseteq \partial_{\triangle}\mathcal{X}$ is *visible* if there is a geodesic ray γ so that $\mathcal{H}(\gamma)$ represents h (i.e. h is the boundary simplex corresponding to the equivalence class of $\mathcal{H}(\gamma)$). In this case, we say that γ *represents* h. The pair h, h' is *visible* if there is a bi-infinite geodesic γ in $\mathcal{X}$ which is the union of two rays with bounded intersection that respectively represent h, h'. The cube complex $\mathcal{X}$ is *fully visible* if every simplex is visible.

Remark 1.12. Given a subcomplex $A \subseteq \partial_{\triangle}\mathcal{X}$, one could say A is "visible" if there is a subspace $\mathcal{Y} \subseteq \mathcal{X}$ whose convex hull has boundary A. I haven't thought about this general notion.

Full visibility, and the failure thereof, is not fully understood; see Problem 1.3.1.(5).

Proposition 1.13. *Every maximal simplex of $\partial_{\triangle}\mathcal{X}$ is visible.*

Sketch. Let m be a maximal simplex of $\partial_{\triangle}\mathcal{X}$, so that m is represented by a boundary set $\mathcal{M}$. We will attempt to build a ray γ with $\mathcal{H}(\gamma) = \mathcal{M}$; if we fail, it's going to reflect that m lies in some larger simplex. First, let $\mathcal{M} = \{H_i\}_{i\geq 0}$ be numbered so that $i < j < k$ whenever H_j separates H_i from H_k. For each i, find a geodesic segment γ_i so that the set of hyperplanes crossing γ_i is exactly $\{H_0, \ldots, H_i\}$. This is possible because $\mathcal{M}$ is closed under separation and because of our numbering scheme.

Now let $\mathcal{Y}_i$ be the convex hull of the union of all possible choices of γ_i. Check that either we can find $\gamma_0 \subset \gamma_1 \subset \ldots$, and thus build a limiting ray, or the following happens: for arbitrarily large i, j, we can find a hyperplane H separating $\mathcal{Y}_i, \mathcal{Y}_j$. In this way, we build a sequence of hyperplanes H_i so that each $M \in \mathcal{M}$ crosses all but finitely many H_i. Result: $\mathcal{M}$ is contained in, but not equivalent to, a large boundary set, contradicting maximality of m. $\square$

Proposition 1.14. *Let h, h' be visible simplices of $\partial_{\triangle}\mathcal{X}$. Then h, h' is a visible pair if and only if $h \cap h' = \emptyset$.*

Sketch. Represent h, h' by rays γ, γ' with a common initial point. If γ, γ' only have finitely many common hyperplanes (i.e. if $h \cap h' = \emptyset$), then there is some folding and truncation one can do to build a bi-infinite geodesic. Otherwise, $h \cap h' \neq \emptyset$. $\square$

1.3. **Comparing the boundaries.** The definition of $\partial_{\triangle}\mathcal{X}$ looks more like the definition of the Roller boundary, but in practice, $\partial_{\triangle}\mathcal{X}$ looks more like the Tits boundary. The following exercises discuss these comparisons. (This section is a series of exercises rather than real exposition since it's interesting but not relevant to my propaganda mission about Question A.)

1.3.1. *The simplicial boundary and the Tits boundary.* Assume that $\mathcal{X}$ is fully visible.

(1) Let v be a simplex of $\partial_{\triangle}\mathcal{X}$, spanned by 0–simplices $v_0, \ldots, v_k$. Show that for each i, we can choose a combinatorial geodesic ray γ_i representing v_i so that there is a cubical isometric embedding $\prod_i C_i \to \mathcal{X}$, where C_i denotes the convex hull of γ_i.

(2) Prove that for each i, there is a CAT(0) geodesic ray α_i, with the same initial point as γ_i, that crosses exactly those hyperplanes crossed by γ_i. Deduce that, somewhere inside $\prod_i C_i$, there is a isometrically embedded (in the CAT(0) sense) Euclidean orthant $\prod_i \alpha_i$.

(3) Produce an embedding $\partial_{\triangle}\mathcal{X} \to \partial_T\mathcal{X}$ (the Tits boundary for the usual CAT(0) metric) sending each simplex to a right-angled spherical simplex.

∗ (4) What is the failure of your embedding to be surjective? Explain how the various C_i "parameterize" the embeddings $\partial_{\triangle}\mathcal{X} \to \partial_T\mathcal{X}$ of the type you just constructed.

∗ (5) Later, we'll see a condition that not only guarantees full visibility but also ensures that the above embeddings are bijective (and are in fact isometries if you regard $\partial_{\triangle}\mathcal{X}$ as a CAT(1) space made of right-angled spherical simplices). For the moment: let $\{H_n\}_{n\geq 1}$ be the set of hyperplanes of C_i (so, intersections of hyperplanes of $\mathcal{X}$ with C_i), numbered so that, if $m < n$, then H_n does not separate H_m from the initial point of γ_i. Suppose that the following holds: consider all subcomplexes you can build by examining the image of $\mathcal{N}(H_m)$ under combinatorial closest-point projection to $\mathcal{N}(H_n)$, as m, n vary. Suppose that the resulting collection of subcomplexes has bounded multiplicity. Then, if this holds for each 0–simplex of $\partial_\triangle \mathcal{X}$, the embedding $\partial_\triangle \mathcal{X} \to \partial_T \mathcal{X}$ you constructed above is bijective. (To experiment with this part of the exercise, you might think about "cubical sectors" in the standard tiling of $\mathbb{R}^2$ by cubes. If your sector has "bottlenecks" in the appropriate sense, then you win: there is a unique (up to asymptoty) CAT(0) geodesic ray, and you win. Otherwise, you have a choice in how to embed the simplicial boundary in the Tits boundary. See Figure 7.)

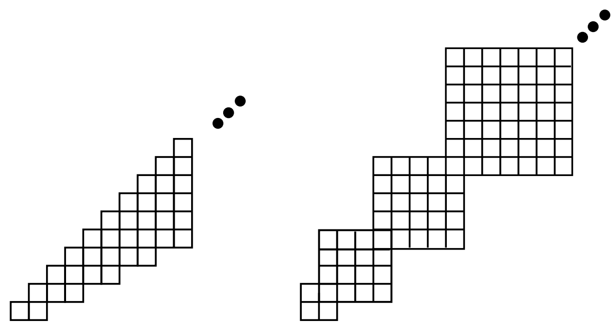

FIGURE 7. $\partial_\triangle \mathcal{X} \to \partial_T \mathcal{X}$ can fail to be surjective (left picture). The bottlenecks in the right picture disallow that sort of behaviour.

1.4. More exercises and problems from Part I.

(1) Prove Proposition 1.5. (Idea: given a boundary set $\mathcal{U}$, use the fact that sets of pairwise-crossing hyperplanes are finite to produce a sequence $H_0, H_1, \ldots \in \mathcal{U}$ with H_i separating $H_{i\pm1}$ for $i \geq 1$. Now close this set under separation...)

(2) Let γ be a combinatorial geodesic ray with convex hull C. Then $\partial_{\triangle} C \subseteq \partial_{\triangle} \mathcal{X}$ is a single simplex, represented by γ.

* (3) Why doesn't the bizarre phenomenon exhibited by cliffs happen when $\mathcal{X}$ is fully visible? What can be said about the structure of the convex hull of a geodesic ray when $\mathcal{X}$ is fully visible?

* (4) Just like the Tits boundary of a CAT(0) space can detect splittings of the CAT(0) space as a metric product, the simplicial boundary can detect splittings as a cubical product. One direction was seen in Example 1.9. Conversely, let $\mathcal{X}$ be a fully visible CAT(0) cube complex which is *essential* (i.e. every halfspace contains a hyperplane). Suppose that $\partial_{\triangle} \mathcal{X} \cong A \star B$ for nonempty subcomplexes A, B. Prove that $\mathcal{X} \cong \mathcal{X}_A \times \mathcal{X}_B$, where $\mathcal{X}_A$ and $\mathcal{X}_B$ are unbounded convex subcomplexes. (Hint: fix a base vertex $x \in \mathcal{X}$ and, for each 0–simplex v of $\partial_{\triangle} \mathcal{X}$ lying in A, use full visibility to represent v by a ray γ_v emanating from x. Let $\mathcal{X}_A$ be the convex hull of the union of all of these rays, and define $\mathcal{X}_B$ analogously...). Also: exhibit examples showing that the essentiality and full visibility hypotheses are necessary.

(5) It is known [Hag14b] that, if $\mathcal{X}$ is uniformly locally finite, then $\mathcal{X}^{(1)}$ is hyperbolic exactly when the following holds for some integer q: for any sets $\mathcal{V}, \mathcal{H}$ of hyperplanes with every element of $\mathcal{V}$ crossing every element of $\mathcal{H}$, we have $\min\{|\mathcal{V}|, |\mathcal{H}|\} \leq q$. Deduce from this that $\partial_{\triangle} \mathcal{X}$ is a discrete set when $\mathcal{X}$ is hyperbolic.

* (6) Give "nice" conditions on $\mathcal{X}$ ensuring that it is fully visible. In particular: given $\mathcal{X}$ a CAT(0) cube complex with a group G acting properly and cocompactly, is $\mathcal{X}$ fully visible? (Behrstock and I guess "yes", in [BH, Conjecture 2.8]. A positive answer to this question would follow from a positive answer to a further-reaching question about cubical groups that we'll discuss later.)

(7) Do the "folding and truncation" from Proposition 1.14, as suggested by Figure 8.

