May 2015

The Markov-Dubins Problem with Free Terminal Direction in a Nonpositively Curved Cube Complex

Jason Thomson La Corte
University of Wisconsin-Milwaukee

The Markov-Dubins Problem with Free Terminal Direction in a Nonpositively Curved Cube Complex

by

Jason Thomson La Corte

A Dissertation Submitted in

Partial Fulfillment of the

Requirements for the Degree of

Doctor of Philosophy

in

Mathematics

at

The University of Wisconsin–Milwaukee

May 2015

ABSTRACT

THE MARKOV-DUBINS PROBLEM WITH FREE TERMINAL DIRECTION IN A NONPOSITIVELY CURVED CUBE COMPLEX

by

Jason Thomson La Corte

The University of Wisconsin–Milwaukee, 2015
Under the Supervision of Professor Craig Guilbault

State complexes are nonpositively curved cube complexes that model the state spaces of reconfigurable systems. The problem of determining a strategy for reconfiguring the system from a given initial state to a given goal state is equivalent to that of finding a path between two points in the state complex. The additional requirement that allowable paths must have a prescribed initial direction and minimal turning radius determines a Markov-Dubins problem with free terminal direction (MDPFTD).

Given a nonpositively curved, locally finite cube complex X, we consider the set of unit-speed paths which satisfy a certain smoothness condition in addition to the boundary conditions and curvature constraint that define a MDPFTD. We show that this set either contains a path of minimal length, or is empty.

We then focus on the case that X is a surface with a nonpositively curved cubical structure. We show that any solution to a MDPFTD in X must consist of finitely many geodesic segments and arcs of constant curvature, and we give an algorithm for determining those solutions to the MDPFTD in X which are CL paths, that is, made up of an arc of constant curvature followed by a geodesic segment. Finally, under the assumption that the 1-skeleton of X is d-regular, we give sufficient conditions for a topological ray in X of constant curvature to be a rose curve or a proper ray.

Dedicated to my father,
who always believed.

Table of Contents

List of Figures

Acknowledgments

To my parents, my grandparents, and my brother, whose support throughout the years gave me the comfort and confidence I needed to carry on.

To my wife Diana, who brightened the darkest of hours with enduring love, hope, and patience.

To Carolle Bendle, whose bold intervention led to all the educational opportunities which followed. You're still my hero, after all this time.

To the man she reached out to, whose generosity enabled me to experience worlds I would otherwise never have known.

To David Snyder, for his integrity and for the guidance he provided as my first graduate advisor, which I have so often recalled in the intervening years.

To Craig Guilbault, for his invaluable academic and professional counsel, and for always encouraging me to pursue my interests.

To all the faculty of the Mathematical Sciences Department at UW–Milwaukee, for creating and sustaining a supportive, inclusive departmental culture of which its members can be proud.

And to all the individuals without whom my journey through higher education would not have been possible—thank you.

List of Symbols

$\mathbb{N}$	set of positive integers
$\mathbb{E}^2$	Euclidean plane $\mathbb{R}^2$ with canonical cubical structure
$\#A$	cardinality of a set A
$\approx$	is homeomorphic to
$\mathrm{B}(x;r)$	metric ball with center x and radius r
$\ell(\gamma)$	length of a path γ
$\overline{\gamma}$	reverse path of γ
$\alpha * \beta$	concatenation of paths α and β
$[xy]$	geodesic from x to y
$x * A$	union of geodesics from a point x to points $y \in A$
$\tilde{\angle}$	comparison angle
$\hat{\angle}$	upper angle
$\angle$	Alexandrov angle
o	little-o Landau symbol
d_∞	uniform metric
$\mathcal{C}(X,Y)$	set of continuous functions $X \to Y$
$\rho(\sigma)$	total rotation of a piecewise geodesic σ
$\tau(\gamma)$	total curvature of a path γ
$\mathrm{Flag}(\mathcal{P})$	flag complex of a poset $\mathcal{P}$
$\lvert\mathcal{P}\rvert$	geometric realization of a poset $\mathrm{Flag}(\mathcal{P})$
$L(W,S)$	nerve of a Coxeter system (W,S)
$\Sigma(W,S)$	Davis complex of (W,S)
$X(W,S)$	Davis complex as piecewise Euclidean cube complex
$S(E)$	label of an edge E in a Cayley graph
$H(E)$	1-dimensional hyperplane dual to an edge E
$H \,)(\, H'$	hyperplanes H and H' osculate
a-$\mathcal{C}^{1,1}$	a-Lipschitz differentiable
a-$\mathcal{C}^{1,1}(\mathcal{Q},\mathcal{T})$	a-$\mathcal{C}^{1,1}$ with respect to a cube path $\mathcal{Q}$ and breakpoint sequence $\mathcal{T}$
a-$\mathcal{C}^{1,1}(M)$	a-$\mathcal{C}^{1,1}$ with at most M breakpoints
$\mathcal{U}(\mathcal{Q})$	unfolding complex for a square path $\mathcal{Q}$
$\mathcal{X}_d$	d-plane

Chapter 1

Introduction

1.1 Motivation

The current work is devoted to the problem of finding the shortest curvature-constrained path between two points with prescribed initial direction in a nonpositively curved cube complex X, which we will call the *Markov-Dubins problem with free terminal direction*. Special attention is paid to the case dim $X = 2$.

This problem arises naturally in connection with *reconfigurable systems*, a class of controllable dynamical systems introduced by Ghrist in [1]. Ghrist's initial examples of reconfigurable systems were metamorphic robots—aggregates capable of changing shape through the independent motion of their constituent cells—but many of the industrial robots which can already be found in the factories of the present day may be formally regarded as reconfigurable systems, and further examples outside the field of robotics abound [17]. The state space of a reconfigurable system can be constructed as a nonpositively curved cube complex called a *state complex*. Paths in the state complex correspond to strategies for reconfiguring the system from one state to another. If we require our paths to have bounded curvature, we impose a dynamic constraint which not only rules out physically unrealistic motion, but also limits the amount of strain placed on the moving parts of the system.

Finding a computationally feasible method for determining shortest curvature-constrained paths with prescribed initial direction is an important first step toward the development of fault-tolerant control algorithms able to seamlessly adopt new reconfiguration strategies in response to a dynamically changing environment. Indeed, the task which must be carried out by such an algorithm can be viewed as a game of pursuit and evasion in which a human or computer controller chases a moving goal state in the state complex, and it is partly in response to the request for additional study of dynamically constrained pursuit-evasion problems in CAT(0) domains made in [2] that the current work has been undertaken.

1.2 Reconfigurable systems and their state spaces

Roughly speaking, a reconfigurable system consists of a graph $\mathcal{G}$ with labeled vertices, called the *workspace graph*, and a collection of partial relabelings, or *generators*. A labeling of the workspace graph represents a state of the underlying system, e.g. the angle of all the joints in a robotic arm. Each generator φ consists of a pair of labelings of the vertices in a given subgraph $\mathcal{H}$ of $\mathcal{G}$. If the current labeling of $\operatorname{Vert}\mathcal{H} \subseteq \operatorname{Vert}\mathcal{G}$ matches one of the labelings in φ, then φ *acts* on $\mathcal{G}$ by relabeling $\operatorname{Vert}\mathcal{H}$ so that it matches the other labeling in φ. In this way, each generator represents an invertible elementary motion of the underlying system.

The formal definition [17] of a reconfigurable system is as follows.

Definition 1.2.1. Fix a set $\mathcal{A}$ of labels, and a graph $\mathcal{G}$ called the **workspace graph**. A **state** is a function $\operatorname{Vert}(\mathcal{G}) \to \mathcal{A}$. A **generator** φ is an ordered triple

$$(\operatorname{supp}(\varphi),\ \operatorname{tr}(\varphi),\ \{u_0^{\mathrm{loc}}, u_1^{\mathrm{loc}}\}),$$

where

- $\mathrm{supp}(\varphi)$ is a subgraph of $\mathcal{G}$, called the **support** of φ;
- $\mathrm{tr}(\varphi)$ is a subgraph of $\mathrm{supp}(\varphi)$, called the **trace** of φ; and
- $\{u_0^{\mathrm{loc}}, u_1^{\mathrm{loc}}\}$ is an unordered pair of functions

$$u_0^{\mathrm{loc}}, u_1^{\mathrm{loc}} : \mathrm{Vert}(\mathrm{supp}\,\varphi) \to \mathcal{A},$$

called **local states**, such that

$$u_0^{\mathrm{loc}} \neq u_1^{\mathrm{loc}}$$

and

$$u_0^{\mathrm{loc}}\big|_{\mathrm{supp}(\varphi)\smallsetminus\mathrm{tr}(\varphi)} = u_1^{\mathrm{loc}}\big|_{\mathrm{supp}(\varphi)\smallsetminus\mathrm{tr}(\varphi)}.$$

A generator φ is said to be **admissible** at a state u if

$$u\big|_{\mathrm{supp}(\varphi)} = u_0^{\mathrm{loc}}.$$

If a generator φ is admissible at a state u, we define the **action** of φ on $\mathcal{G}$ to be

$$\varphi[u] = \begin{cases} u & \text{on } \mathrm{Vert}\left(\mathcal{G} \smallsetminus \mathrm{supp}(\varphi)\right) \\ u_1^{\mathrm{loc}} & \text{on } \mathrm{supp}(\varphi). \end{cases}$$

A **reconfigurable system** on $\mathcal{G}$ is a collection of generators and a collection of states closed under all possible actions. A reconfigurable system is **locally finite** if the number of generators admissible at u, as u ranges over all possible states, is bounded above.

For example, consider a pair of robots free to slide along two tracks in a factory floor, one a line segment, and one a circle. To obtain the workspace graph, we discretize the two tracks. (See Figure 1.1.)

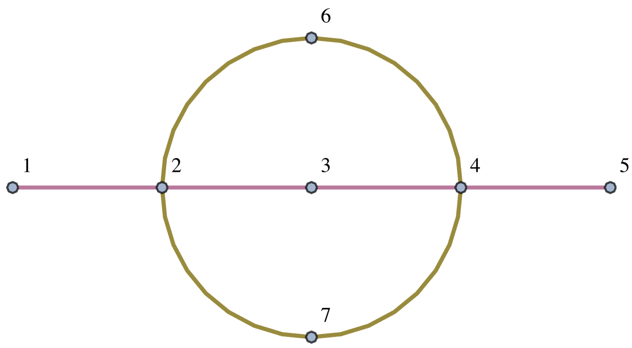

Figure 1.1: Workspace graph.

Representing the two robots as dots of different colors, a state is a function

$$u : \{1, 2, 3, 4, 5, 6, 7\} \to \{\circ, \bullet, \bullet\}.$$

A generator φ is defined by the following pair of moves:

$$\begin{aligned} &\text{If vertex } 1 \text{ is occupied by } \bullet \text{ and vertex } 2 \text{ is unoccupied,} \\ &\qquad\qquad\qquad\qquad \text{move } \bullet \text{ to vertex } 2. \\ &\text{If vertex } 2 \text{ is occupied by } \bullet \text{ and vertex } 1 \text{ is unoccupied,} \\ &\qquad\qquad\qquad\qquad \text{move } \bullet \text{ to vertex } 1. \end{aligned}$$

Here, $\text{supp}(\varphi) = \{1, 2\}$, $\text{tr}(\varphi) = \varnothing$, and

$$u_0^{\text{loc}} : 1 \mapsto \bullet,\ 2 \mapsto \circ; \qquad\qquad u_1^{\text{loc}} : 1 \mapsto \circ,\ 2 \mapsto \bullet.$$

Definition 1.2.2. Given a reconfigurable system, we define its **transition graph** $\mathcal{T}$ to be the graph whose vertices are possible states, and join a pair of states by an edge if there exists a generator that toggles the system between them (Figure 1.2).

The transition graph $\mathcal{T}$ is analogous to the Cayley graph of a group, but need not be homogeneous: the number of generators that can be applied need not be the same at every state.

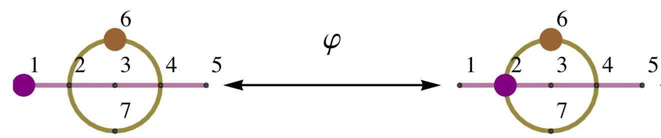

Figure 1.2: An edge in the transition graph.

A path in the transition graph determines a strategy for reconfiguring the system from one state to another by sequentially carrying out elementary moves. But when robots can move to adjacent positions on their tracks without collisions, it is more efficient to move them at the same time.

Definition 1.2.3. We say that a set of k generators φ_i, $i \in \{1, \ldots, k\}$, **commute** if

$$\mathrm{tr}(\varphi_i) \cap \mathrm{supp}(\varphi_j) = \varnothing$$

whenever $i \neq j$. Wherever the 1-skeleton $Q^{(1)}$ of a k-cube appears in $\mathcal{T}$, attach a k-cube if for each vertex u of $Q^{(1)}$, the generators corresponding to the edges incident with u commute at u (Figure 1.3). The resulting cube complex is the **state complex** of the system.

Physically, generators commute at a state u iff they can be applied to u simultaneously, and if the resulting configuration is independent of the order in which they are carried out.

Interior points of a cube are intermediate stages of a transition between states. A path along a k-cube's diagonal represents the simultaneous execution of the k commuting generators corresponding to the k parallelism classes of the cube's edges.

In a real world environment, changing circumstances in the physical workspace may intervene to make a reconfiguration strategy that is already in progress impossible to complete. If a goal state has been prescribed, but an obstruction prevents us from attaining it, a new goal state in the state complex may then be prescribed. It is impractical to bring the moving parts of the system to a dead stop and then instantaneously follow the new strategy, and it is inefficient to slow them to a halt whenever a new strategy is required. We therefore

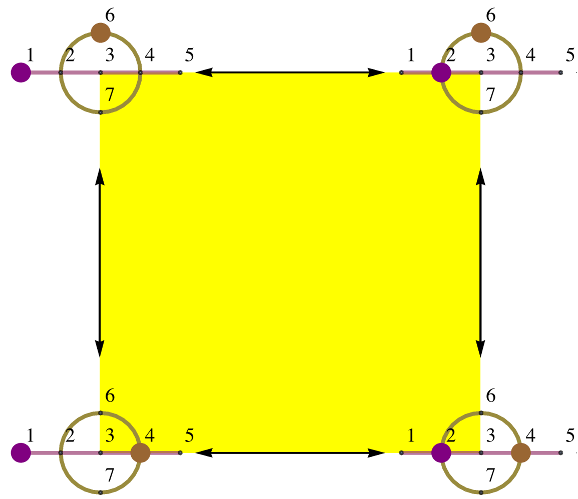

Figure 1.3: A 2-cell in the state complex.

seek a solution to the problem of finding a shortest path in the state complex with a given initial direction.

We make the assumption that sudden, drastic changes in direction by a path in the state space correspond to jerky movements that put stress on the physical components of the system. Taking an extreme example, a 180° change in direction represents a rotation of at least one joint instantaneously followed by a rotation in the opposite direction. Such a movement is neither desirable nor practical, and we would like to rule it out.

If our state space was Euclidean space, the "jerkiness" of such movements could be limited simply by constraining the integral curvature of our paths. But generic state spaces do not sufficiently resemble manifolds for classical analysis to be applicable without substantial augmentation. For starters, they need not be of uniform dimension. Consider, for example, the positive articulated robot arm described in [17]. It consists of several joints that bend in the same plane, constrained so that the arm can only extend up and to the right. Its state space (Figure 5 in [17]) is not a manifold. Examples like this show that state complexes need not be amenable to traditional theories of flow.

We also note that the appearance of sharp corners in paths when the state complex is

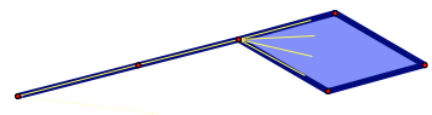

Figure 1.4: Geodesics cannot be uniquely extended through the indicated vertex.

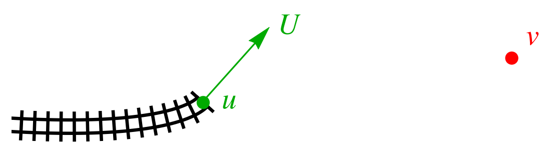

Figure 1.5: Markov's problem.

visualized as a subset of Euclidean space can be misleading. In the state complex pictured, a geodesic beginning at the leftmost point may take on any of several directions when it meets the leftmost square, because no matter what direction is chosen, the metric on this complex forces the angle at the point of entry into the square to be 0 (Figure 1.4). Any approach to finding a curvature-constrained path in such a space must account for this phenomenon.

Our search for a suitable formulation of the problem of finding length-minimal paths of bounded curvature begins with a problem posed by Markov in the nineteenth century, and solved by Dubins nearly seventy years later.

1.3 Dubins problems for Riemannian manifolds

In an 1889 paper published in the Russian-language journal *Communications of the Kharkov Mathematical Society*, Markov considered the problem of joining an existing section of railroad track to a given destination v using as little new track as possible, with turning radius nowhere less than a given bound $r > 0$. The heading U of the existing track is given, but any heading at the terminal point of the new track is acceptable (Figure 1.5).

A formal statement of Markov's problem, which we will now give, appears in [24]. Given two points $u, v \in \mathbb{R}^2$, a number $r > 0$, and a unit tangent vector U at u, find the shortest smooth curve satisfying the following properties:

- The curve's radius of curvature is nowhere less than r.
- The curve has initial position u, initial direction U, and terminal position v.
- The curve consists of finitely many convex arcs, which need not be convex in the same sense.

Markov showed that the solution to this problem always exists, and consists of a line segment beginning at u followed by an arc of a circle C_v of radius r, where C_v is a translate of a circle tangent to (U, u) along the direction U. In the same 1889 paper, Markov discussed the case where a terminal vector V at v is also specified, but, as noted by Kreĭn and Nudel'man [24], did not provide a correct solution.

The latter problem was taken up in a 1957 paper by Dubins [15] which begins by asking which definition of "smoothness" ought to be used in its formulation. At first ignoring Markov's curvature constraint, Dubins observes that it is easy to choose u, v, U, V such that the infimum of lengths of $\mathcal{C}^1$ paths with initial tangent vector (U, u) and terminal tangent vector (V, v) is not attained by any such path. With only slightly more work, Dubins shows that the same is true if we consider $\mathcal{C}^2$ paths. The main achievements of the paper are to show (*i*) that the collection

$$\mathcal{C}_{\text{Dubins}} := \{\Gamma \in A(u, U, v, V) : \gamma \text{ is differentiable and } \gamma' \text{ is } (1/r)\text{-Lipschitz}\},$$

where $A(u, U, v, V)$ is the set of paths satisfying the boundary conditions, must contain a path of minimal length for any choice of u, v, U, V, and (*ii*) that such a path is made up of at most three pieces, each a line segment or an arc of a circle of radius r.

Several authors have replicated Dubins' results and extended them to spaces other than the Euclidean plane. Using nothing more than advanced calculus and plane geometry, Reeds and Shepp gave a new proof of Dubins' existence and characterization results, and extended Dubins' results by allowing paths to reverse directions, thus allowing optimal arcs to contain cusps [33]. Thomaschewski also gives a brief proof of Dubins' characterization using elementary methods [37]. Boissonat, Cérézo, and Leblond re-proved Dubins' results using Pontryagin's maximum principle [7], while Sussmann and Tang gave an alternate proof of Reeds and Shepp's results which combined Pontryagin's maximum principle with Lie-algebraic techniques [36]. Sussmann later adapted these techniques to solve Dubins' problem in Euclidean 3-space [35]. Again using a combination of Pontryagin's maximum principle and Lie algebra, Monroy and Mittenhuber gave sufficient and necessary conditions for the existence of a solution to Dubins' problem in the 2-sphere [30] and in the hyperbolic plane [28], respectively, and in a pair of papers, Mittenhuber subsequently did the same for the higher-dimensional manifolds of constant curvature $\mathbb{H}^n$, $\mathbb{R}^n$, and S^n [27, 29]. Finally, Chitour and Sigalotti gave sufficient and necessary conditions for the existence of a solution to Dubins' problem in a connected, oriented, complete Riemannian 2-manifold in [12, 34, 13].

While the techniques employed by these authors are effective for studying Markov-Dubins problems on a Riemannian manifold, they can not easily be adapted to the cases of non-manifolds and manifolds which do not admit a Riemannian metric. Our approach in the current work is to apply the theory of comparison geometry. This theory is equally well-suited to Riemannian manifolds, non-Riemannian manifolds, and non-manifolds, and is typically no more difficult to apply to higher-dimensional spaces than to lower-dimensional ones. We are particularly interested in square complexes, but our existence result applies to a very broad range of nonpositively curved cube complexes.

The application we have in mind is to change from a reconfiguration strategy that is already in process to a new reconfiguration, possibly with a different goal state. We are interested only in attaining the goal state, and we assume that the underlying reconfigurable system will cease all motion once it is attained. Therefore, in the current work, the terminal direction of the path in the state space will be regarded as meaningless, and we focus on the Markov-Dubins problem with free terminal direction.

1.4 Algorithms for finding geodesics in a cube complex

In order to numerically determine solutions to the Markov-Dubins problem in a cube complex, it will be necessary to employ algorithms which compute the geodesic between two points. As we have concentrated on dimension-2 complexes, the “funnel algorithm” of Lee and Preparata [25] will be adequate. Chepoi and Maftuleac have shown how the funnel algorithm can be adapted for use in any nonpositively curved square complex [11], and it is the latter adaptation which has been used in all our numerical experiments.

An algorithm for use in nonpositively curved complexes of arbitrary dimension has been given by Ardila, Owen, and Sullivant [5]. We look forward to employing an implementation of their algorithm in the future; however, their algorithm will not be discussed in the present work, which will not treat the numerical solution of Markov-Dubins problems in cube complexes of dimension $\geqslant 3$.

1.5 Summary of results in this work

Our first significant result is to show that a certain collection of paths is guaranteed to contain a length-minimizing solution to the Markov-Dubins problem with free terminal direction in a locally finite nonpositively curved cube complex X, provided that it is possible to satisfy the boundary conditions with some admissible (not necessarily length-minimal) path. That is, given two points $u, v \in X$, a direction U at u, and a curvature bound $a > 0$, we define a set $\mathcal{C}$ of “smooth” paths from u to v with initial direction U and curvature nowhere greater than a, and show that this set either is empty, or contains a path of minimum length (Chapter 3).

Our smoothness condition is twofold, having a piecewise component and a condition on behavior at breakpoints. First, we require our paths to be *a-Lipschitz differentiable* with respect to some cube path. Such a path γ is differentiable with a-Lipschitz derivative in each cube in some cube path for γ. At breakpoints, which need not have Euclidean neighborhoods, we require our paths to have zero turning angle. The direction of a path at a point will be identified with a geodesic segment, and a pair of geodesic segments with common initial point p will be identified with the class of curves having prescribed approach and departure directions at p (Definition 3.3.3). These pairs, which we call *hinges*, give us a concrete representation of the directional behavior of a path which is flexible enough to accommodate the phenomenon of bifurcating geodesics.

In Section 4.1, we characterize the elements of $\mathcal{C}$ which have minimal length as being made up of finitely many geodesic segments and arcs of circles of radius a. It is clear that no bound on the number of such segments exists, due to the variety of spaces which can be given the cell structure of a nonpositively curved cube complex. (For example, any winding hallway which does not intersect itself can be regarded as a CAT(0) complex.)

We then give an algorithm for numerically determining a large class of length-minimizing paths in nonpositively curved square complexes of bounded curvature, and verify its effectiveness (Section 4.2). This algorithm requires a procedure for finding geodesics in a simple polygon and a procedure for finding geodesics in a nonpositively curved square complex. We use the classical funnel algorithm of Lee and Preparata [25] for the former, and a technique due to Chepoi and Maftuleac [11] for the latter. We have implemented our algorithm in Mathematica for the spaces $\mathcal{X}_d$, $d \geqslant 4$ (Definition 5.0.8).

Chapter 5 focuses on the behavior of topological rays (see Definition 2.1.1) of constant curvature in the family of spaces $X = \mathcal{X}_d$, $d \geqslant 5$. These spaces, which are the Davis complexes of certain right-angled Coxeter systems (Lemma 5.1.1), are homeomorphic to planes. We find that rays of constant curvature may be properly embedded half-lines, homeomorphic to circles, or self-intersecting rose curves, and give sufficient conditions for each case (§5.2.2 and §5.2.3).

In §5.2.2, we develop a procedure for transferring a *stack*—that is, a sequence of successively osculating hyperplanes—between cube complexes. (We define *hyperplanes* in a cube complex in Definition 2.4.3, and what it means for two hyperplanes to *osculate* in Definition 2.4.9.) Our strategy for showing that a ray $\gamma : [t_1, \infty) \to \mathcal{X}_d$ is proper is to construct a sequence of nested halfspaces $H_k^+ \supset H_{k+1}^+$ containing the tails of the ray: that is, such that

$$\gamma\big([t_k, \infty)\big) \subset H_k^+ \tag{1.1}$$

for some $t_1 < t_2 < \cdots$. Finding such a sequence of halfspaces is equivalent to determining a stack $(H_k)_{k=1}^{\infty}$. That the hyperplanes H_k bound nested halfspaces can be deduced from the group presentation of (W, S) (Lemma 2.4.11) and the fact that $\mathcal{X}_d$ is a special cube complex.

In order to construct a suitable sequence of hyperplanes in $\mathcal{X}_d$, we transfer the ray γ and its carrier to $\mathbb{E}^2$ using a procedure analogous to analytic continuation. In the case under consideration, it is easy to find (finite) sequences of of osculating hyperplanes in $\mathbb{E}^2$—parallel lines which are distance 1 apart—through which the ray passes. These sequences are then transferred back to $\mathcal{X}_d$, and we verify (Lemma 5.2.5) that the result is an (infinite) sequence of osculating hyperplanes through which γ passes.

When transferring stacks from $\mathbb{E}^2$ to $\mathcal{X}_d$, choices must be made, as the carriers of a pair of osculating hyperplanes of $\mathbb{E}^2$ meet along infinitely many pairs of adjacent 2-cells, and each selection of such a pair determines a different pair of hyperplanes in $\mathcal{X}_d$. We keep track of our choices using a *scaffold* (Figure 2.8). A scaffold for a given stack is a sequence of pairs of adjacent edges such that the pairs in each edge are respectively dual to a pair of osculating hyperplanes. Given a suitable scaffold in $\mathbb{E}^2$, mapping the edges of the scaffold from $\mathbb{E}^2$ to $\mathcal{X}_d$ yields a stack in $\mathcal{X}_d$ such that γ satisfies condition (1.1) (Theorem 5.2.2).

Chapter 2

Preliminaries

2.1 Metric geometry

2.1.1 Length spaces

Definition 2.1.1. Let X be a topological space. A map $\gamma : J \to X$, where $J \subset \mathbb{R}$ is an interval of the form $[a,b]$, $[a,b)$, or $[a,\infty)$, is a **curve** in X with **initial point** $\gamma(a)$. If $J = [a,b]$, γ is a **path from** $\gamma(a)$ **to** $\gamma(b)$. If $J = [a,\infty)$, γ is a **(topological) ray**. A curve is **degenerate** if its image is a point.

Definition 2.1.2. The **reverse path** of a path $\gamma : [a,b] \to X$ in a topological space X is

$$\overline{\gamma}(t) = \alpha\big(b + (b-t)\big), \qquad b \leqslant t \leqslant b + (b-a).$$

Definition 2.1.3. Let X be a topological space. The **concatenation** $\alpha * \beta$ of paths $\alpha : [a,b] \to X$ and $\beta : [b,c] \to X$ such that $\alpha(b) = \beta(b)$ is

$$(\alpha * \beta)(t) = \begin{cases} \alpha(t) & a \leqslant t \leqslant b, \\ \beta(t) & b \leqslant t \leqslant c. \end{cases}$$

Definition 2.1.4. A **change of variable** is a nondecreasing continuous surjection between compact intervals.

Definition 2.1.5. The **length** $\ell(\gamma)$ of a path $\gamma : [a, b] \to M$ in a metric space (M, d) is

$$\ell(\gamma) = \sup \left\{ \sum_{k=1}^{m} d(\gamma(t_{k-1}), \gamma(t_k)) : a = t_0 < t_1 < \cdots < t_m = b \right\}.$$

Lemma 2.1.6 ([9]). If $\gamma : [a, b] \to X$ is a path in a topological space X and $L : [c, d] \to [a, b]$ is a change of variable, then $\ell(\gamma L) = \ell(\gamma)$.

Definition 2.1.7. A metric space (X, d) is a **length space** if for each $x, y \in X$,

$$d(x, y) = \inf \{\ell(\gamma) : \gamma \text{ is a path from } x \text{ to } y\}.$$

Definition 2.1.8 ([6, 8]). Let (X, d) be a length space. A path $\gamma : [a, b] \to X$ such that for some $c > 0$,

$$d(\gamma(s), \gamma(t)) = c|s - t|, \qquad \forall\, s, t \in [a, b],$$

is a **geodesic (path)** from $u = \gamma(a)$ to $v = \gamma(b) \neq u$. Its image $[uv]$ is a **geodesic segment**. If $[uv]$ is given, we call γ a **parametrization** of $[uv]$, and we say γ has **constant speed** c. A **geodesic ray** $\gamma : [a, \infty) \to X$ is defined analogously.

Remark. A parametrization of a geodesic segment in a length space is a homeomorphism, being a continuous bijection from a compact interval to a Hausdorff space.

Definition 2.1.9. A length space X is a **[uniquely] geodesic space** if for any two points $x, y \in X$ there is a [unique] geodesic in X from x to y.

Angles

Definition 2.1.10. Let x, y, z be three points in a metric space (X, d) such that $x \neq y \neq z$. The **comparison angle** $\tilde{\angle} xyz$ is

$$\tilde{\angle} xyz = \arccos \frac{d(x, y)^2 + d(y, z)^2 - d(x, z)^2}{2d(x, y)d(y, z)}. \tag{2.1}$$

Lemma 2.1.11. Let x, y, z be three points in a metric space (X, d). If $x \neq y \neq z$, then $\tilde{\angle} xyz$ exists.

Proof. By a proposition of Euclid, a triangle Δabc in $\mathbb{R}^2$ with

$$|a - b| = d(x, y),$$

$$|b - c| = d(y, z),$$

$$|a - c| = d(x, z)$$

can be constructed, since the sum of any two of these quantities is no less than the third ([21], Bk. I, Prop. 22). By hypothesis, $d(x, y)\, d(y, z) \neq 0$. The Law of Cosines now yields

$$\frac{d(x, y)^2 + d(y, z)^2 - d(x, z)^2}{2d(x, y)\, d(y, z)} \in [-1, 1].$$

□

Definition 2.1.12. Let $J = [a, b]$ or $[a, \infty)$. A curve $\gamma : J \to X$ in a space X is **initially injective** if there exists a $\delta > 0$ such that γ is injective on $[a, a + \delta)$.

Definition 2.1.13. Let (X, d) be a metric space. Let $\alpha : [a, c) \to X$ and $\beta : [b, d) \to X$ be initially injective curves with $\alpha(a) = \beta(b)$. The **upper angle** of α and β at $\alpha(a)$ is

$$\hat{\angle}(\alpha, \beta) = \limsup_{\substack{s \to a^+ \\ t \to b^+}} \tilde{\angle} \alpha(s)\alpha(a)\beta(t) = \lim_{\substack{s \to a^+ \\ t \to b^+}} \Big[\sup_{\substack{\sigma \geqslant s \\ \tau \geqslant t}} \tilde{\angle} \alpha(\sigma)\alpha(a)\beta(\tau) \Big],$$

and the **(Alexandrov) angle** of α and β at $\alpha(a)$ is

$$\angle(\alpha, \beta) = \lim_{\substack{s \to a^+ \\ t \to b^+}} \tilde{\angle}\alpha(s)\alpha(a)\beta(t).$$

Remark. Observe that the angle of two curves is independent of a choice of an initially injective parametrization for each. We therefore define the angle of two nondegenerate geodesic paths $[a, b]$ and $[a, c]$ in the obvious way, as the angle of their (constant speed) parametrizations.

Lemma 2.1.14. Let (X, d) be a metric space. Let $\alpha : [a, c) \to X$ and $\beta : [b, d) \to X$ be initially injective curves with $\alpha(a) = \beta(b)$. Then $\hat{\angle}(\alpha, \beta)$ exists.

Proof. By hypothesis, $d\big(\alpha(s), \alpha(a)\big) \neq 0$ and $d\big(\beta(b), \beta(t)\big) \neq 0$ for s near a and t near b. Thus by Lemma 2.1.11, the comparison angle $\tilde{\angle}\alpha(s)\alpha(a)\beta(t)$ is defined whenever s and t are sufficiently near a and b respectively. □

Definition 2.1.15. A length space X is **nonpositively curved** if for each point in X, there exists a neighborhood $\mathcal{N}$ of x such that

$$\angle yxz \leqslant \tilde{\angle} yxz, \qquad \angle xyz \leqslant \tilde{\angle} xyz, \qquad \angle xzy \leqslant \tilde{\angle} xzy$$

for all $x, y, z \in \mathcal{N}$ such that $x \neq y \neq z$. We say X is **CAT(0)** if in addition X is simply connected.

Lemma 2.1.16 ([8])**.** CAT(0) spaces are contractible and uniquely geodesic.

Lemma 2.1.17 ([9])**.** If a sequence of paths γ_n in a nonpositively curved space X converges uniformly to a path γ, then

$$\ell(\gamma) \leqslant \liminf_{n \to \infty} \ell(\gamma_n).$$

Lemma 2.1.18 ([8])**.** Let α and β be two initially injective rays in a length space X with common initial point. If X is nonpositively curved, then $\angle(\alpha, \beta) \in [0, \pi]$ exists.

Lemma 2.1.19 ([9], Proposition 3.6.27.1 and Exercise 3.6.28)**.** Let $\alpha : [0, \varepsilon] \to \mathbb{R}^d$ and $\beta : [0, \varepsilon] \to \mathbb{R}^d$ be two paths parametrized by arclength. Suppose $\alpha(0) = \beta(0)$. If α and β are differentiable at 0, then

$$\cos \angle(\alpha, \beta) = \langle \alpha'(0), \beta'(0) \rangle.$$

Proof. The angle $\angle(\alpha, \beta)$ exists and is equal to $\theta \in [0, \pi]$ if

$$|\alpha(s) - \beta(t)|^2 = s^2 + t^2 - 2st \cos\theta + o(st).$$

Since α and β are differentiable at 0,

$$\alpha(s) = \alpha(0) + s\alpha'(0) + o(s) \qquad \text{and} \qquad \beta(t) = \beta(0) + t\beta'(0) + o(t).$$

Noting that $|\alpha'(0)| = |\beta'(0)| = 1$ by hypothesis, we have

$$\begin{aligned} |\alpha(s) - \beta(t)|^2 &= \left|s\alpha'(0) - t\beta'(0) + o(s) - o(t)\right|^2 \\ &= s^2|\alpha'(0)|^2 + t^2|\beta'(0)|^2 - 2st \langle \alpha'(0), \beta'(0)\rangle + o(s^2) + o(t^2) + o(st) \\ &= s^2 + t^2 - 2st \langle \alpha'(0), \beta'(0)\rangle + o(st). \end{aligned}$$

□

Lemma 2.1.20 (Triangle inequality for angles: [9], Theorem 3.6.34)**.** Let $\alpha, \beta, \gamma : [0, \varepsilon] \to X$ be paths in a metric space (X, d) with common initial point $\alpha(0) = \beta(0) = \gamma(0) = p$. If each of $\angle(\alpha, \beta)$, $\angle(\beta, \gamma)$, and $\angle(\beta, \gamma)$ exist, then $\angle(\alpha, \gamma) \leqslant \angle(\alpha, \beta) + \angle(\beta, \gamma)$.

2.1.2 Piecewise geodesics and turning angle of a curve

Definition 2.1.21. Let $(x_k)_{k=0}^m$ be a finite sequence of points in a geodesic space such that $x_{k-1} \neq x_k$ for each $k \in \{1, \ldots, m\}$. Let γ_k be a parametrization of $[x_{k-1}x_k]$ for each

$k \in \{1, \ldots, m\}$. Then the concatenation $\gamma_1 * \cdots * \gamma_m$ is a **piecewise geodesic path** with **breakpoints** $(x_k)_{k=1}^m$.

Approximation by piecewise geodesics

Definition 2.1.22. Let X be a metric space. If $\gamma : [a, b] \to X$ is a path,

$$a = t_0 < \cdots < t_m = b,$$

and

$$\gamma(t_{k-1}) \neq \gamma(t_k)$$

for $k \in \{1, \ldots, m\}$, we say a piecewise geodesic path σ with breakpoints $\big(\gamma(t_k)\big)_{k=0}^m$ is **inscribed** in γ.

Lemma 2.1.23 ([16])**.** If X is a geodesic space, then for any nondegenerate rectifiable path $\gamma : [a, b] \to X$ and any $\varepsilon > 0$, there is a piecewise geodesic path $\alpha : [a, b] \to X$ inscribed in γ such that $d_\infty(\alpha, \gamma) < \varepsilon$, where d_∞ is the sup metric on $\mathcal{C}\big([a, b], X\big)$.

Proof. Let

$$a = t_0 < \cdots < t_m = b$$

be a partition of $[a, b]$ such that $\gamma(t_{k-1}) \neq \gamma(t_k)$ for $k \in \{1, \ldots, m\}$. Let

$$\alpha_k : [t_{k-1}, t_k] \to X$$

be a constant-speed parametrization of $[\gamma(t_{k-1})\gamma(t_k)]$ for $k \in \{1, \ldots, m\}$. Since γ is uniformly continuous, as $[a, b]$ is compact, there exists a function

$$\omega_\gamma : [0, b - a] \to [0, \infty)$$

such that

$$\omega_\gamma(\delta) \xrightarrow{\delta \to 0} 0 = \omega_\gamma(0)$$

and

$$d\big(\gamma(s), \gamma(t)\big) \leqslant \omega_\gamma\big(|s - t|\big)$$

for all $s, t \in [a, b]$. For each $k \in \{1, \dots, m\}$ and each $t \in [t_{k-1}, t_k]$, since

$$\begin{aligned} d\big(\gamma(t_k), \alpha_k(t)\big) = d\big(\alpha_k(t_k), \alpha_k(t)\big) &= \frac{t_k - t}{t_k - t_{k-1}} d\big(\alpha_k(t_k), \alpha_k(t_{k-1})\big) \\ &\leqslant d\big(\alpha_k(t_k), \alpha_k(t_{k-1})\big) = d\big(\gamma(t_k), \gamma(t_{k-1})\big), \end{aligned}$$

we have

$$\begin{aligned} d\big(\gamma(t), \alpha_k(t)\big) &\leqslant d\big(\gamma(t), \gamma(t_k)\big) + d\big(\gamma(t_k), \gamma(t_{k-1})\big) \\ &\leqslant \omega_\gamma(t - t_k) + \omega_\gamma(t_k - t_{k-1}) \xrightarrow{\text{mesh}(t_j) \to 0} 0. \end{aligned}$$

Thus, given $\varepsilon > 0$, we can take $\text{mesh}(t_j)$ so small that

$$\sup_{t_{k-1} \leqslant t \leqslant t_k} d\big(\gamma(t), \alpha_k(t)\big) < \varepsilon$$

for each $k \in \{1, \dots, m\}$. □

Turning angle

Definition 2.1.24. The **turning angle** at x_k of a piecewise geodesic with breakpoints (x_i) is

$$\pi - \angle\big([x_k x_{k-1}], [x_k x_{k+1}]\big).$$

Definition 2.1.25. Let X be a metric space. The **turning angle** at t of a curve $\gamma : J \to X$ is

$$\pi - \angle(\overline{\gamma|_{J_t^-}}, \gamma|_{J_t^+}),$$

where $J_t^- = J \cap (-\infty, t]$, $J_t^+ = J \cap [t, \infty)$, and the overbar indicates the reverse path.

Lemma 2.1.26 ([8])**.** For any three points $x \neq y \neq z$, the piecewise geodesic $[xy] * [yz]$ is a geodesic iff its turning angle at y is 0.

Convergence behavior of angles

Definition 2.1.27. Let $(x_n)_{n=1}^{\infty}$ and $(y_n)_{n=1}^{\infty}$ be sequences in a geodesic space. The geodesics $[x_n y_n]$ **converge** to a geodesic $[xy]$, and we write $[x_n y_n] \to [xy]$, if there exist parametrizations of $[x_n y_n]$ that uniformly converge to a parametrization of $[xy]$.

Theorem 2.1.28 (**Upper semicontinuity of angles**: [9], Theorem 4.3.11)**.** Let $(x_n)_{n=1}^{\infty}$, $(y_n)_{n=1}^{\infty}$, and $(z_n)_{n=1}^{\infty}$ be sequences in a nonpositively curved geodesic space X with $x_n \neq y_n \neq z_n$ distinct for each n. If $[x_n y_n] \to [xy]$ and $[y_n z_n] \to [yz]$ for some $x, y, z \in X$ such that $x \neq y \neq z$, then $\limsup_{n\to\infty} \angle x_n y_n z_n \leqslant \angle xyz$.

Corollary 2.1.29. Let $(x_n)_{n=1}^{\infty}$, $(y_n)_{n=1}^{\infty}$, and $(z_n)_{n=1}^{\infty}$ be sequences in a nonpositively curved geodesic space X with $x_n \neq y_n \neq z_n$ distinct for each n, and suppose $[x_n y_n] \to [xy]$ and $[y_n z_n] \to [yz]$ for some $x, y, z \in X$. If $[x_n y_n] * [y_n z_n]$ is a geodesic for each n, then $[xy] * [yz]$ is a geodesic.

Proof. We have $\angle xyz \in [0, \pi]$ from Lemma 2.1.18, $\angle([y_n x_n], [y_n z_n]) \equiv \pi$ from Lemma 2.1.26, and $\angle xyz \geqslant \pi$ by Theorem 2.1.28, so $[xy] * [yz]$ is a geodesic by Lemma 2.1.26. □

2.2 Total curvature

2.2.1 Total rotation and total curvature

Remark. A piecewise geodesic's *total rotation* is the discrete analogue of the integral curvature of a regular curve in a smooth manifold. It is defined (below) as the sum of the change in direction at each interior vertex. The idea that piecewise geodesics may be used to approximate the *total curvature* of a path in $\mathbb{R}^d$ that need not be differentiable, much less regular, appears in a 1950 paper by Milnor [26], and, as Alexandrov showed [4], generalizes nicely to paths in certain non-Euclidean spaces.

Definition 2.2.1. Let X be a metric space. Let $\sigma : [a, b] \to X$ be a piecewise geodesic path with breakpoints $\left(\sigma(t_k)\right)_{k=0}^{m}$, where $a = t_0 < t_1 < \cdots < t_m = b$. The **angle** of σ at its interior vertex $\sigma(t_k)$ is

$$\rho_k := \angle\left([\sigma(t_k)\sigma(t_{k-1})], [\sigma(t_k)\sigma(t_{k+1})]\right).$$

If the angle of σ exists at each of its interior vertices, the **total rotation** of σ is

$$\rho(\sigma) := \sum_{k=1}^{m-1} (\pi - \rho_k).$$

Lemma 2.2.2. The total rotation of any piecewise geodesic in a nonpositively curved space exists.

Proof. This follows immediately from Lemma 2.1.18. □

Definition 2.2.3 ([26])**.** The **total curvature** $\tau(\gamma)$ of a nondegenerate path $\gamma : [a, b] \to X$ in a metric space X is defined by

$$\tau(\gamma) := \sup\left\{\rho(\sigma) : \sigma \text{ is a piecewise geodesic path inscribed in } \gamma\right\}.$$

We set $\tau(\gamma) := 0$ if γ is degenerate.

2.2.2 Properties of total curvature

Lemma 2.2.4. If $\gamma : I \to X$ is a path in a length space X, and $L : I' \to I$ is a change of variable, then $\tau(\gamma L) = \tau(\gamma)$.

Proof. Write $\tau = \tau(\gamma)$ and $\tau' = \tau(\gamma L)$. We show $\tau' \leqslant \tau$, and omit the proof of the reverse inequality.

Let $\varepsilon > 0$. There exists a piecewise geodesic path $c : [a'', b''] \to X$ inscribed in $\gamma L : [a', b'] \to X$ such that $\tau' - \varepsilon < \rho(c)$. Then there exist partitions $(u_0)_{k=0}^m$ of $[a'', b'']$ and $(s_0)_{k=0}^m$ of $[a', b']$ such that $c(u_{k-1}) \neq c(u_k)$ for $k \in \{1, \ldots, m\}$, $c(u_k) = \gamma L(s_k)$ for $k \in \{0, \ldots, m\}$, and

$$\begin{aligned}\rho(c) &= \sum_{k=1}^{m-1} \pi - \angle\big([c(u_k)c(u_{k-1})], [c(u_k)c(u_{k+1})]\big) \\ &= \sum_{k=1}^{m-1} \pi - \angle\big([\gamma(t_k)\gamma(t_{k-1})], [\gamma(t_k)\gamma(t_{k+1})]\big),\end{aligned}$$

where $L(s_k) = t_k$.

Noting that

$$\gamma(t_{k-1}) = c(u_{k-1}) \neq c(u_k) = \gamma(t_k),$$

it suffices to show $(t_k)_{k=0}^m$ is a partition of $[a, b]$. (Then c is inscribed in γ, so $\tau \geqslant \rho(c) > \tau' - \varepsilon$, and letting $\varepsilon \to 0$ yields $\tau' \geqslant \tau$.) This is elementary: since L is nondecreasing and $\gamma(t_{i-1}) \neq \gamma(t_i)$, we have $i < j \Rightarrow s_i < s_j \Rightarrow t_i < t_j$, and since L is a continuous surjection, $t_0 = a$ and $t_m = b$. □

Theorem 2.2.5 ([23, 3])**.** Let X be a CAT(0) space. If a sequence of piecewise geodesic paths

$\sigma_n : [a, b] \to X$ converges uniformly to a path α, then

$$\tau(\alpha) \leqslant \liminf_{n\to\infty} \rho(\sigma_n).$$

Furthermore, if the σ_n are inscribed in α, and have mesh approaching 0, then

$$\tau(\alpha) = \lim_{n\to\infty} \rho(\sigma_n).$$

Corollary 2.2.6. Let α and β be paths whose concatenation $\alpha * \beta$ is defined. Then

$$\tau(\alpha * \beta) = \tau(\alpha) + \tau(\beta) + \pi - \angle(\overline{\alpha}, \beta).$$

Proof. The conclusion clearly holds when α and β are piecewise geodesics. Let $(\rho_n)_{n=1}^{\infty}$ and $(\sigma_n)_{n=1}^{\infty}$ be sequences of piecewise geodesics respectively inscribed in α and β, each with mesh approaching 0. Then $\tau(\overline{\rho_n}) \to \tau(\overline{\alpha}) = \tau(\alpha)$, $\tau(\sigma_n) \to \tau(\beta)$, and $\tau(\overline{\rho_n} * \sigma_n) \to \tau(\alpha * \beta)$ by Theorem 2.2.5, and $\angle(\overline{\rho_n}, \sigma_n) = \angle(\overline{\alpha}, \beta)$ by definition. Thus

$$\tau(\alpha*\beta) = \lim_{n\to\infty} \tau(\overline{\rho_n}*\sigma_n) = \lim_{n\to\infty} \big(\tau(\overline{\rho_n})+\tau(\sigma_n)+\pi-\angle(\overline{\rho_n}, \sigma_n)\big) = \tau(\alpha)+\tau(\beta)+\pi-\angle(\overline{\alpha}, \beta).$$

□

Corollary 2.2.7 (Lower semicontinuity of total curvature for paths). If a sequence of nondegenerate paths γ_n in a CAT(0) space X converges uniformly to a path γ, then

$$\tau(\gamma) \leqslant \liminf_{n\to\infty} \tau(\gamma_n).$$

Proof. By Lemma 2.1.23, for each n, there exists a piecewise geodesic path σ_n such that $d_\infty(\sigma_n, \gamma_n) \leqslant \varepsilon/n$. Since $\gamma_n \to \gamma$,

$$d_\infty(\sigma_n, \gamma) \leqslant \varepsilon/n + d_\infty(\gamma_n, \gamma) \xrightarrow{n\to\infty} 0,$$

so $\sigma_n \to \gamma$. By definition of total curvature, $\rho(\sigma_n) \leqslant \tau(\gamma_n)$ for each n. By Theorem 2.2.5,

$$\tau(\gamma) \leqslant \liminf_{n\to\infty} \rho(\sigma_n) \leqslant \liminf_{n\to\infty} \tau(\gamma_n).$$

□

Lemma 2.2.8. If $\gamma : [a, b] \to \mathbb{R}^d$ is $\mathcal{C}^1$ with Lipschitz derivative, then $\tau(\gamma) = \int |\gamma''|$.

Remark. Lemma 2.2.8 is well known for twice-differentiable γ ([4], §5.3). To see that it holds for a $\mathcal{C}^1$ path with Lipschitz derivative requires only an application of the following result.

Theorem 2.2.9 (Rademacher [9]**).** A Lipschitz map $\mathbb{R}^n \to \mathbb{R}^m$ is differentiable almost everywhere.

2.3 Cell complexes

2.3.1 CW complexes

Definition 2.3.1. Write $I = [0, 1]$. Let X be a topological space. A ***k*-dimensional closed cell**, or **closed *k*-cell in** X ($k \in \mathbb{N}_0$) is the image e of a continuous map $c : I^k \to X$, called its **characteristic map**, such that the restriction of c to $\text{int}(I^k)$ is a homeomorphism. The corresponding **open cell** $\mathring{e}$ is the image of int I^k under c.

Definition 2.3.2. Let X be a topological space. Let $\mathcal{E} = (e_\lambda)_{\lambda\in\Lambda}$ be an indexed collection of k-dimensional closed cells in X with characteristic maps $c_\lambda : I^{k(\lambda)} \to e_\lambda$. If

- every point in X lies in exactly one open cell $\mathring{e}_\lambda$, and
- each $c_\lambda(\partial I^{k(\lambda)})$ lies in finitely many open cells, each of dimension less than $k(\lambda)$,

then X is a **CW complex** with **cell structure** $\mathcal{E}$ and **attaching maps** $a_\lambda = c_\lambda|_{\partial I^k}$, and we topologize X by declaring

$$A \subset X \text{ closed} \iff A \cap e_\lambda \text{ is closed in } e_\lambda \text{ for each } \lambda.$$

The elements of X are the **cells of** X. If X is **finite-dimensional**, i.e. if

$$\dim(X) := \sup\{\text{dimension of cells in } \mathcal{E}\} < \infty,$$

the topology on X is the quotient topology.

Example 2.3.3. For each $(m, n) \in \mathbb{Z}^2$, define $c_{(m,n)} : I \times I \to \mathbb{R}^2$ by

$$c_{(m,n)}(s, t) = (m, n) + (s, t),$$

and let

$$\mathcal{E} = \{\text{Image } c_{(m,n)}\big|_F : (m, n) \in \mathbb{Z}^2 \text{ and } F \text{ is a face of } I \times I\}.$$

Then $\mathbb{R}^2$ is a CW complex with cell structure $\mathcal{E}$. We will denote this CW complex by $\mathbb{E}^2$.

Definition 2.3.4. Let X be a CW complex with cell structure $\mathcal{E}$. A subset Y of X is a **subcomplex** of X if Y is a CW complex with cell structure $\{e \in \mathcal{E} : e \subset Y\}$. The ***k*-skeleton** $X^{(k)}$ of X is the subcomplex of X consisting of all j-cells with $j \leqslant k$. The **underlying graph** of X is its 1-skeleton. A cell $e \in \mathcal{E}$ is a **vertex** of X if $\dim e = 0$, and an **edge** if $\dim e = 1$.

2.3.2 Simplicial complexes

Definition 2.3.5. Let $\mathcal{V} = \{v_0, \ldots, v_k\}$ be a set of $k+1$ points in $\mathbb{R}^d$ such that $v_1 - v_0, \ldots, v_n - v_0$ are linearly independent. Then the convex hull of $\mathcal{V}$, $\text{conv}(\mathcal{V})$, is called a ***k*-simplex** in $\mathbb{R}^d$ with vertex set $\mathcal{V}$. If $\mathcal{U}$ is a nonempty subset of $\mathcal{V}$, we call $\text{conv}(\mathcal{U})$ a ***j*-face** of K, where $j = \#\mathcal{U}$.

Definition 2.3.6. The union of a finite set S of simplices in $\mathbb{R}^d$ is a **simplicial complex** K if each face of a member of S is again a member of S, and if the intersection of any two pair of members of S is a face of each. The ***k*-skeleton** of a simplicial complex K is the union of all j-faces of simplices in K with $j \leqslant k$.

2.3.3 Cube complexes

Our definition of a cube complex deviates from that of some other authors insofar as we require that each closed cell must be embedded (condition (1) in Definition 2.3.7), that the link of each vertex must be simplicial (condition (3) in Definition 2.3.7), and that the dimension of a cube complex must be finite. None of the following CW complexes satisfies our definition of a cube complex, as each violates the condition indicated in parentheses:

- a circle made up of a single edge with its endpoints glued together (1);
- a torus constructed from a single square cell by identifying each pair of opposite edges (1);
- a pair of squares glued together along adjacent edges (3);
- a pair of 3-cubes glued together along adjacent faces (3).

On the other hand, our definition does allow a pair of cubes to meet along two or more (nonadjacent) faces. For example, a cylinder obtained by gluing two opposite sides of a square to two opposite sides of a second square, respectively, qualifies as a cube complex under Definition 2.3.7.

Our reason for incorporating conditions (1) and (3) into our definition of a cube complex is that we are primarily interested in cube complexes that arise as the state complexes of reconfigurable systems, and all state complexes satisfy these two conditions (see [32], Proposition 3.1.2 and the proof of Proposition 3.2.11). We will not have occasion to discuss other species of CW complexes built from Euclidean cubes, which include (*combinato-*

rial) *cubical complexes*, *cubings*, *cubed complexes*, and *cubulations*, but we refer the interested reader to the concise summary given in [32] of the menagerie hinted at here, noting only that the objects described by Definition 2.3.7 would be described as *simple cube complexes* in that summary.

Definition 2.3.7. We define a **cube complex** to be a connected finite-dimensional CW complex X such that

(1) the characteristic map of each closed cell is an isometric embedding,

(2) the restriction of each attaching map $I^k \to X$ to each $(k-1)$-face of I^k is an isometry onto I^{k-1} postcomposed with the characteristic map of some closed $(k-1)$-cell of X, and

(3) the link of each vertex is a simplicial complex.

Definition 2.3.8. A **square complex** is a cube complex that is the union of its 2-cells.

Definition 2.3.9. We will always regard a cube complex X as being assembled from a collection of disjoint cubes C_λ in $\mathbb{R}^d$ for some $d \in \mathbb{N}$. The map

$$p : \coprod C_\lambda \to X$$

whose restriction to each C_λ is an isometry $C_\lambda \to I^k$ postcomposed with the characteristic map $I^k \to X$ is a **defining projection** for the cube complex X.

Definition 2.3.10. Let X be a cube complex. The **carrier** of a subset A of X is the smallest closed subcomplex of X containing A, and will be denoted by $\mathrm{Carrier}(A)$. We say X is **locally finite** if each point of X has a neighborhood that meets only finitely many cells of X.

Definition 2.3.11. The **link** of a vertex v of a cube complex X, which we denote by $\mathrm{link}(v)$, is the abstract simplicial complex whose k-cells E^k are the $(k+1)$-cells in X incident with v, with j-dimensional faces identified iff the corresponding $(j+1)$-dimensional faces are

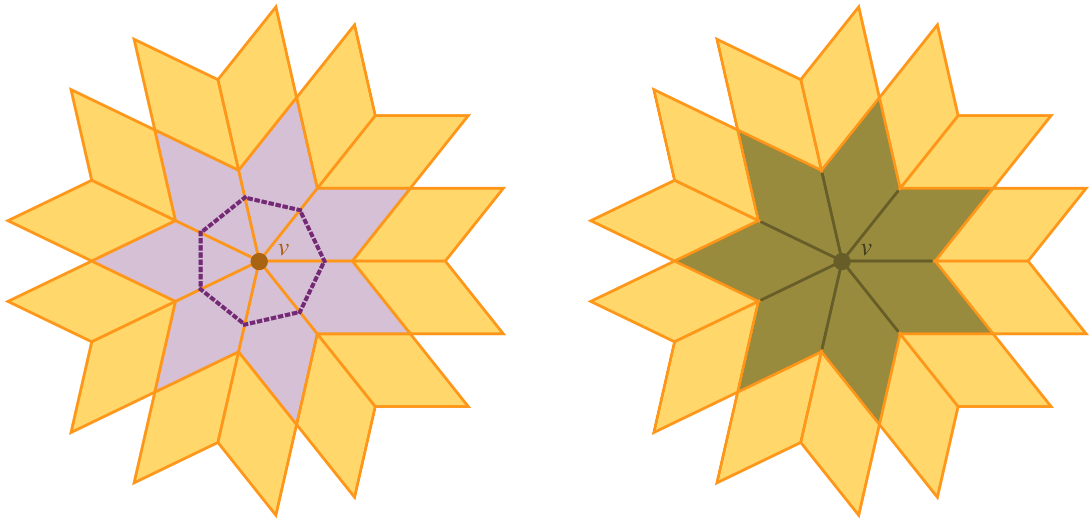

Figure 2.1: LEFT: The link of a vertex v in a square complex (dotted purple lines). RIGHT: The open star of v (green).

identified in X (Figure 2.1).

Definition 2.3.12. The **open star** of a vertex v of a cube complex X, written $\text{star}(v)$, is the union of the interiors of cells incident with v (Figure 2.1). The **closed star** of v, written $\overline{\text{star}}(v)$, is the subcomplex of X whose cells are incident with v.

Remark. The link of a vertex v in a cube complex can equivalently be defined as the spherical simplicial complex whose underlying space is a small metric sphere (boundary of a ball) about v. See [8, 14] for details.

Definition 2.3.13. A simplicial complex K is **flag** if for every 1-skeleton of a k-simplex in K, the corresponding k-simplex appears in K.

Definition 2.3.14. A cube complex X is **nonpositively curved** if the link of each vertex is flag.

Definition 2.3.15. A cube complex X is **CAT(0)** if X is nonpositively curved and simply connected.

Theorem 2.3.16 (Gromov: [18], p. 122). Let X be a cube complex. If the link of each vertex in X is flag, then X is CAT(0).

2.3.4 The cube complex associated with a Coxeter system

The definitions and much of the notation in this subsection are taken from [14].

Definition 2.3.17. Let $S = \{s_1, \ldots, s_n\}$ be a finite set. Let $M = [m_{ij}]_{i,j=1}^n$ be a symmetric matrix with entries in $\mathbb{N} \cup \{\infty\}$ such that $m_{ii} = 1$ and $m_{ij} \geqslant 2$ if $i \neq j$. Set

$$W = \langle S \mid (s_i s_j)^{m_{ij}} = 1 : m_{ij} < \infty \rangle.$$

The pair (W, S) is a **Coxeter system** (of **finite rank**) with matrix M, and W is a **Coxeter group**. If each m_{ij} is either 2 or ∞, we say (W, S) is **right-angled**.

Definition 2.3.18. The **Cayley graph** Cayley(W, S) of the Coxeter system (W, S) is defined by taking Vert$(\Gamma) = W$, and attaching for each $(w, s) \in W \times S$ an undirected edge

$$E = E(w, s)$$

with $\partial E = \{w, ws\}$. We say such an edge is **labeled** by

$$S(E) := s.$$

Definition 2.3.19. The **flag complex** Flag$(\mathcal{P})$ of a poset P is the abstract simplicial complex whose simplices are the finite chains in P. The **geometric realization** $|P|$ of a poset is the geometric realization of its flag complex Flag(P).

Example 2.3.20 ([14], Example A.4.6)**.** Write bX for the barycentric subdivision of a simplicial complex X. If K is an abstract simplicial complex, and $X = $ Geom(K) is its geometric realization, then the geometric realization of K is

$$|K| = \mathrm{Flag}(K) = \text{cone on } bX = \text{join of } bX \text{ and a point corresponding to } \varnothing.$$

If P is a poset of cells in a convex cell complex Λ, then $\mathrm{Flag}(P)$ is the barycentric subdivision of Λ.

Definition 2.3.21. Let (W, S) be a right-angled Coxeter system. A **special subgroup** W_T of W is the subgroup generated by $T \subseteq S$. As a special case, we take $W_\varnothing$ to be the trivial group $\{e\}$. A subset T of S is **spherical** if W_T is finite. A **spherical coset** of (W, S) is a coset of a special subgroup W_T for spherical $T \subseteq S$. Write

$$\mathcal{S} = \mathcal{S}(W, S) = \{T \subseteq S \text{ spherical}\},$$

$$\mathcal{S}_{>\varnothing}(W, S) = \mathcal{S}(W, S) \setminus \{\varnothing\}.$$

The **nerve** $L(W, S)$ of (W, S) is the abstract simplicial complex $(\mathcal{S}_{>\varnothing}(W, S), \subseteq)$. The geometric realization of the poset $(W\mathcal{S}, \subseteq)$, where

$$W\mathcal{S} = \bigcup_{T \in \mathcal{S}(W,S)} W/W_T,$$

the poset of special cosets of (W, S) under subset inclusion, is a simplicial complex called the **Davis complex**

$$\Sigma(W, S) = |W\mathcal{S}|$$

of (W, S). A cell structure $\mathcal{E}_\square$ on $\Sigma(W, S)$ that makes $\Sigma(W, S)$ a cube complex is obtained by taking as cells the geometric realizations of all posets of the form

$$(W\mathcal{S})_{\leqslant wW_T} = \{C \in W\mathcal{S} : C \subseteq wW_T\}$$

for $w \in W$ and $T \in \mathcal{S}(W, S)$ ([14], p. 233). The Davis complex with the cubical cell structure

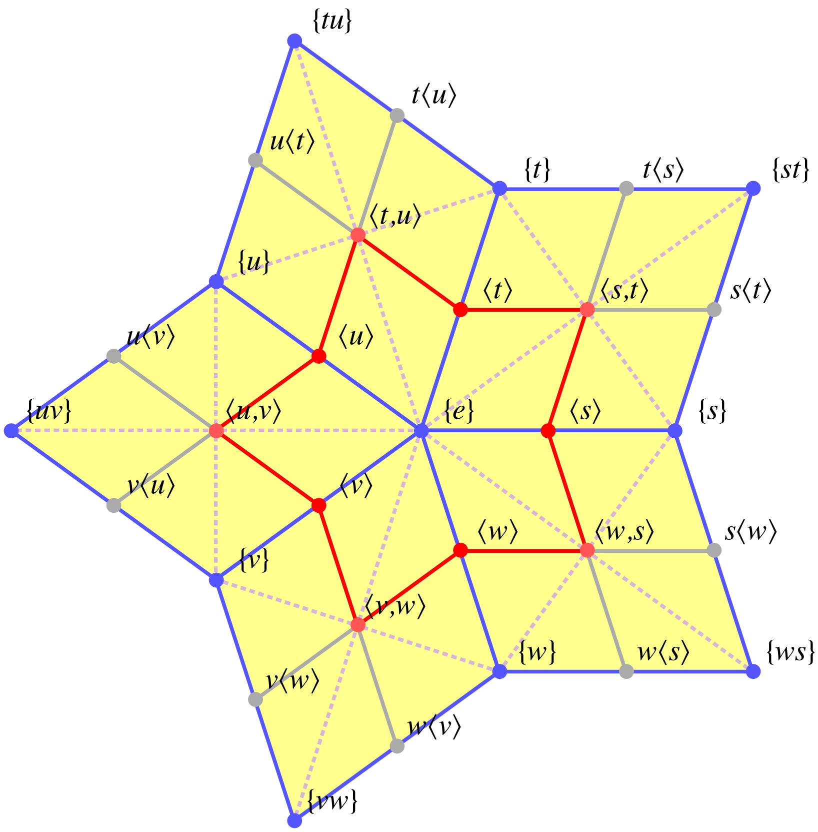

Figure 2.2: The carrier (yellow and blue) in the square complex $X(W,S)$ of the nerve $L(W_5, S_5)$ (red), where W_5 is the right-angled Coxeter group with generating set $S_5 = \{s, t, u, v, w\}$ and nontrivial relators $(st)^2$, $(tu)^2$, $(uv)^2$, $(vw)^2$, $(ws)^2$. The dotted lines indicate the simplicial structure of $\Sigma(W_5, S_5)$.

$\mathcal{E}_\square$ and the piecewise Euclidean metric will be denoted by $X(W, S)$. (See Figure 2.2.)

Remark. The inclusion $\mathcal{S}_{>\varnothing}(W, S) \hookrightarrow W\mathcal{S}$ defined by $T \mapsto W_T$ induces a simplicial embedding of the nerve $L(W, S)$ into $\Sigma(W, S)$.

The cardinality of $T \in \mathcal{S}(W, S)$ is the dimension of the cell $(W\mathcal{S})_{\leqslant wW_T}$ in $\mathcal{E}_\square$. For example, when $T = \varnothing$, so that wW_T is the coset of $W_\varnothing = \{e\}$ consisting of the single element $w \in W$, we have a vertex. When $T = \{s\}$, we get a coset of the form $w\langle s\rangle$, which contains the cosets $\{w\}$ and $\{ws\}$ of $W_\varnothing = \{e\}$, and so on. In Figure 2.2, the 2-cells of $\mathcal{E}_\square$ are shown in yellow. The nerve L is a barycentrically subdivided pentagon, shown in red; it is isomorphic to the link of each vertex in $\mathcal{E}_\square$.

We note that W acts geometrically (properly, cocompactly, and by isometries) on $X(W, S)$, although we will not have occasion to use this fact. The fundamental domain is the coned-off nerve, $K = e * L$. As Davis puts it, 'the elements of S act as "reflections" across the "mirrors" of K' [14]. In our illustration, the mirrors of K are the unsubdivided edges of the pentagon L with vertices $\langle s, t\rangle$, $\langle t, u\rangle$, $\langle u, v\rangle$, $\langle v, w\rangle$, and $\langle w, s\rangle$. For example, $(W\mathcal{S})_{\leqslant\langle s\rangle}$ is reflected in the mirror with endpoints $\langle s, t\rangle$ and $\langle w, s\rangle$ via the isometry of $X(W, S)$ induced by left multiplication by s.

Lemma 2.3.22. Let (W, S) be a right-angled Coxeter system. If every $s \in S$ commutes with exactly one $t \in S \setminus \{s\}$, then $X(W, S)$ is a square complex.

Proof. For a right-angled Coxeter system (W, S), all nontrivial relators are of the form $(st)^2$, so the cardinality of T for each special subgroup W_T is at most 2. □

Remark. The product $(st)^2$ is a relator for a right-angled Coxeter system (W, S) if and only if every vertex in $X(W, S)$ is incident with a square with edges alternately labeled by distinct $s, t \in S$. The opposite edges E and E' of such a square have the same labels, $S(E) = S(E')$.

Theorem 2.3.23 ([14], Theorem 12.2.1)**.** Let (W, S) be a Coxeter system. If (W, S) is right-angled, then $X(W, S)$ is CAT(0).

Remark. As we are concerned only with cube complexes, we will not need Moussong's well-known generalization ([14], Theorem 12.2.3) of Theorem 2.3.23 to arbitrary Coxeter systems.

2.4 Hyperplanes

2.4.1 Basic definitions

Throughout this section, (W, S) denotes a right-angled Coxeter system.

Definition 2.4.1. For two edges E and E' in a square complex X, write $E \parallel E'$ iff

$$\text{Vert}(Q) = \partial E \cup \partial E'$$

for some 2-cell Q of X, and let $\sim$ be the transitive closure of the relation $\parallel$. That is, $E \sim E'$ iff there exist edges $E_1, \ldots, E_n$ ($n \geqslant 2$) such that

$$E = E_1 \parallel E_2 \parallel \cdots \parallel E_n = E'.$$

Then $\sim$ is an equivalence relation on

$$\mathcal{F} := \{\text{1-dimensional faces of 2-cells}\}.$$

Since a square complex is by definition the union of its 2-cells, $\mathcal{F} = \text{Edges}(X)$. An equivalence class $[E] \in \mathcal{F}/\sim$ is a **parallelism class** of edges in X.

Definition 2.4.2 ([17])**.** Each k-cell Q in a cube complex X ($k \in \mathbb{N}$) is isometric to the product space I^k by definition. Such an isometry $\varphi_Q : I^k \to Q$ is a **coordinate map** on Q. Given $k \in \mathbb{N}$, a k-cell E of X, and $j \in \{1, \ldots, k\}$, the set

$$\{\varphi(t_1, \ldots, t_k) : t_j = \tfrac{1}{2} \text{ and } t_i \in I \text{ for } i \neq j\}$$

is a **midplane** of Q. A midplane m is **dual** to an edge E if $m \cap E \neq \varnothing$.

Definition 2.4.3. Two midplanes m and m' respectively, are **midplane-equivalent**, and we write $m \simeq m'$, if $m \cap m'$ is a midplane. Let $\cong$ be the equivalence relation obtained by taking the transitive closure of the relation $\simeq$ on midplanes of cubes in X. The equivalence classes of $\cong$ are **hyperplanes** of X. The union of all midplanes in such an equivalence class is also called a **hyperplane** of X. If H is a hyperplane (in the latter sense) and m is a midplane contained in H which is dual to an edge E, then H is **dual** to E; we also say in this case that H is **dual** to the edge's label $S(E)$. If a hyperplane H contains a midplane m' that is perpendicular in some cube of X to a midplane m, we say that H is **dual** to m.

Remark. The k-dimensional hyperplanes ($k \in \mathbb{N}$) of a square complex are in one-to-one correspondence with the parallelism classes $[E] \in \mathcal{F}/\sim$. This fact follows from the observation that any 1-cube m that is a midplane meets opposite edges E and E' of the unique 2-cell Q that contains it, so

$$m \text{ midplane of } X \quad \Longleftrightarrow \quad \partial E \cup \partial E' = \text{Vert}(Q).$$

We will alternately regard a hyperplane as a collection of midplanes, the (connected) union of such midplanes, or the parallelism class of dual edges, whichever is most convenient. For instance, identifying a hyperplane with the parallelism class of edges $[E]$ makes it plain that each hyperplane H is uniquely identified by a choice of edge,

$$H = H(E) = [E].$$

If X is the square complex associated with a Coxeter system (W, S), each edge of X is determined by a group element and a generator, so each $(w, s) \in W \times S$ determines a hyperplane,

$$H = H(w, s) := H\big(E(w, s)\big).$$

Definition 2.4.4. Two hyperplanes H and H' (not necessarily distinct) are **orthogonal**, and we write $H \perp H'$, if H is dual to edges E_1 and E_2, H' is dual to edges E_1' and E_2' such that for some 2-cell Q,

$$\partial E_1 \cup \partial E_2 = \text{Vert}(Q) = \partial E_1' \cup \partial E_2'.$$

Remark. If two midplanes m and m' meet in the interior of some 2-cell Q, their respective midplane-equivalence classes H and H' are obviously orthogonal.