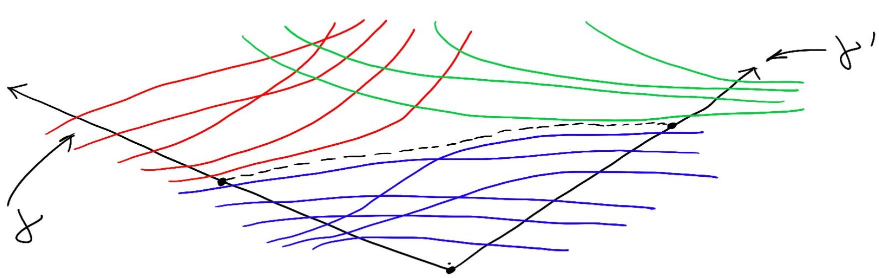

FIGURE 8. Folding and truncating to make a bi-infinite geodesic: "go beyond the last common hyperplane".

(8) If γ is a combinatorial geodesic ray whose convex hull is hyperbolic, and γ' is a geodesic ray representing the same boundary simplex, then γ' and γ fellow-travel.

(9) Let G act geometrically on $\mathcal{X}$ and let $\mathcal{Y} \subseteq \mathcal{X}$ be the G–essential core. Prove that $\partial_{\triangle} \mathcal{X} = \partial_{\triangle} \mathcal{Y}$.

2. Part II: Quasi-trees, rank-one stuff, and "hierarchies"

2.1. **The contact graph.** Given a CAT(0) cube complex $\mathcal{X}$, the *contact graph* $\mathcal{CX}$ is the graph with a vertex for each hyperplane and an edge joining two vertices if and only if the corresponding hyperplanes contact. In other words, $\mathcal{CX}$ is the intersection graph of the hyperplane-carriers. The contact graph is quasi-isometric to the space obtained from $\mathcal{X}^{(1)}$ by coning off the 1–skeleton of each hyperplane carrier.

Remark 2.1 (Mapping class group digression)**.** In many ways, the contact graph is to the cube complex as the curve graph of a surface is to the marking complex/mapping class group, so, if you're familiar with the Masur-Minsky approach to the mapping class group [MM99, MM00], it might help to have the curve graph in the back of your mind when thinking about the contact graph. In fact, the relationship between the contact graph and the simplicial boundary allows one to build a boundary for the mapping class group [BHS14, BHS15, DHS15].

Remark 2.2 (Contact graph facts)**.** Some basic facts about the contact graph:

(1) If $\mathcal{Y} \subseteq \mathcal{X}$ is a convex subcomplex, then $\mathcal{CY}$ is an induced subgraph of $\mathcal{CX}$ in the obvious way: each hyperplane $H \cap \mathcal{Y}$ of $\mathcal{Y}$ (representing a vertex of $\mathcal{CY}$) corresponds to the hyperplane H of $\mathcal{X}$ (and thus to the corresponding vertex of $\mathcal{CX}$).
(2) If $\mathcal{A}, \mathcal{B}$ are CAT(0) cube complexes, then $\mathcal{C}(\mathcal{A} \times \mathcal{B}) = \mathcal{CA} \star \mathcal{CB}$, where $\star$ denotes the join. Indeed, every hyperplane crossing one factor crosses every hyperplanes crossing the other factor. The CAT(0) cube complexes we have in mind have horrifically locally infinite contact graphs, in general, because they can have unbounded hyperplanes, but this example shows that the contact graph can nonetheless be bounded.
(3) Each 0–cube $x \in \mathcal{X}$ corresponds to a clique in $\mathcal{CX}$ whose vertices correspond to the hyperplanes whose carriers contain x. Hence the clique number of $\mathcal{CX}$ is equal to the maximum degree of a vertex in $\mathcal{X}$.

Theorem 2.3 ([Hag14b])**.** *$\mathcal{CX}$ is quasi-isometric to a tree.*

In view of Theorem 2.3 and Remark 2.2.(2), it's natural to ask when $\mathcal{CX}$ is a quasi-point, i.e. a quasi-tree by virtue of being bounded. In fact, this question motivated the definition of the simplicial boundary, and it's closely related to rank-rigidity.

2.2. **Rank-one things, the contact graph, and the boundary.** Let γ be a combinatorial geodesic ray in $\mathcal{X}$. What does γ look like? More precisely, what does the convex hull of γ look like (intrinsically in $\mathcal{X}$ and from the point of view of the simplicial boundary)?

For $i \geq 1$, let H_i be the hyperplane dual to the i^{th} 1–cube of γ. The sequence $H_1, H_2, H_3, \ldots$ defines an edge-path in $\mathcal{CX}$, since $H_i \perp\!\!\!\!\smile H_{i+1}$ for each i. Here is the main question:

- When is $\{H_i\}$ unbounded? In other words, when does projecting γ to $\mathcal{CX}$ prove that $\mathcal{CX}$ is an unbounded quasi-tree? Even better, when is $\gamma(i) \mapsto H_i$ a quasi-isometric embedding?

The following is proved in Section 2 of [Hag13]. We'll skip the proof here, and instead focus on a more quantitative version of the same statement, which is more obviously related to what we'll do in the next section.

Theorem 2.4 (The projection trichotomy)**.** *Let γ be a combinatorial geodesic ray in $\mathcal{X}$. Then one of the following holds:*

(1) γ lies in a uniform neighborhood of some hyperplane of $\mathcal{X}$;
(2) γ lies in an isometrically embedded staircase in $\mathcal{X}$;
(3) the set of hyperplanes intersecting γ span an unbounded subset of $\mathcal{CX}$.

(The first two conclusions can both hold.) In particular, γ projects to an unbounded subset in $\mathcal{CX}$ only if it represents an isolated 0–simplex in $\partial_\triangle \mathcal{X}$.

In fact, Theorem 2.2 of [Hag13] is stronger; paraphrased, it says:

Theorem 2.5 (The quantitative projection trichotomy [Hag13])**.** *Let γ be a combinatorial geodesic ray in $\mathcal{X}$. Suppose that there exists M so that for all hyperplanes H of $\mathcal{X}$, there are at most M hyperplanes U such that $U \cap H \neq \emptyset$ and $U \cap \gamma \neq \emptyset$. Then the subgraph of $\mathcal{C}\mathcal{X}$ spanned by the vertices corresponding to the set of hyperplanes crossing γ is quasi-isometric to γ.*

Proof. Here's slicker version of the proof than you'll find in [Hag13], because I'm older and wiser (well, only about this theorem). Let C be the convex hull of γ in $\mathcal{X}$, and let Λ be the subgraph of $\mathcal{C}\mathcal{X}$ spanned by the set of hyperplanes crossing C (which is the same as the set of hyperplanes crossing γ). Let H, H' cross C, and let $H = H_0, \ldots, H_n = H'$ be a geodesic sequence in $\mathcal{C}\mathcal{X}$ joining them. Choose $x \in \mathcal{N}(H_0) \cap C, y \in \mathcal{N}(H_n) \cap C$ to be 0–cubes, and let $\alpha_0\alpha_1 \cdots \alpha_{n-1}\alpha_n$ be a combinatorial path from x to y, where each α_i is a combinatorial geodesic in $\mathcal{N}(H_i)$. Let β be a combinatorial geodesic from x to y (so β lies in C – for example, we could take $x, y \in \gamma$ and have β be a subpath of γ, but it doesn't matter). See Figure 9.

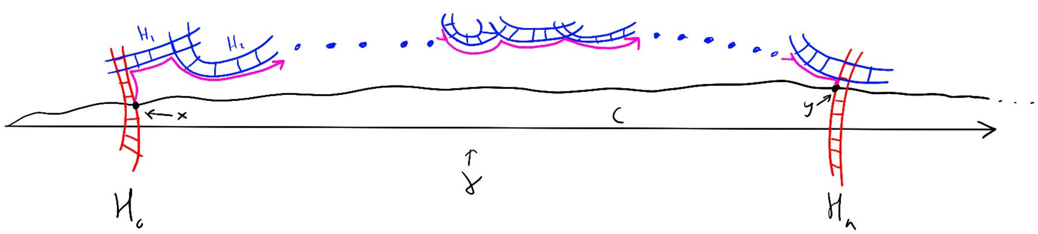

FIGURE 9. The proof of Theorem 2.5.

Since $\beta^{-1}\alpha_0 \cdots \alpha_n$ is a closed path, it bounds a disc diagram $D \to \mathcal{X}$. Let all the choices above be made so that, when we pull D tight, its area is as small as possible (so, as few squares as possible and as few backtracks in the boundary path as possible). Let K be a dual curve in D starting on some α_i. If K ends on α_j with $|i-j| > 2$, then we contradicted that $H_i, H_{i+1}, \ldots, H_j$ (or whatever) is a $\mathcal{C}\mathcal{X}$–geodesic. If $i = j$, we contradict that α_i is a geodesic. If $|i - j| = 1$, then we can do some folding and remove a backtrack. If $|i - j| = 2$, then we can replace H_{i+1} (say) by the hyperplane to which K maps and get a lower-area D. Hence all dual curves travel between $\alpha_0 \cdots \alpha_n$ and β, i.e. $\alpha_0 \cdots \alpha_n$ is a geodesic starting and ending on C and thus lying in C. In particular, each hyperplane crossing each α_i crosses γ, so $|\alpha_i| \leq M$ for all i. Hence n grows linearly in $|\beta|$, and we're done. □

Remark 2.6. A similar argument yields the notion of *hierarchy paths in cube complexes* – i.e. geodesics of $\mathcal{X}$ that track geodesics in $\mathcal{C}\mathcal{X}$. It's shown in [BHS14, Section 3] that any two vertices are joined by such a geodesic, using an argument basically identical to the proof of Theorem 2.5.