We have defined a square complex so that all its k-cells are isometric to I^k. By examining cases, we see that for any two hyperplanes H and H' in a square complex, either $d(H, H') \geqslant 1$ or $d(H, H') = 0$, and if the latter is true, then H and H' must meet in the interior of some k-cell ($k \in \mathbb{N}$). If H and H' meet in the interior of a 1-cell E, then $H = [E] = H'$. If H and H' meet in the interior of a 2-cell, then $H \perp H'$. This proves

Lemma 2.4.5. Let H and H' be distinct hyperplanes of a square complex that are not orthogonal. Then $d(H, H') \geqslant 1$.

Definition 2.4.6. A hyperplane H is **two-sided** if it disconnects its carrier, i.e. if Carrier$(H) \smallsetminus H$ is disconnected.

Remark. Suppose H is a two-sided hyperplane, and E is an edge dual to H. All the edges dual to H can be consistently oriented by choosing an orientation for E and declaring all edges in the parallelism class $[E]$ to be oriented in the same direction [19].

Arguments establishing the following well-known elementary properties of hyperplanes in CAT(0) cube complexes can be found in [38].

Lemma 2.4.7 ([38])**.** Let H be a hyperplane in a CAT(0) cube complex X. Then:

(1) H is two-sided.

(2) $X \smallsetminus H$ has exactly two connected components.

(3) H is not orthogonal to itself.

(4) H is simply connected.

Definition 2.4.8. If H is a hyperplane such that $X \smallsetminus H$ has two connected components $H^{\pm}$, the sets $H^{\pm}$ are called **halfspaces**.

Remark. In general, a two-sided hyperplane in a square complex X need not separate X. For example, identifying the opposite edges of $[0, 2] \times [0, 2] \subset \mathbb{E}^2$ yields a torus containing four hyperplanes of two midplanes each, and cutting the torus along any one of these hyperplanes disconnects its carrier, but not the space.

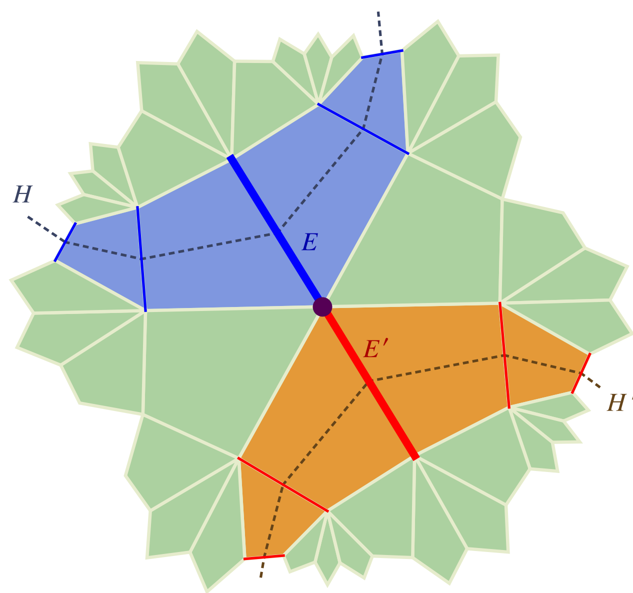

Figure 2.3: Osculating hyperplanes.

2.4.2 Osculating hyperplanes

Definition 2.4.9 ([19]). Let H and H' be hyperplanes (not necessarily distinct). We say H and H' **osculate**, and we write $H \)\(H'$, if there exist adjacent edges E and E' such that H is dual to E, H' is dual to E', and E and E' are not both contained in any 2-cell. (See Figure 2.3.)

Counterexample. Give $X = \partial(I^2)$ a cell structure by taking each face of the cube I^3 as a 2-cell. Then there are no osculating hyperplanes in X.

Definition 2.4.10. For a subset A and an element x of a group G, write

$$C_A(x) = \{a \in A : ax = xa\}.$$

Remark. Let $E = E(w, s)$ and $E(w, s')$ be adjacent edges of $X(W, S)$ with common vertex w, where $s, s' \in S$. Let $H = H(w, s)$ be the hyperplane dual to E, and let $H' = H(w, s')$ be the hyperplane dual to E'. Then H and H' osculate iff $s \notin C_S(s')$.

Lemma 2.4.11. The distance between osculating hyperplanes in $X = X(W, S)$ is 1.

Proof. Let $w \in \text{Vert}(X)$, and let $E = E(w, s)$ and $E' = E(w, s')$ be adjacent edges meeting at w, where $s, s' \in S$. Then the distance between $H = [E]$ and $H' = [E']$ is no more than 1. Suppose H and H' osculate. Then $s \notin C_S(s')$. In particular, $s \neq s'$, and this implies

$$H = [E(w, s)] \neq [E(w, s')] = H'.$$

In principle, it is possible that two osculating hyperplanes are orthogonal: see Figure 2.4 (bottom right). But since s and s' do not commute, there is no square in $X(W, S)$ with boundary edges cyclically labeled s, s', s, s'. Therefore, H and H' are not orthogonal. Since H and H' are distinct and not orthogonal, $d(H, H') \geqslant 1$ by Lemma 2.4.5. □

Lemma 2.4.12. Let H and H' be hyperplanes in a square complex X. Suppose H and H' are disjoint. Then Carrier(H') lies in a component of $X \setminus H$.

Proof. If Carrier(H') meets H, then $H = H'$ or $H \perp H'$; in either case, H and H' are not disjoint. Thus Carrier$(H') \subset X \setminus H$. Since H' is connected, so is its cubical neighborhood Carrier(H'). The conclusion now follows. □

Lemma 2.4.13. Let H and H' be hyperplanes in $X = X(W, S)$. If H and H' are osculating hyperplanes, then Carrier(H') lies in a halfspace H^- of X such that $\partial H^- = H$.

Proof. Follows from Lemmas 2.4.7, 2.4.11, and 2.4.12. □

2.4.3 Special cube complexes

The definitions in this subsection are taken from [19]. Illustrations of the hyperplane configurations discussed in Definitions 2.4.14, 2.4.15, and 2.4.16 can be found in Figure 2.4.

Definition 2.4.14. Let H be a two-sided hyperplane; then there are two possible ways that all edges dual to H may be consistently oriented. We say H **self-osculates** if for one of

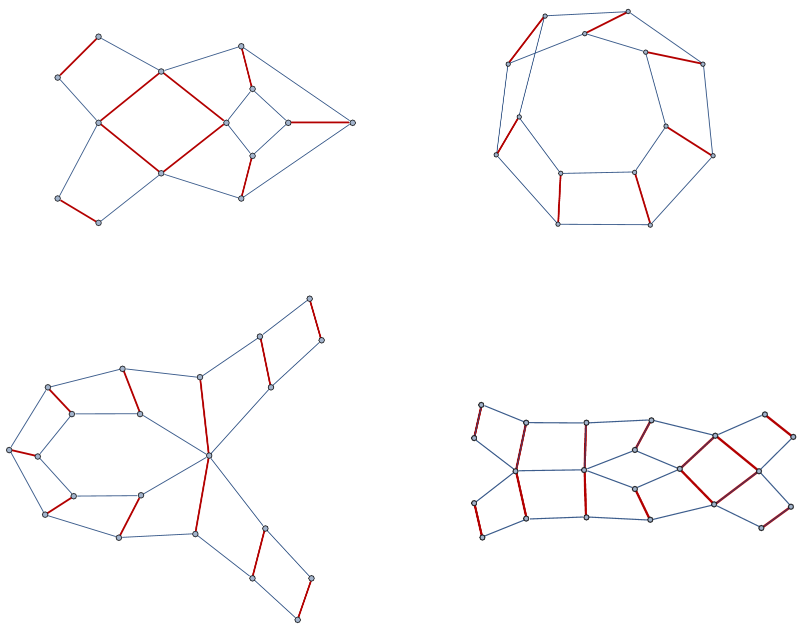

Figure 2.4: Hyperplane configurations that are forbidden in a special cube complex. CLOCKWISE FROM TOP LEFT: a self-orthogonal hyperplane, a hyperplane that is not two-sided, a self-osculating hyperplane, and a pair of interosculating hyperplanes.

these choices of orientation, the carrier of H contains a vertex which is the initial vertex of two distinct edges dual to H.

Definition 2.4.15. A pair of distinct hyperplanes **interosculate** if they osculate and are orthogonal.

Definition 2.4.16. A cube complex X is **special** if

(1) no hyperplane of X is orthogonal to itself,

(2) every hyperplane of X is two-sided,

(3) no hyperplane of X self-osculates, and

(4) no two hyperplanes of X interosculate.

The pathological cases which cannot occur in a special cube complex are illustrated in Figure 2.4.

Theorem 2.4.17 ([20])**.** A CAT(0) cube complex is special.

Theorem 2.4.18. The cube complex associated with a right-angled Coxeter system is special.

Proof. Immediate from Theorems 2.3.23 and 2.4.17. □

2.5 Some constructions for square complexes

2.5.1 Stacks

In Chapter 5, we will classify the topological rays $\gamma : [t_1, \infty) \to X$ of constant curvature in a certain family of CAT(0) square complexes X. (Recall that we define a topological ray simply to be a continuous map of a half-line into a space.) One possibility is that γ is **proper**, that is, such that $\gamma^{-1}(K)$ is compact for every compact $K \subset X$. In this subsection, we give a simple sufficient condition for γ to be proper. This condition is stated in terms of a sequence of successively osculating hyperplanes

$$H_1 \mathrel{)\!(} H_2 \mathrel{)\!(} H_3 \mathrel{)\!(} \cdots$$

which bound a sequence of nested halfspaces. When X is CAT(0), the additional properties we will require of the sequence $(H_k)_{k=1}^{\infty}$ follow from the fact that its successive terms osculate.

Definition 2.5.1. A **stack** in a cube complex X is a sequence $(H_k)_{k=1}^{\infty}$ of successively osculating hyperplanes, $H_k \mathrel{)\!(} H_{k+1}$. We say a stack $(H_k)_{k=1}^{\infty}$ is **oriented** if $X \smallsetminus H_k$ has two connected components for each k, and if the components $H_k^{\pm}$ of $X \smallsetminus H_k$ are labeled so that $H_k^- \subset H_{k+1}^-$ and $H_k^+ \supset H_{k+1}^+$ for each k.

Lemma 2.5.2. Every stack in a CAT(0) cube complex X can be oriented.

Proof. It suffices to show that if H_1 and H_2 are osculating hyperplanes, then the respective components of $X \setminus H_1$ and $X \setminus H_2$ can be labeled as $H_1^{\pm}$ and $H_2^{\pm}$ so that $H_1^- \subset H_2^-$ and $H_1^+ \supset H_2^+$.

By Lemma 2.4.7, each hyperplane H of X partitions $X \setminus H$ into two connected components. Since X is CAT(0), hence special (Lemma 2.4.17), osculating hyperplanes of X cannot be orthogonal (Definitions 2.4.16 (4) and 2.4.15). In particular, osculating hyperplanes cannot intersect (see the Remark following Definition 2.4.4).

Let H_1 and H_2 be osculating hyperplanes, and let H_1', H_1'' (H_2', H_2'') be the two distinct components of $X \setminus H_1$ ($X \setminus H_2$, respectively). Since H_1 is connected and does not meet H_2, it must lie in one of the two connected components of $X \setminus H_2$. Assume $H_1 \subset H_2''$ without loss of generality. Since $H_2'' \subset X \setminus H_2'$, we have $H_2' \subset X \setminus H_1$. If there exist $x \in H_2' \cap H_1'$ and $y \in H_2' \cap H_1''$, then x and y can be joined by a path in $H_2' \subset X \setminus H_1$. But the two distinct components H_1' and H_1'' of $X \setminus H_1$ cannot be joined by a path in $X \setminus H_1$: contradiction. Thus $H_2' \subset (X \setminus H_1') \cap (X \setminus H_1) = H_1''$ or $H_2' \subset (X \setminus H_1'') \cap (X \setminus H_1) = H_1'$. $\square$

It is easy to see that the hyperplanes $H_1 \,)(\, H_2 \,)(\, H_3 \,)(\, \cdots$ in an oriented stack "approach infinity," in the sense that $d(H_1, H_n) \geqslant n - 1$ for each n (use Lemma 2.4.5 and connectedness). In a similar sense, a ray γ approaches infinity if there exists a stack $(H_k)_{k=1}^{\infty}$ of hyperplanes which chop γ up into initial segments $\gamma\big([t_1, t_k)\big) \subset H_k^-$ and tails $\gamma\big((t_k, \infty)\big) \subset H_k^+$ for some choice of parameter values $t_1 < t_2 < t_3 < \cdots$. To show that γ is a proper ray in this situation requires only minimal additional work, as we will now show.

Definition 2.5.3. A ray $\gamma : [t_1, \infty) \to X$ is **properly segmented** by an oriented stack $(H_k)_{k=1}^{\infty}$ if for any $k \in \mathbb{N}$, there exists a $t_k > t_1$ such that $\gamma\big([t_1, t_k)\big) \subset H_k^-$ and $\gamma\big((t_k, \infty)\big) \subset H_k^+$. (See Figure 2.5.)

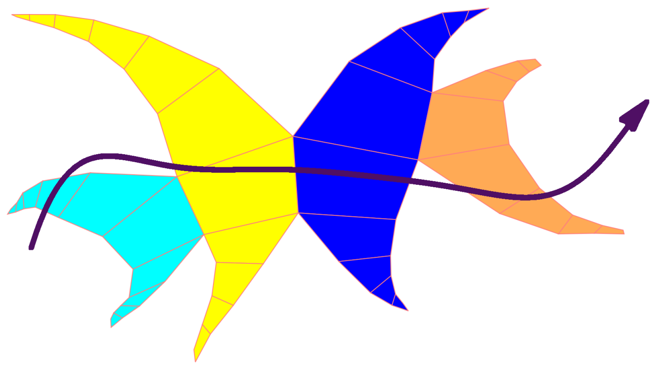

Figure 2.5: A ray that is properly segmented by a stack of hyperplanes in a square complex.

Lemma 2.5.4. If a ray γ in a cube complex X is properly segmented by an oriented stack of hyperplanes, then γ is proper.

Proof. Let $K \subset X$ be compact. Then $\gamma^{-1}(K)$ is closed, and $K \subset \bigcup_{k=1}^{N} H_k^- = H_N^-$ for some N. Since

$$\gamma_N^+ = \gamma\big((t_N, \infty)\big) \subset H_N^+ \subset X \setminus H_N^- \subset X \setminus K,$$

we have $\gamma^{-1}(K) \subset [t_1, t_N]$. □

2.5.2 Cube paths

As we have just seen, the hyperplanes of a cube complex X can be used to provide a simple, intuitive description of the behavior of a ray in X. On the other hand, the determination of the precise intersection of a given ray with a given hyperplane would require tedious calculations based on the ray's local coordinates in each cube. Indeed, many of the figures that appear in the present work were produced by Mathematica code which carries out just such calculations. But it is more natural and, in certain special cases, far less computationally intensive to determine the interaction of hyperplanes with the sequence of cells through which the ray passes, rather with than the ray itself. We therefore focus on *square paths* rather than rays in the sequel, noting that—for example—it is trivial to show that a ray γ is proper given sufficient information about the behavior of a square path for

γ as it passes through each hyperplane in a given stack.

Definition 2.5.5. An increasing sequence $(t_k)_{k=1}^m$ in $\mathbb{R}$ is a **subdivision** of $[t_1, t_m]$. An unbounded increasing sequence $(t_k)_{k=1}^\infty$ is a **subdivision** of $[t_1, \infty)$. The t_k are the **breakpoints** of the subdivision.

Definition 2.5.6. Let $\gamma : [a, b] \to X$ be a path in X, and let $a = t_1 < t_2 < \cdots < t_m = b$ be a subdivision of $[a, b]$. A sequence $(Q_k)_{k=1}^m$ of cells in X is a **cube path for γ with breakpoints** $(t_k)_{k=1}^m$ if $\gamma([t_k, t_{k+1}]) \subset Q_k$, $Q_k \not\subset Q_{k+1}$, and $Q_k \not\supset Q_{k+1}$ for all suitable k. A cube path for a curve $\gamma : [a, \infty) \to X$ is defined analogously, taking $m = \infty$. We write

$$\mathcal{Q}\big(\gamma, (t_k)_{k=1}^m\big) = \big\{\text{cube paths for } \gamma \text{ with breakpoints } (t_k)_{k=1}^m\big\}.$$

If each Q_k has dimension 2, $\mathcal{Q}$ is a **square path**. For an interval I in $\mathbb{N}$,

$$\mathcal{Q}I := \bigcup \{Q_k : k \in I\}.$$

A cube path $\mathcal{Q}$ is **edgewise** if $Q_k \cap Q_{k+1} \in \mathrm{Edges}(X)$ for all k. A square path $\mathcal{Q}$ is **locally monotone edgewise** if $\mathcal{Q}$ is edgewise and $Q_k \cap Q_{k+1} \neq Q_{k+1} \cap Q_{k+2}$ for all k.

2.5.3 Properly segmented square paths

We now modify Definition 2.5.3 by considering a square path for γ instead of γ itself. The new definition of a properly segmented ray can still be used to show that γ is proper, but does not require the calculation of the precise intersection of γ with the various hyperplanes involved.

Definition 2.5.7. Let $\mathcal{Q} := (Q_n)_{n=1}^\infty$ be an edgewise square path. $\mathcal{Q}$ is **properly segmented** by an oriented stack $(H_k)_{k=1}^\infty$ if there exist integers $1 \leqslant a_1 \leqslant b_1 < a_2 \leqslant b_2 < \cdots$ such that for each $k \in \mathbb{N}$,

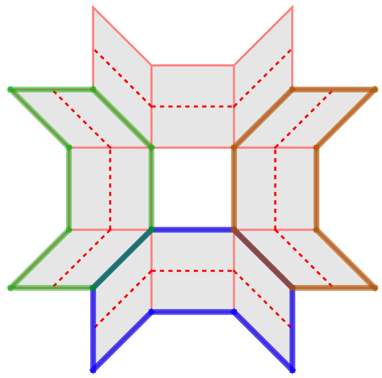

Condition (3) of Definition 2.5.7 violated.

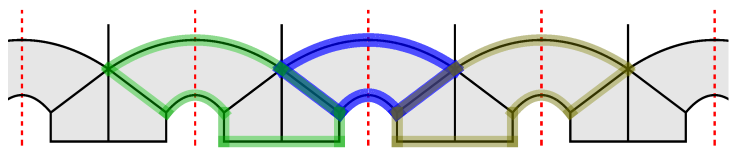

Condition (3) of Definition 2.5.7 satisfied.

Figure 2.6: A square path that is properly segmented with respect to a stack passes through the carrier of each hyperplane in the stack.

(1) $\mathcal{Q}[a_k, b_k] \subset \text{Carrier}(H_k)$,

(2) $\mathcal{Q}(b_k, a_{k+1}) \subset X \setminus (H_k \cup H_{k+1})$, and

(3) $\mathcal{Q}[a_{k-1}, a_k)$ and $\mathcal{Q}(b_k, b_{k+1}]$ meet distinct components of $\mathcal{Q}[a_k, b_k] \setminus H_k$ (see Figure 2.6).

Lemma 2.5.8. Let X be a square complex. If $\mathcal{Q} = (Q_n)_{n=1}^{\infty}$ is a square path for a ray $\gamma : [t_1, \infty) \to X$, and $\mathcal{Q}$ is properly segmented by an oriented stack $(H_k)_{k=1}^{\infty}$, then γ is proper.

Proof. Let $K \subset X$ be compact; then $\gamma^{-1}(K)$ is closed. Let $(a_k)_{k=1}^{\infty}$ and $(b_k)_{k=1}^{\infty}$ be sequences of integers as in the previous definition. By compactness,

$$K \cap \left(\bigcup \mathcal{Q}\right) \subset \mathcal{Q}[a_1, a_N]$$

for some $N \in \mathbb{N}$. Since

$$\begin{aligned}\gamma\big((t_{b_{N+1}}, \infty)\big) \subset \mathcal{Q}[b_{N+1}, \infty) \subset H^+_{N+1} \subset X \setminus H^-_{N+1} &\subset X \setminus \mathcal{Q}[a_1, a_{N+1}) \\ &\subset X \setminus \mathcal{Q}[a_1, a_N] \\ &\subset X \setminus \Big(K \cap \big(\bigcup \mathcal{Q}\big)\Big),\end{aligned}$$

we have $\gamma^{-1}(K) \subset [a, t_{b_{N+1}}]$. □

2.5.4 Stable cube paths

The material in this subsection will be used in Chapter 3, in which we will extract a convergent subsequence β_{n_k} from a given sequence of paths β_n using the Arzelà-Ascoli Theorem. Here, we show how to construct a common cube path for all the β_{n_k}.

Definition 2.5.9. Let $(\gamma^n)_{n=1}^{\infty}$ be a sequence of curves in X with respective parameter domains $J^{(n)}$. Let $\big(t_k^{(n)}\big)$ be a subdivision of $J^{(n)}$ for each n. A sequence (Q_k) of cells in X is a **stable cube path for $(\gamma^n)_{n=1}^{\infty}$ with breakpoints $\big(t_k^{(n)}\big)$** if

$$(Q_k) \in \bigcap_{n=1}^{\infty} \mathcal{Q}\big(\gamma^n, \big(t_k^{(n)}\big)\big).$$

Lemma 2.5.10 (Construction of a stable cube path). For each $n \in \mathbb{N}$, let $\gamma^n : [0, b_n] \to X$ be a path, let

$$\mathcal{T}^n = \big(t_k^{(n)}\big)_{k=1}^{m_n}$$

be a subdivision of $[0, b_n]$, and let

$$(Q_j^n)_{j=1}^{m_n - 1} \in \mathcal{Q}(\gamma^n, \mathcal{T}^n).$$

If $M := \sup_n m_n < \infty$ and X is locally finite, then there exist integers $m \leqslant M$, integers

$n_1 < n_2 < n_3 < \cdots$, and a stable cube path

$$(Q_j)_{j=1}^{m-1} \in \bigcap_{i=1}^{\infty} \mathcal{Q}(\gamma^{n_i}, \mathcal{T}^{n_i}) \tag{2.2}$$

for the subsequence $(\gamma^{n_i})_{i=1}^{\infty}$ with breakpoints $(t_k^{(n_i)})_{k=1}^{m_{n_i}}$.

Proof. We construct $(n_i)_{i=1}^{\infty}$ satisfying (2.2) by an induction of finitely many steps. Let $\#E(k, n)$ be the number of paths γ^p whose cube paths agree with the cube path of γ^n in the kth slot,

$$\#E(k, n) = \#\{p \in \mathbb{N} : Q_k^p = Q_k^n\}.$$

Since X is locally finite, $\#E(1, n_1) = \infty$ for some $n_1 \in \mathbb{N}$. Then there exist integers $n_1 < n_2 < n_3 < \cdots$ such that

$$(Q_k^{n_i})_{k=1}^{m_{n_i}-1} \in \mathcal{Q}(\gamma^{n_i}, \mathcal{T}^{n_i})$$

and $Q_1^{n_i} = Q_1^{n_1}$ for all $i \in \mathbb{N}$. Reindex all sequences by the rule $i :\leftarrow \square_i$.

Let $k \geqslant 2$, and assume for induction that $Q_j^{n_j} = Q_j^n$ for each $n \in \mathbb{N}$ and each $j \in \{1, \ldots, k-1\}$. If $\#\{n \in \mathbb{N} : m_n = k\} = \infty$, there exist integers $n_1 < n_2 < n_3 < \cdots$ such that for all $i \in \mathbb{N}$, $m_{n_i} = k$ and

$$(Q_j^{n_j})_{j=1}^{k-1} \overset{\left(\substack{\text{I.H.,}\\ n=n_i}\right)}{=} (Q_j^{n_i})_{j=1}^{m_{n_i}-1} \overset{\text{(hyp)}}{\in} \mathcal{Q}(\gamma^{n_i}, \mathcal{T}^{n_i})$$

so that

$$(Q_j^{n_j})_{j=1}^{k-1} \in \bigcap_{i=1}^{\infty} \mathcal{Q}(\gamma^{n_i}, \mathcal{T}^{n_i})$$

is as required.

Suppose $\#\{n \in \mathbb{N} : m_n = k\} < \infty$. For each $n \in \mathbb{N}$,

$$(Q_1^{n_1}, \ldots, Q_k^{n_k}, Q_{k+1}^n, \ldots, Q_{m_n-1}^n) \overset{\text{(I.H.)}}{=} (Q_j^n)_{j=1}^{m_n-1} \overset{\text{(hyp)}}{\in} \mathcal{Q}(\gamma^n, \mathcal{T}^n).$$

As before, $\#E(k+1, n_{k+1}) = \infty$ for some $n_{k+1} \in \mathbb{N}$ by local finiteness, so there exist integers

$n_1 < n_2 < n_3 < \cdots$ such that $Q_{k+1}^{n_i} = Q_{k+1}^{n_{k+1}}$ for all $i \in \mathbb{N}$. Reindex all sequences, $i :\leftarrow \square_i$.

The inductive step in this construction repeats for each incremental value of k with

$$\#\{n \in \mathbb{N} : m_n = k\} < \infty,$$

and halts in finitely many steps because $M = \sup_n m_n < \infty$. $\square$

2.5.5 Unfolding complexes, folding maps, and scaffolds

The *unfolding complex* and *folding maps* defined in this section are adaptations of the ideas presented by Chepoi and Maftuleac in [11]. In the present work, Chepoi and Maftuleac's "unfolding" map of a subcomplex of a CAT(0) box complex into $\mathbb{R}^d$ appears as the embedding whose existence is guaranteed by Theorem 4.2.1.

A key technique which will be used in Chapter 5 is the transfer of rays and their square paths between a given square complex X and some better understood square complex X'. (In this work, X' will always be $\mathbb{E}^2$.) We will now give a construction which will serve as an intermediary during the transfer, called an *unfolding complex*.

The unfolding complex can be roughly understood as the of a square path $\mathcal{Q}$ for a ray γ in X to some common universal cover of X and X' which need not be specified or defined. The point is simply that the unfolding complex is the simplest possible square complex $\mathcal{U}(\mathcal{Q})$ which records the adjacency data of successive squares in $\mathcal{Q}$.

The unfolding complex comes with natural projections, called *folding maps*, onto $\mathcal{Q}$ and onto some square path in X' to be determined. By lifting γ to $\mathcal{U}(\mathcal{Q})$ and projecting the result into X', we obtain a curve $\breve{\gamma}$ in X'. As we will explain in the Remark at the end of this subsection, these folding maps can be used to determine the behavior of the ray γ in X by using information about the interaction between $\breve{\gamma}$ and the hyperplanes of X'.

Definition 2.5.11. Let X be a square complex, and let $\mathcal{Q} := (Q_n)_{n=1}^m$ ($m \leqslant \infty$) be a locally monotone edgewise square path. The **unfolding complex** $\mathcal{U} = \mathcal{U}(\mathcal{Q})$ (or the **unfolding of** $\mathcal{Q}$) is the surface with boundary

$$\mathcal{U} = \left(\coprod_{n=1}^{m} (Q_n \times \{n\}) \right) / \sim,$$

where $\sim$ glues copies of successive squares in $\mathcal{Q}$ along their shared edge. The **lift** of a ray $\gamma : [t_1, \infty) \to X$ to $\mathcal{U}$ is the unique ray $\hat{\gamma} = \tilde{\gamma} : [t_1, \infty) \to \mathcal{U}$ defined in the obvious manner, by requiring $\text{proj}_{Q_1}\big(\hat{\gamma}(t_1)\big) = \gamma(t_1)$.

Remark. We require $\mathcal{Q}$ to be locally monotone edgewise in Definition 2.5.11 in order to guarantee that $\mathcal{U}$ is locally Euclidean. If, say, $Q_1 \cap Q_2 = Q_2 \cap Q_3 \in \text{Edges}(X)$, then $\bigcup_{n=1}^{3} Q_n \times \{n\} \subset \mathcal{U}$ would be homeomorphic to $\mathbf{Y} \times I$, where $\mathbf{Y}$ is the cone on three points.

Definition 2.5.12. A **folding map** in a cube complex X for an unfolding complex $\mathcal{U}$ is a cellular map $\varphi : \mathcal{U} \to X$ whose restriction to $\text{int}(\mathcal{U})$ is a local isometry.

Definition 2.5.13. Let $\mathcal{Q} = (Q_n)_{n=1}^{\infty}$ be a square path in a square complex X with defining projection p, say with $p(C_n) = Q_n$ for disjoint cubes $C_n \subset \mathbb{R}^d$. If $\mathcal{U}$ is the unfolding complex of $\mathcal{Q}$, the **natural folding** of $\mathcal{U}$ (Figure 2.7) is the map $\psi : \mathcal{U} \to \bigcup \mathcal{Q}$ defined by

$$\overline{(p(x), n)} \mapsto p(x) \qquad (x \in C_n).$$

Given a square path $\mathcal{Q}$ in a square complex X, and a folding map $\varphi : \mathcal{U}(\mathcal{Q}) \to \mathbb{E}^2$, we would like to establish a formal procedure for mapping a stack of hyperplanes in $\mathbb{E}^2$ whose carriers cover $\varphi(\mathcal{U}(\mathcal{Q}))$ onto a stack of hyperplanes in X whose carriers cover $\bigcup \mathcal{Q}$. The remainder of this subsection standardizes the bookkeeping involved in this procedure, which will be used in §5.2.2.

Lemma 2.5.14. Let $\mathcal{Q} = (Q_n)_{n=1}^{\infty}$ be a square path in a square complex X. Let $\mathcal{U}$ be the unfolding of $\mathcal{Q}$. If $\varphi : \mathcal{U} \to Y$ is a folding map into a square complex Y, and $\psi : \mathcal{U} \to X$ is the natural folding of $\mathcal{U}$, then there exist isometries $\chi_n : \varphi\big(Q_n \times \{n\}\big) \to Q_n$ making the

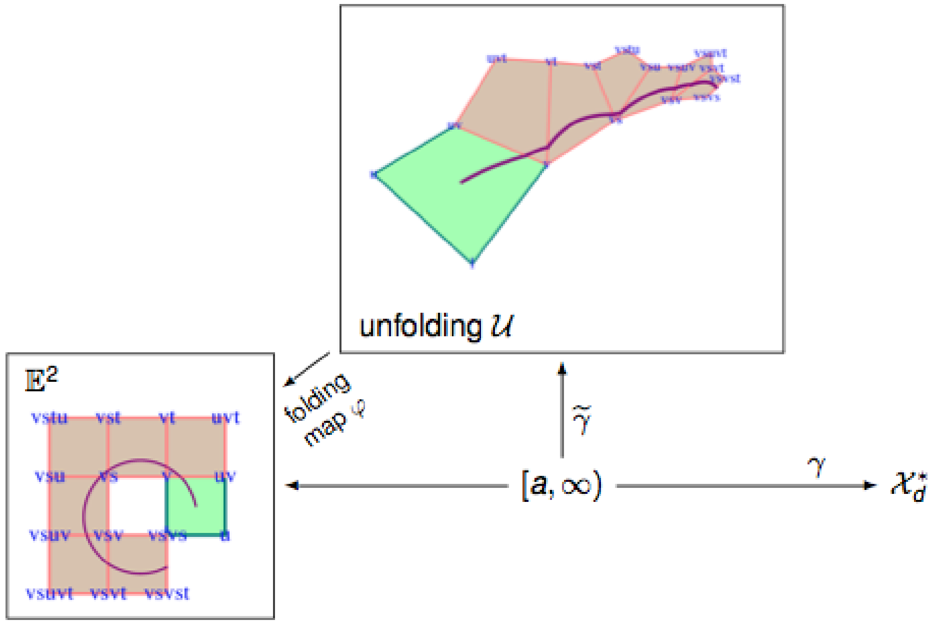

Figure 2.7: Unfolding complex and natural folding.

following diagram commute for each $n \in \mathbb{N}$.

$$\begin{array}{ccc} & Q_n \times \{n\} & \\ \swarrow \varphi & & \psi \searrow \\ \varphi(Q_n \times \{n\}) & \xrightarrow{\chi_n} & Q_n \end{array}$$

Proof. Let p be a defining projection for X, say with $p(C_n) = Q_n$ for disjoint cubes $C_n \subset \mathbb{R}^d$. Choose $\chi_n = \psi \circ \left(\varphi|_{Q_n \times \{n\}}\right)^{-1}$, so that χ_n takes the generic element

$$\varphi\left(\overline{(p(x), n)}\right) \qquad (x \in C_n)$$

of $\varphi(Q_n \times \{n\})$ to $p(x)$. □

Definition 2.5.15. Let $\mathcal{Q} = (Q_n)_{n=1}^m$ $(m < \infty)$ be a square path in a square complex X that is properly segmented with respect to a stack $\mathcal{H} = (H_k)_{k=1}^p$ $(p < \infty)$ of hyperplanes of X.

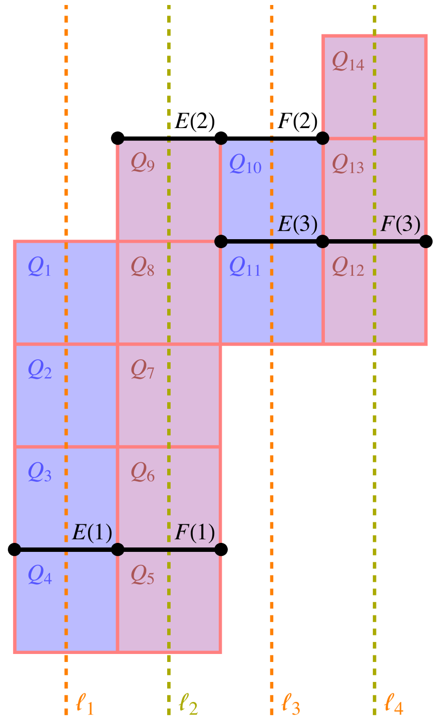

Figure 2.8: A scaffold in $\mathbb{E}^2$.

A pair

$$\big((E_k)_{k=1}^{p-1}, (F_k)_{k=1}^{p-1}\big),$$

where each E_k and F_k is an edge of X, is a **scaffold** for $\mathcal{Q}$ with respect to $\mathcal{H}$ (Figure 2.8) if there exist integers

$$1 = a_1 \leqslant b_1 < a_2 \leqslant b_2 < \cdots \leqslant b_p = m$$

as in Definition 2.5.7 satisfying the additional assumptions that, for each $k \in \{1, \ldots, p-1\}$,

- $E_k \in \text{Edges}(Q_{b_k})$ and is dual to H_k,
- $F_k \in \text{Edges}(Q_{a_{k+1}})$ and is dual to H_{k+1}, and
- E_k and F_k are adjacent.

Remark. Scaffolds will be used to transfer stacks between square complexes that are the codomains of folding maps. The proof of Theorem 5.2.2 will show how scaffolds enable us to identify a sequence of osculating hyperplanes with respect to which a given square path is properly segmented, given a sequence of parallel hyperplanes (lines) in $\mathbb{E}^2$, each a distance 1 from its successor. We will then be able to show that a ray is proper by mapping the unfolding complex of its square path $\mathcal{Q}$ into $\mathbb{E}^2$, finding a stack of parallel lines with respect to which the resulting square path $\check{\mathcal{Q}}$ in $\mathbb{E}^2$ is properly segmented, and then applying Theorem 5.2.2 to a suitable scaffold for $\check{\mathcal{Q}}$. Before proceeding, we record in the following lemma, essential to the proof of Theorem 5.2.2, that it is easy to find a scaffold for a properly segmented square path in $\mathbb{E}^2$.

Lemma 2.5.16 (Construction of scaffolds in $\mathbb{E}^2$). If $\mathcal{Q} = (Q_n)_{n=1}^m$ is an edgewise square path in $\mathbb{E}^2$ that is properly segmented with respect to a stack $\mathcal{H} = (H_k)_{k=1}^p$ of hyperplanes of $\mathbb{E}^2$, then there exists a scaffold for $\mathcal{Q}$ with respect to $\mathcal{H}$.

Proof. Let

$$1 = a_1 \leqslant b_1 < a_2 \leqslant b_2 < \cdots \leqslant b_p = m$$

be as in Definition 2.5.7. We claim that $a_{k+1} = b_k + 1$ for each $k \in \{1, \ldots, p-1\}$. Then $Q_{a_{k+1}}$ and Q_{b_k} are adjacent, so

$$D_k := Q_{b_k} \cap Q_{a_{k+1}}$$

is an edge of $\mathbb{E}^2$. Furthermore, D_k is parallel to H_k and H_{k+1} (that is, parallel to the midplane in Q_k of H_k, and to the midplane in Q_{k+1} of H_{k+1}), because $H_{k_1} \,)\!(\, H_k$ and $\mathbb{E}^2$ contains no interosculating hyperplanes. Let $x_k \in \partial D_k$ for $k \in \{1, \ldots, p-1\}$. Since D_k is parallel to each of H_k and H_{k+1}, there exist $E_k \in \text{Edges}(Q_{b_k})$ and $F_k \in \text{Edges}(Q_{a_{k+1}})$, each incident with x_k, such that E_k is dual to H_k and F_{k+1} is dual to H_{k+1}.

To verify the claim, let $r \in \mathbb{N}$ such that $a_{k+1} = b_k + r$, and assume for a contradiction that $r \geqslant 2$. Suppose $Q_{b_k+1} \subset H_k^-$. Since $\mathcal{Q}[b_k+1, a_{k+1}) \cap H_k = \varnothing$, we have $\mathcal{Q}[b_k+1, a_{k+1}) \subset H_k^-$. But since $\mathcal{Q}[a_{k+1}, m] \subset H_k^+$, the union of squares in $\mathcal{Q}$ must be disconnected, contradicting the definition of a square path. Thus $Q_{b_k+1} \subset H_k^+$. But the only 2-cells of $\mathbb{E}^2$ in H_k^+ adjacent to Carrier H_k lie in Carrier H_{k+1}, so $\mathcal{Q}(b_k, a_{k+1}) \cap H_{k+1} \neq \varnothing$, which contradicts that $\mathcal{Q}$ is properly segmented with respect to $\mathcal{H}$. □

Chapter 3

Existence proof for the Markov-Dubins problem with free terminal direction in a locally finite nonpositively curved cube complex

3.1 Piecewise Lipschitz differentiable paths in a cube complex

Throughout this section, X is a locally finite cube complex, p is a defining projection $\coprod C_\lambda \to X$ for X (see Definition 2.3.9), and when a sequence $(Q_k)_{k=1}^{\infty}$ of cubes of X is given, we write

$$p_k := p\big|_{C_k},$$

where $p(C_k) = Q_k$.

Definition 3.1.1. Let $\gamma : J \to \mathbb{R}^d$ be a curve, where J is a closed interval bounded below. For $t \in \partial J$, let $\gamma'(t)$ be the righthand (lefthand) derivative if $t = \min J$ ($t = \max J$, respectively), and for $t \in \operatorname{int} J$, let $\gamma'(t)$ be the two-sided derivative. For $a > 0$, the curve γ is ***a*-$\mathcal{C}^{1,1}$** (or ***a*-Lipschitz differentiable**) if $\gamma' : J \to \mathbb{R}^d$ is defined and a-Lipschitz.

Definition 3.1.2. Let $a > 0$, and let $\mathcal{Q} := (Q_k)$ be a cube path for a curve $\gamma : J \to X$ with breakpoints $\mathcal{T} := (t_k)_{k=1}^{m}$ ($m \leqslant \infty$), where J is a closed interval (so $m < \infty$ and $t_m = \max J$ if $\sup J < \infty$, and $t_1 = \min J$: see Definitions 2.5.5 and 2.5.6). If

$$p_k^{-1}\gamma_k : [t_k, t_{k+1}] \to \mathbb{R}^d$$

is a-$\mathcal{C}^{1,1}$ for all suitable k, then γ is ***a*-$\mathcal{C}^{1,1}(\mathcal{T}, \mathcal{Q})$** (or ***a*-Lipschitz differentiable with respect to $\mathcal{T}$ and $\mathcal{Q}$**).

Definition 3.1.3. Let J be a compact interval, and fix $M \in \{2, 3, 4, \ldots\}$. A curve $\gamma : J \to X$ is ***a*-$\mathcal{C}^{1,1}(M)$** (or ***a*-Lipschitz differentiable with at most M breakpoints**) if there exists a subdivision $\mathcal{T} := (t_k)_{k=1}^{m}$ of J with $m \leqslant M$ and a cube path $\mathcal{Q} := (Q_k)_{k=1}^{m-1}$ for γ with breakpoints $\mathcal{T}$ such that γ is a-$\mathcal{C}^{1,1}(\mathcal{T}, \mathcal{Q})$.

Example 3.1.4. There exist curves η, ω, ζ in $\mathbb{R}^2$ made up of smooth segments such that η has a cube path but is not a-$\mathcal{C}^{1,1}$ with respect to any cube path, ω is a-$\mathcal{C}^{1,1}$ but not with respect to any finite cube path, and ζ has no cube path at all. Let $X := [0, 1] \times [-1, 1]$ have 2-cells $Q_1 = [0, 1]^2$ and $Q_0 = \operatorname{cl}(X \smallsetminus Q_1)$. Write $e^0 = Q_0 \cap Q_1 = [0, 1] \times \{0\}$.

For $0 \leqslant t < 1$, let

$$f(t) = t^2 \sin \tfrac{1}{t}, \qquad \eta(t) = (t, f(1 - t)).$$

Let $t_0 = 0$ and $t_k = 1 - \frac{1}{k\pi}$ for $k \in \mathbb{N}$. Write $J_k = [t_k, t_{k+1}]$. Then $\eta : [0, 1) \to X$ has cube path $Q_0, Q_1, Q_0, \ldots$ with respect to the subdivision $(J_k)_{k=0}^{\infty}$. (See Figure 3.1.)

- η is not a-$\mathcal{C}^{1,1}$ for any $a \geqslant 0$.

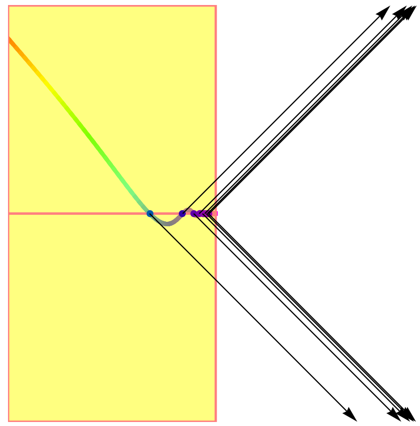

Figure 3.1: A path which is not a-$\mathcal{C}^{1,1}$ with respect to any cube path. (Tangent vectors to the curve η shown in black.)

Proof. Let $a \geqslant 0$. For all $k \in \mathbb{N}$,

$$|\eta'(t_k) - \eta'(t_{k-1})| = |(1,\ \cos k\pi) - (1,\ \cos(k-1)\pi)| = 2.$$

Thus for k so large that $a|t_k - t_{k-1}| < 2$, η' is not a-Lipschitz on $[t_k, t_{k-1}]$. □

Replace each segment $\eta|_{J_k}$ by a parametrization $J_k \to Q_k$ of an arc of a circle of radius $R \geqslant 1/2\pi$, with endpoints $\eta(\partial J_k) \subset e^0$. Assuming each arc is oriented with increasing first coordinate, let ω be the concatenation of all these arcs.

- ω is $(1/R)$-$\mathcal{C}^{1,1}$, being piecewise $\mathcal{C}^\infty$ with

$$|\omega'(s) - \omega'(t)| = \int_s^t |\omega''(u)|\, du = \tfrac{1}{R}(t-s)$$

for any $k \in \mathbb{N}$ and $[s,t] \subseteq J_k$, but clearly there does not exist a finite cube path for ω, as for all $t \in [0,1)$, $\omega([t,1]) \nsubseteq (Q_0 \smallsetminus Q_1) \cup (Q_1 \smallsetminus Q_0)$.

Define $\zeta : [0,1] \to X$ by

$$\zeta(t) = \begin{cases} (t,\ f(1-2t)) & \text{for } 0 \leqslant t \leqslant \frac{1}{2}, \\ (t,\ f(2t-1)) & \text{for } \frac{1}{2} \leqslant t \leqslant 1. \end{cases}$$

Figure 3.2: A path with no cube path.

Then ζ is the concatenation of two pathological curves glued together at their singular points (Figure 3.2).

- ζ has no cube path, because for any increasing sequence $(t_k)_{k=0}^{\infty}$ in $[0, 1]$, for some k, we have $\zeta([t_k, t_{k+1}]) \not\subset (Q_0 \setminus Q_1) \cup (Q_1 \setminus Q_0)$.

Remark. It is easy to see that any one of the pathological behaviors discussed in Example 3.1.4 is exhibited by some limit of finitely a-$\mathcal{C}^{1,1}$ paths. (For example, for each $n \in \mathbb{N}$, modify η to obtain η_n by replacing $\eta\big|_{[t_n, 1]}$ by a regular parametrization of $[t_n, 1]$. Then $d_\infty(\eta_n, \eta) \to 0$.) For this reason, we will work only with sequences of paths that are a-$\mathcal{C}^{1,1}$ with respect to m-cube paths for some fixed $m < \infty$.

Lemma 3.1.5 ([2]). The total curvature of an a-$\mathcal{C}^{1,1}$ path γ in $\mathbb{R}^d$ is $\int |\gamma''(t)|\, dt$.

3.2 Outline of existence proof

As before, X is a locally finite cube complex whose cells Q_λ are the image under a defining map p of disjoint cubes $C_\lambda \subset \mathbb{R}^d$ for some fixed $d \in \mathbb{N}$, and we write $p_\lambda := p\big|_{C_\lambda}$.

Definition 3.2.1. For u, v in X such that $u \neq v$, a unit tangent vector U to $p(C_\lambda)$ at $p_\lambda(u)$

for some λ, $a > 0$, and $M \in \mathbb{N}$, write

$$\mathcal{A}(u, v, U) = \{\gamma : [0, b_\gamma] \to X \,|\, \gamma \text{ is a unit-speed path, } \gamma(0) = u, \gamma(b_\gamma) = v, \gamma'(0) = U\},$$

$$\mathcal{C}(a, M, u, v, U) = \{\gamma \in \mathcal{A}(u, v, U) : \gamma \text{ is } a\text{-}\mathcal{C}^{1,1}(M)\}.$$

In this and the following subsection, we prove:

Theorem 3.2.2. If $\mathcal{C} = \mathcal{C}(a, M, u, v, U)$ is nonempty, then $\mathcal{C}$ contains a path β of minimal length, i.e. $\ell(\beta) = \inf_{c \in \mathcal{C}} \ell(c)$.

We now outline the proof, whose structure parallels that of Dubins' existence proof in [15]. The following diagram indicates the various maps involved.

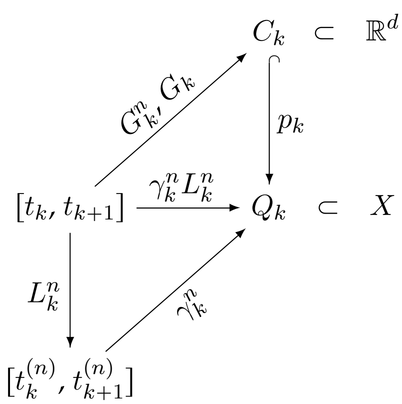

(i) Let $(\gamma^n)_{n=1}^\infty$ be a sequence in $\mathcal{C}(a, M, u, v, U)$ such that $\ell(\gamma^n) \searrow \inf_{c \in \mathcal{C}} \ell(c)$.

(ii) Pass to a subsequence of $(\gamma^n)_{n=1}^\infty$ for which there exists a stable finite cube path $(Q_k)_{k=1}^m$, and let

$$L_k^n : [t_k, t_{k+1}] \to [t_k^{(n)}, t_{k+1}^{(n)}]$$

be nonincreasing changes of variable for some fixed choice of $t_1 < t_2 < \cdots < t_m$. We now have a-$\mathcal{C}^{1,1}$ paths

$$G_k^n = p_k^{-1} \gamma_k^n L_k^n$$

in $C_k \subset \mathbb{R}^d$ with common domain for all $n \in \mathbb{N}$ and $k \in \{1, \ldots, m-1\}$.

(iii) Using Arzelà-Ascoli, recursively pass to subsequences so that

$$G_k^n \xrightarrow{n\to\infty} G_k \quad \text{and} \quad G_k^{n\prime} \xrightarrow{n\to\infty} g_k$$

for some $G_k : [t_k, t_{k+1}] \to C_k \subset \mathbb{R}^d$ and a-Lipschitz $g_k : [t_k, t_{k+1}] \to \mathbb{R}^d$ such that $G_k' = g_k$.

(iv) Check that the uniform limit β of the concatenations

$$\beta^n = pG_1 * \cdots * pG_m$$

is indeed in $\mathcal{C}$.

The main construction (i)–(iii) is carried out in §3.2.1. The verification (iv) of the constructed path's suitability appears in §3.2.2.

Finally, in §3.3, we show that a certain smoothness condition is preserved under the operation of taking uniform limits. It then follows that if we modify our choice of $\mathcal{C}$ to include only paths satisfying the smoothness condition, then the same construction yields a "smooth" path whose length is minimal over all such paths. We will restate our result as Corollary 3.3.11 at the end of this Chapter.

3.2.1 Construction of the length-minimizer β

Assume that $\mathcal{C} = \mathcal{C}(a, M, u, v, U)$ is nonempty. Then $\inf_{c\in\mathcal{C}} \ell(c) < \infty$, and there exists a sequence in $\mathcal{C}$ of paths $\gamma^n : [0, b_n] \to X$ such that

$$\ell(\gamma^n) \searrow \inf_{c\in\mathcal{C}} \ell(c)$$

as $n \to \infty$. In this subsection, we will construct a path β whose length attains this infimum.

For each $\gamma^n : [0, b_n] \to X$, there exists a cube path $(Q_k^n)_{k=1}^{m_n-1}$, $m_n \leqslant M$, for γ^n with breakpoints $(t_k^{(n)})_{k=1}^{m_n}$. Since the $b_n = \ell(\gamma^n)$ are decreasing, we have $b_n \leqslant \ell(\gamma^1) < \infty$.

By Lemma 2.5.10, there exists a stable cube path $(Q_k)_{k=1}^{m-1}$, $m \leqslant M$, for $(\gamma^n)_{n=1}^{\infty}$, so that

$$\gamma^n\big([t_k^{(n)}, t_{k+1}^{(n)}]\big) \subset Q_k$$

for all $k \in \{1, \ldots, m-1\}$ and $n \in \mathbb{N}$. As is spelled out by Lemma 3.2.4, since the b_n are bounded above, there exist $0 = t_1 < \cdots < t_m < \infty$ and changes of variable $L_k^n : [t_k, t_{k+1}] \to [t_k^{(n)}, t_{k+1}^{(n)}]$ such that $|L_k^{n\prime}| \leqslant 1$.

Each Q_k is the image of some cube $C_k \subset \mathbb{R}^d$ under the quotient map p defining X. Write

$$\gamma_k^n = \gamma^n\big|_{[t_k^{(n)}, t_{k+1}^{(n)}]}, \qquad G_k^n = p_k^{-1}\gamma_k^n L_k^n. \tag{3.1}$$

Claim 1. The set $\displaystyle\bigcup_{n=1}^{\infty}\bigcup_{k=1}^{m-1} \text{Image}(G_k^{n\prime}) \subset \mathbb{R}^d$ is bounded.

Proof of Claim 1. Write $\rho = \min_{k=1}^{m-1}(t_{k+1} - t_k)$, and observe that

$$|G_k^{n\prime}(t)| \leqslant \lim_{s \to t} \frac{|G_k^n(s) - G_k^n(t)|}{|s-t|} \leqslant \frac{\ell(\gamma^1)}{\rho}$$

because for all $s, t \in [t_k, t_{k+1}]$, we have

$$\begin{aligned} |G_k^n(s) - G_k^n(t)| = d(\gamma_k^n L_k^n(s), \gamma_k^n L_k^n(t)) \leqslant \ell(\gamma_k^n L_k^n\big|_{[s,\,t]}) &= \frac{\ell(\gamma_k^n L_k^n)}{t_{k+1} - t_k}|s-t| \\ &\leqslant \frac{\ell(\gamma^1)}{\rho}|s-t|. \end{aligned}$$

□

Now, for each $k \in \{1, \ldots, m-1\}$, the collection of functions $\{G_1^n{}' : n \in \mathbb{N}\}$ is uniformly bounded, and by Lemma 3.2.5 it is equicontinuous, so by the corollary to the Arzelà-Ascoli Theorem (3.2.6), we can recursively pass to subsequences so that for each k, the uniform limit of $G_k^n{}'$ exists; call it g_k.

Since G_1^n is differentiable, $G_1^n(0) \xrightarrow{\equiv} p_1^{-1}(u)$, and $G_1^n{}' \to g_1$, by Theorem 3.2.8, the G_1^n have a uniform limit $G_1 : [t_1, t_2] \to C_1$, and $G_1{}' = g_1$.

<u>Claim 2</u>. $(G_1^n(t_1))$ converges in C_1.

Proof of Claim 2. Write

$$P_n := G_1^n(t_1) \in \partial C_1, \qquad P := G_1(t_1), \qquad R_n := G_2^n(t_1) \in \partial C_2,$$

and

$$\overline{x} = p(x) \text{ for } x \in C_1 \sqcup C_2.$$

We have

$$pG_1^n(t_2) = \overline{P_n} \in \partial Q_1 \quad \text{and} \quad pG_2^n(s_2) = \overline{R_n} \in \partial Q_2,$$

so $\partial Q_1 \cap \partial Q_2 \neq \varnothing$. It follows from property (2) of Definition 2.3.7 that there exists a shared face F of ∂Q_1 and ∂Q_2 which is the image under p of faces E_k of C_k $(k = 1, 2)$, and an isometry $\varphi : E_1 \to E_2$ such that $\overline{x} = \overline{y} \Leftrightarrow \varphi(x) = y$ for $x \in E_1$ and $y \in E_2$. Since $\overline{P_n} = \overline{R_n}$, we have $R_n = \varphi(P_n)$, and we know $P_n \to P$ because $G_1^n \to G_1$. Thus by continuity,

$$\lim_{n\to\infty} R_n = \lim_{n\to\infty} \varphi(P_n) = \varphi(P).$$

□

Applying Theorem 3.2.8 again yields a differentiable function $G_2 = \lim_{n\to\infty} G_2^n$ with $G_2{}' = g_2$.

<u>Claim 3</u>. The concatenation $pG_1 * pG_2$ is a path in $Q_1 \cup Q_2$.

Proof of Claim 3.

$$pG_1(t_1) = \lim_{n\to\infty} pG_1^n(t_1) = \lim_{n\to\infty} \overline{P_n} = \lim_{n\to\infty} \overline{R_n} = \lim_{n\to\infty} pG_2^n(t_1) = pG_2(t_1).$$

□

Repeating these steps, we obtain a path

$$\beta := p_k G_1 * \cdots * p_k G_m : [t_1, t_m] \to X, \qquad \beta|_{[t_k,\, t_{k+1}]} = p_k G_k,$$

and $\lim_{n\to\infty} \gamma^n = \beta$ by construction.

3.2.2 Verification that β is a length-minimizer

We now establish Theorem 3.2.2 by proving first that $\ell(\beta)$ is minimal, and then that $\beta \in \mathcal{C}$.

Since

$$\begin{aligned}
\left|\ell(\gamma_k^n) - \ell(\beta|_{[t_k,\, t_{k+1}]})\right| &= |\ell(p_k G_k^n) - \ell(p_k G_k)| && \text{((3.1) and Lemma 2.1.6)} \\
&= |\ell(G_k^n) - \ell(G_k)| && \text{(each } p_k \text{ is an isometry)} \\
&= \left| \int_{[t_k,\, t_{k+1}]} |G_k^{n\prime}| - \int_{[t_k,\, t_{k+1}]} |G_k{}'| \right| \\
&\leqslant \int_{[t_k,\, t_{k+1}]} |G_k^{n\prime} - G_k'| \\
&\leqslant (t_{k+1} - t_k)\, d_\infty(G_k^{n\prime}, G_k') \overset{n\to\infty}{\longrightarrow} 0
\end{aligned}$$

for each $k \in \{1, \ldots, m-1\}$, we have

$$\ell(\beta) = \sum_{k=1}^{m-1} \ell(\beta|_{[t_k,\, t_{k+1}]}) = \lim_{n\to\infty} \sum_{k=1}^{m-1} \ell(\gamma_k^n) = \lim_{n\to\infty} \ell(\gamma^n) = \inf_{c\in\mathcal{C}} \ell(c)$$

After reparametrizing so that β is unit speed, the verification that $\beta \in \mathcal{C}$ is routine:

- Since $\gamma^n(0) = u$, $\gamma^n(b_n) = v$, and $\gamma_n'(0) = U$ for each $n \in \mathbb{N}$, we have $\beta(t_1) = u$, $\beta(t_m) = v$, and $\beta'(t_1) = U$.

- β is a-$\mathcal{C}^{1,1}$ with respect to $(t_k)_{k=1}^m$ and $(Q_k)_{k=1}^{m-1}$:

  To verify that $\beta\big([t_k, t_{k+1}]\big) \subseteq Q_k$ for each $k \in \{1, \ldots, m-1\}$, observe that for each $t \in [t_k, t_{k+1}]$, since $\beta(t) = \lim_{n\to\infty} pG_k^n(t)$ and $G_k^n(t) \in C^k$, we have $\beta(t) \in \overline{Q_k} = Q_k$.

  We now show that each ${G_k}'$ is a-Lipschitz. By Lemma 3.2.9, it is enough to show that each $G_k^{n\prime} = (p_k^{-1}\gamma_k^n L_k^n)'$ is a-Lipschitz, which follows immediately from the inequality

$$\begin{aligned}
\big|G_k^{n\prime}(s) - G_k^{n\prime}(t)\big| &= \big|\big((p_k^{-1}\gamma_k^n)'L_k^n\big)(s)\; L_k^{n\prime}(s) - \big((p_k^{-1}\gamma_k^n)'L_k^n\big)(t)\; L_k^{n\prime}(t)\big| \\
&= |L_k^{n\prime}|\,\big|(p_k^{-1}\gamma_k^n)'\big(L_k^n(s)\big) - (p_k^{-1}\gamma_k^n)'\big(L_k^n(t)\big)\big| \\
&\leqslant \big|(p_k^{-1}\gamma_k^n)'\big(L_k^n(s)\big) - (p_k^{-1}\gamma_k^n)'\big(L_k^n(t)\big)\big| \\
&\leqslant a|L_k^n(s) - L_k^n(t)| \\
&\leqslant a|s-t|.
\end{aligned}$$

This concludes the proof of Theorem 3.2.2.

3.2.3 Lemmas used in existence proof

Lemma 3.2.3. A nonexpanding linear map $L : [a, b] \to [c, d]$ is Lipschitz with constant $|L'| \leqslant 1$.

Lemma 3.2.4. For each $n \in \mathbb{N}$, let $\big(t_k^{(n)}\big)_{k=1}^m$ be a subdivision of $\big[0, t_m^{(n)}\big]$, where $t_m^{(n)} < S$ for some fixed $S < \infty$. Then there exist $T < mS$, a subdivision $(t_k)_{k=1}^m$ of $[0, T]$, and nonexpanding linear surjections $L_k^n : [t_1, t_k] \to [t_k^{(n)}, t_{k+1}^{(n)}]$.

Proof. Let

$$\psi^k = \sup_{n\in\mathbb{N}} \left(t_{k+1}^{(n)} - t_k^{(n)}\right) \in (0, S]$$

for $k = 0, \ldots, m$. Write $s_{-1} = 0$ and $s_k = \sum_{i=0}^{k} \psi^i$ for $k = 0, \ldots, m$. Take

$$L_k^n(t) = (t - s_k)\frac{t_{k+1}^{(n)} - t_k^{(n)}}{\psi^k} + t_k^{(n)}.$$

Then

$$\left|L_n^k(t) - L_n^k(u)\right| \leqslant |t - u|$$

for each $t, u \in [s_k, s_{k+1}]$ and $L_n^{k\prime} \leqslant 1$ for each k and n. □

Lemma 3.2.5. Let $a > 0$. A collection of a-Lipschitz functions from a metric space X to a metric space Y is equicontinuous.

Proof. For any $f \in F$, $x \in X$, and $\varepsilon > 0$,

$$d(f(x), f(y)) \leqslant a\, d(x, y) < \varepsilon$$

whenever $d(x, y) < \varepsilon/a$. □

Theorem 3.2.6 (Arzelà-Ascoli: [31], Theorem 45.4)**.** Let X be a compact space. Then $F \subseteq \mathcal{C}(X, \mathbb{R}^d)$ has compact closure under the uniform metric d_∞ iff F is equicontinuous and uniformly bounded.

Corollary 3.2.7 ([31], Exercise 45.3)**.** If $F \subseteq \mathcal{C}(X, \mathbb{R}^d)$ is equicontinuous and pointwise bounded under the uniform metric d_∞, then every sequence in F has a subsequence that converges in $\left(\mathcal{C}(X, \mathbb{R}^d), d_\infty\right)$.

Proof. Assume there is a sequence $E := (f_n)_{n=1}^{\infty}$ in F that has no convergent subsequence. Then E has no limit points, otherwise there would exist some $f \in \mathcal{C}(X, \mathbb{R}^d)$ such that $f_{n_k} \to f$ for some subsequence $(f_{n_k})_{k=1}^{\infty}$. Since each f_n is therefore an isolated point of E, there exist $\varepsilon_n > 0$ for each $n \in \mathbb{N}$ such that

$$E \cap \mathrm{B}(f_n; \varepsilon_n) = \{f_n\}. \tag{3.2}$$

The set E is closed in $\big(\mathcal{C}(X, \mathbb{R}^d), d_\infty\big)$, hence compact. Then finitely many $\mathrm{B}(f_n; \varepsilon_n)$ suffice to cover E, so one of these balls contains infinitely many f_n, which contradicts (3.2). □

Theorem 3.2.8. Let $(f_n)_{n=1}^{\infty}$ be a sequence of $\mathcal{C}^1$ paths $[a, b] \to \mathbb{R}^d$ such that $\big(f_n(t_0)\big)_{n=1}^{\infty}$ converges for some $t_0 \in [a, b]$. If $(f_n')_{n=1}^{\infty}$ converges uniformly, then $(f_n)_{n=1}^{\infty}$ converges uniformly and $\big(\lim_{n\to\infty} f_n\big)' = \lim_{n\to\infty} f_n'$.

Lemma 3.2.9. The uniform limit of a-Lipschitz functions is a-Lipschitz.

Proof. If f is the uniform limit of functions $f_n : X \to Y$ between metric spaces X, Y, and each f_n is a-Lipschitz, then

$$d(f(s), f(t)) \leqslant d(f(s), f_n(s)) + d(f_n(s), f_n(t)) + d(f_n(t), f(t)) \leqslant a\, d(s, t) + \varepsilon$$

for all $\varepsilon > 0$, $s, t \in X$, and n sufficiently large. □

3.3 Smoothness at breakpoints

In the above argument, admissible paths are required to be piecewise smooth, but may have arbitrary angle where two cubes meet. We now require our paths to be smooth at their breakpoints, in the sense of having zero turning angle (Definition 2.1.25) where two cubes meet. The key fact needed to apply the same argument as above in this case is that the property of having zero turning angle at a point is preserved under uniform limits.

Remark. The link at an arbitrary point p in a geodesic space X can be regarded as the collection of geodesic segments issuing from p modulo subset containment [8]. In this sense, the link at p is the *space of directions* at p, and endowing link(p) with the angle metric makes it a metric space. Indeed, we have already recorded the triangle inequality for this space as Lemma 2.1.20. We have chosen to omit the definition of link(p) given in [8] for arbitrary p, and will present only those facts which are essential to our arguments. The remarks in this subsection are purely explanatory, and our proofs do not depend on them.

Any ray or path γ with initial point p for which there exists a geodesic segment in X making angle 0 with γ is represented by an element of link(p). We may then say that γ *has an initial direction*. Two rays or paths issuing from p, not necessarily geodesic, will be said to have the *same direction* if the angle between them exists and is 0. (For simplicity, we will refer only to "paths," rather than "rays or paths," from this point forward.) The relation of having the same direction is an equivalence relation on the set of paths issuing from p that have a direction. The equivalence classes under this relation will be called *directions*.

If γ is differentiable with nonzero derivative in coordinates at its initial point, in the sense that $(\varphi\gamma)'(a) \neq 0$ for some isometric map φ into $\mathbb{R}^N$ whose domain contains the image of γ, then γ clearly has a direction. A vector $V \in \mathbb{R}^N$ with direction the same as that of $\varphi\gamma$, where γ is a path, and φ is an isometry of a cell containing γ, will be called a *pre-direction* for γ. Clearly, a pre-direction for γ depends on the choice of φ. But if φ is fixed, and γ has the same direction as some geodesic in X, then

(*i*) the pre-directions for a given path γ are determined up to a positive scalar multiple, and

(*ii*) measuring angles between a pair of paths in a cell of X is equivalent to measuring the angles between a corresponding pair of pre-directions under φ.

Definition 3.3.1. Let X be a geodesic space, and let α and β be curves in X with common initial point x. If $\angle(\alpha, \beta) = 0$, the curves α and β have the **same (initial) direction**. For

$y \in X \setminus \{x\}$, the curve α and the geodesic segment $[xy]$ have the **same (initial) direction** if there exists a parametrization λ of $[xy]$ such that $\angle(\alpha, \lambda) = 0$. If $X = \mathbb{R}^N$ and $V \in \mathbb{R}^N \setminus \{0\}$, the curve α and the vector V have the **same (initial) direction** provided that there exists a parametrization ν of the path with direction V issuing from x such that $\angle(\alpha, \nu) = 0$. Define $\angle([xy], \beta)$, $\angle(V, \beta)$, etc., analogously.

Definition 3.3.2. Let Q_λ be a cube in a cube complex with defining projection p, say with $p(C_\lambda) = Q_\lambda$ for disjoint cubes $C_\lambda \subset \mathbb{R}^N$, and write $\varphi = p_\lambda^{-1}$. A vector $V \in \mathbb{R}^N \setminus \{0\}$ is an **(initial) pre-direction** for a path $\gamma : [a, b] \to Q$ if the path $\varphi\gamma$ is differentiable at a, and $\varphi\gamma$ and V have the same direction.

Remark. A path γ in a geodesic space X that passes through a point $\gamma(t_0) = p$ can be regarded as the concatenation of the path $\gamma|_{[t_0-\varepsilon,\, t_0]}$, whose reverse path we denote by γ^-, with the path

$$\gamma^+ := \gamma|_{[t_0,\, t_0+\varepsilon]}.$$

If γ is a smooth path and X is a smooth manifold, then for any $w \in X$ such that

$$\angle(\gamma^-, [pw]) = \pi,$$

i.e., the turning angle of the concatentation $\gamma|_{[t_0-\varepsilon,\, t_0]} * [pw]$ is 0, we have

$$\angle([pw], \gamma^+) = 0.$$

This implication does not hold in a general geodesic space, even if γ is a geodesic. For example, if γ is a geodesic path passing through a point p in a square complex such that the link at p has total angle measure m such that $2\pi < m < 3\pi$, then there exists a w such that

$$\angle(\gamma^-, [pw]) = \pi$$

and

$$\angle([pw], \gamma^+) = \theta$$

for any $\theta \in [0, m - 2\pi]$.

When discussing the directional behavior of a path passing through a point in a cube complex, we will therefore keep track of two geodesic segments, one making angle 0 with γ^-, and one making angle 0 with γ^+. This convention enables us to distinguish between paths which are so-called *bifurcating geodesics*, that is, geodesic paths that have the same initial direction, but distinct tails (e.g., Figure 1.4). Such a pair of geodesic segments—a special case of what we will call a *hinge*—acts as a canonical choice of *integral curve*, i.e. a path having prescribed approach and departure directions at a point p.

Definition 3.3.3. The concatenation $\alpha * \beta$ of two nondegenerate paths $\alpha : [a, b] \to X$ and $\beta : [b, c] \to X$ in a length space X is a **hinge** with **hinge point** $p := \alpha(b) = \beta(b)$. Its **(hinge) angle** is $\angle(\overline{\alpha}, \beta)$.

Definition 3.3.4. Let $\alpha, \alpha_1, \alpha_2, \cdots : [a, b] \to X$ and $\beta, \beta_1, \beta_2, \cdots : [b, c] \to X$ be paths such that each of $\alpha * \beta, \alpha_1 * \beta_1, \alpha_2 * \beta_2, \ldots$ is a hinge. If $\alpha_n \to \alpha$ and $\beta_n \to \beta$, where α and β are nondegenerate paths, we say the sequence of hinges $\alpha_n * \beta_n$ **converges** to the hinge $\alpha * \beta$, and we write $\alpha_n * \beta_n \xrightarrow{*} \alpha * \beta$. (It is automatic that $\alpha(b) = \beta(b)$, so that $\alpha * \beta$ is indeed a hinge.)

We now turn to the task of showing that if the angles of hinges $\alpha_n * \beta_n$ are each π, and if the hinges converge, $\alpha_n * \beta_n \xrightarrow{*} \alpha * \beta$, then the angle of $\alpha * \beta$ is also π. This is trivial when the hinge points lie in the interior of a single cube. A bit more work is required for the case we are concerned with, namely, the case that all hinge points are breakpoints of the cube complex X, with the α_n all lying in one cube Q_1, and the β_n all lying in another cube Q_2.