In the next section, we'll see that Theorem 2.5 is the tip of the iceberg in terms of understanding points in $\partial_{\triangle}\mathcal{X}$ in terms of projections to quasi-trees. First, we can state some related results about isometries (note that $\text{Aut}(\mathcal{X})$ acts by isometries on $\mathcal{C}\mathcal{X}$). Recall from [Hag07] that, after passing if necessary to the first cubical subdivision of $\mathcal{X}$ (which does nothing to the boundary and nothing serious to the contact graph), each $g \in \text{Aut}(\mathcal{X})$ either fixes a 0–cube (elliptic) or has a combinatorial geodesic axis (hyperbolic).

Corollary 2.7. *Let $\mathcal{X}$ be a CAT(0) cube complex with no infinite set of pairwise-crossing hyperplanes. Let $g \in \mathrm{Aut}(\mathcal{X})$. Then one of the following holds:*

1. *g is elliptic;*
2. *g stabilizes a clique in $\mathcal{CX}$. In particular, if $\mathcal{X}$ is locally finite, then g^N stabilizes a hyperplane for some $N \geq 0$;*
3. *any axis of g bounds a combinatorial half-flat, and there is a $\langle g \rangle$–orbit in $\mathcal{CX}$ of diameter at most 3;*
4. *g acts loxodromically on $\mathcal{CX}$.*

In the last case, g is a rank-one isometry of $\mathcal{X}$.

Sketch. Suppose g is hyperbolic and let A be a combinatorial axis for g in $\mathcal{X}$. If A projects to a quasi-geodesic in $\mathcal{CX}$ (i.e. if the set of hyperplanes crossing A corresponds to a set of $\mathcal{CX}$–vertices spanning a quasi-line), then the last option holds. In this case, g must be rank-one since, since half-flats have uniformly bounded projection to the contact graph. The rest of the claim follows Theorem 2.4 by translating staircases or hyperplanes by powers of g. □

Remark 2.8. If you're into mapping class groups, and the Corollary looks like an unsatisfying analogue of the Nielsen-Thurston classification, fear not, for soon we'll pass from the contact graph to a more refined analogue of the curve graph which really will give a sense in which the second two conclusions describe "reducible" elements.

Finally, we note that the Caprace-Sageev Irreducibility Criterion [CS11] combines with the above classification to show that if G acts geometrically and essentially on $\mathcal{X}$, then the following are equivalent:

1. $\mathcal{X} = A \times B$ for unbounded subcomplexes A, B;
2. $\mathcal{CX}$ is bounded;
3. $\mathcal{CX} = \mathcal{C}A \star \mathcal{C}B$;
4. $\partial_{\triangle}\mathcal{X}$ is connected;
5. $\partial_{\triangle}\mathcal{X}$ has no isolated 0–simplex;
6. $\partial_{\triangle}\mathcal{X} = \partial_{\triangle}A \star \partial_{\triangle}B$;
7. G contains no rank-one isometry of $\mathcal{X}$;
8. G contains no loxodromic isometry of $\mathcal{CX}$.

Later, we'll discuss how to prove this (in many interesting cases) without recourse to the Irreducibility Criterion.

2.3. **Factor-systems and the factored contact graph.** This part of the lectures will prepare us for an alternative definition of $\partial_{\triangle}\mathcal{X}$ corresponding to a more sophisticated version of projections of geodesic rays to quasi-trees, and the resulting compactification of $\partial_{\triangle}\mathcal{X}$. Most of the following comes from [BHS14]. Given a convex subcomplex $C \subseteq \mathcal{X}$, there is a closest-point projection map $\mathfrak{g}_C : \mathcal{X} \to C$: given $x \in \mathcal{X}^{(0)}$, there is a unique closest 0–cube $\mathfrak{g}_C(x)$ of C (exercise), and this extends in an obvious way to higher-dimensional cubes, so that they either map cubically, or collapse onto faces and then map cubically.

The preceding discussion showed that a geodesic γ either projects quasi-geodesically to $\mathcal{CX}$, or there is some hyperplane H so that $\mathfrak{g}_{\mathcal{N}(H)}(\gamma)$ is large. Presumably, there should be a way to proceed inductively, and conclude that the distance along γ is coarsely the same as the sum of a bunch of distances along hyperplanes, etc. (If you know about subsurface projections, this should remind you of the Masur-Minsky distance formula...)

This plan looks doomed in the infinite-dimensional case, and in fact the story is a bit more complicated than the above, so we have to impose some additional restrictions on $\mathcal{X}$.

Definition 2.9 (Factor system)**.** A *factor system* $\mathfrak{F}$ is a set of convex subcomplexes so that:

1. $\mathcal{X} \in \mathfrak{F}$ and $H \in \mathfrak{F}$ for each combinatorial hyperplane H;

(2) there exists N so that for all $x \in \mathcal{X}^{(0)}$, we have $|\{F \in \mathfrak{F} : x \in F\}| \leq N$;
(3) whenever $F, F' \in \mathfrak{F}$, either $\mathfrak{g}_F(F') \in \mathfrak{F}$ or $\mathfrak{g}_F(F')$ is a single vertex, in which case we do *not* include it (for a technical reason in the proof of Theorem 3.15; everything works just fine if we allow single points in $\mathfrak{F}$).

So, for example, if $\mathcal{X}$ is a tree, then $\mathfrak{F} = \{\mathcal{X}\}$. If $\mathcal{X}$ is the standard tiling of $\mathbb{R}^2$ by 2–cubes, then $\mathfrak{F}$ consists of $\mathcal{X}$, the hyperplanes (two parallelism classes of lines), and nothing else. If $\mathcal{X}$ is a staircase, then it can't contain a factor-system (exercise; easy and important!).

Remark 2.10. The first and second conditions imply that $\mathcal{X}$ has bounded degree, and in particular finite dimension. The third condition can actually by relaxed somewhat, so that $\mathfrak{g}_F(F')$ need only be included if its diameter exceeds some predetermined threshold.

Remark 2.11. Lots of cube complexes have factor systems: if $\mathcal{X}$ has a factor system, then so does any convex subcomplex of $\mathcal{X}$, and moreover universal covers of Salvetti complexes of RAAGs have factor systems (exercises). In particular, if G is a virtually compact special group, then G acts geometrically on a CAT(0) cube complex with a factor system. The exact scope of the definition is the subject of Question A'!

Definition 2.12 (Parallel)**.** Two convex subcomplexes $F, F' \subseteq \mathcal{X}$ are *parallel* if they intersect the same hyperplanes, i.e. if $\mathcal{C}F, \mathcal{C}F'$ are the same subgraph of $\mathcal{C}\mathcal{X}$. Given a factor-system $\mathfrak{F}$, let $\overline{\mathfrak{F}}$ denote $\mathfrak{F}$ mod parallelism (or, sometimes, a set containing exactly one element F in each parallelism class $[F]$).

Definition 2.13 (Factored contact graph, projection)**.** For each $F \in \mathfrak{F}$, let $\mathfrak{F}_F$ be the set of parallelism classes $[E] \in \overline{\mathfrak{F}}$ so that E is parallel to a proper subcomplex of F. Then define $\overline{\mathcal{F}}_F$ to be the set of parallelism classes in $\mathfrak{F}_F$. To F (well, its parallelism class) we can now associate the *factored contact graph* $\widehat{\mathcal{C}}F$ which is obtained from $\mathcal{C}F$ by coning off $\mathcal{C}E \subset \mathcal{C}F$ for each $E \in \mathfrak{F}_F$. There is a coarse projection $\pi_F : \mathcal{X} \to 2^{\widehat{\mathcal{C}}F}$: for each 0–cube or open cube c of $\mathcal{X}$, let $\pi_F(c)$ be the clique in $\mathcal{C}F \subset \widehat{\mathcal{C}}F$ whose vertices are the hyperplanes whose carriers contain $\mathfrak{g}_F(c)$.

Theorem 2.14. *There exists $\lambda \geq 1$ (independent of $\mathcal{X}$) so that $\widehat{\mathcal{C}}F$ is (λ, λ)–quasi-isometric to a tree for each $F \in \mathfrak{F}$.*

Remark 2.15. The factor system and the projections make $\mathcal{X}$ a *hierarchically hyperbolic space* (in fact, one of the main examples) in the sense of [BHS15].

Here's the most important consequence of the existence of factor systems.

Theorem 2.16 (Distance formula)**.** *For each sufficiently large s, there exists C so that for all $x, y \in \mathcal{X}^{(0)}$,*

$$\mathsf{d}_{\mathcal{X}}(x, y) \asymp_C \sum_{\substack{F \in \overline{\mathfrak{F}}, \\ \mathsf{d}_{\widehat{\mathcal{C}}F}(\pi_F(x), \pi_F(y)) \geq s}} \mathsf{d}_{\widehat{\mathcal{C}}F}(\pi_F(x), \pi_F(y)).$$

Proof. This uses some disc diagram considerations and induction on the multiplicity of the factor system (passing from $\mathfrak{F}$ to some $\mathfrak{F}_F$ lowers multiplicity). It's beyond the scope of these notes, but most of the work lies in defining factor systems and factored contact graphs correctly; the details are in [BHS14]. □

2.4. **Problems on Part 2.**

(1) If $C \subseteq \mathcal{X}$ is a convex subcomplex, and $x \in \mathcal{X}^{(0)}$, then there is a unique closest 0–cube of C to x.
(2) Let Γ be a finite graph and let S_Γ be the Salvetti complex of the corresponding right-angled Artin group A_Γ. Find at least two different (i.e. in principle different; maybe they agree for some specific Γ) factor-systems on $\widetilde{S}_\Gamma$ that you can describe in terms of Γ and A_Γ.
(3) Let $\mathcal{X}$ be a cube complex with a factor system $\mathfrak{F}$ and let $\mathcal{Y} \subset \mathcal{X}$ be a convex subcomplex. Prove that $\{F \cap \mathcal{Y} : F \in \mathfrak{F}\}$ is a factor system.
∗ (4) Prove Theorem 2.14. (Starred because tricky.)
(5) Give a nice combinatorial description of the convex hull of a geodesic ray.
∗ (6) Using the description of the convex hull of a geodesic ray, prove that, if $\mathcal{X}$ contains a factor system, then it is fully visible. (Starred because interesting; this is sort of like the bottleneck thing above, in the sense that you should start with a maximal simplex, represent it by a ray (maximal simplices are visible), and then use the fact that factor systems on $\mathcal{X}$ induce factor systems on convex subcomplexes to study the structure of the convex hull of your ray...)