Our primary tool for analyzing the behavior of angles under limits is Theorem 2.1.28,

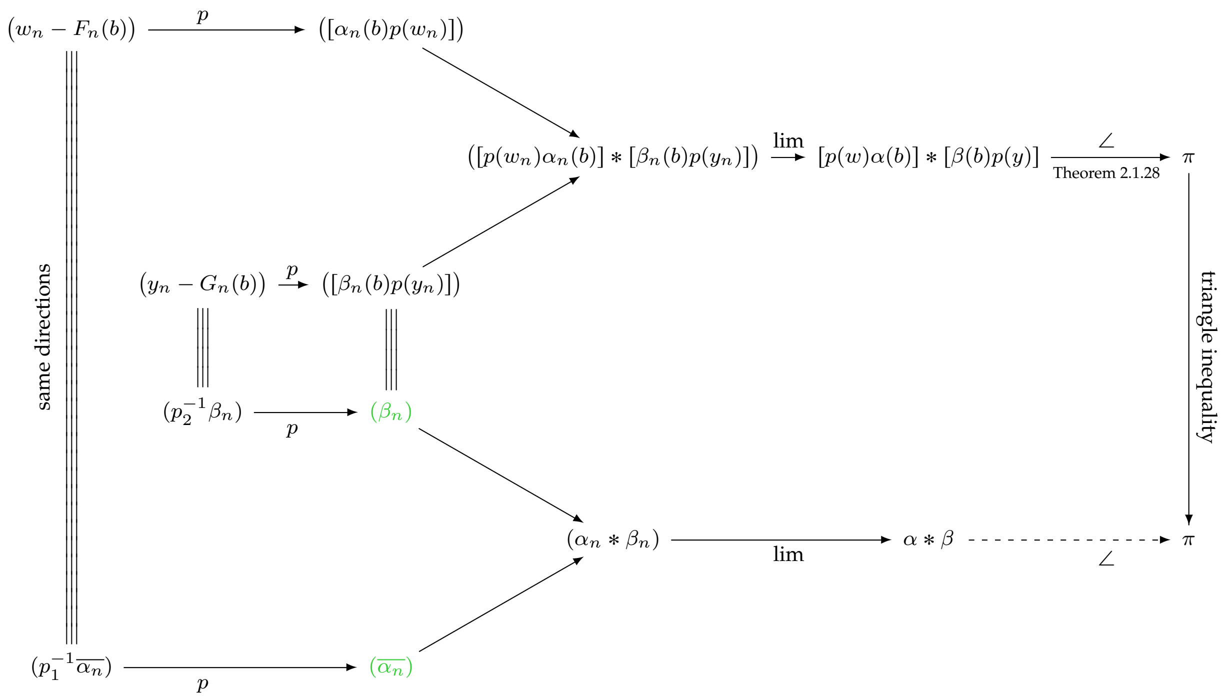

Figure 3.3: Schematic for Lemma 3.3.7. (The proof begins with the green nodes.)

which says that if $[a_n b_n] * [b_n c_n] \xrightarrow{*} [ab] * [bc]$, then the hinge angle of $[ab] * [bc]$ is no less than the $\limsup$ of the hinge angles of $[a_n b_n] * [b_n c_n]$. In order to apply this to generic, not necessarily geodesic hinges, we will pull back to pre-directions, which can be projected back into our cube complex to obtain directions for $\overline{\alpha_n}$ and β_n. Theorem 2.1.28 can then be applied. The schematic diagram given in Figure 3.3 may serve to clarify the process.

It is in choosing pre-directions that a certain technical issue must be shown to be avoidable: we must make sure that it is possible to choose *nondegenerate* pre-directions for the "legs" of the various hinges, otherwise the angle between the corresponding directions will be undefined. We write $p_k = p|_{C_k}$ as usual, where p is the defining projection for X, and consider the pullbacks $F_n = p_1^{-1}\overline{\alpha_n}$ and $G_n = p_2^{-1}\beta_n$ ($n \in \mathbb{N}$), along with the limits

$$F = \lim_{n\to\infty} F_n, \qquad G = \lim_{n\to\infty} G_n, \qquad \alpha = \lim_{n\to\infty} \alpha_n, \qquad \beta = \lim_{n\to\infty} \beta_n.$$

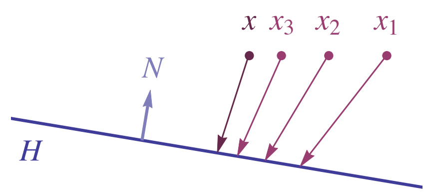

Figure 3.4: Illustration of the hypotheses of Lemma 3.3.5.

In Lemma 3.3.6, we establish that a convergent sequence of nondegenerate geodesic predirections for the $\overline{\alpha_n}$ (and β_n) can be constructed if it is known that each of the $F_n'(b)$ (and $G_n'(b)$) is nonzero. As a preliminary to Lemma 3.3.6, we prove a technical geometric convergence result, Lemma 3.3.5.

Lemma 3.3.5. Let H be a codimension-1 hyperplane in $\mathbb{R}^m$, i.e. a set of the form $a + P$, where $a \in \mathbb{R}^m$ and P is a $(m-1)$-dimensional linear subspace of $\mathbb{R}^m$ ($m \geqslant 2$). Let $(x_n)_{n=1}^{\infty}$ be a sequence in an open halfspace of $\mathbb{R}^m$ bounded by H such that $x_n \to x \notin H$. Let $(v_n)_{n=1}^{\infty}$ be a sequence of vectors in $\mathbb{R}^m$ such that $v_n \to v$, $x + sv \in H$ for some $s > 0$, and $x_n + s_n v_n \in H$ for some $s_n > 0$ ($n \in \mathbb{N}$). Then the distances $s_n/|v_n|$ from x_n to H along v_n converge to the distance $s/|v|$ from x to H along v.

Proof. Let N be a unit normal vector for H, and let $c \in \mathbb{R}$ such that $H = \{z \in \mathbb{R}^m : z \cdot N + c = 0\}$. For simplicity, assume $c = 0$.

We know x lies in the same open halfspace bounded by H as do the x_n ($n \in \mathbb{N}$) since $x_n \to x \notin H$. Since $s > 0$ and $s_n > 0$ ($n \in \mathbb{N}$),

$$\langle v, N \rangle = -\tfrac{1}{s}\langle x, N \rangle \qquad \text{and} \qquad \langle v_n, N \rangle = -\tfrac{1}{s_n}\langle x_n, N \rangle \;\; (n \in \mathbb{N})$$

all have the same sign.

The distance from an arbitrary point $y \in \mathbb{R}^m$ to H along an arbitrary vector $u \in \mathbb{R}^m$ is

$$D(y, u) = \frac{\langle y, N \rangle}{\langle u, N \rangle}.$$

This function is continuous on each component of $\mathbb{R}^m \times (\mathbb{R}^m \setminus \{\langle \cdot, N\rangle = 0\})$. Thus $D(x_n, v_n) \to D(x, v)$. □

Lemma 3.3.6. Let $C \subset \mathbb{R}^N$ be a cube. Let $F, F_1, F_2, \cdots : [a, b] \to C \subset \mathbb{R}^N$ be paths that are differentiable at a with nonzero derivative at a. If $F_n(a) \to F(a)$ and $F_n'(a) \to F'(a)$ for each n, then there exist $w, w_1, w_2, \cdots \in C$ such that

$$w_n \to w, \qquad F_n(a) \neq w_n, \qquad F(a) \neq w,$$

$w_n - F_n(a)$ has the same direction as F_n, and $w - F(a)$ has the same direction as F.

Proof.

Claim: The (nonzero) distances r_n from $F_n(a)$ to $\partial C \setminus E_n$ along the rays $\{F_n(a) + sF_n'(a) : s \geqslant 0\}$ converge to the nonzero distance r from $F(a)$ to $\partial C \setminus E$ along the ray $\{F(a) + sF'(a) : s \geqslant 0\}$.

> *Proof of Claim.* Let E be the union of open faces $\operatorname{int} S$ of C such that S contains $F(a)$, and let E_n be the union of open faces $\operatorname{int} S$ of C such that S contains $F_n(a)$. If $F_n(a) \notin \operatorname{Vert} C$, then $F_n(a) \in \operatorname{int} S$ for some k-dimensional ($k \geqslant 1$) face S of C, and $0 < d(F_n(a), \partial S) = d(F_n(a), \partial C \setminus E_n)$. If $F_n(a) \in \operatorname{Vert} C$, then $1 = d(F_n(a), \partial C \setminus E_n)$. In either case, $d(F_n(a), \partial C \setminus E_n) > 0$. Similarly, $d(F(a), \partial C \setminus E) > 0$. Thus the distances along the specified rays to $\partial C \setminus E_n$ ($\partial C \setminus E$ respectively) can be no smaller. Take $r_n = D_n / |F_n'(a)|$, where $D_n > 0$ is the distance along $F_n'(a)$ from $F_n(a)$ to $\partial C \setminus E_n$, and $r = D/|F'(a)|$, where $D > 0$ is the distance along $F'(a)$ from $F(a)$ to $\partial C \setminus E$. Now by Lemma 3.3.5, $r_n \to r$. □

Choose

$$w = F(a) + rF'(a), \qquad w_n = F_n(a) + r_n F_n'(a).$$

Then

$$\angle\big(w - F(a), F\big) = \angle\big(rF'(a), F'(a)\big) = 0,$$
$$\angle\big(w_n - F_n(a), F_n\big) = \angle\big(r_n F_n'(a), F_n'(a)\big) = 0$$

by Lemma 2.1.19, and $w_n \to w \neq F(a)$. □

Lemma 3.3.7 (Preservation of smoothness at breakpoints under uniform limits). Let Q_1, Q_2 be cubes in a nonpositively curved cube complex X, and let $p_i : \mathbb{R}^N \supset C_i \to Q_i$ be an isometry ($i = 1, 2$). Let $\alpha, \alpha_1, \alpha_2, \cdots : [a, b] \to Q_1$ and $\beta, \beta_1, \beta_2, \cdots : [b, c] \to Q_2$ be paths parametrized at constant speed. Suppose:

- each of $\alpha * \beta, \alpha_1 * \beta_1, \alpha_2 * \beta_2, \ldots$ is a hinge,
- $\alpha_n * \beta_n \xrightarrow{*} \alpha * \beta$,
- $\angle(\overline{\alpha_n}, \beta_n) = \pi$ for all n,
- each of $F := p_1^{-1}\overline{\alpha}$, $F_n := p_1^{-1}\overline{\alpha_n}$, $G := p_2^{-1}\beta$, $G_n := p_2^{-1}\beta_n$ is differentiable at b with nonzero derivative at b for all n,
- $F_n'(b) \to F'(b)$, and $G_n'(b) \to G'(b)$

Then $\angle(\overline{\alpha}, \beta) = \pi$.

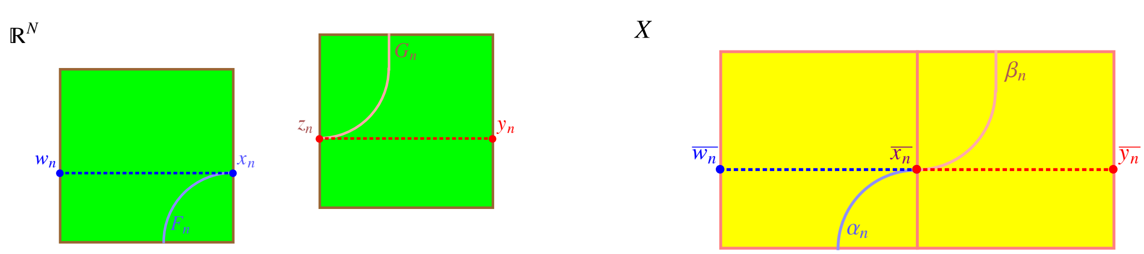

Proof. Using Lemma 3.3.6 and the fact that each of p_1 and p_2 is an isometry (hence angle-preserving), we obtain pre-directions $w - F(b), w_1 - F_1(b), w_2 - F_2(b), \cdots \in C_1$ for $F, F_1, F_2, \ldots$, and pre-directions $y - G(b), y_1 - G_1(b), y_2 - G_2(b), \cdots \in C_2$ for $G, G_1, G_2, \ldots$, respectively, such that $w_n \to w \neq F(b)$ and $y_n \to y \neq G(b)$. By Lemma 2.1.20,

$$\begin{aligned} &\angle\big([\alpha_n(b)\,p_1(w_n)], [\beta_n(b)\,p_2(y_n)]\big) - \angle\big(\overline{\alpha_n}, \beta_n\big) \\ &\qquad \leqslant \angle\big([\alpha_n(b)\,p_1(w_n)], \overline{\alpha_n}\big) + \angle\big(\beta_n, [\beta_n(b)\,p_2(y_n)]\big) \\ &\qquad = 0, \end{aligned}$$

so

$$\angle\big([\alpha_n(b)\,p_1(w_n)], [\beta_n(b)\,p_2(y_n)]\big) = \angle(\overline{\alpha_n}, \beta_n) = \pi.$$

Noting that

$$[\alpha_n(b)\,p_1(w_n)] \to [\alpha(b)\,p_1(w)] \neq \{\alpha(b)\}, \qquad [\beta_n(b)\,p_2(y_n)] \to [\beta(b)\,p_2(y)] \neq \{\beta(b)\},$$

and that $[\alpha_n(b)\,p_1(w_n)] * [\beta(b_n)\,p_2(y_n)]$ is a hinge for each n, we apply Theorem 2.1.28 to get

$$\angle\big([\alpha(b)\,p_1(w)], [\beta(b)\,p_2(y)]\big) = \pi.$$

Now Lemma 2.1.20 yields

$$\begin{aligned} \angle\big([\alpha(b)\,p_1(w)], [\beta(b)\,p_2(y)]\big) - \angle\big(\overline{\alpha}, \beta\big) &\leqslant \angle\big([\alpha(b)\,p_1(w)], \overline{\alpha}\big) + \angle\big(\beta, [\beta(b)\,p_2(y)]\big) \\ &= 0, \end{aligned}$$

so $\angle(\overline{\alpha}, \beta) = \pi$. □

Definition 3.3.8. We say an a-$\mathcal{C}^{1,1}(M)$ path γ is **smooth at breakpoints** if there exists a cube path (Q_k) for γ with breakpoints (t_k) such that γ has zero turning angle at $\gamma(t_k)$ for all k.

Definition 3.3.9. For u, v in X such that $u \neq v$, $U \in \text{link}(u)$, $a > 0$, and $M \in \mathbb{N}$, write

$$\mathcal{C}_0(a, M, u, v, U) = \{\gamma \in \mathcal{A}(u, v, U) : \gamma \text{ is } a\text{-}\mathcal{C}^{1,1}(M) \text{ and smooth at breakpoints}\}.$$

Theorem 3.3.10. For any sequence $(\beta^n)_{n=1}^{\infty}$ of elements of $\mathcal{C}_0 = \mathcal{C}_0(a, M, u, v, U)$ with uniformly bounded lengths, there exist reparametrizations $\widehat{\beta}^n$ of the β^n with a common domain such that some subsequence $(\widehat{\beta}^{n_i})_{i=1}^{\infty}$ has uniform limit in $\mathcal{C}_0$.

Corollary 3.3.11. If $\mathcal{C}_0$ is nonempty, then $\mathcal{C}_0$ contains a path β of minimal length, i.e.

$$\ell(\beta) = \inf_{c \in \mathcal{C}_0} \ell(c).$$

Proof of Theorem 3.3.10. Let $(\beta^n)_{n=1}^{\infty}$ in $\mathcal{C}_0$ be a sequence with uniformly bounded lengths $\ell(\beta^n)$. By Lemma 2.5.10, there exists a stable cube path $\mathcal{Q} := (Q_k)_{k=1}^{m-1}$, $m \leqslant M$, for $(\beta^n)_{n=1}^{\infty}$. As in subsection 3.2.1, we may reparametrize the β^n to obtain a common domain $[0, t_m]$, $t_m < \infty$ (the uniform boundedness of the lengths $\ell(\beta^n)$ is used here), and a subdivision $0 = t_1 < \cdots < t_m$ such that $\mathcal{Q}$ is a stable cube path for the reparametrized β^n with breakpoints $(t_k)_{k=1}^{m}$. Write $\beta_k^n := \beta^n|_{[t_k, t_{k+1}]}$ and $\beta_k := \beta|_{[t_k, t_{k+1}]}$ for $k \in \{1, \ldots, m-1\}$. We verify that β is smooth at breakpoints.

If $\lim_{n\to\infty} \ell(\beta_1^n * \cdots * \beta_k^n) \to 0$, say $\beta^n\big([t_1, t_{k+1}]\big) = \{p\}$, then we may reparametrize β so that $\beta^{-1}(p) = \{0\}$ and replace $\mathcal{Q}$ with the cube path $(Q_i)_{i=k+1}^{m-1}$. Then p is not a breakpoint for the reparametrized β, and can be ignored in the present verification. The case that $\lim_{n\to\infty} \ell(\beta_k^n * \cdots * \beta_{m-1}^n) \to 0$ can be dealt with similarly.

For each k such that $\lim_{n\to\infty} \ell(\beta_k^n) \neq 0$ we have $\ell(\beta_k) \neq 0$ by the Dominated Convergence Theorem and Lemma 2.2.8, and we can pass to a subsequence so that $\ell(\beta_k^n) \neq 0$ for each n. If $\lim_{n\to\infty} \ell(\beta_k^n) \neq 0$ and $\lim_{n\to\infty} \ell(\beta_{k+1}^n) \neq 0$ then, possibly after passing to subsequences, $\beta_k * \beta_{k+1}$ and each $\beta_k^n * \beta_{k+1}^n$ is a hinge (in particular, each of β_k^n, β_{k+1}^n, β_k, β_{k+1} is nondegenerate). Again as in subsection 3.2.1, we may pass to subsequences so that the derivatives in coordinates of the β_k^n converge uniformly to the derivative in coordinates of β_k, and by Lemma 3.3.7, $\angle(\overline{\beta_k}, \beta_{k+1}) = \pi$.

We claim that $\beta_k * \cdots * \beta_{k+1}$ is smooth at breakpoints for each k such that $\ell(\beta_k) \neq 0$, $\ell(\beta_{k+1}) \to 0$, and $\ell(\beta_{k+2}) \neq 0$. Write

$$\begin{aligned} \tau[t_i, t_j] &:= \tau(\beta_i * \cdots * \beta_j), \\ \tau^n[t_i, t_j] &:= \tau(\beta_i^n * \cdots * \beta_j^n), \\ \tau_i &:= \tau(\beta_i), \\ \tau_i^n &:= \tau(\beta_i^n). \end{aligned}$$

Using Lemma 3.1.5 and the Dominated Convergence Theorem (note that $|G_k^{n\prime\prime}(t)| \leqslant a$ almost everywhere by Rademacher's Theorem, 2.2.9), we have

$$\tau_k^n = \int_{t_k}^{t_{k+1}} |G_k^{n\prime\prime}| \xrightarrow{n\to\infty} \int_{t_k}^{t_{k+1}} |G_k{}''| = \tau_k,$$

where $G_k^n := p_k^{-1}\beta_k^n$. Since $\ell(\beta_{k+1}) \to 0$, we have

$$\tau_{k+1}^n \xrightarrow{n\to\infty} 0.$$

Using Corollary 2.2.6 and the fact that each $\beta^n := \beta_1^n * \cdots * \beta_m^n$ is smooth at breakpoints,

$$\begin{aligned} \tau^n[t_k, t_{k+3}] &= \tau_k^n + \left(\pi - \angle(\overline{\beta_k^n}, \beta_{k+1}^n)\right) + \tau_{k+1}^n + \left(\pi - \angle(\overline{\beta_{k+1}^n}, \beta_{k+2}^n)\right) + \tau_{k+2}^n \\ &= \tau_k^n + \tau_{k+1}^n + \tau_{k+2}^n \\ &\xrightarrow{n\to\infty} \tau_k + \tau_{k+2}. \end{aligned}$$

By lower semicontinuity of total curvature (Lemma 2.2.5),

$$\tau_k + (\pi - \angle(\overline{\gamma_k}, \gamma_{k+2})) + \tau_{k+2} = \tau[t_k, t_{k+3}]$$
$$\leqslant \lim_{n\to\infty} \tau^n[t_k, t_{k+3}] = \tau_k + \tau_{k+2}.$$

Thus $\angle(\overline{\gamma_k}, \gamma_{k+2}) = \pi$. A similar argument shows that if $\ell(\beta_k) \neq 0$, $\ell(\beta_{k+1} * \cdots * \beta_{k+j}) = 0$, and $\ell(\beta_{k+j+1}) \neq 0$, then $\angle(\overline{\gamma_k}, \gamma_{k+j+1}) = \pi$. $\square$

Chapter 4

Numerical results

Having established the existence of a solution to our version of the Markov-Dubins problem with free terminal direction for a broad category of cube complexes X, we will focus in the remainder of this work on the case that X is a square complex.

In the first section of this Chapter, we generalize a result due to Jacobs and Canny [22] which characterizes solutions to the Markov-Dubins problem with specified initial and terminal direction in a planar region with polygonal obstacles. (Note that a planar region from which polygonal regions have been deleted is indeed a nonpositively curved square complex, which of course fails to be CAT(0) if the resulting region is not simply connected.) Perhaps unsurprisingly, in light of Dubins' characterization, Jacobs and Canny found that solutions are necessarily made up of geodesic segments and arcs of circles. We find that the same is true in an arbitrary nonpositively curved square complex. It is then immediate that the same characterization holds for solutions of the Markov-Dubins problem with free terminal direction, since any solution to the problem without specified terminal direction is a solution to some problem with the terminal direction prescribed.

Section 4.2 gives the details for an algorithm which numerically determines the solution in the case that it is a *CL path*, that is, made up of a path of constant curvature $\kappa > 0$, followed by a geodesic path. As Kreĭn and Nudel'man note in [24], the solution in the plane to the problem with free terminal direction is always of this form. We do not address the question of determining when and if there exist solutions which are not CL paths. However, it is plain that for certain choices of nonpositively curved square complex X, boundary conditions u, v, U, and curvature constant $\kappa > 0$, no admissible path whatsoever exists,[1] and in Chapter 5, we will encounter spaces homeomorphic to $\mathbb{R}^2$ which contain unbounded regions inaccessible by CL paths with prescribed u, v, U, and κ. (Recall that our existence result claims only that a length-minimizer exists when the set $\mathcal{C}$ of admissible paths is nonempty.) We will return to the problem of determining a sufficient condition for the existence of an arbitrary admissible path in future work.

4.1 Characterization of length-minimizers

Theorem 4.1.1 ([22], Corollary to Theorem 2)**.** Let $\Omega \subset \mathbb{E}^2$ be a finite union of polygons. Let $u, v \in \mathbb{E}^2$, $U, V \in S^1$, and $a > 0$. Let

$$\mathcal{C}_{\mathrm{FP}} := \mathcal{C}_{\mathrm{FP}}(a, \Omega, u, v, u, v)$$

be the collection of paths $\gamma : [0, b_\gamma] \to \Omega$ such that $|\gamma'| \equiv 1$, γ' is a-Lipschitz, $\gamma(0) = u$, $\gamma(b_\gamma) = v$, $\gamma'(0) = U$, and $\gamma'(b_\gamma) = V$. Then

1) $\mathcal{C}_{\mathrm{FP}}$ is empty or contains a length-minimizing element.

2) Every length-minimizer is $\mathcal{C}^1$ and is made up of finitely many pieces, each a line segment or an arc of a circle of radius $1/a$, and meets $\partial\Omega$ in finitely many points and/or line segments.

Theorem 4.1.2. Let u, v be distinct points of a nonpositively curved square complex X, let $U \in \mathrm{link}(u)$, $a > 0$, and $M \in \mathbb{Z}_{\geqslant 0}$, Any length-minimizing element of $\mathcal{C} = \mathcal{C}(a, M, u, v, U)$

[1] E.g., retrieving one's car from a parking garage without using the reverse gear.

is made up of finitely many geodesic segments and/or arcs of a metric circle of radius $1/a$ (i.e., the image of the arc under a distance-preserving map into $\mathbb{E}^2$ is an arc of a circle of radius $1/a$), and meets $X^{(1)}$ in finitely many points and/or geodesic segments.

Proof. Suppose $\gamma \in \mathcal{C}$ is a length-minimizer, i.e. $\ell(\gamma) = \inf_{c\in\mathcal{C}} \ell(c)$, and let $(t_i)_{i=1}^{m}$, $(Q_i)_{i=1}^{m-1}$ be a subdivision and cube path as in the definition of $\mathcal{C}$. Without loss of generality, we may assume each Q_i has dimension 2. (For example, a 1-cell Q_i followed by a 2-cell Q_{i+1} can be replaced by a 2-cell $R^i \neq Q_{i+1}$ since $Q_{i+1} \not\supset Q_i$ and X is the union of its 2-cells.)

For $i \in \{1, \ldots, m-1\}$, suppose $\gamma_i := \gamma|_{[t_i, t_{i+1}]}$ has initial position $\gamma_i(s_i) = u_i$ and terminal position $\gamma_i(t_i) = v_i$. We show that γ_1 is as claimed. Let e be an isometry from Q_0 onto a square in $\mathbb{E}^2$. It is clear that

$$e\gamma_1 \in \mathcal{C}_{\text{FP}} := \mathcal{C}_{\text{FP}}(a, e(Q_1), e(u_1), e(v_1), (e\gamma_1)'(s_1), (e\gamma_1)'(t_1)).$$

Since $\mathcal{C}_{\text{FP}} \neq \varnothing$, by Theorem 4.1, $\mathcal{C}_{\text{FP}}$ contains a length-minimizing element β_1. Take the domain of β_1 to be $H_1 = [\tilde{t}_1, \tilde{t}_2]$, where $\tilde{t}_2 = t_2$, and set $\tilde{t}_i = t_i$ for $i \in \{3, \ldots, m\}$. Assume for a contradiction that $\ell(e\gamma_1) > \ell(\beta_1)$, i.e. $e\gamma_1$ is not a length-minimizing element of $\mathcal{C}_{\text{FP}}$.

Observe that

- $\tilde{\gamma} := e^{-1}\beta_1 * \gamma_2 * \cdots * \gamma_{m-1}$ is a-$\mathcal{C}^{1,1}$ with respect to $(\tilde{t}_i)$ and (Q_i).
- $e^{-1}\beta_1$ has the same boundary conditions (position of, and direction at endpoints) as γ_1.
- $\angle(-e^{-1}\beta_1, \gamma_2 * \cdots * \gamma_{m-1}) = \angle(-\gamma_1, \gamma_2 * \cdots * \gamma_{m-1}) = \pi$ since $-\beta_1$ and $-e\gamma_0$ have the same direction at $e(v_1)$.

It follows that $\tilde{\gamma} \in \mathcal{C}$. But then, since γ is a length-minimizing element of $\mathcal{C}$,

$$\ell(e\beta_1) + \ell(\gamma_2 * \cdots * \gamma_{m-1}) = \ell(\tilde{\gamma}) \geqslant \ell(\gamma) = \ell(\gamma_1) + \ell(\ell(\gamma_2 * \cdots * \gamma_{m-1}),$$

which contradicts that $\ell(e\beta_1) < \ell(\gamma_1)$. Thus $e\gamma_1$ is indeed a length-minimizing element of $\mathcal{C}_{\text{FP}}$, so by Theorem 4.1, $e\gamma_1$ is made up of finitely many line segments and arcs of circles of radius $1/a$. Since e^{-1} takes line segments to geodesics in X, the path γ_1 is as required, and similar reasoning shows that each of $\gamma_2, \ldots, \gamma_{m-1}$ is as well. □

4.2 Computation of length-minimal CL paths

By a **CL path** in a geodesic space, we mean the concatenation of a path of constant curvature followed by a geodesic path, with turning angle 0 at the breakpoint. This terminology is based on Dubins' 1957 paper, in which a piecewise $\mathcal{C}^2$ path made up of circular arcs and line segments is described by a word in the letters C (for circle) and L for line segment. In this section, we will state and justify an algorithm for approximating the shortest CL path in a nonpositively curved square complex with prescribed endpoints and initial direction. The key idea is to reduce the problem to that of finding a root of a continuous function of one variable.

Although the shortest curvature-constrained path satisfying such boundary conditions need not be a CL path, and indeed may not exist at all, our algorithm provides a numerical solution to the Markov-Dubins problem with free terminal direction in a nonpositively curved square complex in the case that the solution is a CL path.

Let us outline the algorithm. The verification of its effectiveness will be given as Corollary 4.2.6.

- Let u and v be the desired initial and terminal position, and let U be the desired initial direction. Let Σ be one of the two oriented curves of constant curvature $a > 0$ tangent to U at u. Let s be the length of an arc of Σ beginning at u and following the orientation of Σ. Let $P(s)$ be the endpoint of this arc, and let $T(s)$ be the tangent to Σ at $P(s)$. (See Figure 4.1.)

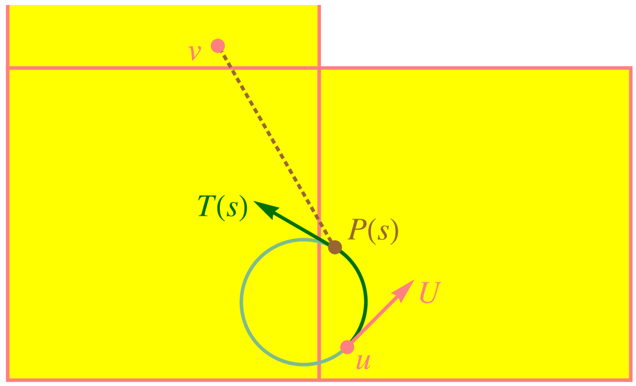

Figure 4.1: Determining a CL path by finding a root of Θ.

- The directed angle $\Theta(s)$ from the tangent $T(s)$ to the geodesic $[P(s)v]$ in X can be numerically computed by using an algorithm due to Chepoi and Maftuleac [11] to determine a decomposition of $[P(s)v]$ into line segments, each contained in a single 2-cell of X,

$$[P(s)v] = [P(s)w_1] * \cdots * [w_m v], \qquad w_k = w_k(s) \in \text{Vert}(X). \tag{4.1}$$

  Then the directed angle from $T(s)$ to $[P(s)v]$ is equal to the directed angle from $T(s)$ to $[P(s)w_1]$.

- If there exists a CL path $\gamma : [a, b] \to X$ from u to v with initial direction U, we must have

$$\Theta(s) = 0$$

  for some $s \in [a, b]$. Since Θ is a continuous, piecewise monotone function (Theorem 4.2.3), we can find all such values of s on a finite segment of Σ by repeatedly applying any algorithm guaranteed to find roots of a continuous function on an interval. Noting that, in some cases, a curve Σ of constant curvature in a nonpositively curved square complex is a proper ray, we must restrict our attention to an initial segment of Σ having finite length.

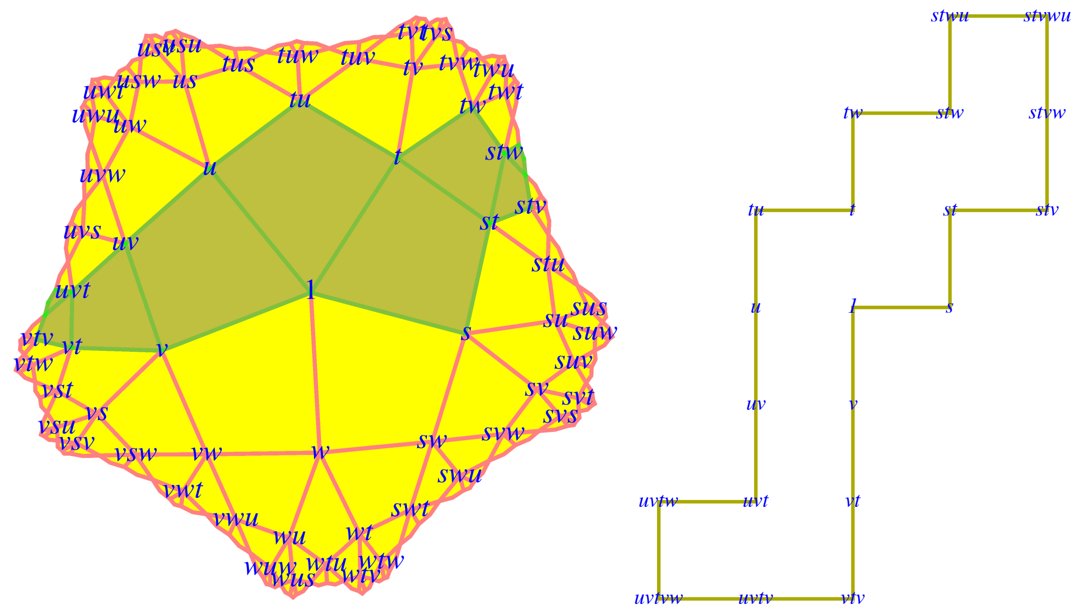

Figure 4.2: (LEFT:) A subcomplex (tan and green) of the space $\mathcal{X}_5$ (yellow and pink) described in Chapter 5, and (RIGHT:) its embedding as a polygon P in $\mathbb{E}^2$.

The fact that a geodesic in X can be decomposed as claimed in (4.1) can be seen as follows. Given the endpoints of a geodesic $[xy]$ to be determined, Chepoi and Maftuleac show (Theorem 4.2.1) how to determine a geodesic $[xy]$ in an NPC square complex between two given endpoints by first constructing an isometry that maps a subcomplex of X containing $[xy]$ onto a planar polygon P isometrically embedded in $\mathbb{E}^2$ (Figure 4.2), and then applying Lee and Preparata's funnel algorithm to the images of x and y in P (Figures 4.3–4.5). Indeed, P is a *monotone polygonal chain*, and consequently can be triangulated so that the dual graph of the triangulation is homeomorphic to an interval. Thus Lemma 4.2.2 immediately implies the existence of a decomposition as in (4.1).

Theorem 4.2.1 ([11])**.** The geodesic between two points u, v of a CAT(0) square complex X lies in a subcomplex K of X that embeds isometrically in $\mathbb{R}^2$ as a chain of monotone polygons. Moreover, K depends only on the choice of 2-cells containing u and v.

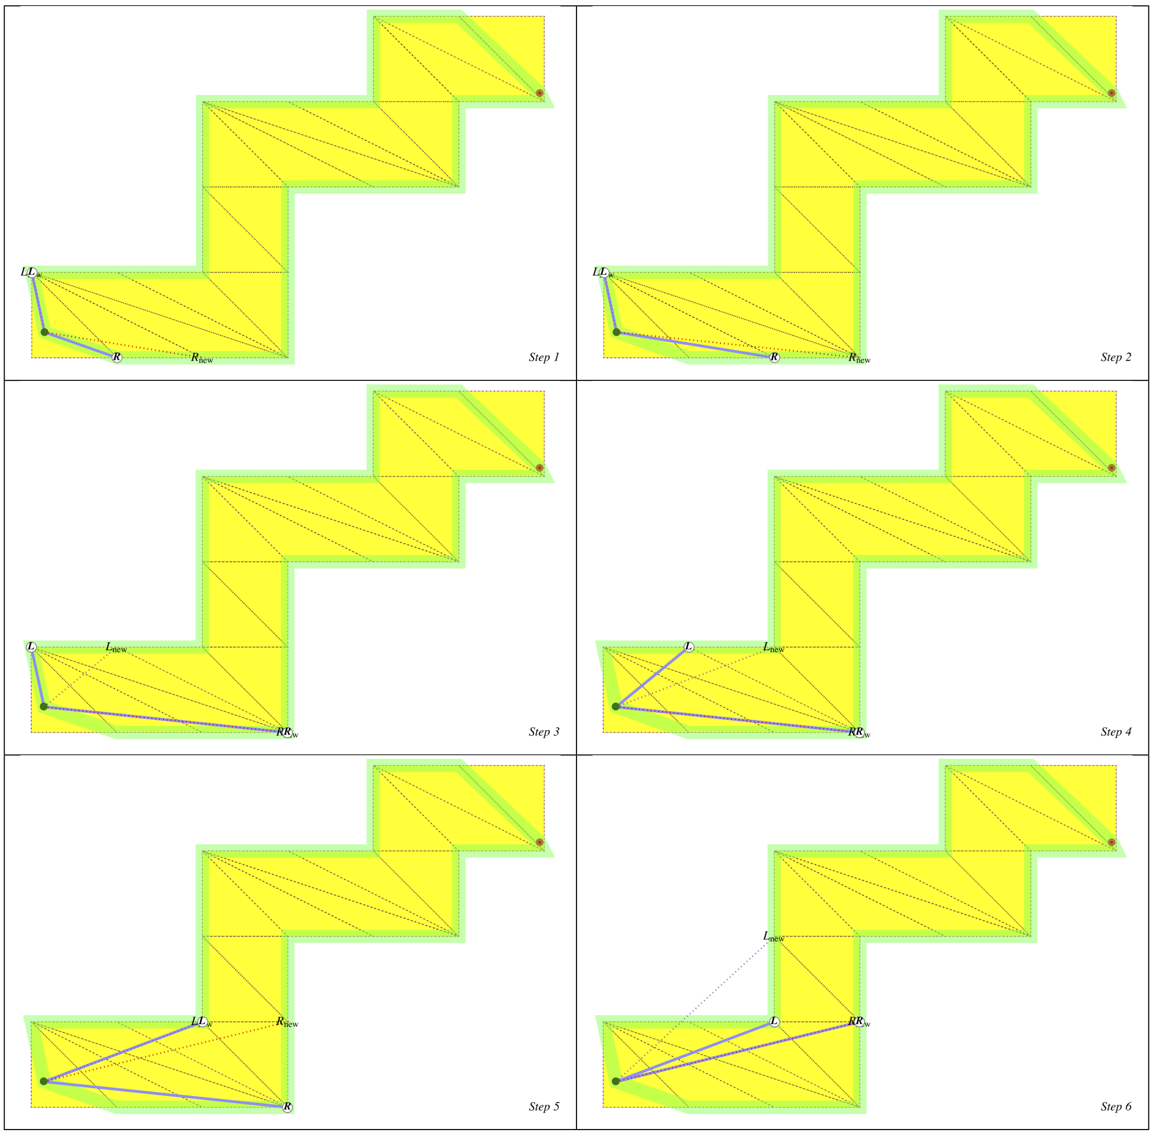

Figure 4.3: An illustration of Lee and Preparata's funnel algorithm for finding the shortest path between two points in a planar polygon.

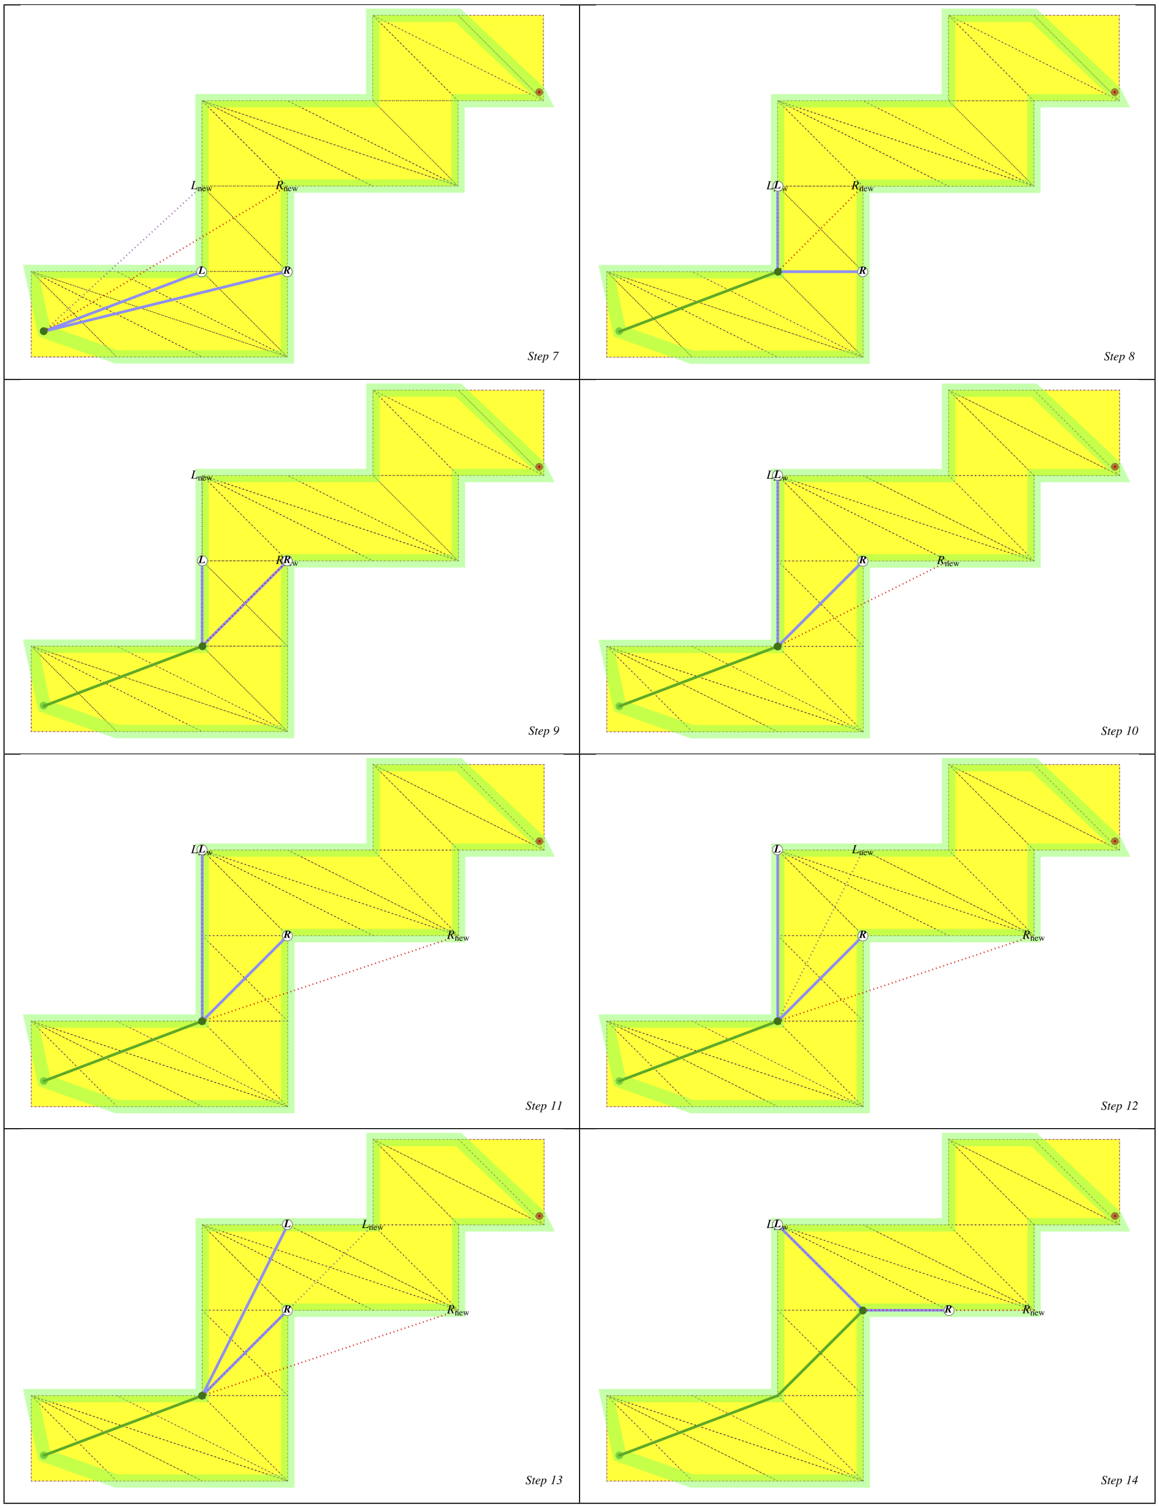

Figure 4.4: Continuation of previous figure.

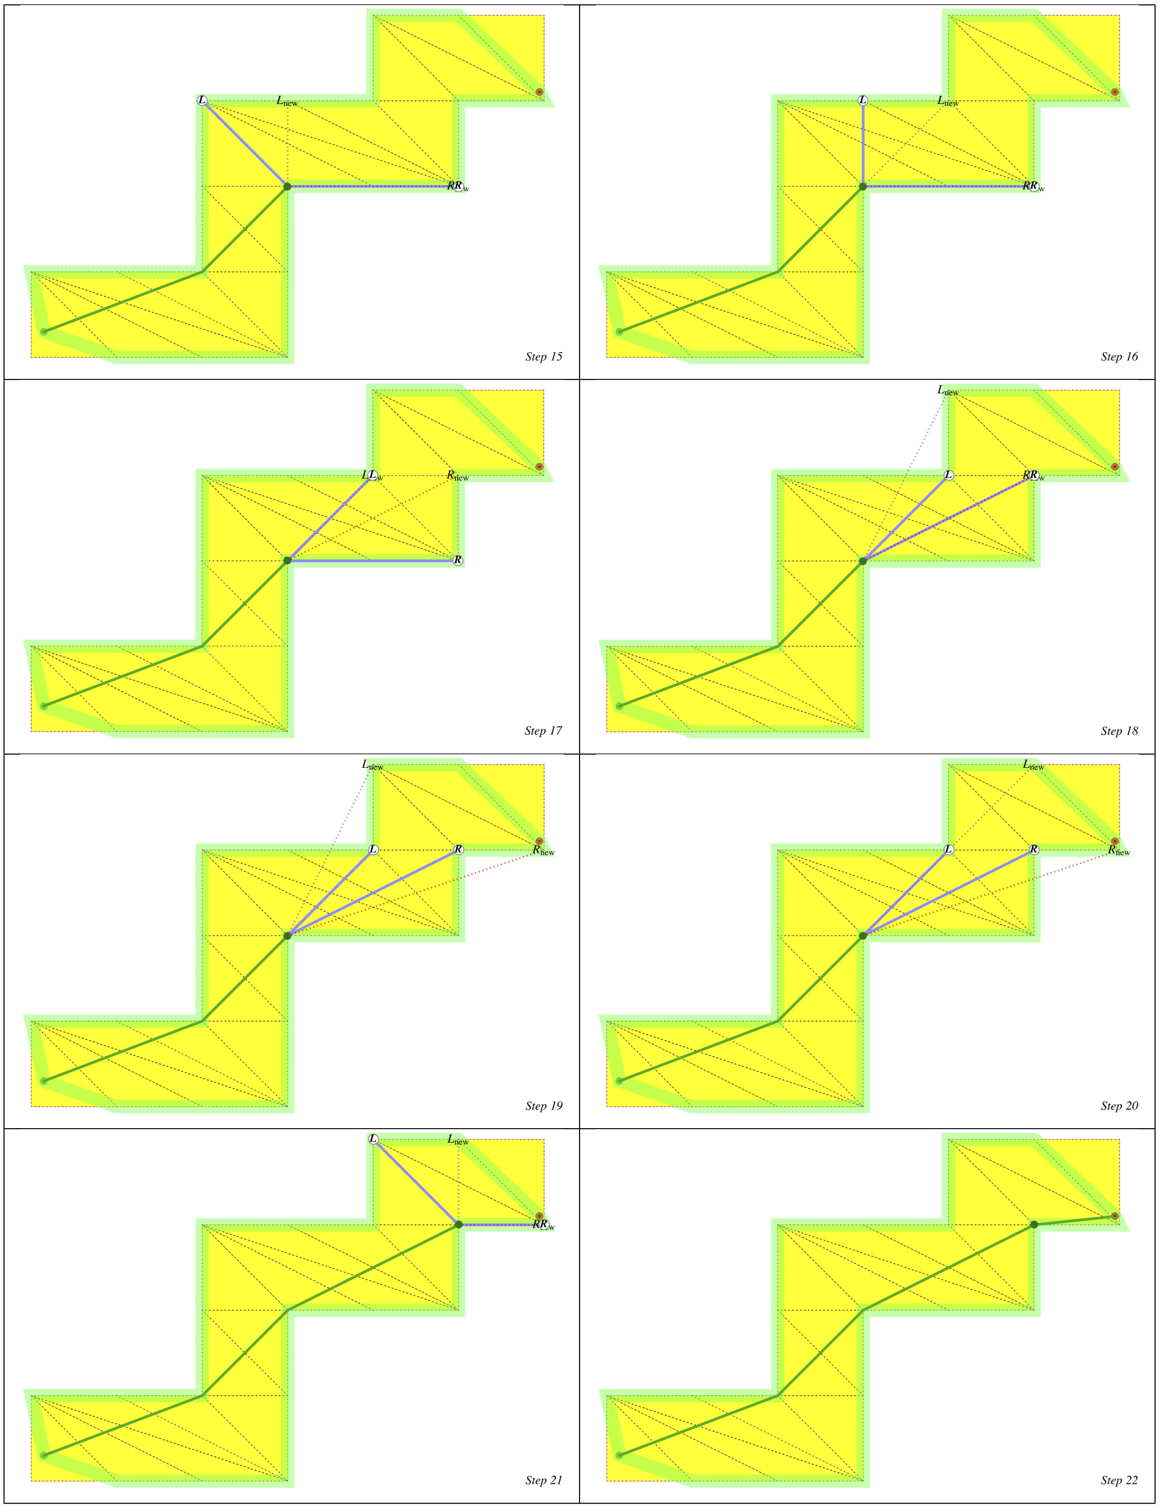

Figure 4.5: Continuation of previous figure.

Lemma 4.2.2 ([25])**.** Let P be a simple polygon, and let u, v be distinct points of P. Let $\mathcal{T}$ be a triangulation of P with $\text{Vert}(\mathcal{T}) \subset \partial P$, and assume that the dual graph of $\mathcal{T}$ is an interval. The geodesic $[uv]$ in P is made up of line segments each of whose endpoints is u, v, or a vertex of $\mathcal{T}$.

A bit more work is needed to verify the elementary fact that the directed angle function described above is continuous and piecewise monotone. (The proof of Theorem 4.2.3 will be momentarily deferred.)

Theorem 4.2.3. Let X be a nonpositively curved square complex and let $\gamma : [a, b] \to X \smallsetminus \text{Vert}(X)$ be a unit-speed path of constant positive curvature. Let $v \in X \smallsetminus \text{Image}(\gamma)$. The directed angle function $\Theta(t)$ from the tangent $\gamma'(t)$ to the geodesic $[\gamma(t)v]$ is a continuous function $[a, b] \to (-2\pi, 2\pi]$. Moreover, there exists a subdivision $a = s_1 < s_2 < \cdots < s_{r+1} = b$ such that $\Theta|_{[s_k, s_{k+1}]}$ is monotone for each k.

From Theorem 4.2.3 we conclude that for any $S > a$, the bisection method can be applied finitely many times to find all zeros of $\Theta|_{[a, S]}$, since $[s_k, s_{k+1}]$ contains exactly one zero of Θ if and only if $\Theta(s_k)\Theta(s_{k+1}) \leqslant 0$. For if $\Theta(s_k)\Theta(s_{k+1}) < 0$, then $\Theta(s_k)$ and $\Theta(s_{k+1})$ have opposite signs, and $[s_k, s_{k+1}]$ contains exactly one zero of Θ by monotonicity and continuity, while if $\Theta(s_k)\Theta(s_{k+1}) > 0$, then $\Theta(s_k)$ and $\Theta(s_{k+1})$ have the same signs, in which case $[s_k, s_{k+1}]$ contains no zeroes of Θ, as Θ is monotone.

Lemma 4.2.4. Let $\sigma : [a, a + 2\pi] \to \mathbb{R}^2$ be a unit-speed parametrization of a circle. Let $v \in \mathbb{R}^2 \smallsetminus \text{Image}\,\sigma$. There exists a continuous monotone function $\theta : [a, a + 2\pi] \to [-2\pi, 2\pi]$ such that $\theta(t)$ is the directed angle from $\sigma'(t)$ to the vector $\sigma(t) - v$, and $\theta(a) \cdot \theta(a + 2\pi) < 0$.

Proof. Let $\texttt{ATan2} : \mathbb{R}^2 \setminus \{(0,0)\} \to (-\pi, \pi]$ be the computer scientist's arctangent function,

$$\texttt{ATan2}(x, y) = \operatorname{Arg}(x + iy) = \begin{cases} \arctan(y/x) & x > 0 \\ \arctan(y/x) + \pi & x < 0,\ y \geqslant 0 \\ \arctan(y/x) - \pi & x < 0,\ y < 0 \\ \pi/2 & x = 0,\ y > 0 \\ -\pi/2 & x = 0,\ y < 0. \end{cases}$$

Suppose σ parametrizes the circle counterclockwise. Then

$$\alpha(t) = \begin{cases} \texttt{ATan2}(\sigma'(t)) & \text{if } \texttt{ATan2}(\sigma'(t)) \geqslant \texttt{ATan2}(\sigma'(a)) \\ \texttt{ATan2}(\sigma'(t)) + 2\pi & \text{otherwise} \end{cases}$$

defines a monotone increasing function $[a, a + 2\pi] \to (-\pi, 3\pi]$ such that $\alpha(t) \pmod{2\pi} \in (-\pi, \pi]$ is the directed angle from the x-axis to the vector $\sigma'(t)$.

Define $\beta : [a, a + 2\pi] \to [0, 2\pi]$ by

$$\beta(t) = \begin{cases} \texttt{ATan2}(\sigma(t) - v) & \text{if } \texttt{ATan2}(\sigma(t) - v) \geqslant 0 \\ \texttt{ATan2}(\sigma(t) - v) + 2\pi & \text{otherwise,} \end{cases}$$

so that $\beta(t)$ is the angle measured counterclockwise from the x-axis to the vector $\sigma(t) - v$.

Let x and y be the two points of the circle through which the line meeting v is tangent to the circle, Recalling that the angle between a tangent vector and a chord of a circle measures one-half the angle of the arc subtended by the chord on the side of the tangent vector, we see that

$$\theta(t) := \alpha(t) - \beta(t)$$

is continuous and monotonically increasing in t, reaching a maximum at one of the two points x, y and a minimum at the other (Figure 4.6).

To ensure that $\theta(a) \cdot \theta(a + 2\pi) < 0$, modify θ by adding 2π if $\alpha(a) - \beta(a) < -2\pi$, or -2π if

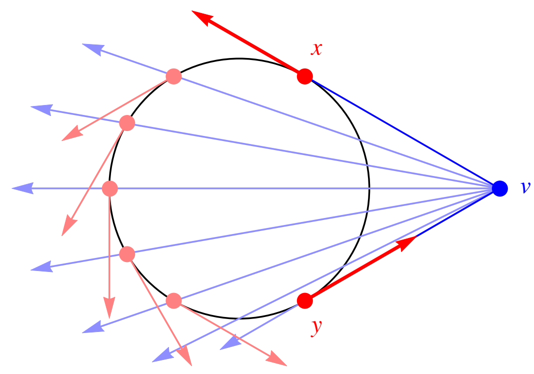

Figure 4.6: $\theta(t)$ is the directed angle between a tangent $\sigma'(t)$ to the circle and $\sigma(t) - v$.

$\alpha(a) - \beta(a) > 0$. □

Lemma 4.2.5. Let $P \subset \mathbb{R}^2$ be a simply connected simple polygon. Let $\sigma : [a, b] \to P$ be a unit-speed parametrization of an arc of a circle. Let $v \in P \setminus \text{Image}(\sigma)$. There exist $n \in \mathbb{N}$, $z_1, \ldots, z_{n-1} \in \{v\} \cup \text{Vert}(P)$, and $a = s_1 < s_2 < \cdots < s_n = b$ such that for $k \in \{1, \ldots, n\}$, the directed angle between $\sigma'(t)$ and the geodesic in P from $\sigma(t)$ to v is equal to the directed angle between $\sigma'(t)$ and the geodesic in $\mathbb{R}^2$ from $\sigma(t)$ to z_k for $t \in [s_k, s_{k+1}]$.

Proof. Consider the application of the Lee-Preparata funnel algorithm to determine the geodesic in P with initial point v and terminal point $\sigma(t)$. In particular, recall that a "funnel" is the subpolygon of P obtained by cutting off each ray in the convex cone spanned by the rays $\overrightarrow{z\ell}$ and $\overrightarrow{zr}$ where it meets ∂P, with $z \in \{v\} \cup \text{Vert}(P)$ and $\ell, r \in \text{Vert}(P)$ distinct. (See Figures 4.3–4.5.) At each step, the algorithm halts if and only if the current funnel F, say with apex z, meets σ in a nondegenerate arc containing $\sigma(t)$. (Note that, if F meets σ in one or two isolated points, then F cannot be the final funnel, for in that case the intersection of $[z\sigma(t)] \setminus \{z, \sigma(t)\}$ and ∂P must contain some vertex of P which must be processed by the algorithm before halting.) The algorithm outputs the desired geodesic $[v\sigma(t)]$ as a chain of line segments, the last of which is $[z\sigma(t)]$.

We now show that the same z can be used for all $w \in F \cap \sigma$: that is, if F is the final funnel computed by the algorithm for some geodesic $[v\sigma(t)]$, then for each point $w \in A := F \cap \sigma$, the apex z of the final funnel computed by the algorithm for $[vw]$ depends only on the choice of A, not on the choice of $w \in A$. Assume for a contradiction that there exist w, $w' \in A$ such that the final funnel F computed for $[vw]$ has apex $z \in \{v\} \cup \mathrm{Vert}(P)$, and the final funnel F' computed for $[vw']$ has apex $z' \in \{v\} \cup \mathrm{Vert}(P)$. Since F is the union of Euclidean shortest paths from z to points of A, $[zw]$ is a geodesic in P. By the same reasoning applied to F', $[z'w]$ is a geodesic in P. Since P is uniquely geodesic, we must have $[zw] \subset [z'w] \subset [vw]$ or $[z'w] \subset [zw] \subset [vw]$, otherwise there would exist two distinct geodesics from v to w, one containing z and not z', and the other containing z' and not z. It follows that $z = z'$, since each point of $[vw] \cap \big(\{v\} \cup \mathrm{Vert}(P)\big)$ is the apex of some funnel that is processed by the algorithm when computing $[vw]$, and only one point of $\{v\} \cup \mathrm{Vert}(P)$ can be the apex of the final such funnel. Thus, for each $A = F \cap \sigma$, where F is a terminal funnel with apex z, and for every $w \in A$, the last in the chain of line segments making up the geodesic $[vw]$ in P is $[zw]$, where z depends only on the choice of A.

Let $\mathcal{F}$ be the set of final funnels computed by the algorithm for a geodesic $[vw]$ with $w \in \sigma$. Since there are only finitely many $\leqslant \binom{\#\mathrm{Vert}(P)+1}{3}$ possible funnels when computing $[vw]$ as w ranges over σ, the partition of σ determined by the endpoints of the (at most two) components of $\sigma \cap F$, as F ranges over $\mathcal{F}$, has only finitely many breakpoints $a = s_1 < s_2 < \cdots < s_n = b$. For $k \in \{1, \ldots, n-1\}$, let z_k be the apex of the final funnel F_k computed by the algorithm for a geodesic $[v\sigma(t)]$ whose intersection with σ is $\sigma\big([s_k, s_{k+1}]\big)$. The result now follows. $\square$

Theorem 4.2.3 can now be established.

Proof of Theorem 4.2.3. Given $\gamma : [a, \infty) \to X$, identify an initial segment $\widehat{\gamma} = \gamma\big|_{[a,\,b]}$ and a square path $(Q_k)_{k=1}^{n}$ with breakpoints $(t_k)_{k=1}^{n+1}$ for $\widehat{\gamma}$. Write $\gamma_k = \gamma\big|_{[t_k,\,t_{k+1}]}$. After applying the Chepoi-Maftuleac embedding, Lemma 4.2.5 can be applied to partition the domain of each γ_k,

$$t_k = s_{k,1} < \cdots < s_{k,m_k} = t_{k+1},$$

so that each of the directed angle functions $\Theta_{k,j} : [s_{k,j}, s_{k,j+1}] \to [-2\pi, 2\pi]$ for a restriction γ_k to $[s_{k,j}, s_{k,j+1}]$ is continuous and monotone. Reindex the $s_{k,j}$, ordered lexicographically by their indices (k, j), as $s_1, s_2, \ldots, s_{r+1}$, and reindex the corresponding directed angle functions $\Theta_{k,j}$ accordingly. Since these functions agree at the endpoints of their domain intervals, the concatenation

$$\Theta = \Theta_1 * \cdots * \Theta_r$$

is continuous. □

It is now trivial to verify a sufficient condition for the existence of a CL path with prescribed endpoints and initial direction. Moreover, if this condition holds, and if there exists such a path of length $\leqslant \mu$ for some $\mu > 0$, then our algorithm is guaranteed to find the shortest such path in finitely many steps.

Corollary 4.2.6. Let $\gamma : [a, \infty) \to X \setminus \mathrm{Vert}(X)$ be a ray of constant positive curvature in a nonpositively curved square complex X. Let $v \in X \setminus \mathrm{Image}(\gamma)$. For each $s > 0$, let $\Theta(s)$ be the directed angle from $\gamma'(s)$ to $[\gamma(s)v]$ at $\gamma(s)$. There exists a CL path in X with initial point $\gamma(a)$, terminal point v, and initial direction $\gamma'(a)$ if and only if there exist $s > a$ and $t > a$ such that $\Theta(s)\Theta(t) \leqslant 0$. Furthermore, for fixed $\mu > 0$, all such CL paths of length $\leqslant \mu$ can be computed numerically in finite time.

Proof. We prove sufficiency; the reverse implication follows from the definition of a CL path, since we must have $\Theta(s) = 0$ at the breakpoint $\gamma(s)$. If $\Theta(s)\Theta(t) \leqslant 0$, we can apply the bisection method finitely many times to find all zeros of $\Theta|_{[s,t]}$. Since all CL paths have zero turning angle at their breakpoints by definition, this process determines all CL paths as required. □

Chapter 5

The *d*-plane

Throughout this chapter, we will focus on a family of square complexes we call *d-planes* ($d \in \{4, 5, 6, \dots\}$). They are homeomorphic to the Euclidean plane, but not isometric for any $d \geqslant 5$.

Definition 5.0.7. A graph is ***d*-regular** if every vertex has degree d, i.e. is incident with exactly d edges.

Definition 5.0.8. The ***d*-plane** $\mathcal{X}_d$ ($d \geqslant 4$) is a simply connected surface without boundary that is a piecewise Euclidean square complex with a d-regular graph as its 1-skeleton.

Thus the neighborhood of each vertex v of the d-plane consists of d copies of the product space $I \times I$ arranged cyclically around v. In §5.1, we verify that the d-plane is the Davis complex of a right-angled Coxeter system. Figure 2.2 illustrates the simplicial and cubical cell structures of the neighborhood of a vertex in $\mathcal{X}_5$ when the latter is regarded as a Davis complex. We then exhibit a reconfigurable system whose state space has the 5-plane as its universal cover. Finally, in §5.2, we examine the shape of curves of constant curvature in the d-plane.

5.1 Basic properties

Lemma 5.1.1. Let $d \geqslant 4$, and let W_d be the group with generating set $S_d = \{s_1, \ldots, s_d\}$ and relators

$$s_1^2 = \cdots = s_d^2 = (s_1 s_2)^2 = (s_2 s_3)^2 = \cdots (s_d s_1)^2 = 1.$$

(a) $X(W_d, S_d)$ is a topological 2-manifold without boundary.

(b) The 1-skeleton $X^{(1)}$ of $X(W_d, S_d)$ is d-regular.

(c) The link of each vertex in $X(W_d, S_d)$ is a d-gon.

(d) $X(W_d, S_d)$ is CAT(0), hence special.

(e) $\mathcal{X}_d = X(W_d, S_d)$.

Proof. We will prove (a) directly, verifying (b) and (c) along the way.

We know by Lemma 2.3.22 that $X(W_d, S_d)$ is a square complex. To show that $X(W_d, S_d)$ is locally Euclidean, it suffices to show that each point of $X^{(1)}$ has a neighborhood homeomorphic to a disk, since all other points of $X(W_d, S_d)$ lie in the interiors of 2-cells.

The cells of $X(W_d, S_d)$ are the geometric realizations of the posets under subset inclusion

$$\alpha = (W\mathcal{S})_{\leqslant x\langle T\rangle} = \{x'\langle T'\rangle \subseteq x\langle T\rangle : x' \in W_d, T' \subseteq S, \#\langle T'\rangle < \infty\}$$

for $x \in W_d$ and $T \subset S$ such that $\#\langle T\rangle < \infty$ (see Definition 2.3.21). Recall that $\dim \alpha = \#T$, and write

$$e^{\#T}(x\langle T\rangle) = |\alpha|.$$

Suppose $|\alpha|$ is a vertex. We show that star $|\alpha|$ is homeomorphic to a disk. Since $\dim |\alpha| = 0$, we have $|\alpha| = e^0(\{x\})$ for some $x \in W_d$. Then $|\alpha|$ is incident exactly with the edges $e^1(x\langle s\rangle)$, $s \in S_d$—from which (b) follows—and with the 2-cells $e^2(x\langle s_i, s_j\rangle)$, $(s_i s_j)^2 = 1$.

Two such 2-cells meet in an edge if they have a common generator, and are disjoint if not:

$$e^2(x\langle s_i, s_j\rangle) \cap e^2(x\langle s_k, s_\ell\rangle) = \begin{cases} e^1(x\langle s_j\rangle) & \text{if } j = k,\, i \neq \ell \\ \varnothing & \text{if } \{i, j\} \cap \{k, \ell\} = \varnothing. \end{cases} \qquad (*)$$

Thus the 2-cells incident with $|\alpha|$ are arranged cyclically, giving (c), and their union

$$\mathcal{N}_{\{x\}} := \overline{\text{star}}\,|\alpha|$$

is homeomorphic to a disk.

Suppose $|\alpha|$ is an edge. Then $|\alpha| = e^1(x\langle s_j\rangle)$ for some $s_j \in S_d$, and is the intersection of exactly two 2-cells, as in (*). The union of these 2-cells is a neighborhood for any interior point of $|\alpha|$ which is homeomorphic to a disk. We conclude that every point of the 1-skeleton $X^{(1)}$ has a neighborhood homeomorphic to a disk. Since $X(W_d, S_d)$ is the union of neighborhoods $\mathcal{N}_{\{x\}}$ ($x \in W_d$) as above, and $\partial\mathcal{N}_{\{x\}} \subset X^{(1)}$, it follows that $X(W_d, S_d)$ has no boundary. Now (a) has been shown.

Part (d) is simply an application of Theorems 2.3.23 and 2.4.17. In particular, $X(W_d, S_d)$ is simply connected, and (e) now follows. □

5.1.1 The 5-plane is the universal cover of the state complex for a reconfigurable system

In this subsection, we exhibit a reconfigurable system whose state complex is covered by the 5-plane.

The workspace graph $\mathcal{W}$ of this system is a 5-gon with vertices labeled distinctly. The generators are transpositions of adjacent labels. We might imagine that these transpositions are carried out by an industrial robot, say, by a rotating pincer-shaped arm able to pick up and swap distinctly labeled checkers initially placed on the vertices. (See Figure 5.1.)

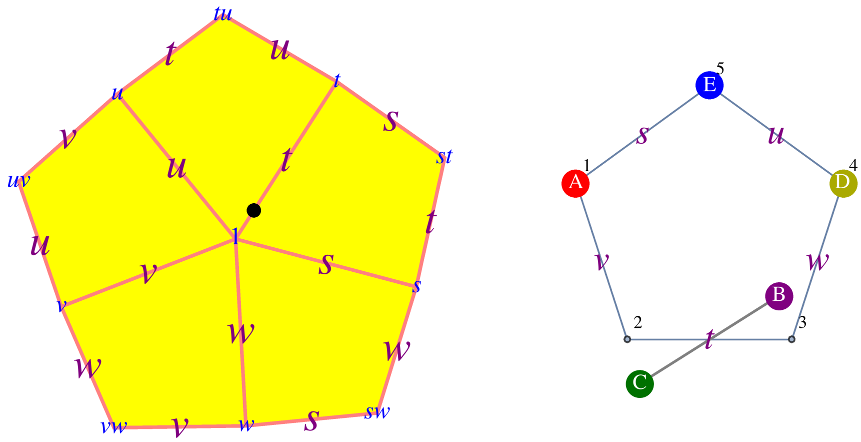

Figure 5.1: (RIGHT:) A transposition of adjacent labels in progress, and (LEFT:) a point in $\mathcal{X}_5$ corresponding to the system's state.

Observe that generators commute if and only if the corresponding edges are disjoint. Since each edge is disjoint from two edges, and no set of three edges is pairwise disjoint, we find that each vertex in the state complex is incident with exactly five squares, arranged cyclically.

The state complex is a closed surface of genus 16, as can be verified by an Euler characteristic argument, and is orientable, as can be checked directly [17]. Its universal cover is the 5-plane $\mathcal{X}_5$ (Figure 5.2), that is, the Davis complex of the right-angled Coxeter system with generators s, t, u, v, w and relators

$$(st)^2 = (tu)^2 = (uv)^2 = (vw)^2 = (ws)^2 = 1.$$

The reconfigurable system we have described provides a concrete example of the practical applications of the present work. Suppose the system begins in a given state, and a goal state is given, but a new state is prescribed before the original goal state has been achieved. At the instant when the new goal state is prescribed, we have a Markov-Dubins problem with free terminal direction.

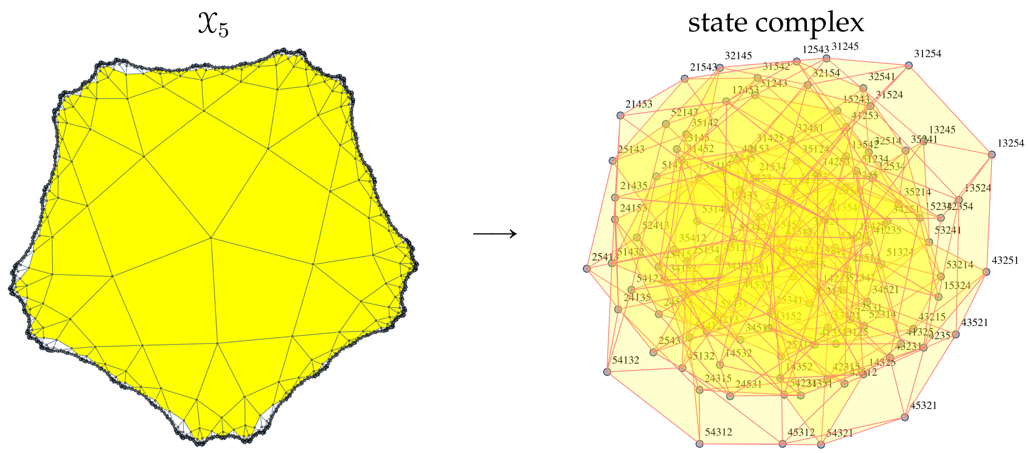

Figure 5.2: The shaded region in $\mathcal{X}_5$ is a minimal region sent by the covering map onto the state complex.

5.2 Curves of constant curvature in the d-plane

5.2.1 Characterization and numerical experiments

As with Dubins' 1957 characterization of the solutions to the Markov-Dubins problem with prescribed terminal direction in $X = \mathbb{R}^2$, the solutions to the problem in the d-plane are made up of finitely many geodesic segments and arcs of constant curvature a. We are thus led to consider the problem of characterizing curves of constant curvature in $\mathcal{X}_d$.