3. Part III: The "compact simplicial boundary", rank-rigidity etc., and Question A

3.1. **Redefining $\partial_\triangle \mathcal{X}$ using factor systems.** We now assume $\mathcal{X}$ is a CAT(0) cube complex with a factor system $\mathfrak{F}$. In particular, $\mathcal{X}$ has bounded degree and is thus finite dimensional (so $\partial_\triangle \mathcal{X}$ is defined). We can and shall assume that $\mathfrak{F}$ is the minimal factor system, i.e. it contains $\mathcal{X}$, every combinatorial hyperplane, and every other subcomplex needed to make it closed under projection, but nothing else. As before, $\overline{\mathfrak{F}}$ is the set of parallelism classes, which indexes the set of factored contact graphs $\widehat{\mathcal{C}}F$.

In the previous exercises, it was proved that:

Proposition 3.1. *If $\mathcal{X}$ has a factor system, then $\mathcal{X}$ is fully visible.*

We're going to apply the distance formula to re-imagine the simplicial boundary in a way that explicitly involves the factored contact graphs. The goal is to build a compact version of the simplicial boundary. As an application, we'll give a new proof of rank-rigidity, under the assumption that factor-systems exist. Rank-rigidity, its consequences, and related things are awesome, so this will complete the mission to motivate Question A!

The following discussion appears in a different form in the forthcoming paper [DHS15], but the one below is more streamlined; things are easier the second time one does them.

Definition 3.2 (Nesting, orthogonality, complement)**.** Let $F, F' \in \mathfrak{F}$. We write $F \sqsubseteq F'$ if F is parallel to a subcomplex of F'. We write $F \perp F'$ if $F, F' \hookrightarrow \mathcal{X}$ extends to a convex embedding $F \times F' \to \mathcal{X}$. One can check that if $\mathcal{P}$ is a set of pairwise-parallel elements of $\mathfrak{F}$, then there is a convex subcomplex E and a convex embedding $F \times E \to \mathcal{X}$ so that each element of $\mathcal{P}$ has the form $F \times \{e\}$ for some $e \in E$. If $\mathcal{P}$ is an entire parallelism class, represented by some $F \in \mathfrak{F}$, then the complex $E = E_F$ is the *complement* of F.

Definition 3.3 (γ–relevant)**.** Given a combinatorial geodesic ray γ of $\mathcal{X}$ and $F \in \mathfrak{F}$, let $\pi_F(\gamma)$ be the subgraph of $\widehat{\mathcal{C}}F$ spanned by the hyperplanes of F that cross γ. We say that $F \in \mathfrak{F}$ is *γ–relevant* if $\pi_F(\gamma)$ is unbounded.

Proposition 3.4. *Let $F \in R(\gamma)$. Then $\pi_F(\gamma)$ is a quasigeodesic ray in $\widehat{\mathcal{C}}F$, and hence picks out a unique point $p^F_\gamma \in \partial\widehat{\mathcal{C}}F$ (the Gromov boundary). Moreover, $R(\gamma) \neq \emptyset$.*

Proof. Distance formula! □

Proposition 3.5. *Let γ be a combinatorial geodesic ray. Then the set $R(\gamma)$ of relevant elements of $\mathfrak{F}$ is pairwise-orthogonal and hence contains at most $n+1$ elements, where γ represents an n–simplex of $\partial_{\triangle}\mathcal{X}$.*

Proof. Let $v = [v_0, \ldots, v_n]$ be the simplex of $\partial_{\triangle}\mathcal{X}$ represented by γ. We will prove the claim by induction on n. You showed in an exercise that, since $\mathcal{X}$ has a factor system, it is fully visible. Moreover, it is clear that the claim depends only on the asymptoty class of γ, so we may modify γ in this way if needed. Hence, as shown in Section 3 of [Hag13], for each i we can choose a geodesic ray γ_i representing v_i so that there is a combinatorial isometric embedding $P = \prod_i \gamma_i \to \mathcal{X}$ whose image contains γ as a ray emanating from the basepoint and crossing every hyperplane crossing P.

Principle I: Let $\mathcal{H}(\gamma)$ be the set of hyperplanes crossing γ. If $F, F' \in R(\gamma)$, then $F \cap \mathcal{H}(\gamma), F' \cap \mathcal{H}(\gamma)$ are both infinite. If these sets have infinite intersection, then either $F = F' = \mathfrak{g}_F(F') = \mathfrak{g}_{F'}(F)$, or $F'' = \mathfrak{g}_F(F')$ is a proper subcomplex of F crossed by infinitely many hyperplanes crossing γ. So the projection of γ to $\widehat{\mathcal{C}}F$ is a subgraph consisting of vertices all adjacent to the cone-point over $\widehat{\mathcal{C}}_{F''} \subset \widehat{\mathcal{C}}F$, contradicting that $F \in R(\gamma)$. Thus if $F, F' \in R(\gamma)$ are distinct, then each of F, F' can cross only finitely many hyperplanes crossing γ.

The base case $n = 0$: If $F, F' \in R(\gamma)$ and $n = 0$, then $\mathcal{H}(\gamma)$ is a minimal boundary set. In particular, if infinitely many elements of $\mathcal{H}(\gamma)$ cross F, and infinitely many cross F', then infinitely many cross F and F'. So Principle I tells us $F = F'$. In particular, $R(\gamma)$ is a pairwise-orthogonal set with $\leq n+1$ elements.

The inductive step: Write $P' = \prod_{i=0}^{n-1} \gamma_i$, so that $P = P' \times \gamma_n$. Choose a "diagonal" ray γ'' in P' crossing every hyperplane, so that P contains $\gamma'' \times \gamma_n$ and so that $R(\gamma) = R(\gamma_n) \sqcup R(\gamma'')$ (since $R(\gamma)$ just depends on the set of hyperplanes crossing γ). By induction, $R(\gamma'')$ is a pairwise-orthogonal set with $\leq n$ elements, and $R(\gamma_n)$ is a singleton since v_n is a 0–simplex. Let F_n be the unique element of $R(\gamma_n)$ and let F'' be the product of the elements of $R(\gamma'')$, which is a convex subcomplex.

Let H be a hyperplane crossing γ'' (hence F''). Then H crosses each hyperplane crossing γ_n, whence $\mathfrak{g}_H(F_n) \sqsubset F_n$ contains the projection of γ, so $\mathfrak{g}_H(F_n) = F_n$, i.e. $F_n \sqsubset H$, by Principle I. This holds for any H crossing γ''. Similarly, for any V crossing γ_n and any $F' \in R(\gamma'')$, we have $F' \sqsubset V$. Another application of Principle I now shows that every hyperplane crossing F'' crosses every hyperplane crossing F_n. Applying Proposition 2.5 of [CS11] completes the proof. □

Remark 3.6 ("Nielsen-Thurston classification")**.** This is a digression, but we can now give the promised "Nielsen-Thurston classification. If $\mathcal{X}$ is a CAT(0) cube complex with a factor system $\mathfrak{F}$, and $g \in \mathrm{Aut}(\mathcal{X})$, then on of the following holds:

- g is elliptic;
- g is *reducible axial*: there exists $N > 0$ and $F \in \mathfrak{F} - \{\mathcal{X}\}$ so that g^N stabilizes F and acts loxodromically on $\widehat{\mathcal{C}}F$;
- g is *irreducible axial*: g acts loxodromically on $\widehat{\mathcal{C}}\mathcal{X}$.

The proof is an exercise and the digression is over.

We now define a simplicial complex $\Delta\mathcal{X}$ as follows: for each set $R = \{F_0, \ldots, F_n\}$ of pairwise-orthogonal elements of $\mathfrak{F}$, and each choice $(p_0, \ldots, p_n) \in \prod_i \partial\widehat{\mathcal{C}}F_i$, we put an n–simplex, which we regard as the set of formal sums $\sum_{i=0}^n a_i p_i$, where each $a_i \in [0,1]$ and $\sum_i a_i = 1$. (This makes the face relation obvious, and clearly $\Delta\mathcal{X}$ is a flag complex.)

Proposition 3.7. *$\Delta\mathcal{X}$ is isomorphic to $\partial_{\triangle}\mathcal{X}$.*

Sketch. There are a few steps. A detailed proof is carried out in Section 10 of [DHS15].

Check that $\mathcal{X}$ is fully visible: This was an exercise!

Represent boundary points with combinatorial rays: For each 0–simplex v of $\partial_{\triangle}\mathcal{X}$, let γ be a combinatorial geodesic ray representing v and let $R(\gamma)$ be the γ–relevant elements of $\mathfrak{F}$. From the definition of a minimal boundary set, it follows that $R(\gamma)$ has a unique element, say $F = F(\gamma)$. Indeed, if $F, F' \in R(\gamma)$, then they are orthogonal, so infinitely many hyperplanes crossing γ cross F (but not F') and infinitely many cross F' (but not F). These infinite sets contain sub-boundary sets of $\mathcal{H}(\gamma)$, contradicting that $\dim v = 0$. Let $\pi_F(\gamma)$ be the (uniformly quasigeodesic) projection of γ to $\widehat{\mathcal{C}}F$, representing a point $p_v \in \partial\widehat{\mathcal{C}}F$.