One consequence of the nonpositive curvature of $\mathcal{X}_d$ is that, when $n \geqslant 5$, the circumference of a arc of sufficiently large constant curvature $a \geqslant 1$ centered at a vertex is $(\pi d/2)R > 2\pi R$, where $R = 1/a$. This situation recalls that of the hyperbolic plane $\mathbb{H}^2$ (that is, the Riemann surface of constant curvature -1), in which the circumference of a circle of radius R is $2\pi \sinh(R) > 2\pi R$. But the behavior of curves of constant curvature exhibits not only similarities, but also dramatic differences between the geometries of the square complex $\mathcal{X}_d$ and the Riemannian 2-manifolds of constant nonpositive curvature. A ray $\gamma : [0, \infty) \to \mathbb{H}^2$ of constant curvature $a \geqslant 0$ has image either an embedded circle, a hypercycle, or a horocycle, depending on whether $a > 1$, $a < 1$, or $a = 1$ [10]. In any case, the image of γ is a smooth 1-manifold, as is the case for the images of constant-curvature rays in $\mathbb{R}^2$.

In $\mathcal{X}_d$, on the other hand, a ray of constant curvature may self-intersect. Furthermore, numerical experiments show that the behavior of a curve γ of constant curvature a in $\mathcal{X}_d$ is not a function of a, but depends instead on the combinatorics of the sequence of square cells γ visits. Such a curve γ may be a proper topological ray that attains an arbitrarily large distance from its initial point (Figure 5.3), a rose curve made up of $M = \text{lcm}\{4, d\}$ circular arcs (Figure 5.10), or an embedded (metric) circle. It will be convenient to examine only those rays which do not meet Vert($\mathcal{X}_d$), because an arc in $\mathcal{X}_d$ ($d \geqslant 5$) which is smooth at breakpoints (Definition 3.3.8) may have any of infinitely many exit directions as it passes through a breakpoint. Furthermore, since the shape of a ray of constant curvature depends in part on its initial direction, the classification of a ray of constant curvature which pass through vertices of $\mathcal{X}_d$ is not a well-posed problem unless an exit direction is specified at each vertex it meets. The restriction of our focus to rays in $\mathcal{X}_d^*$ is not a significant loss of generality, however, since any arc in $\mathcal{X}_d$ which meets Vert($\mathcal{X}_d$) can be decomposed into subarcs, whose interiors do not.

5.2.2 Annulus Condition

In the square complex $\mathbb{E}^2$, non-intersecting hyperplanes are parallel lines separated by a distance of at least 1. Consequently, in order to prove that a ray $\gamma : [t_1, \infty) \to \mathbb{E}^2$ with square path $\mathcal{Q} = (Q_n)_{n=1}^{\infty}$ is proper, it would be enough to show that there exist subpaths $\mathcal{Q}[a_k, b_k] = \bigcup_{i=a_k}^{b_k} Q_i$ of $\mathcal{Q}$ and halfspaces H_k^+ respectively bounded by hyperplanes H_k of $\mathbb{E}^2$ such that $H_1^+ \supset H_2^+ \supset H_3^+ \supset \cdots$ and $\mathcal{Q}[a_{k+1}, b_{k+1}] \subset H_k^+$ for all $k \in \mathbb{N}$ (Figure 5.4).

A similar argument can be used to prove that a ray in a generic square complex X is proper if one can construct a sequence of hyperplanes H_k bounding sequences of nested halfspaces as in the previous subsection (see Lemma 2.5.4). When X is the square complex

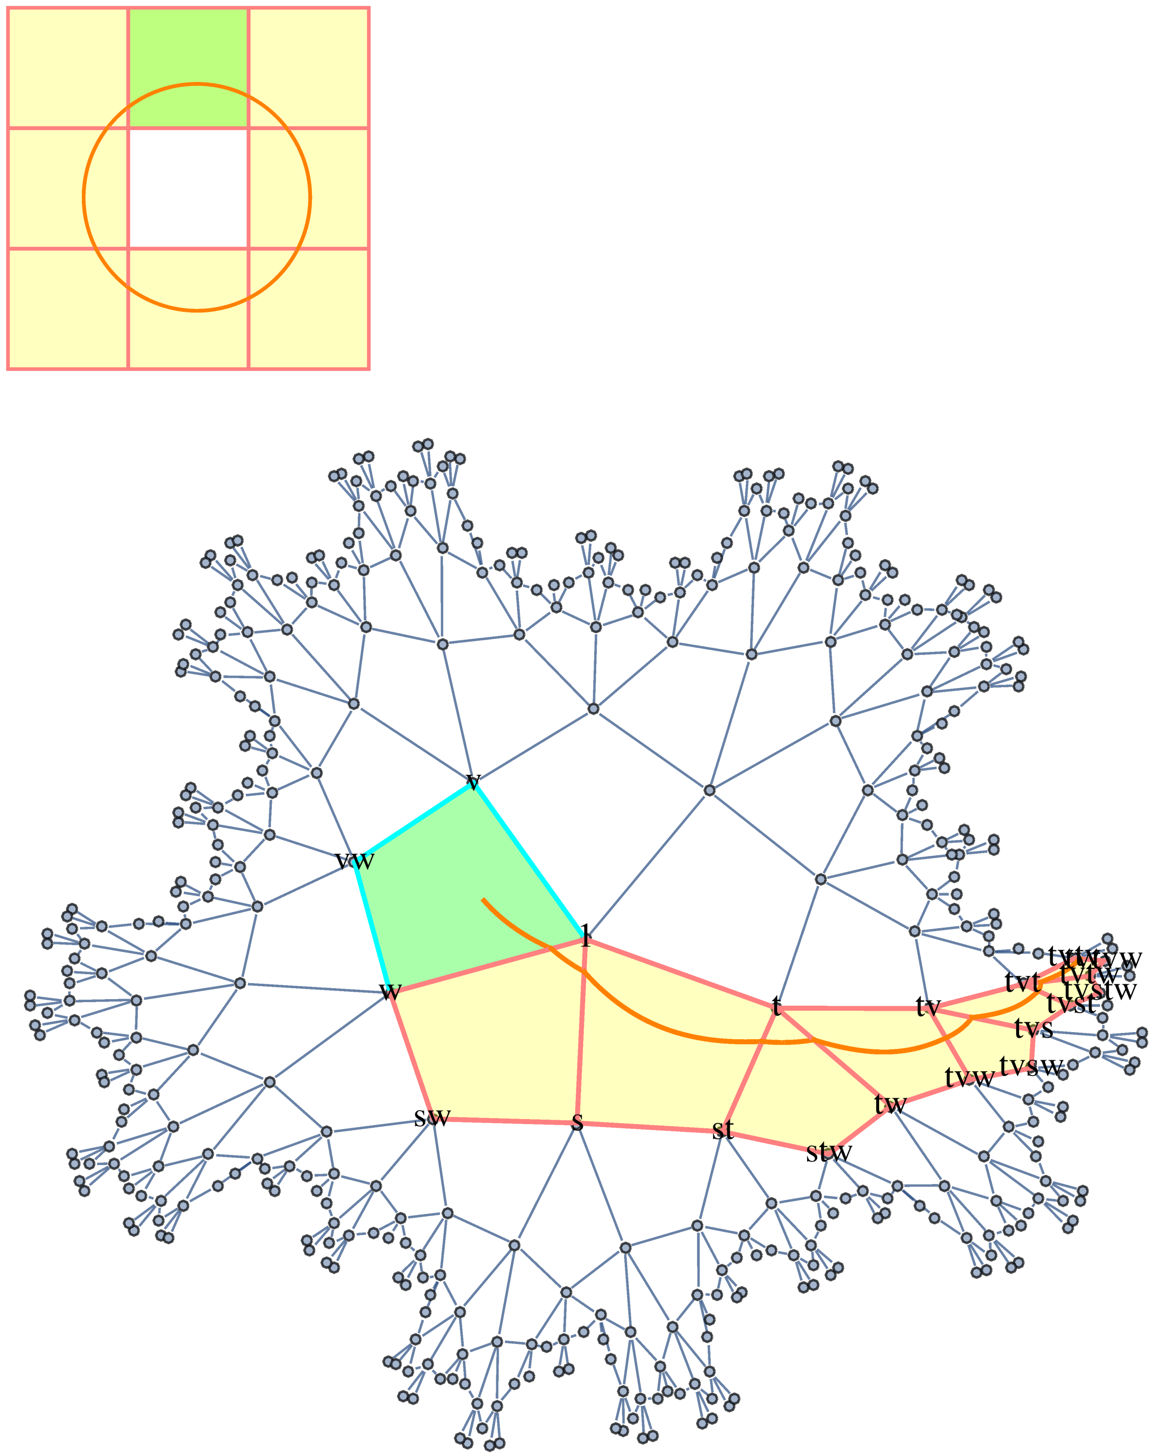

Figure 5.3: BOTTOM: A proper ray γ of constant curvature in $\mathcal{X}_5^*$. TOP: The corresponding annular carrier in $\mathbb{E}^2$.

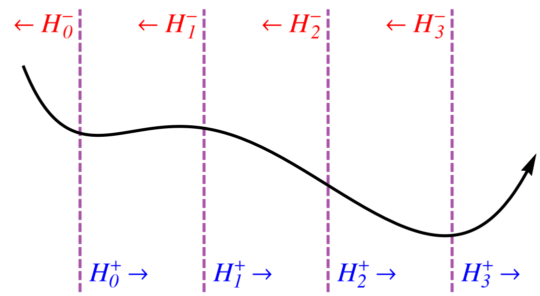

Figure 5.4: Showing a ray in $\mathbb{E}^2$ is proper with nested halfspaces.

associated with a right-angled Coxeter system (W, S), we can establish the existence of such a sequence. Indeed, if two hyperplanes H and H' osculate (Definition 2.4.9), then H and H' behave analogously to hyperparallel lines in $\mathbb{H}^2$, and have distance 1 between them (Lemma 2.4.11).

To produce the desired sequence of nested halfspaces for a ray γ of constant curvature $a > 0$ in $\mathcal{X}_d$, we proceed as follows. (That these steps can in fact be carried out will be verified in the proofs of this subsection.)

- Let $\mathcal{Q}$ be a square path for γ. Let $\varphi : \mathcal{U}(\mathcal{Q}) \to \mathbb{E}^2$ be a folding map, where $\mathcal{U}(\mathcal{Q})$ is the unfolding of $\mathcal{Q}$, and let Σ be the image under φ of the lift of γ to $\mathcal{U}(\mathcal{Q})$. (See Figure 2.7.) We assume in the current subsection that Carrier Σ is an annulus, i.e. homeomorphic to $S^1 \times I$.
- Taking $k = 1$, construct two stacks $\breve{\mathcal{H}}_k$ and $\breve{\mathcal{H}}_{k+1}$ of hyperplanes of $\mathbb{E}^2$ such that
  - (i) $(\breve{Q}_n)_{n=1}^{n_k+1}$ is properly segmented with respect to $\breve{\mathcal{H}}_k$ for some $n_k > 1$,
  - (ii) $(\breve{Q}_n)_{n=n_k-1}^{n_{k+1}+1}$ is properly segmented with respect to $\breve{\mathcal{H}}_{k+1}$ for some $n_{k+1} > n_k + 1$, and
  - (iii) the hyperplanes of $\breve{\mathcal{H}}_{k+1}$ are perpendicular to those of $\breve{\mathcal{H}}_k$

  (see the proof of Theorem 5.2.7.) The hyperplanes of $\breve{\mathcal{H}}_k$ and $\breve{\mathcal{H}}_{k+1}$ lift to hyperplanes in $\mathcal{X}_d$ (Theorem 5.2.2) which form a stack with respect to which $(Q_n)_{n=1}^{n_{k+1}+1}$ is prop-

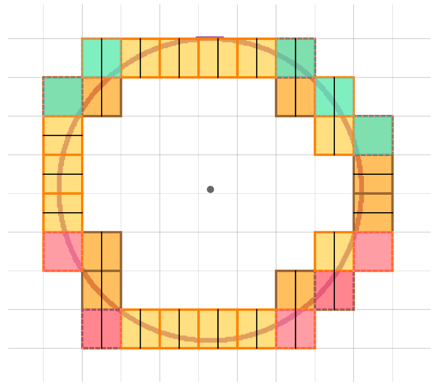

Figure 5.5: The image of a folding $\varphi : \mathcal{U}(\mathcal{Q}) \to \mathbb{E}^2$ that satisfies the Annulus Condition.

erly segmented (Elbow Lemma, 5.2.4).

- Increment k and repeat the previous step, constructing a stack $\breve{\mathcal{H}}_k$ in $\mathbb{E}^2$ such that (ii) and (iii) are satsified. Repeat this until every square of Carrier Σ meets the carrier of a hyperplane of one of the stacks so constructed, with the hyperplanes of successive stacks being perpendicular. This can be done with exactly four stacks repeated periodically, $\breve{\mathcal{H}}_5 = \breve{\mathcal{H}}_1, \breve{\mathcal{H}}_6 = \breve{\mathcal{H}}_2, \ldots$.
- By induction, $(Q_n)_{n=1}^{\infty}$ is properly segmented with respect to lifts of the hyperplanes of $\breve{\mathcal{H}}_1, \breve{\mathcal{H}}_2, \breve{\mathcal{H}}_3, \ldots$ to $\mathcal{X}_d$ (Lemma 5.2.5).

The following conventions will be in effect for the remainder of this section.

Notation. Given a square path $\mathcal{Q} = (Q_n)_{n=1}^{\infty}$ in $\mathcal{X}_d$, a folding map $\varphi : \mathcal{U}(\mathcal{Q}) \to \mathbb{E}^2$, and the natural folding $\psi : \mathcal{U}(\mathcal{Q}) \to \mathcal{X}_d$ (Definition 2.5.13), we write

$$\breve{Q}_n = \varphi(\hat{Q}_n) \subset \mathbb{E}^2, \qquad \chi_n = \psi \circ \left(\varphi|_{\hat{Q}_n}\right)^{-1}.$$

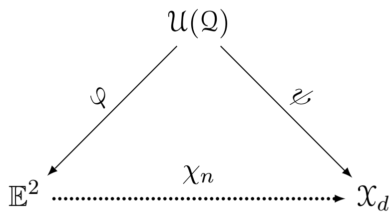

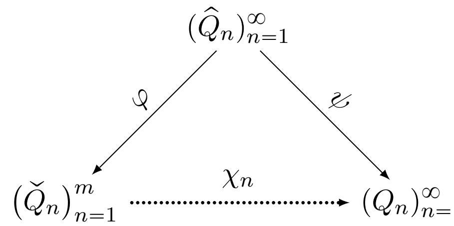

A square path $\mathcal{Q}$ in $\mathcal{X}_d$ can be transferred to the Euclidean plane by constructing its unfolding complex $\mathcal{U}(\mathcal{Q})$ (Definition 2.5.11) and sending it into $\mathbb{E}^2$ via the natural folding, a process which may be visualized as a cellular analogue of analytic continuation. In the next lemma, we reverse the direction of transfer, and show how a sequence of nested halfspaces in $\mathbb{E}^2$ can be used to determine a sequence of nested halfspaces in $\mathcal{X}_d$.

Though our goal is to use these nested halfspaces to prove that certain rays are proper, we will work in terms of square paths for the rays in question. We therefore do not directly refer to nested halfspaces, but instead carry out our discussion in terms of properly segmented square paths.

In Theorem 5.2.2, we establish a technique which uses scaffolds to transfer finite stacks of hyperplanes from $\mathbb{E}^2$ to stacks in $\mathcal{X}_d$, and we show that, for a finite square path $\mathcal{Q}$ to be properly segmented in $\mathcal{X}_d$, it is enough to construct a stack in $\mathbb{E}^2$ with respect to which its counterpart under a folding map $\mathcal{U}(\mathcal{Q}) \to \mathbb{E}^2$ is properly segmented. But first, we make a simple observation.

Lemma 5.2.1. Let D be an edge in a nonpositively curved square complex. If E_1, E_2 are adjacent edges of X such that E_1 and D are adjacent sides of some square Q_1 of X, and E_2 and D are adjacent sides of some square Q_2 of X, then the hyperplanes $H(E_1)$ and $H(E_2)$ osculate. (See Figure 5.6.)

Proof. Let v the vertex where E_1 and E_2 meet. If there were some 2-cell Q in X with adjacent sides E_1 and E_2, then the graph link(v) would contain a cycle of three edges, one for each of the 2-cells Q_1, Q_2, Q, and hence would not be flag, in violation of the definition

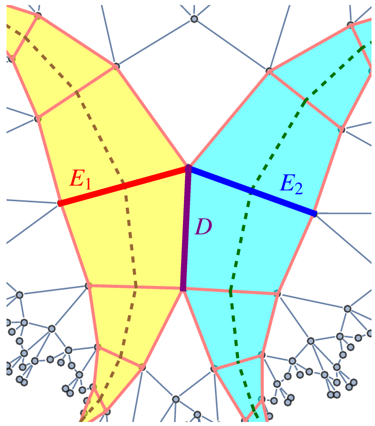

Figure 5.6: Illustration for Lemma 5.2.1.

of a nonpositively curved cube complex. Therefore, the adjacent edges E_1 and E_2 are not contained in any single 2-cell in X, so the hyperplanes $H(E_1)$ and $H(E_2)$ osculate. □

Theorem 5.2.2 (Transfer of a finite stack). Let $\mathcal{Q} := (Q_n)_{n=1}^m$ be an edgewise square path in a special square complex X. Let $\varphi : \mathcal{U}(\mathcal{Q}) \to \mathbb{E}^2$ be a folding map. Suppose the square path $(\breve{Q}_n)_{n=1}^m$ is properly segmented with respect to some stack $\breve{\mathcal{H}} = (\breve{H}_k)_{k=1}^p$ of hyperplanes of $\mathbb{E}^2$. Then there exist integers

$$1 = a_1 \leqslant b_1 < a_2 \leqslant b_2 < \cdots \leqslant b_p = m$$

and a scaffold $\big((E_k)_{k=1}^{p-1}, (F_k)_{k=1}^{p-1}\big)$ for $\breve{\mathcal{Q}}$ with respect to $\breve{\mathcal{H}}$ such that $\mathcal{H} := (H_k)_{k=1}^p$ is a stack, where

$$\begin{aligned} H_1 &:= H(\chi_{b_1}(E_1)), \\ H_2 &:= H(\chi_{a_2}(F_1)) = H(\chi_{b_2}(E_2)), \\ &\vdots \\ H_{p-1} &:= H(\chi_{a_{p-1}}(F_{p-2})) = H(\chi_{b_{p-1}}(E_{p-1})), \\ H_p &:= H(\chi_{a_p}(F_{p-1})), \end{aligned}$$

and $\mathcal{Q}$ is properly segmented with respect to $\mathcal{H}$.

Proof. Let

$$1 = a_1 \leqslant b_1 < a_2 \leqslant b_2 < \cdots \leqslant b_p = m$$

be as in Definition 2.5.7. As in the proof of Lemma 2.5.16, we have $a_{k+1} = b_k + 1$ for each $k \in \{1, \ldots, p-1\}$, and there exist $E_k \in \text{Edges}(Q_{b_k})$ and $F_k \in \text{Edges}(Q_{a_{k+1}})$ for $k \in \{1, \ldots, p-1\}$ such that

$$\left((E_k)_{k=1}^{p-1}, (F_k)_{k=1}^{p-1}\right)$$

is a scaffold for $\breve{\mathcal{Q}}$ with respect to $\breve{\mathcal{H}}$. Note that, for each k, E_k and F_k are adjacent, $H(E_k) = \breve{H}_k$, and $H(F_k) = \breve{H}_{k+1}$ by Definition 2.5.15.

To see that $\mathcal{H}$ is a stack, we verify that $H(\chi_{b_k}(E_k))$ and $H(\chi_{a_{k+1}}(F_k))$ osculate for each $k \in \{1, \ldots, p-1\}$. First observe that

$$\chi_{b_k}(\breve{Q}_{b_k}) = Q_{b_k}, \qquad \chi_{a_{k+1}}(\breve{Q}_{a_{k+1}}) = Q_{a_{k+1}},$$

and χ_{b_k} and $\chi_{a_{k+1}}$ agree on

$$\breve{D} := \breve{Q}_{b_k} \cap \breve{Q}_{a_{k+1}}.$$

Then the edges $\chi_{b_k}(E_k)$ and $\chi_{a_{k+1}}(F_k)$ are adjacent, as $E_k \cap F_k \subset \breve{Q}_{b_k} \cap \breve{Q}_{a_{k+1}}$. Since each of χ_{b_k} and $\chi_{a_{k+1}}$ map adjacent sides to adjacent sides, we know

$$D := \chi_{b_k}(\breve{D}) = \chi_{a_{k+1}}(\breve{D})$$

is adjacent to $\chi_{b_k}(E_k)$ in Q_{b_k}, and D is adjacent to $\chi_{a_{k+1}}(F_k)$ in $Q_{a_{k+1}}$. Now Lemma 5.2.1 can be applied to the edges D, $\chi_{b_k}(E_k)$, and $\chi_{a_{k+1}}(F_k)$, and we conclude that $H(\chi_{b_k}(E_k))$ and $H(\chi_{a_{k+1}}(F_k))$ osculate.

It is now obvious that Definition 2.5.7 is satisfied. We will spell out this routine verification.

First, we verify condition (1) of Definition 2.5.7, $\mathcal{Q}[a_k, b_k] \subset$ Carrier H_k, for $k \in \{1, \ldots, p-1\}$. The continuous map

$$\mathfrak{X}_k : \bigcup_{n=a_k}^{b_k} \check{Q}_n \to \bigcup_{n=a_k}^{b_k} Q_n \qquad \mathfrak{X}_k|_{\check{Q}_n} = \chi_n,$$

takes midplane-equivalent midplanes of $\check{\mathcal{Q}}[a_k, b_k]$ to midplane-equivalent midplanes of $\mathcal{Q}[a_k, b_k]$. For $n \in \{a_k, \ldots, b_k\}$, since $\check{Q}_n \cap \check{H}_k \subset H(E_k)$, we therefore have

$$\mathfrak{X}_k(\check{Q}_n \cap \check{H}_k) \subset H(\mathfrak{X}_k(E_k)) = H(\chi_{b_k}(E_k)) = H_k.$$

Thus

$$Q_n = \text{Carrier}\, \mathfrak{X}_k(\check{Q}_n \cap \check{H}_k) \subset \text{Carrier}\, H_k.$$

The proof of condition (1) for $k = p$ is the same, mutatis mutandis (replace E_k above by $F_{p-1} \subset \check{H}_p$).

Since $a_{k+1} = b_k + 1$ for each $k \in \{1, \ldots, p-1\}$, condition (2) of Definition 2.5.7 is satisfied vacuously.

The fact that

$$Q[a_{k-1}, a_k) = Q[a_{k-1}, b_{k-1}] \subset \text{Carrier}\, H_{k-1}$$

and

$$Q(b_k, b_{k+1}] = Q[a_{k+1}, b_{k+1}] \subset \text{Carrier}\, H_{k+1}$$

meet different components of Carrier$(H_k) \setminus H_k$, which is condition (3) of Definition 2.5.7, is also easily checked, using connectedness and the fact that, by Definition 2.4.16, H_k is two-sided. To see this, observe that by continuity, $\mathfrak{X}_k$ maps the connected components of Carrier$(\check{H}_k) \setminus \check{H}_k$ into connected components of Carrier$(H_k) \setminus H_k$, and that the latter two components are distinct since Carrier H_k is two-sided and not self-orthogonal. $\square$

The next two technical lemmas set forth sufficient conditions under which we can obtain a stack $\mathcal{H}$ of hyperplanes with respect to which a given square path $\mathcal{Q}$ in $\mathcal{X}_d$ is properly segmented, given a partition of $\mathcal{Q}$ into two pieces $\mathcal{Q}_1$ and $\mathcal{Q}_2$ that are properly segmented with respect to given stacks $\mathcal{H}_1$ and $\mathcal{H}_2$ respectively. We subsequently extend to the case (Lemma 5.2.5) that $\mathcal{Q}$ is partitioned into infinitely many pieces $\mathcal{Q}_i$ ($i \in \mathbb{N}$).

In the first of these lemmas (Lemma 5.2.3), which is stated for a general square complex, we require that the sequences $\mathcal{H}_1$ and $\mathcal{H}_2$ overlap, with the last two hyperplanes of $\mathcal{H}_1$ matching the first two of $\mathcal{H}_2$; then it is just a matter of reindexing the hyperplanes and checking definitions.

In the second (Lemma 5.2.4), we take $\mathcal{Q}$ to be a square path in $\mathcal{X}_d$ for $d \geqslant 5$, push $\mathcal{Q}$ into $\mathbb{E}^2$ via a folding map $\mathcal{U}(\mathcal{Q}) \to \mathbb{E}^2$, and establish conditions under which two stacks of hyperplanes in $\mathbb{E}^2$ can be transferred back to $\mathcal{X}_d$ in such a way that the original square path $\mathcal{Q}$ is properly segmented with respect to their concatenation. The key requirements are that the square paths in $\mathbb{E}^2$ corresponding to $\mathcal{Q}_1$ and $\mathcal{Q}_2$ are joined by three squares that form an "elbow" shape that overlaps the tail of $\mathcal{Q}_1$ and the head of $\mathcal{Q}_2$, and that the hyperplanes meeting this elbow shape do so in a suitable configuration.

Lemma 5.2.3. Let $(Q_n)_{n=1}^{q}$ be a locally monotone edgewise square path in a square complex X. Let $p \in \{2, \ldots, q-1\}$, and suppose that each of the three square paths

$$(Q_n)_{n=1}^{p+1}, \quad (Q_n)_{n=p-1}^{q}, \quad (Q_{p-1}, Q_p, Q_{p+1}),$$

is properly segmented with respect to stacks

$$(H_k)_{k=1}^{r+1}, \quad (H_k)_{k=r}^{t}, \quad (H_r, H_{r+1})$$

of hyperplanes, respectively. Then $(Q_n)_{n=1}^{q}$ is properly segmented with respect to $(H_k)_{k=1}^{t}$.

Proof. There are three cases:

(i) both Q_{p-1} and Q_p lie in Carrier H_r, and Q_{p+1} lies in Carrier H_{r+1},

(ii) Q_{p-1} lies in Carrier H_r, Q_{p+1} lies in Carrier H_{r+1}, and Q_p lies in neither, and

(iii) Q_{p-1} lies in Carrier H_r, and both Q_p and Q_{p+1} lie in Carrier H_{r+1}.

Assume the former; the remaining cases are similar.

By hypothesis, there exist integers

$$1 = a_1^{(1)} \leqslant b_1^{(1)} < \cdots < a_{r+1}^{(1)} \leqslant b_{r+1}^{(1)}$$

and

$$a_r^{(2)} \leqslant b_r^{(2)} < \cdots < a_t^{(2)} \leqslant b_t^{(2)} = q$$

such that

(1) $\mathcal{Q}[a_k^{(1)}, b_k^{(1)}] \subset \text{Carrier}(H_k)$ for $k \in \{1, \ldots, r+1\}$ and $\mathcal{Q}[a_k^{(2)}, b_k^{(2)}] \subset \text{Carrier}(H_k)$ for $k \in \{r, \ldots, t\}$,

(2) $\mathcal{Q}(b_k^{(1)}, a_{k+1}^{(1)}) \subset X \setminus (H_k \cup H_{k+1})$ for $k \in \{1, \ldots, r\}$ and $\mathcal{Q}(b_k^{(2)}, a_{k+1}^{(2)}) \subset X \setminus (H_k \cup H_{k+1})$ for $k \in \{r, \ldots, t-1\}$,

(3) $\mathcal{Q}[a_{k-1}^{(1)}, a_k^{(1)})$ and $\mathcal{Q}(b_k^{(1)}, b_{k+1}^{(1)}]$ meet distinct components of $\mathcal{Q}[a_k^{(1)}, b_k^{(1)}] \setminus H_k$ for $k \in \{1, \ldots, r+1\}$ and $\mathcal{Q}[b_{k-1}^{(2)}, b_k^{(2)})$ and $\mathcal{Q}(b_k^{(2)}, b_{k+1}^{(2)}]$ meet distinct components of $\mathcal{Q}[a_k^{(2)}, b_k^{(2)}] \setminus H_k$ for $k \in \{r, \ldots, t-1\}$.

Choose $a'_1, \ldots, a'_t$ and $b'_1, \ldots, b'_t$ as follows:

$$1 = \underset{\underset{a'_1}{\|}}{a_1^{(1)}} < \cdots \leqslant \underset{\underset{b'_{r-1}}{\|}}{b_{r-1}^{(1)}} < \underset{\underset{a'_r}{\|}}{a_r^{(1)}} \leqslant \underset{\underset{b'_r}{\|}}{b_r^{(1)}} < \underset{\underset{a'_{r+1}}{\|}}{a_{r+1}^{(2)}} \leqslant \underset{\underset{b'_{r+1}}{\|}}{b_{r+1}^{(2)}} < \cdots < \underset{\underset{a'_t}{\|}}{a_t^{(2)}} \leqslant \underset{\underset{b'_t}{\|}}{b_t^{(2)}} = q.$$

Condition (1) in Definition 2.5.7 is satisfied for all indices of the sequence $(a'_k)_{k=1}^t$. We need to check conditions (2) and (3) only for selected indices near $b'_r = b_r^{(1)} = p$ and $a_{r+1}^{(1)} = b_{r+1}^{(1)} = p + 1 = a_{r+1}^{(2)} = a'_{r+1}$:

(2) $\mathcal{Q}(b'_r, a'_{r+1}) = \mathcal{Q}(b_r^{(1)}, a_{r+1}^{(1)}) \subset X \smallsetminus (\text{Carrier } H_r \cup \text{Carrier } H_{r+1})$,

(3) $\mathcal{Q}[a'_{r-1}, a'_r) = \mathcal{Q}[a_{r-1}^{(1)}, a_r^{(1)})$ and $\mathcal{Q}(b'_r, a'_{r+1}] = \mathcal{Q}(b_r^{(1)}, a_{r+1}^{(1)}]$ meet distinct components of

$$\mathcal{Q}[a'_r, b'_r] \smallsetminus H_r,$$

and $\mathcal{Q}[a'_r, a'_{r+1}) = \mathcal{Q}[a_r^{(1)}, a_{r+1}^{(1)})$ and $\mathcal{Q}(b'_{r+1}, a'_{r+2}] = \mathcal{Q}(b_{r+1}^{(2)}, a_{r+2}^{(2)}]$ meet distinct components of

$$\mathcal{Q}[a'_{r+1}, b'_{r+1}] \smallsetminus H_{r+1}.$$

□

Lemma 5.2.4 (Elbow Lemma). Let $(Q_n)_{n=1}^q$ be a locally monotone edgewise square path in $\mathcal{X}_d$ ($d \geqslant 5$). Let $p \in \{2, \ldots, q-1\}$, and suppose that each of the two square paths $(\breve{Q}_n)_{n=1}^{p-1}, (\breve{Q}_n)_{n=p+1}^q$ in $\mathbb{E}^2$ is properly segmented with respect to stacks $(\breve{H}_k)_{k=1}^{r-1}, (\breve{H}_k)_{k=r}^t$, respectively, of hyperplanes in $\mathbb{E}^2$. Further assume that

- $\breve{H}_{r-1} \perp \breve{H}_r$,
- $\breve{H}_{r-1} \cap \breve{H}_r \subset \mathbb{E}^2 \smallsetminus \breve{Q}_p$,
- $\breve{H}_{r-1}$ contains a midplane of $\breve{Q}_{p-1}$, and
- $\breve{H}_r$ contains a midplane of $\breve{Q}_{p+1}$.

Then $(Q_n)_{n=1}^q$ is properly segmented with respect to some stack $(H_k)_{k=1}^t$ of hyperplanes in $\mathcal{X}_d$.

Proof. The bulleted assumptions on $\breve{H}_{r-1}$ and $\breve{H}_r$ force the square path $(\breve{Q}_{p-1}, \breve{Q}_p, \breve{Q}_{p+1})$ to form an "elbow" shape as shown in Figure 5.7. Since $d \geqslant 5$, the pair of hyperplanes in $\mathcal{X}_d$ that correspond to $\breve{H}_{r-1}$ and $\breve{H}_r$ osculate, and thus constitute a stack in $\mathcal{X}_d$ with respect to which the square path $(\breve{Q}_{p-1}, \breve{Q}_p, \breve{Q}_{p+1})$ is properly segmented. The result now follows from Lemma 5.2.3 and Theorem 5.2.2. □

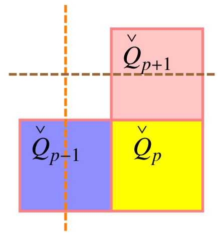

Figure 5.7: Illustration for Lemma 5.2.4.

Lemma 5.2.5. Let $(Q_n)_{n=1}^{\infty}$ be a locally monotone edgewise square path in $\mathcal{X}_d$ ($d \geqslant 4$). Let $(p_i)_{i=1}^{\infty}$ and $(r_i)_{i=1}^{\infty}$ be two increasing sequences of integers with $p_1 > 1$ and $r_1 > 1$. Suppose that for each $i \in \mathbb{N}$ that each of the square paths

$$(\check{Q}_n)_{n=p_{i-1}}^{p_i+1}, \quad (\check{Q}_n)_{n=p_i-1}^{p_{i+1}}, \quad (\check{Q}_{p_i-1}, \check{Q}_{p_i}, \check{Q}_{p_i+1}),$$

in $\mathbb{E}^2$ is properly segmented with respect to stacks

$$(\check{H}_k)_{k=r_{i-1}}^{r_i+1}, \quad (\check{H}_k)_{k=r_i}^{r_{i+1}}, \quad (\check{H}_{r_i}, \check{H}_{r_i+1}),$$

respectively, of hyperplanes in $\mathbb{E}^2$. Then $(Q_n)_{n=p_1-1}^{\infty}$ is properly segmented with respect to some stack $(H_k)_{k=1}^{\infty}$ of hyperplanes in $\mathcal{X}_d$.

Proof. This follows by induction from Lemma 5.2.4. □

It remains only to show how to produce a stack of hyperplanes in $\mathbb{E}^2$ with respect to which Carrier Σ is properly segmented. Then Lemma 5.2.5 applies to show that the original ray γ in $\mathcal{X}_d$ is properly segmented by some stack of hyperplanes, and hence is proper.

Lemma 5.2.6. Let $\breve{\gamma} : [t_1, \infty) \to \mathbb{E}^2$ be a unit-speed ray of constant curvature $|\breve{\gamma}''| \equiv a > 0$ with image Σ and an edgewise square path $(\check{Q}_n)_{n=1}^{\infty}$. Assume that Carrier Σ is homeomorphic to an annulus.

(a) There exists a p such that $\check{Q}_{n+p} = \check{Q}$ for all n.

(b) There exist horizontal hyperplanes $\breve{H}$, $\breve{H}'$ and $n_1, n_2, n_3, n_4 \in \mathbb{N}$ such that

$$n_1 + 1 < n_2, \qquad n_2 + 1 < n_3, \qquad n_3 + 1 < n_4 < n_1 + p,$$

$$\text{Carrier}(\breve{H}) \cap \text{Carrier}(\Sigma) = \bigcup_{n=n_1}^{n_2} \breve{Q}_n,$$

and

$$\text{Carrier}(\breve{H}') \cap \text{Carrier}(\Sigma) = \bigcup_{n=n_3}^{n_4} \breve{Q}_n.$$

(c) There exist stacks $\breve{\mathcal{H}}_1$, $\breve{\mathcal{H}}_3$ of vertical hyperplanes of $\mathbb{E}^2$ and stacks $\breve{\mathcal{H}}_2$, $\breve{\mathcal{H}}_4$ of horizontal hyperplanes such that $(\breve{Q}_n)_{n=n_i+1}^{n_{i+1}-1}$ is properly segmented with respect to $\breve{\mathcal{H}}_i$ $(i \in \{1, 2, 3\})$ and $(\breve{Q}_n)_{n=n_4+1}^{n_1+p-1}$ is properly segmented with respect to $\breve{\mathcal{H}}_4$.