The isomorphism I: Define $I : \partial_{\triangle}\mathcal{X} \to \Delta\mathcal{X}$ as follows. For each 0–simplex v, let $I(v) = p_v$, where $p(v)$ is the 0–simplex of $\Delta\mathcal{X}$ defined above. Let v, v' be adjacent 0–simplices of $\partial_{\triangle}\mathcal{X}$. Then, since they're visible, we have (up to changing the basepoints of the rays) a cubical isometric embedding $\gamma \times \gamma' \to \mathcal{X}$, where γ, γ' represent v, v', as shown in [Hag13]. Hence there is a "diagonal" geodesic ray γ'' representing the simplex $v \star v'$, lying in $\gamma \times \gamma'$, and having the same projection as γ (resp. γ') on $F(\gamma)$ (resp. $F(\gamma')$) and its factored contact graph. Thus $F(\gamma), F(\gamma')$ are in $R(\gamma'')$ and thus orthogonal by Proposition 3.5. Thus there is a 1–simplex $[p_v, p_{v'}]$ in $\Delta\mathcal{X}$, and we let $I([v, v']) = [p_v, p_{v'}]$. Higher simplices can be handled similarly, or you can use that $\Delta\mathcal{X}$ and $\partial_{\triangle}\mathcal{X}$ are both flag complexes, or you can induct, or whatever. □

We've now understood points in the simplicial boundary in terms of the Gromov boundaries of the factored contact graphs. Since the factored contact graphs are (locally infinite, in general) quasi-trees, their boundaries have interesting but understandable topology, which we're going to exploit to build a sort of "contact graph cone topology" on $\partial_{\triangle}\mathcal{X}$. There's an extra wrinkle introduced by the orthogonality relations/the many non-asymptotic rays that in general represent the same simplex, i.e. are associated to the same collection of rays in factored contact graphs, namely that different rays can traverse their shadows on the factored contact graphs at different speeds (or, if you prefer the "old-school simplicial boundary" interpretation, they can punch through their different minimal boundary sets in different orders). This is part of the reason for the horrendousness of the definition of the topology that we now define.

Theorem 3.8 (The HHS boundary)**.** *Let $\mathcal{X}$ be a CAT(0) cube complex with a factor system $\mathfrak{F}$. Then there is a space $\partial_f\mathcal{X}$ with the following properties:*

(1) $\partial_f\mathcal{X} \cup \mathcal{X}$ is a compact, Hausdorff, separable space;
(2) $\mathcal{X}$ is dense in $\partial_f\mathcal{X} \cup \mathcal{X}$;
(3) there is a continuous bijection $\partial_{\triangle}\mathcal{X} \to \partial_f\mathcal{X}$ which is an embedding on each simplex;
(4) there is an embedding $\partial\widehat{\mathcal{C}}F \to \partial_f\mathcal{X}$ for each $F \in \mathfrak{F}$;
(5) if $\partial\widehat{\mathcal{C}}\mathcal{X}$ is nonempty (i.e. $\widehat{\mathcal{C}}\mathcal{X}$ is unbounded), and there is a group acting geometrically on $\mathcal{X}$, then $\partial\widehat{\mathcal{C}}\mathcal{X}$ is dense in $\partial_f\mathcal{X}$.

Remark 3.9. It follows easily from the definitions that $\widehat{\mathcal{C}}\mathcal{X}$ is quasi-isometric to $\mathcal{C}\mathcal{X}$ – the former is obtained by coning off bounded subgraphs of the latter – so the latter part of the theorem is related to the question of when the contact graph is bounded.

Remark 3.10. In fact, the boundary mentioned in the theorem is a special case of the *boundary of a hierarchically hyperbolic space (HHS)* introduced by Durham, the author, and Sisto in [DHS15]. This same construction also yields a bordification of the mapping class group, a new bordification of Teichmüller space, universal covers of non-geometric 3–manifolds, etc. From the HHS point of view, the cubical case is a particularly nice test case, both because the hyperbolic spaces out of whose Gromov boundaries we're building the boundary are particularly simple (quasi-trees) and because there's an alternative interpretation of the boundary (well, not topologically, but more than just as a set), namely the original definition of $\partial_{\triangle}\mathcal{X}$.

Remark 3.11 (Bounded Geodesic Image)**.** In one spot in the following discussion of Theorem 3.8, we need the "cubical Bounded Geodesic Image theorem" which we now briefly discuss. Let $H, F \in \mathfrak{F}$ and suppose $F \sqsubset H$. For concreteness, since we'll be working at the level of the factored contact graphs, we can just assume that $F \subsetneq H$. Now, for any $t \in \widehat{\mathcal{C}}H$, we can find $x \in \mathcal{X}$ that projects uniformly close to t, and then send it to $\widehat{\mathcal{C}}F$ by considering $\pi_F(x)$. This gives a (coarsely surjective) coarse map $\widehat{\mathcal{C}}H \to \widehat{\mathcal{C}}F$ (there's a map the other way, too, which is less interesting: just send everything in $\widehat{\mathcal{C}}F$ to the cone-point in $\widehat{\mathcal{C}}H$ over $\widehat{\mathcal{C}}F$, or use the inclusion, or whatever).

Now, suppose that γ is a (quasi)geodesic in $\widehat{\mathcal{C}}H$ that does not pass through the subgraph $\widehat{\mathcal{C}}F$. Then for any points $t, t' \in \gamma$, the corresponding points $x, x' \in \mathcal{X}$ can be chosen so that each hyperplane separating them (or, depending on the exact situation, all but uniformly many hyperplanes separating them) fails to cross F. Thus the image of γ under $\widehat{\mathcal{C}}H \to \widehat{\mathcal{C}}F$ is uniformly bounded. It follows that if γ is a (quasi)geodesic ray originating near the cone-point over $\widehat{\mathcal{C}}F$, then there is a point $t \in \gamma$ so that $\widehat{\mathcal{C}}H \to \widehat{\mathcal{C}}F$ is coarsely constant on $\gamma([t, \infty))$. This technical point plays a fairly major role in this business, and is a factor system analogue of the bounded geodesic image theorem in the mapping class group [MM00], but in the present text it only plays a minor technical role, in the definition of the topology on $\partial_f \mathcal{X}$ below.

Discussion of the proof of Theorem 3.8. As a set, $\partial_f \mathcal{X}$ is just $\Delta\mathcal{X}$. The most illuminating part is the definition of the topology on $\Delta\mathcal{X}$ yielding the boundary $\partial_f \mathcal{X}$, which is explained below. Proposition 3.7 provides the bijection $\partial_\triangle \mathcal{X} \to \partial_f \mathcal{X}$; it's just the identity (one must check it's an embedding on simplices). The definition of $\Delta\mathcal{X}$ provides the inclusions $\partial\widehat{\mathcal{C}}F \hookrightarrow \partial_f \mathcal{X}$ (which one must check are embeddings). The definition of the topology will make it obvious that $\partial_f \mathcal{X}$ is Hausdorff and separable.

The details of the proof are in [DHS15] and are more or less technical applications of the definitions, although the proof of compactness is quite involved. In [DHS15], we work in the context of "hierarchically hyperbolic spaces". As explained in [BHS14], the factor system $\mathfrak{F}$ and the projections prove that $\mathcal{X}$ is an HHS, so that the arguments in [DHS15] apply. Section 10 of [DHS15] also discusses the relationship between $\partial_\triangle \mathcal{X}$ and $\partial_f \mathcal{X}$, in a slightly different way.

We now discuss how to define the topology. Specifically, we'll build a neighborhood basis for a topology on $\mathcal{X} \cup \partial_f \mathcal{X}$ consisting of open balls in $\mathcal{X}$ together with sets that we now define.

- Let $p = \sum_{F \in R} a_F p_F$, where R is a set of parallelism classes in $\mathfrak{F}$ with pairwise-orthogonal representatives, and each $a_F \in (0,1)$, and each $p_F \in \partial\widehat{\mathcal{C}}F$. Let $\epsilon > 0$, and let $\mathcal{O}_F$ be a neighborhood of p_F in $\widehat{\mathcal{C}}F \cup \partial\widehat{\mathcal{C}}F$ (with the cone topology). We will associate a set $\mathcal{N}_{\epsilon,\{\mathcal{O}_F\}}(p) \subset \mathcal{X} \cup \Delta\mathcal{X}$ to this data, and the set of all such sets will be our neighborhood basis.
- Let's first decide what it means for an interior point $x \in \mathcal{X}$ to be in $\mathcal{N}_{\epsilon,\{\mathcal{O}_F\}}$. To be close to p, it seems reasonable, in view of the distance formula, that we should require x to project close to p in each relevant factored contact graph, i.e. $\pi_F(x) \in \mathcal{O}_F$ for each $F \in R$. That's not enough, though: there's a whole orthant in $\prod_{F \in R} F$ with the given property; we also need to account for the slope of the ray representing p. Accordingly, fixing once and for all a basepoint $x_0 \in \mathcal{X}$, we require that

$$\left| \frac{a_F}{a_{F'}} - \frac{\mathsf{d}_{\widehat{\mathcal{C}}F}(\pi_F(x_0), \pi_F(x))}{\mathsf{d}(\pi_{F'}(x_0), \pi_{F'}(x))} \right| < \epsilon$$

  for $F, F' \in R$. Finally, we don't want some $H \in \mathfrak{F}$, orthogonal to every $F \in R$, so that x has some large H–coordinate dragging it out of the orthant corresponding to the p_F.

Therefore, we require that for all such H and all $F \in R$,

$$\frac{\mathsf{d}_{\widehat{\mathcal{C}}H}(\pi_H(x_0), \pi_H(x))}{\mathsf{d}_{\widehat{\mathcal{C}}F}(\pi_F(x_0), \pi_F(x))} < \epsilon.$$

This defines the *interior* part of our basic set $\mathcal{N}_{\epsilon, \{\mathcal{O}_F\}}$.

- Let $q = \sum_{H \in R'} b_H q_H$ be a boundary point, with $q_H \in \partial\widehat{\mathcal{C}}H$ and $b_H \in (0,1)$. To test whether $q \in \mathcal{N}_{\epsilon,\{\mathcal{O}_F\}}(p)$, we have to test whether q lies in approximately the same direction as p as measured in $\widehat{\mathcal{C}}F, F \in R$, and that q does not make much progress away from p elsewhere. To formalize this, we have to consider two types of q:
  - **Non-generic:** Here, either $R \cap R' \neq \emptyset$, or there exists $H \in R'$ which is orthogonal to each $F \in R$. In this case, let $R'' = R \cap R'$. In order to include q in our neighborhood, we require that:
    * $q_H \in \mathcal{O}_H$ for all $H \in R''$;
    * $\sum_{H \in R' - R''} b_H < \epsilon$ (i.e. most of the action is in directions q has in common with p);
    * $|a_H - b_H| < \epsilon$ for all $H \in R''$.
  - **Generic:** The final case is that where $R \cap R' = \emptyset$ and, moreover, each $F \in R$ is **not** orthogonal to some $H_F \in R'$. First we require that

$$\sum_H b_H < \epsilon$$

    , where the sum is taken over the $H \in R'$ that are orthogonal to every $F \in R$. Next, we are going to define a projection $\pi_F(p) \in \widehat{\mathcal{C}}F$ for each $F \in R$ and require that q project close to the p–direction in the p–relevant factored contact graphs, with the right contribution from each one, i.e.:
    * $\pi_F(q) \in \mathcal{O}_F$ for each $F \in R$;
    * for all $F, F' \in R$,

$$\left| \frac{\mathsf{d}_{\widehat{\mathcal{C}}F}(\pi_F(x_0), \pi_F(q))}{\mathsf{d}_{\widehat{\mathcal{C}}F'}(\pi_{F'}(x_0), \pi_{F'}(q))} - \frac{a_F}{a_{F'}} \right| < \epsilon.$$

    How to define π_F? Well, by definition we have some $H \in R'$ so that H is neither orthogonal to F nor equal to F. One can check that this provides a uniformly bounded subset ρ of $\widehat{\mathcal{C}}F$ obtained by closest-point projecting H to F and then projecting to $\widehat{\mathcal{C}}F$ unless $F \sqsubset H$. So, in this case, we let $\pi_F(q) = \rho$; this is forced on us: we'd better be projecting the boundary of H to the same place in $\widehat{\mathcal{C}}F$ we're projecting the rest of H!

    Otherwise, $F \sqsubset H$. In this case, let γ_H be a quasigeodesic ray in $\widehat{\mathcal{C}}H$ from the cone-point over $\widehat{\mathcal{C}}F \subset \widehat{\mathcal{C}}H$ to $q_H \in \partial\widehat{\mathcal{C}}H$. Bounded Geodesic Image provides a first point $t \in \gamma_H$ so that γ_H is coarsely constant after t. Choose $x \in \mathcal{X}$ so that $\pi_H(x)$ coarsely coincides with t, then let $\pi_F(q) = \pi_F(x)$.

This completes the description of the topology. □

Example 3.12 (Examples)**.** Two simple examples:

(1) If $\mathcal{X}$ is a hyperbolic CAT(0) cube complex with a proper cocompact group action, then Agol's virtual specialness theorem combines with results of [BHS14] to show that $\mathcal{X}$ has a factor system. As one would hope and expect, $\partial_f\mathcal{X}$ turns out to be homeomorphic to the Gromov boundary in this case, as shown in [DHS15].

(2) If $\mathcal{X}$ decomposes as a product of unbounded subcomplexes, then $\partial_f\mathcal{X}$ is homeomorphic to $\partial_\triangle\mathcal{X}$, which (recall) is a join.

3.2. **Stationary measures on the boundary, rank-rigidity, and the like.** We'll focus on uses of $\partial_f\mathcal{X}$ that involve the next few lemmas. Everything in this section works in the more general context of *hierarchically hyperbolic spaces* (see [BHS14, BHS15, DHS15]), but the cubical case is more concrete but does not hide any of the essential ideas. Also, one can relax cocompactness in various ways. (Besides, these are cubical lectures.)

In this section, we work in the setting of a group G acting geometrically on $\mathcal{X}$, which is a CAT(0) cube complex equipped with a factor-system $\mathfrak{F}$. Equip G with a Borel probability measure μ whose support generates G (use that G is finitely-generated if you want).

Using compactness of $\partial_f\mathcal{X}$, one can construct a μ–stationary measure ν on $\mathcal{X}\cup\partial_f\mathcal{X}$, i.e. a Borel probability measure such that

$$\nu(E)=\sum_{g\in G}\mu(g)\nu(g^{-1}E)$$

for all ν–measurable subsets $E\subseteq\mathcal{X}\cup\partial_f\mathcal{X}$.

Consider the action of G on the set $\overline{\mathfrak{F}}$ induced by the action on $\mathfrak{F}$ and the fact that G preserves parallelism. Note that $\mathcal{X}$ is fixed.

Lemma 3.13. *If G has no finite orbit in $\overline{\mathfrak{F}}-\{[\mathcal{X}]\}$ and no finite orbit in $\partial\widehat{\mathcal{C}}\mathcal{X}$, then ν is supported on $\partial\widehat{\mathcal{C}}\mathcal{X}\subset\partial_f\mathcal{X}$.*

Hence if $\mathrm{diam}(\widehat{\mathcal{C}}\mathcal{X})<\infty$*, there exists $F\in\mathfrak{F}-\{\mathcal{X}\}$ and $G'\leq_{f.i.}G$ so that gF is parallel to F for all $g\in G'$.*

Proof. The second assertion follows from the first since ν can't be supported on $\emptyset$, so it remains to prove the first assertion.

Let D be the set of finite subsets of $\overline{\mathfrak{F}}$, so that G has no finite orbit in D other than $G\cdot\{[\mathcal{X}]\}$ and $G\cdot\emptyset$. We can assume there exists $F\in\mathfrak{F}-\{\mathcal{X}\}$; otherwise $\partial_f\mathcal{X}=\partial\widehat{\mathcal{C}}\mathcal{X}$ and we're done.

We'll define a G–equivariant map $\mathcal{O}:\mathcal{X}\cup\partial_f\mathcal{X}\to D$ which is measurable when D is endowed with the probability measure $\tilde{\nu}$ given by $\tilde{\nu}(A)=\nu(\mathcal{O}^{-1}(A))$ for all $A\subseteq D$. We will also check that $E=\mathcal{X}\cup\partial_f\mathcal{X}-\partial\widehat{\mathcal{C}}\mathcal{X}$ is measurable and $\mathcal{O}(E)$ contains no finite G–orbit. It follows from e.g. [KM96, Lemma 2.2.2],[Woe89],[Bal89],[Hor14] that $\nu(E)=0$, and we're done.

So: if $p\in\partial_f\mathcal{X}$ is represented by some ray γ, let $R(\gamma)$ be the set of relevant factor-system elements, and let $\mathcal{O}(p)=R(\gamma)$. This is G–equivariant, and $\mathcal{O}(p)=\{[\mathcal{X}]\}$ if and only if $p\in\partial\widehat{\mathcal{C}}\mathcal{X}$. To define the map on $\mathcal{X}$, let $C\subset\mathcal{X}$ be a subset which contains exactly one open cube of $\mathcal{X}$ in each G–orbit, and exactly one 0–cube in each G–orbit. Since G acts geometrically, C is a finite union of Borel subsets of $\mathcal{X}\cup\partial_f\mathcal{X}$, so C is measurable. Fix a set $A\in D-\{\emptyset,[\mathcal{X}]\}$, which exists because we can use $\{F\}$. Let $\mathcal{O}(x)=A$ for each $x\in C$, and extend G–equivariantly to all of $GC=\mathcal{X}$. From the definition, one can easily check that $\mathcal{O}$ is Borel (use that G is countable).

Since $\mathcal{X}$ is an open subset of $\mathcal{X}\cup\partial_f\mathcal{X}$, it follows from Lemma 3.14 that E is measurable. Finally, if $p\in\mathcal{X}$, then $G\cdot\mathcal{O}(p)=G\cdot A$, which is infinite. If $p\in\partial_f-\partial\mathcal{X}$, then $G\mathcal{O}(p)$ must be infinite, since otherwise it would be a finite union of subsets of $\overline{\mathfrak{F}}-\{[\mathcal{X}]\}$, which cannot be G–invariant by hypothesis. □

Lemma 3.14. *For any finite subset R of $\overline{\mathfrak{F}}$, the set of $p\in\partial_f\mathcal{X}$ whose set of relevant factored contact graphs is R is a Borel set.*

Proof. This is a horrible chase (it's 12 lines in [DHS15], but not fun lines) through the definition of the topology on the boundary; exercise. □

We can now reprove a special case of the Caprace-Sageev rank-rigidity theorem [CS11], as a test of our tool.

Theorem 3.15 (Rank-rigidity assuming factor systems)**.** *Let $\mathcal{X}$ be a CAT(0) cube complex with a factor system $\mathfrak{F}$. Suppose that some group G acts properly and cocompactly on $\mathcal{X}$. Then one* *of the following holds provided G acts* essentially *on $\mathcal{X}$ in the sense that every halfspace contains G–orbit points arbitrarily far from the associated hyperplane:*

(1) $\mathcal{C}\mathcal{X}$ is unbounded, and G contains a rank-one isometry of $\mathcal{X}$ acting loxodromically on $\mathcal{C}\mathcal{X}$.