(d) Each pair of stacks $\breve{\mathcal{H}}_i$, $\breve{\mathcal{H}}_{i+1}$, taking indices mod 4, satisfies the bulleted hypotheses of Lemma 5.2.4.

Proof. (a) is obvious.

(b): Let x be the point of the circle Σ with minimum vertical coordinate x_2, let $\breve{H}$ be the horizontal hyperplane of greatest vertical coordinate whose carrier contains x, and write

$$\lceil\lceil x_2 \rceil\rceil = \begin{cases} \lceil x_2 \rceil & \text{if } x_2 \notin \mathbb{Z}, \\ \lceil x_2 \rceil + 1 & \text{if } x_2 \in \mathbb{Z}. \end{cases}$$

Clearly $C := \text{Carrier}\left\{\xi \in \Sigma : \xi_2 \leqslant \lceil\lceil x_2 \rceil\rceil\right\}$ is connected and contained in Carrier $\breve{H}$.

We show that C consists of at least three successively adjacent 2-cells. Assume that Σ is oriented counterclockwise without loss of generality. Let n_1 be the minimal index such that $\breve{Q}_{n_1}$ is the leftmost 2-cell of C, and let n_2 be the minimal index no less than n_1 such that $\breve{Q}_{n_2}$ is the rightmost square of C.

Let $w = (w_1, w_2)$ and $y = (y_1, y_2)$ be the leftmost and rightmost points of Σ, respectively.

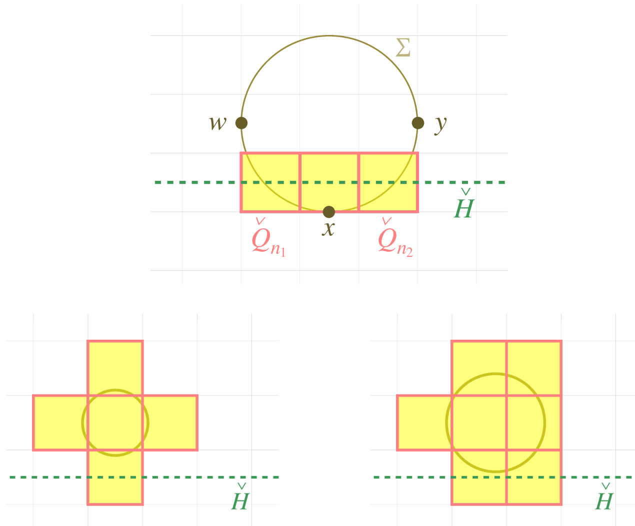

Figure 5.8: C must contain more than two successively adjacent 2-cells, as Carrier $\Sigma \approx S^1 \times I$ and Σ meets no vertex of $\mathbb{E}^2$.

If $w \in C$ (or $y \in C$), then by symmetry, the circle Σ is contained in

$$\{\xi : [[x_2]] - 1 \leqslant \xi_2 \leqslant [[x_2]] + 1\}.$$

But this union of two adjacent horizontal hyperplane carriers contains no subcomplex homeomorphic to an annulus, contrary to our hypothesis. If w and y lie in adjacent 2-cells of $\mathbb{E}^2$, then the left and right semicircles of Σ lie in adjacent vertical hyperplane carriers, which is also impossible. Thus $w_2 = y_2 > [[x_2]]$, and w and y do not lie in adjacent 2-cells of $\mathbb{E}^2$. Examining the remaining cases shows that, since Carrier Σ is an edgewise square path, C must contain at least three successively adjacent 2-cells (see Figure 5.8). Thus $n_1 + 1 < n_2$. We choose $\check{H}'$, n_2, and n_3, and show that $n_2 + 1 < n_3$, $n_3 + 1 < n_4$, and $n_4 < n_1 + p$, in an analogous fashion.

(c): Let M_n be the vertical midplane of $\check{Q}_n$ for each

$$n \in \{n_1 + 1, \ldots, n_2 - 1\} \cup \{n_3 + 1, \ldots, n_4 - 1\},$$

let $\check{\mathcal{H}}_1 = (\check{H}_k)_{k=1}^{p_1}$ be the stack determined by the midplanes M_n, $n_1 < n < n_2$, and let $\check{\mathcal{H}}_3 =$

$(\check{J}_k)_{k=1}^{p_3}$ be determined by the midplanes M_n, $n_3 < n < n_4$. (Note that the hyperplanes $H(M_n)$ need not be distinct for distinct values of n.) Let M_n be the horizontal midplane of $\check{Q}_n$ for

$$n \in \{n_2 + 1, \ldots, n_3 - 1\} \cup \{n_4 + 1, \ldots, n_1 + p - 1\},$$

let $\check{\mathcal{H}}_2 = (\check{I}_k)_{k=1}^{p_2}$ be the stack determined by M_n, $n_2 < n < n_3$, and let $\check{\mathcal{H}}_4 = (\check{K}_k)_{k=1}^{p_4}$ be the stack determined by M_n, $n_4 < n < n_1 + p$. Clearly (c) follows.

(d): Consider $\check{Q}_{n_2}$. Its neighbor $\check{Q}_{n_2-1}$ lies to its left, and by our choice of n_2, we know $\check{Q}_{n_2+1}$ lies above it. It follows that $\check{I}_{p_2} = H(M_{n_2-1})$ and $\check{J}_1 = H(M_{n_2+1})$ are perpendicular. It is now obvious that the bulleted hypotheses of the Elbow Lemma (5.2.4) hold for the square paths $(\check{Q}_n)_{n=n_1+1}^{n_2-1}$, $(\check{Q}_n)_{n_2+1}^{n_3-1}$ and the stacks $\check{\mathcal{H}}_2$ and $\check{\mathcal{H}}_3$, and the remainder of (d) can be verified by similar arguments. □

Theorem 5.2.7 (Annulus Condition). Let $\gamma : [t_1, \infty) \to \mathcal{X}_d^*$ ($d \geqslant 5$) be a ray of constant curvature $a > 0$ with locally monotone edgewise square path $\mathcal{Q} = (Q_n)_{n=1}^{\infty}$. Let $\varphi : \mathcal{U}(\mathcal{Q}) \to \mathbb{E}^2$ be a folding map. If Image φ is an annulus, i.e. homeomorphic to $S^1 \times I$, then γ is a proper ray. (See Figure 5.3.)

Proof. Let $\hat{\gamma}$ be the lift of γ to $\mathcal{U}(\mathcal{Q})$, and write $(\hat{Q}_n)_{n=1}^{\infty}$ for its square path in $\mathcal{U}(\mathcal{Q}) = \bigcup_{n=1}^{\infty} \hat{Q}_n$. Write $\check{\gamma} = \varphi\hat{\gamma}$ and $\check{Q}_n = \varphi(\hat{Q}_n)$ ($n \in \mathbb{N}$). By Lemma 5.2.6, there exist a p such that $\check{Q}_{n+p} = \check{Q}_n$ for all $n \in \mathbb{N}$, a partition of the square path $(\check{Q}_n)_{n=1}^{p}$ into four subpaths $\check{\mathcal{Q}}_i$, $i \in \{1, \ldots, 4\}$, and stacks $\check{\mathcal{H}}_i$, $i \in \{1, \ldots, 4\}$, of hyperplanes such that the hypotheses of Lemma 5.2.4 are satisfied for $\check{\mathcal{Q}}_i$, $\check{\mathcal{Q}}_{i+1}$ and $\check{\mathcal{H}}_i$, $\check{\mathcal{H}}_{i+1}$, $i \in \{1, \ldots, 4\}$, taking indices mod 4. Applying Lemma 5.2.4 yields that $(Q_n)_{n=1}^{p}$ is properly segmented with respect to some stack in $\mathcal{X}_d$.

Using Lemma 5.2.5, the argument just given can be extended to show that the entire square path $\mathcal{Q} = (\hat{Q}_n)_{n=1}^{\infty}$ is properly segmented with respect to some stack of hyperplanes in $\mathcal{X}_d$. Then by Lemma 2.5.8, γ is proper. □

Remark. Numerical experiments suggest that the Annulus Condition is not a necessary condition for a curve of constant curvature in $\mathcal{X}_d$ ($d \geqslant 5$) to be a proper ray (Figure 5.9).

5.2.3 Small Block Condition

We now turn our attention to the case that either the image of the folding map φ is isometric to the square region $[-1,1] \times [-1,1]$.

Let $\gamma : [t_1, \infty) \to \mathcal{X}_d$ ($d \geqslant 5$) be a unit-speed ray of constant nonzero curvature with edgewise square path $\mathcal{Q} = (Q_k)_{k=1}^{\infty}$, and let $\varphi : \mathcal{U}(\mathcal{Q}) \to \mathbb{E}^2$ be a folding map. For simplicity, assume that each $\gamma \cap Q_k$ is a nondegenerate arc.

Before proceeding, let us make the trivial observation that Image φ is isometric to $[-1, 1] \times [-1, 1]$ only if $\Sigma = \varphi\hat{\gamma} : [t_1, \infty) \to \mathbb{E}^2$ contains exactly one vertex in its interior $S := \operatorname{conv}(\Sigma) \smallsetminus \Sigma$, where conv denotes the convex hull. For take Image $\varphi = [-1, 1] \times [-1, 1]$. Clearly $S \subset \operatorname{int}(\operatorname{Image} \varphi)$ contains at most one vertex, the origin 0. Assume for a contradiction that $0 \notin S$. Then by Minkowski's Separating Hyperplane Theorem (see, e.g., [24]), the convex set S lies in an open halfspace L^+ bounded by a line through the origin, say with normal N such that $\langle x, N\rangle \geqslant 0$ for $x \in \overline{S} \supset \Sigma$. Since $\langle e, N\rangle < 0$ for at least one of $\pm e_1$, $\pm e_2$, where e_1, e_2 are the standard unit basis vectors of $\mathbb{E}^2$, we see that $\mathbb{E}^2 \smallsetminus L^+$ contains at least one of the 2-cells of $[-1, 1] \times [-1, 1]$. But then $\sigma \subset L^+$ cannot meet the interior of this 2-cell Q. Since each arc $\gamma \cap Q_k$ is nondegenerate, it follows that Image $\varphi \subset \big([-1, 1] \times [-1, 1]\big) \smallsetminus \operatorname{int} Q$, contrary to our hypothesis. Thus the origin must lie in the interior of Σ.

Theorem 5.2.8 (Small Block Condition). Let $\gamma : [t_1, \infty) \to \mathcal{X}_d$ ($d \geqslant 5$) be a unit-speed ray of constant nonzero curvature with edgewise square path $\mathcal{Q} = (Q_k)_{k=1}^{\infty}$, and let $\varphi : \mathcal{U}(\mathcal{Q}) \to \mathbb{E}^2$ be a folding map. Assume that each $\gamma_k := \gamma \cap Q_k$ is nondegenerate, i.e. not a single point. If Image φ is isometric to $[-1, 1] \times [-1, 1]$, then γ is made up of M distinct circular arcs, where $M = \operatorname{lcm}\{4, d\}$ or $M = d$.

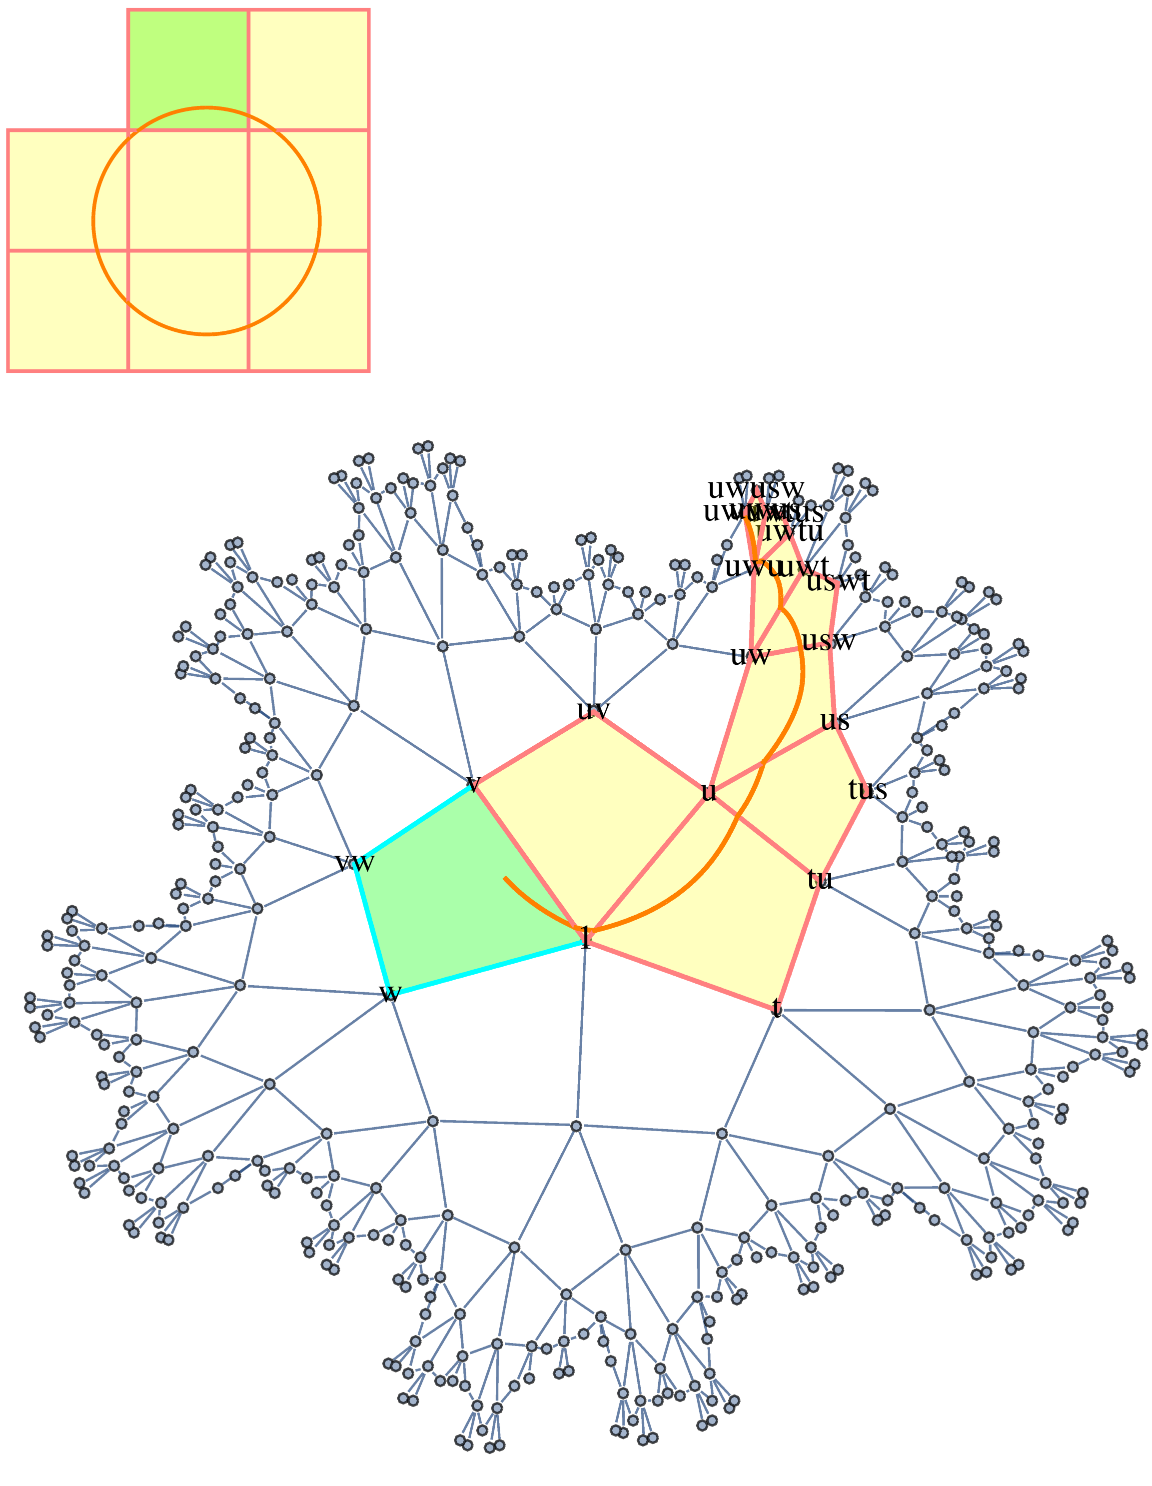

Figure 5.9: BOTTOM: A proper ray γ of constant curvature in $\mathcal{X}_5^*$. TOP: The corresponding non-annular carrier in $\mathbb{E}^2$.

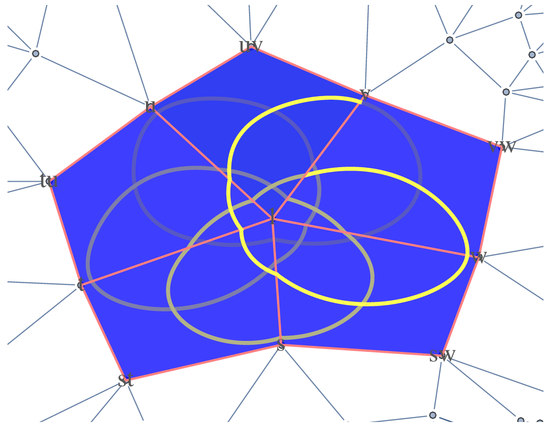

Figure 5.10: A rose curve in $\mathfrak{X}_5$.

Proof. Assume Image $\varphi = [-1, 1] \times [-1, 1]$ without loss of generality. Let $t_1 < t_2 < \cdots$ be a sequence of breakpoints for γ with respect to $\mathcal{Q}$, i.e.

$$\gamma_k := \gamma\big|_{[t_k,\, t_{k+1}]}$$

is contained in Q_k. Write

$$\varphi_k = \varphi\big|_{Q_k} : \mathcal{U}(\mathcal{Q}) \supset Q_k \to [-1, 1] \times [-1, 1], \qquad x_k = \gamma(t_k), \qquad x'_k = \varphi_k(x_k).$$

Since the interior of the circle $\Sigma = \varphi\widehat{\gamma}$ contains exactly one vertex of $\mathbb{E}^2$, the origin, we know that Σ meets the axes of $\mathbb{E}^2$ in exactly four points $x'_k \in \varphi(Q_k)$, $k \in \{1, \ldots, 4\}$. Since Image $\varphi = [-1, 1] \times [-1, 1]$, the 2-cells $Q_k \subset \mathfrak{X}_d$ all share a common vertex $w \in \mathfrak{X}_d$, and $\varphi_k(w) = 0$ for all k. Since each φ_k is an isometric map,

$$d_{\mathfrak{X}_d}(w, x_k) = d(0, x'_k) = |x'_k|.$$

Suppose $|x'_1| = |x'_2|$ and $|x'_3| = |x'_4|$. Since the chord $[x'_1 x'_3]$ is the perpendicular bisector of $[x'_2 x'_4]$ and vice versa, the intersection of these two chords is the center of Σ by a result of Euclid ([21], Bk. III, Porism to Prop. 1). But we know that the intersection of $[x'_1 x'_3]$ and $[x'_2 x'_4]$ is the origin. Since the segments $[0x'_1], \ldots, [0x'_4]$ are thus radii of the circle Σ, and

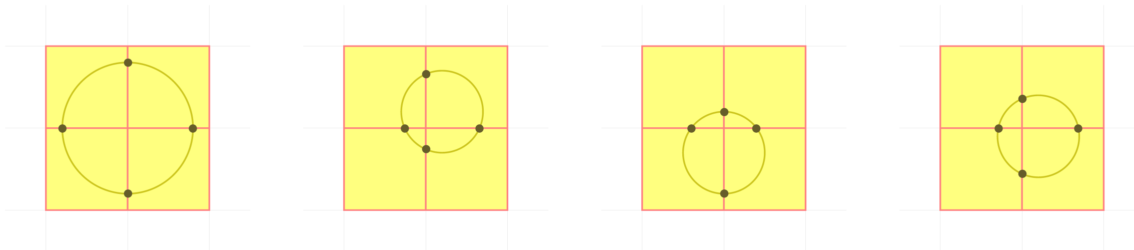

Figure 5.11: Cases for Theorem 5.2.8.

the arcs

$$\wideparen{x'_1 x'_2},\ \wideparen{x'_2 x'_3},\ \wideparen{x'_3 x'_4},\ \wideparen{x'_4 x'_1}$$

are all therefore congruent, so are their preimages

$$\wideparen{x_1 x_2},\ \wideparen{x_2 x_3},\ \wideparen{x_3 x_4},\ \wideparen{x_4 x_5}.$$

It follows that the reflection of $\gamma_1 \subset \mathfrak{X}_d$ in the edge $e_{12} := Q_1 \cap Q_2$ is γ_2, the reflection of γ_2 in $e_{23} := Q_2 \cap Q_3$ is γ_3, and so on, and then that γ is a (metric) circle made up of d arcs.

By the Intersecting Chord Theorem ([21], Bk. III, Prop. 35), we have $|x'_1||x'_3| = |x'_2||x'_4|$. Therefore, if $|x'_1| = |x'_2| = |x'_3|$, we must have $|x'_4| = |x'_1|$, and the conclusions of the previous paragraph obtain.

Up to a cyclic permutation of the labels x'_1, x'_2, x'_3, x'_4, there are three remaining cases:

- $|x'_1| = |x'_2| \neq |x'_3| = |x'_4|$.
- $|x'_1| = |x'_3|$, $|x'_2|$, and $|x'_4|$ are distinct.
- $|x'_1|$, $|x'_2|$, $|x'_3|$, and $|x'_4|$ are all distinct.

Examining each of these possibilities, we see that $|x_{Ln+i}| = |x'_{Ln+i}| = |x'_i|$ for each $i \in \{1, 2, 3, 4\}$ iff $Ln \equiv 0 (\mathrm{mod}\ 4)$. Therefore, γ is made up of $M = \mathrm{lcm}\{4, d\}$ distinct circular arcs. □

Chapter 6

Future directions

As mentioned in the introduction to Chapter 4, the question of whether or not there exist solutions to the Markov-Dubins problem with prescribed initial direction which are not CL paths remains open at the conclusion of the present work. Although the existence of non-CL length-minimal paths of bounded curvature in a square complex seems doubtful, their existence in complexes of dimension $\geqslant 3$ is all but assured by the work of Sussmann and his collaborators (see, e.g. [35]), who find that in $\mathbb{R}^3$, some Markov-Dubins problems have solutions they describe as helicoids. We hope to determine whether non-CL length-minimizers are possible in dimension 2, and to extend the characterization of solutions established by Sussmann and others to cube complexes with dimension $\geqslant 3$.

While preliminary results do indicate that there exists a simple sufficient condition which determines whether or not some admissible path in a nonpositively curved square complex exists, given the boundary conditions which define a Markov-Dubins problem with free terminal vector, we have not presented such a condition in this text. In conjunction with Theorem 3.2.2, such a condition would provide a complete answer to the question of when a solution to a given problem is guaranteed to exist.

It is to be hoped that the results in the current work might be extended to the Markov-

Dubins problem in a nonpositively curved square complex with prescribed in initial *and* terminal vector. The diversity of the possible shapes of solutions in Dubins' original problem is almost entirely due to the specification of the terminal vector, and necessitates the consideration of CLC paths made up of two arcs of a circle joined by a line segment, and CCC (or "teardrop") paths composed of three circular arcs having alternating orientation. (The utility of such a teardrop-shaped path will be clear to any driver who has executed a U-turn in a confined parking lot.) Of course, in a planar region of the type considered by Jacobs and Canny (see §4.1), any number of segments, each a circular arc or a geodesic, is in principle possible for a solution that is required to navigate around several obstacles. The one-dimensional bisection method we have used to find CL paths will not be sufficient to determine shortest paths consisting of three or more segments of varying type, including even the simplest analogue of Dubins' teardrop paths. It remains to be seen whether it is possible, on the basis of information which may be known about a particular square complex, to establish a bound on the number of segments a shortest path may have for any given problem, regardless of the boundary conditions. In the case of the d-plane, it is clear that the explosion of total angle measure at vertices when $d \geqslant 5$ will make the circumnavigation of a vertex dramatically less efficient than it would be for $d = 4$. Future research may reveal that, under circumstances to be determined, the circumnavigation of a vertex can be replaced by a more efficient, but perhaps less intuitively obvious route.

Another area which might be explored is the applicability of our methods to unions and quotients of d-planes. Such spaces arise naturally in connection with right-angled Artin systems, which themselves arise naturally in connection with reconfiguration problems. Indeed, the generalized braid groups G are fundamental to the study of classical configuration spaces. When such a group, whether right-angled or not, acts properly on a CAT(0) cube complex X, we expect that the solution of Markov-Dubins problems in X will be amenable to methods analogous to those we have presented.

A final possible extension of our ideas, proposed informally by Guilbault, would exploit the phenomenon of bifurcating geodesics by associating a variable cost with each of the possible directions by which a geodesic might enter a cube. Given a sufficiently high-dimensional Euclidean space $\mathbb{R}^d$, it should be possible to embed the state complex for a given reconfigurable system in such a way that zero-cost entry vectors correspond to Euclidean geodesics. Geodesic segments under the ambient Euclidean metric might then represent reconfiguration strategies that are deemed ideal with respect to considerations not easily represented in the state complex, e.g., energy consumption, or kinematic constraints on systems external to the one being modeled.

Bibliography

[1] Aaron Abrams and Robert Ghrist, *State complexes for metamorphic robots*, The International Journal of Robotics Research **23** (2004), no. 7-8, 811–826.

[2] Stephanie Alexander, Richard Bishop, and Robert Ghrist, *Pursuit and evasion in non-convex domains of arbitrary dimensions*, Robotics: Science and Systems, Citeseer, 2006.

[3] ______, *Total curvature and simple pursuit on domains of curvature bounded above*, Geometriae Dedicata **149** (2010), no. 1, 275–290.

[4] Aleksandr Alexandrov and Yurii Reshetnyak, *General theory of irregular curves*, Mathematics and its applications (Soviet series), vol. 29, Kluwer Academic Publishers, 1989.

[5] Federico Ardila, Megan Owen, and Seth Sullivant, *Geodesics in CAT(0) cubical complexes*, Advances in Applied Mathematics **48** (2012), no. 1, 142–163.

[6] Werner Ballmann, *Singular spaces of non-positive curvature*, Sur les groupes hyperboliques d'apres Mikhael Gromov, Springer, 1990, pp. 189–201.

[7] Jean-Daniel Boissonat, André Cérézo, and Juliette Leblond, *Shortest paths of bounded curvature in the plane*, Proceedings of the IEEE International Conference on Robotics and Automation, 1992, pp. 2315–2320.

[8] Martin Bridson and André Haefliger, *Metric spaces of non-positive curvature*, vol. 319, Springer Science & Business Media, 1999.

[9] Dmitri Burago, Yuri Burago, and Sergei Ivanov, *A course in metric geometry*, vol. 33, American Mathematical Society Providence, 2001.

[10] Constantin Carathéodory, *Conformal representation*, Courier Corporation, 1998.

[11] Victor Chepoi and Daniela Maftuleac, *Shortest path problem in rectangular complexes of global nonpositive curvature*, Computational Geometry **46** (2013), no. 1, 51–64.

[12] Yacine Chitour and Mario Sigalotti, *Controllability of the Dubins problem on surfaces*, 2005 and 2005 European Control Conference. CDC-ECC '05. 44th IEEE Conference on Decision and Control, IEEE, 2005, pp. 1439–1444.

[13] ———, *Dubins problem on surfaces, I: Nonnegative curvature*, The Journal of Geometric Analysis **15** (2005), no. 4, 565–587.

[14] Michael Davis, *The geometry and topology of Coxeter groups*, vol. 32, Princeton University Press, 2008.

[15] Lester Dubins, *On curves of minimal length with a constraint on average curvature, and with prescribed initial and terminal positions and tangents*, American Journal of Mathematics (1957), 497–516.

[16] Peter Friz and Nicolas Victoir, *Multidimensional stochastic processes as rough paths: theory and applications*, Cambridge University Press, 2010.

[17] Robert Ghrist and Valerie Peterson, *The geometry and topology of reconfiguration*, Advances in Applied Mathematics **38** (2007), no. 3, 302–323.

[18] Mikhael Gromov, *Hyperbolic groups*, in *Essays in group theory* (Steve Gersten, ed.), MSRI Publications, vol. 8, Springer-Verlag, 1987, pp. 75–264.

[19] Frédéric Haglund and Daniel Wise, *Coxeter groups are special*, preprint (2008).

[20] ———, *Special cube complexes*, Geometric and Functional Analysis **17** (2008), no. 5, 1551–1620.

[21] Thomas Heath, *The thirteen books of Euclid's Elements*, Courier Corporation, 1956.

[22] Paul Jacobs and John Canny, *Planning smooth paths for mobile robots*, Nonholonomic Motion Planning, Springer, 1993, pp. 271–342.

[23] Wichitra Karuwannapatana and Chaiwat Maneesawarng, *The lower semi-continuity of total curvature in spaces of curvature bounded above*, East-West Journal of Mathematics **9** (2007), no. 1, 1–9.

[24] Mark Kreĭn and Adolf Nudel'man, *The Markov moment problem and extremal problems*, Translations of mathematical monographs, vol. 50, American Mathematical Society, 1977.

[25] Der-Tsai Lee and Franco Preparata, *Euclidean shortest paths in the presence of rectilinear barriers*, Networks **14** (1984), no. 3, 393–410.

[26] John Milnor, *On the total curvature of knots*, Annals of Mathematics (1950), 248–257.

[27] Dirk Mittenhuber, *Dubins' problem in hyperbolic space*, Conference on Geometric Control and Non-holonomic Mechanics, June 19-21, 1996, Mexico City, Canadian Mathematical Society Conference Proceedings, vol. 25, 1998, pp. 101–114.

[28] ———, *Dubins' problem in the hyperbolic plane using the open disc model*, Conference on Geometric Control and Non-holonomic Mechanics, June 19-21, 1996, Mexico City, Canadian Mathematical Society Conference Proceedings, vol. 25, 1998, pp. 115–152.

[29] ———, *Dubins' problem is intrinsically three-dimensional*, ESAIM: Control, Optimisation and Calculus of Variations **3** (1998), 1–22.

[30] Felipe Monroy-Pérez, *Non-Euclidean Dubins' problem*, Journal of dynamical and control systems **4** (1998), no. 2, 249–272.

[31] James Munkres, *Topology: A first course*, 2nd ed., Prentice-Hall, 2000.

[32] Valerie Peterson, *On state complexes and special cube complexes*, Ph.D. thesis, University of Illinois at Urbana-Champaign, 2000.

[33] James Reeds and Lawrence Shepp, *Optimal paths for a car that goes both forwards and backwards*, Pacific Journal of Mathematics **145** (1990), no. 2, 367–393.

[34] Mario Sigalotti and Yacine Chitour, *Dubins' problem on surfaces, II: Nonpositive curvature*, SIAM Journal on control and optimization **45** (2006), no. 2, 457–482.

[35] Héctor Sussmann, *Shortest 3-dimensional paths with a prescribed curvature bound*, Proceedings of the 34th IEEE Conference on Decision and Control, vol. 4, IEEE, 1995, pp. 3306–3312.

[36] Héctor Sussmann and Guoqing Tang, *Shortest paths for the Reeds-Shepp car: a worked out example of the use of geometric techniques in nonlinear optimal control*, Rutgers Center for Systems and Control Technical Report **10** (1991), 1–71.

[37] Boris Thomaschewski, *Workspaces of continuous robotic manipulators*, Ph.D. thesis, Technische Universität Ilmenau, 2002.

[38] Daniel Wise, *From riches to RAAGs: 3-manifolds, right-angled Artin groups, and cubical geometry*, CBMS Regional Conference Series in Mathematics, vol. 117, American Mathematical Society, 2012.

Curriculum Vitae

Jason Thomson La Corte

Education

M.S., Mathematics. Texas State University–San Marcos, December 2007.

B.A., Liberal Arts. Hampshire College, February 1998.

Presentations

Numerical solution of the Dubins problem with free terminal direction in a state complex. 31st Annual Workshop in Geometric Topology, Milwaukee, WI, June 14, 2014.

Curvature-constrained path planning in a nonpositively curved cube complex. Joint Mathematics Meetings, Baltimore, MD, January 18, 2014.

Classification of curves of constant curvature in the d-plane. Faculty Topology Seminar, University of Wisconsin–Milwaukee, December 9, 2013.

Existence of solutions to the Dubins problem with free terminal direction in a nonpositively curved cube complex. Faculty Topology Seminar, University of Wisconsin–Milwaukee, February–March, 2013.

Awards and Honors

Mark Lawrence Teply Award. University of Wisconsin–Milwaukee, 2014.

Chancellor's Graduate Student Award. University of Wisconsin–Milwaukee, 2013–2015.

GAANN Fellowship. University of Wisconsin–Milwaukee, 2010–2013.

Teaching Experience

Adjunct Faculty, January 2015–May 2015. Saint Mary-of-the-Woods College.

Graduate Teaching Assistant, September 2009–May 2015. University of Wisconsin–Milwaukee.

Adjunct Associate Professor, January 2008–August 2009. Austin Community College.

Graduate Teaching Assistant, August 2005–December 2007. Texas State University–San Marcos.

Saint Mary-of-the-Woods College

Finite Mathematics	Spring 2015	Lab instructor
Fundamentals of Mathematics	Spring 2015	Lab instructor

University of Wisconsin–Milwaukee

Contemporary Mathematics	Spring 2015	Instructor of record
Contemporary Mathematics	Fall 2014	Instructor of record
Calculus and Analytic Geometry I	Summer 2014	Instructor of record
Calculus and Analytic Geometry I	Spring 2014	Instructor of record
Calculus and Analytic Geometry I	Fall 2013	Instructor of record
Intermediate Algebra	Fall 2012	Teaching assistant
Theory of Functions of a Real Variable II	Spring 2012	Grader
Theory of Functions of a Real Variable I	Fall 2011	Grader
Contemporary Mathematics	Fall 2010	Instructor of record
Contemporary Mathematics	Fall 2010	Teaching assistant
Essentials of Algebra	Spring 2010	Instructor of record
Intermediate Algebra	Fall 2009	Instructor of record

Austin Community College

Trigonometry	Summer 2009	Instructor of record
College Algebra	Summer 2009	Instructor of record
College Mathematics	Spring 2009	Instructor of record
College Algebra	Spring 2009	Instructor of record
College Mathematics	Fall 2008	Instructor of record
College Mathematics	Spring 2008	Instructor of record
College Algebra	Spring 2008	Instructor of record

Texas State University–San Marcos

Precalculus	Fall 2007	Lab instructor
Precalculus	Summer 2007	Lab instructor
Pre-College Algebra	Summer 2007	Lab instructor
Calculus II	Spring 2007	Lab instructor
Precalculus	Spring 2007	Lab instructor
Basic Math	Fall 2006	Lab instructor
Basic Math	Spring 2006	Lab instructor
Basic Math	Fall 2005	Lab instructor
Differential Equations	Spring 2005	Grader