(2) $\mathcal{C}\mathcal{X}$ is bounded, and $\mathcal{X}$ splits as the product of two unbounded convex subcomplexes.

The second conclusion holds if and only if $\partial_{\triangle}\mathcal{X}$ decomposes as a nontrivial simplicial join.

Remark 3.16. Caprace-Sageev require a proper, cocompact, essential action *or* an essential action with no fixed point at infinity. Our proof actually works in the second context, too (see [DHS15]), but in either case, we need a factor-system (while [CS11] does not). A positive answer to Question A would allow us to drop that hypothesis in the cocompact setting. More to the point (since rank-rigidity already has a nice proof), it would hopefully allow one to solve other problems using these techniques, in the general setting of cocompactly cubulated groups.

Proof of Theorem 3.15. First suppose that $\mathcal{C}\mathcal{X}$ is bounded, i.e. $\partial\widehat{\mathcal{C}}\mathcal{X} = \emptyset$ (since $\widehat{\mathcal{C}}\mathcal{X}$ and $\mathcal{C}\mathcal{X}$ are q.i.). Lemma 3.13 provides $F \in \mathfrak{F} - \{\mathcal{X}\}$ so that (up to finite index) G preserves the parallelism class of F. Essentiality tells us that $\mathcal{X} = F \times F^{\perp}$. If $F^{\perp}$ is bounded, then essentiality tells us it's trivial, so $\mathcal{X} = F$, a contradiction. Otherwise, F is bounded, so is trivial by essentiality, which we disallowed. Hence the second conclusion holds.

Hence suppose that $\mathcal{C}\mathcal{X}$ is unbounded (i.e. $\widehat{\mathcal{C}}\mathcal{X}$ is unbounded). There are two ways to proceed.

Method I: use Caprace-Sageev double-skewering: Let H, H' be hyperplanes corresponding to vertices at distance at least 10 in $\mathcal{C}\mathcal{X}$. By the Double-Skewering Lemma [CS11], there exists a hyperbolic element $g \in G$ so that H' separates H and gH (i.e. H, H' cut the axis of g). It is easy to check (Exercise 3.3.(4)) that $\langle g \rangle$ has an unbounded orbit in $\mathcal{C}\mathcal{X}$, so our above discussion about projecting geodesic rays to the contact graph, applied to an axis of g, shows that g is loxodromic on $\mathcal{C}\mathcal{X}$ and in particular rank-one.

Method II: use acylindricity: In [BHS14], it is shown that G acts *acylindrically* on $\widehat{\mathcal{C}}\mathcal{X}$. A result of Osin [Osi15] combines with the fact that $\widehat{\mathcal{C}}\mathcal{X}$ is unbounded to show that G contains a loxodromic isometry of $\widehat{\mathcal{C}}\mathcal{X}$, which is necessarily rank-one as above (since half-flats have diameter-≤ 3 projection to $\mathcal{C}\mathcal{X}$). □

The key here (from the point of view of the boundary; we also used other serious tools) was Lemma 3.13, which should provide a template for proving various similar results (in the more general HHS situation, one can use it to prove things like the Tits alternative or a generalization of the Handel-Mosher "omnibus subgroup theorem" for mapping class groups [DHS15]; hopefully in the more restricted setting of cube complexes with factor systems, one can do even more.

We've focused heavily on the case where $\mathcal{X}$ has a factor system; indeed, we've mainly left aside things that one can do with the simplicial boundary in general (see e.g. [?]). Now, the definition of a factor system is not particularly natural-seeming, even if it is what one is led to when trying to write down a Masur-Minsky-style distance formula for right-angled Artin groups (factor systems are not the first attempt at the latter [KK14]). It would be nice if, at least in the case of cube complexes with proper, cocompact group actions, one had access to the factor system technology, the boundary $\partial_f\mathcal{X}$, and the methods we've illustrated above, without having to hypothesize the existence of $\mathfrak{F}$. This is what a positive answer to Question A would enable.

3.3. **Final problems.** We finish with a list of problems related to Part 3, and some other interesting problems related to (or possibly approachable using) the boundary. The latter list is biased toward things I think are very interesting but don't know how to do, or total speculation.

Problems on $\partial_f\mathcal{X}$:

(1) Prove the main property of orthogonal complements: for each $F \in \mathfrak{F}$, there exists a maximal convex subcomplex $F^\perp$ so that $F \to \mathcal{X}$ extends to a convex embedding $F \times F^\perp \to \mathcal{X}$ with the property that each parallel copy of F is the image of $F \times \{e\}$ for some vertex $e \in F^\perp$. (Hint: take the convex hull of the union of all parallel copies of F, and then find two classes of hyperplanes, forming a join in the contact graph, from which you can read off the product decomposition.)

(2) Prove the "Nielsen-Thurston classification" outlined in Remark 3.6.

(3) Prove Lemma 3.14.

(4) Check that if H, H' are hyperplanes at large contact-graph distance, both cutting an axis A of a hyperbolic isometry g of $\mathcal{X}$, then A has unbounded projection to $\mathcal{C}\mathcal{X}$.

(5) Let $\mathbf{R}_2$ denote the standard tiling of $\mathbb{E}^2$ by 2–cubes, and let $\mathcal{X}$ be the obvious tree of $\mathbf{R}_2$s on which $\mathbb{Z}^2 * \mathbb{Z}^2$ acts geometrically. How is $\mathcal{C}\mathcal{X}$ related to the Bass-Serre tree? Describe $\partial_f\mathcal{X}$.

$*$ (6) What else can you prove about cube complexes with factor systems, and groups acting on them, using Lemma 3.13 or related ideas? (E.g. Tits alternative...)

$*$ (7) For interesting non-hyperbolic cube complexes $\mathcal{X}$ with factor systems, arising in nature (e.g. from RAAGs, Coxeter groups, etc.), describe $\partial_f\mathcal{X}$ in a satisfying topological way. (In other words, prove theorems along the lines of various theorems describing Gromov boundaries of hyperbolic groups.)

Problems on relative hyperbolicity, thickness, divergence, quasi-isometries: In the following problems, $\mathcal{X}$ is a CAT(0) cube complex on which the group G acts geometrically (and, say, essentially, although this is only necessary for some parts).

$*$ (8) In [BH], the existence of nontrivial relatively hyperbolic structures on G is characterized in terms of the structure of $\partial_\triangle\mathcal{X}$ and the action of G. Roughly, G is hyperbolic relative to subgroups $\{P_i\}$ if $\partial_\triangle\mathcal{X}$ consists of a G–invariant collection of isolated 0–simplices, together with a bunch of subcomplexes which are the simplicial boundaries of convex subcomplexes stabilized by the various P_i and their conjugates. This can be phrased completely in terms of the action on $\partial_\triangle\mathcal{X}$, although it is fiddly.

In [BDM09], Behrstock-Drutu-Mosher introduced the notion of a *thick space*. Thickness is an obstruction to the existence of a nontrivial relatively hyperbolic structure, and is enjoyed by lots of the usual suspects, like one-ended RAAGs, mapping class groups, $\mathrm{Out}(F_n)$, etc. $\mathcal{X}$ is *thick of order* 0 if it's a product with unbounded factors [Hag13] (there is a more general definition for general metric spaces: no cutpoints in any asymptotic cone). Inductively, a space is thick of order n if there is a collection $\{U_i\}$ of undistorted subspaces, all thick of order $n-1$, so that any two points in the main space are connected through a chain of elements of $\{U_i\}$ so that successive elements have coarsely connected, unbounded intersection. In [BH], we characterised thickness of order 1 for G completely in terms of $\partial_\triangle\mathcal{X}$ and the action of G thereon; roughly, it corresponds to the existence of a connected G–invariant subcomplex, but with the overall boundary being disconnected.

Must G have a relatively hyperbolic structure whose peripheral subgroups have no nontrivial RH structure? Whose peripheral subgroups are thick? Maybe studying $\partial_\triangle\mathcal{X}$ is useful here. Note that this is true for Coxeter groups [BHS], which are cubulated [NR03], though not always cocompactly. This might involve **generalizing the above result on thickness of order 1**: can one identify higher-order thickness of G from $\partial_{\triangle}\mathcal{X}$? Note that for any k,n there is a k–dimensional $\mathcal{X}$ so that G is thick of order exactly n [BH].

Thickness is also related to the divergence function of G: if G is thick of order n, then its divergence is polynomial of order at most $n+1$, see [BD14]. In general, it is hard to prove lower bounds on divergence. **If G is thick of order n, is the divergence function of $\mathcal{X}$ exactly polynomial of order $n+1$?** Must the divergence function of $\mathcal{X}$ be either polynomial or exponential? (A positive answer would follow from a positive answer to the above question on minimal RH structures with thick peripherals, and a positive answer to the question about divergence of thick cubical groups.)

Other problems:

∗ (9) **Graph colouring and Question A.** Let $\mathcal{C}_{\sharp}\mathcal{X}$ be the intersection graph of the hyperplanes in $\mathcal{X}$ (the *crossing graph*). Suppose that $h:\mathcal{C}_{\sharp}\mathcal{X}\to\Gamma$ is a graph homomorphism with the property that $v,w\in\mathcal{C}_{\sharp}\mathcal{X}^{(0)}$ are adjacent whenever $h(v),h(w)$ are adjacent, and suppose that Γ is finite. (Call h a Γ*–colouring of* $\mathcal{X}$.)

Note that if $\mathcal{X}$ admits a Γ–colouring for some finite Γ, then $\mathcal{C}_{\sharp}\mathcal{X}$ has finite chromatic number. In fact, this corresponds to an isometric embedding of $\mathcal{X}$ in a finite product of trees, and this is not always possible even when you remove obvious restrictions by bounding the degree of 0–cubes in $\mathcal{X}$ (see [CH13]). However, if $\mathcal{X}$ has a factor system and a geometric group action and satisfies one additional mild technical hypothesis, then $\mathcal{X}$ admits a very particular type of Γ–colouring corresponding to an *equivariant* embedding in a finite product of trees [BHS14].

So, suppose $\mathcal{X}$ admits a Γ–colouring $h:\mathcal{C}_{\sharp}\mathcal{X}\to\Gamma$. Let S_Γ be the Salvetti complex associated to Γ, with 1–cubes labelled by vertices of Γ and, more generally, n–cubes labelled by cliques in Γ. Define a map $q:\mathcal{X}\to S_\Gamma$ as follows. First, $q(x)$ is the unique 0–cube of S_Γ for each 0–cube $x\in\mathcal{X}$. For each 1–cube e of $\mathcal{X}$, let H be the dual hyperplane, and send e to the 1–cube of S_Γ labelled by $h(H)$. Extend in the obvious way to higher cubes. This map is locally injective because no two 1–cubes with a common endpoint in $\mathcal{X}$ are dual to the same hyperplane. It's a cubical map since h is a homomorphism. It's a local isometry since any cube in S_Γ whose corner is in the image pulls back to $\mathcal{X}$ by the second condition on h. Hence q lifts to a convex embedding $\mathcal{X}\to\widetilde{S}_\Gamma$. In other words, Question A can be rephrased as: if $\mathcal{X}$ admits a proper, cocompact group action, does it admit a Γ–colouring for some finite Γ. Does this help?

∗ (10) **"Restricted triple space".** Let $\mathcal{X}$ be any CAT(0) cube complex for which $\partial_{\triangle}\mathcal{X}$ is defined (i.e. no factor system assumptions). A triple of distinct points $p,q,r\in\partial_{\triangle}\mathcal{X}$ is *visible* if there do not exist distinct $a,b\in\{p,q,r\}$ so that $a\in v,b\in w$ where v,w are simplices with nonempty intersection. Let $\mathfrak{V}$ be the set of visible triples. Using Proposition 1.14 and the relationship between the Tits boundary and $\partial_{\triangle}\mathcal{X}$, along with the fact that $\mathcal{X}^{(1)}$ is a *median graph*, one should be able to construct a (coarse) map $\mathfrak{V}\to\mathcal{X}$ roughly as follows: given a visible triple p,q,r, form an "ideal triangle" of bi-infinite combinatorial geodesics whose endpoints are p,q,r. Then, for each $t\geq 0$, travel distance t in the positive and negative directions on each of these geodesics, to give you three pairs of (uniformly close) points. Take the median (coarsely) and show that this stabilizes as $t\to\infty$. Can $\mathfrak{V}$ be topologized so that this map is useful? Is there some notion of a "uniform convergence action" related to this construction? What if $\mathcal{X}$ has a factor system and you're using triples in $\partial_f\mathcal{X}$ instead?

∗ (11) **Simplicial-ish boundary of a median space?** Define an analogue of the simplicial boundary of a measured wallspace, or for the dual median space (see [CDH10]). What can you do with this object?

* (12) What topological spaces can arise as $\partial_f \mathcal{X}$ where $\mathcal{X}$ is a CAT(0) cube complex with a factor system?
* (13) If $\mathcal{X}$ admits two factor systems, must the associated compactified simplicial boundaries be homeomorphic? (I guess "yes", and maybe this is not so hard, in view of Theorem 3.8.)
* (14) If G acts properly and cocompactly on CAT(0) cube complexes $\mathcal{X}$ and $\mathcal{Y}$, must $\partial_\triangle \mathcal{X}$ and $\partial_\triangle \mathcal{Y}$ be isomorphic simplicial complexes? Homeomorphic? (Croke-Kleiner-type problems [CK00] don't arise because $\partial_\triangle \mathcal{X}$ has no idea which CAT(0) metric you're using. It is shown in [Hag14a] that if G is virtually $\mathbb{Z}^n$, then $\partial_\triangle \mathcal{X}$ has to be the hyperoctahedron you expect, but I don't think the isomorphism $\partial_\triangle \mathcal{X} \to \partial_\triangle \mathcal{Y}$ has to be equivariant in general, when there is an interesting point group. The point of [Hag14a] was to make an argument that generalizes to more interesting cubical groups, but it's not clear that it's the best strategy to follow...)
* (15) Let G act geometrically on $\mathcal{X}$. It is shown in [WW15] that each virtually abelian subgroup $A \leq G$ which is "highest", in the sense that no finite-index subgroup of A is contained in a higher-rank abelian subgroup, stabilizes a convex subcomplex of the form $\prod_i C_i$ where each C_i is the convex hull of a combinatorial line (and lies at finite Hausdorff distance from that line). Hence A stabilizes a hyperoctahedron in the boundary, namely $\partial_\triangle \prod_i C_i$. Can this be used to generalize the main result of [Hag14a] by proving that certain groups cannot admit a cocompact cubulation? Coxeter groups come to mind, here. (Something like: find a virtually abelian subgroup that doesn't act the right way on the required hyperoctahedron (like a $(3,3,3)$ triangle subgroup), then combine the results of [WW15] and [Hag14a].)

References

[Bal89] Werner Ballmann. On the dirichlet problem at infinity for manifolds of nonpositive curvature. In *Forum Mathematicum*, volume 1, pages 201–214, 1989.

[BD14] Jason Behrstock and Cornelia Druţu. Divergence, thick groups, and short conjugators. *Illinois Journal of Mathematics*, 58(4):939–980, 2014.

[BDM09] Jason Behrstock, Cornelia Druţu, and Lee Mosher. Thick metric spaces, relative hyperbolicity, and quasi-isometric rigidity. *Mathematische Annalen*, 344(3):543–595, 2009.

[BH] Jason Behrstock and Mark F. Hagen. Cubulated groups: thickness, relative hyperbolicity, and simplicial boundaries. *To appear in Groups, Geometry, and Dynamics.*

[BHS] J Behrstock, M Hagen, and A Sisto. Thickness, relative hyperbolicity, and randomness in coxeter groups, 2013. *Preprint (with an appendix jointly written with PE. Caprace)*, 12.

[BHS14] Jason Behrstock, Mark F. Hagen, and Alessandro Sisto. Hierarchically hyperbolic spaces I: curve complexes for cubical groups. *arXiv:1412.2171*, 2014.

[BHS15] Jason Behrstock, Mark F. Hagen, and Alessandro Sisto. Hierarchically hyperbolic spaces II: combination theorems and the distance formula. *ArXiv eprint 1509.00632*, 2015.

[CDH10] Indira Chatterji, Cornelia Druţu, and Frédéric Haglund. Kazhdan and Haagerup properties from the median viewpoint. *Advances in Mathematics*, 225(2):882–921, 2010.

[CH13] Victor Chepoi and Mark F Hagen. On embeddings of CAT(0) cube complexes into products of trees via colouring their hyperplanes. *Journal of Combinatorial Theory, Series B*, 103(4):428–467, 2013.

[CK00] Christopher B Croke and Bruce Kleiner. Spaces with nonpositive curvature and their ideal boundaries. *Topology*, 39(3):549–556, 2000.

[CS11] Pierre-Emmanuel Caprace and Michah Sageev. Rank rigidity for CAT(0) cube complexes. *Geom. Funct. Anal.*, 21:851–891, 2011.

[DHS15] Matthew G Durham, Mark F Hagen, and Alessandro Sisto. Boundaries and automorphisms of hierarchically hyperbolic spaces. *In preparation*, pages 1–60, 2015.

[Hag07] Frédéric Haglund. Isometries of CAT(0) cube complexes are semi-simple. *arXiv:0705.3386*, 2007.

[Hag13] Mark F. Hagen. The simplicial boundary of a CAT(0) cube complex. *Alg. Geom. Topol.*, 13:1299–1367, 2013.

[Hag14a] Mark F Hagen. Cocompactly cubulated crystallographic groups. *Journal of the London Mathematical Society*, 90(1):140–166, 2014.

[Hag14b] Mark F Hagen. Weak hyperbolicity of cube complexes and quasi-arboreal groups. *J. Topol.*, 7(2):385–418, 2014.

[Hor14] Camille Horbez. A short proof of Handel and Mosher's alternative for subgroups of $Out(F_N)$. *arXiv preprint arXiv:1404.4626*, 2014.

[KK14] Sang-Hyun Kim and Thomas Koberda. The geometry of the curve graph of a right-angled artin group. *International Journal of Algebra and Computation*, 24(02):121–169, 2014.

[KM96] Vadim A Kaimanovich and Howard Masur. The Poisson boundary of the mapping class group. *Inventiones mathematicae*, 125(2):221–264, 1996.

[MM99] Howard A Masur and Yair N Minsky. Geometry of the complex of curves i: Hyperbolicity. *Inventiones mathematicae*, 138(1):103–149, 1999.

[MM00] Howard A Masur and Yair N Minsky. Geometry of the complex of curves ii: Hierarchical structure. *Geometric and Functional Analysis*, 10(4):902–974, 2000.

[NR03] Graham A Niblo and Lawrence D Reeves. Coxeter groups act on CAT(0) cube complexes. *Journal of Group Theory*, 6(3):399, 2003.

[Osi15] Denis Osin. Acylindrically hyperbolic groups. *Transactions of the American Mathematical Society*, 2015.

[Woe89] Wolfgang Woess. Boundaries of random walks on graphs and groups with infinitely many ends. *Israel Journal of Mathematics*, 68(3):271–301, 1989.

[WW15] Daniel T Wise and Daniel J Woodhouse. A cubical flat torus theorem and the bounded packing property. *arXiv preprint arXiv:1510.00365*, 2015.

DPMMS, University of Cambridge, Cambridge, UK

E-mail address: markfhagen@gmail.com