CHAPTER 7 : More about trees.

7.1. Introduction.

In this chapter, we aim to explore further the "treelike" nature of hyperbolic spaces. The motivation for this chapter comes from [G, Chapter 6].

In Section 3.3, we showed that any finite subset of an almost-hyperbolic space may be spanned by a tree which measures distances up to a certain additive constant. The main aim of this chapter is to prove Theorem 7.6.1, which is a refinement of this result, though some results along the way seem to have some interest in their own right.

By a "tree", T, we mean a connected acyclic finite graph. As usual, we shall identify T, thought of combinatorially, with its realisation as a 1-complex.

Given any closed subset $P \subseteq T$, we write $\mathrm{span}_T P$, or just $\mathrm{span}\, P$ for the subtree of T spanned by P, i.e. the the smallest subtree containing P. Thus, in particular, if x, y are distinct points of T, then $\mathrm{span}(x, y)$ is the unique arc joining x to y.

Given any point $x \in T$, we write $\deg x$ for the degree of x, i.e. the number of connected components of $T \backslash \{x\}$. We write

$$\mathrm{node}\, T = \{x \in T \mid \deg x \geq 3\}$$

for the set of "nodes" of T, and

$$\mathrm{ext}\, T = \{x \in T \mid \deg x = 1\}$$

for the set of "extreme points" of T.

Another way to describe $\mathrm{ext}\, T$ is to say that it is the smallest subset, W, of T such that $T = \mathrm{span}\, W$.

Let V be a finite set. An *(abstract) spanning tree*, $\tau = (T, \sigma, f)$, for V consists of a tree T, together with a path-metric σ on T, and a map $f : V \longrightarrow T$ such that $T = \mathrm{span}\, f(V)$ (or equivalently $\mathrm{ext}\, T \subseteq f(V)$). Clearly (T, σ) is a geodesic space (as defined in Section 1.1). We can alternatively think of σ a finite 1-dimensional measure of full support on T. If $Q \subseteq T$, we write $\sigma(Q)$ for the measure, or "σ-length, of Q. We have $\sigma(x, y) = \sigma(\mathrm{span}(x, y))$ for all $x, y \in T$. Given $v, w \in V$, we write

$$\rho_\tau(v, w) = \sigma(f(v), f(w)).$$

Thus ρ_τ is a pseudometric V, i.e. we allow for the possibility of two distinct points being 0 distance apart. (This extra generality will be convenient in Section 7.3.) Corollary 7.3.2 gives a characterisation of which pseudometrics may be derived in this way.

We inserted the word "abstract" in the definition to distinguish it from from the following notion of "immersed" spanning tree. Let $(\mathcal{S}, d)$ be a geodesic space, and $V \subseteq (\mathcal{S}, d)$ a finite set of points. An *immersed spanning tree*, (τ, g), for V, consists of an abstract spanning tree $\tau = (T, \sigma, f)$, together with a distance non-increasing map $g : (T, \sigma) \longrightarrow (\mathcal{S}, d)$ such that $g \circ f : V \longrightarrow (\mathcal{S}, d)$ is just the inclusion of V into $(\mathcal{S}, d)$. (Thus, in this case, f is injective and ρ_τ is a metric.) We call such a tree an *embedded spanning tree* if g is injective and σ is the path metric on T induced from d. In this case, we may identify T with $g(T) \subseteq \mathcal{S}$, so that the maps f and g become superfluous. In all the cases we deal with, the edges of an embedded spanning tree will be geodesic.

In Section 3.3, we showed that if $(\mathcal{S}, d)$ is k-H1, and $V \subseteq \mathcal{S}$ has $n+1$ elements, then there is an embedded spanning tree $T = T_V$ for V such that for all $v, w \in V$, we have

$$\rho_\tau(v, w) \leq d(v, w) + K(k, n),$$

where $K(k, n)$ depends only on k and n. In fact, we can write $K(k, n)$ in the form $kH(n)$, where $H : \mathbf{N} \longrightarrow [0, \infty)$ is a fixed function of n. This can be seen from a study of the proof, though in fact, it is a consequence of a more general principle, about which we shall say more in Section 7.2. With regard to the dependence on n, however, we were more careless. Our proof would give $H(n)$ exponential in n.

Definition : Given a function $f : \mathbf{N} \longrightarrow [0, \infty)$, and a set V of $n+1$ points in a k-H1 space $(\mathcal{S}, d)$, we shall call an immersed spanning tree, (τ, g), for V, an *$O(f(n))$-approximating tree*, or just *$O(f(n))$-tree*, if for all $v, w \in V$, we have

$$\rho_\tau(v, w) \leq d(v, w) + kO(f(n)).$$

Of course, this is really a property of a method of construction, rather than of a particular tree.

So far, we have shown the existence of embedded exponential approximating trees. Theorem 7.6.1. shows the existence of immersed $O(\log n)$ trees. In fact, this is the best order of growth one could hope for, as the following example shows.

Consider $n+1$ points equally spaced around a large circle in the hyperbolic plane. The best immersed tree in this case is obtained by joining each point to the centre of the circle by a geodesic segment. The angle between two adjacent segments is $O(1/n)$. We see that this gives us an $O(-\log(1/n)) = O(\log n)$-approximating tree.

We shall show (Proposition 7.5.2) that the construction of Section 3.3 gives, in fact, an $O(n)$-approximating tree. Another obvious way to construct an embedded spanning tree is to take a tree of minimal total length spanning the $n+1$ points of V. We call such a tree a *Steiner tree*. We shall show that this also gives an $O(n)$-approximating tree. In both these cases, this is, in general, the best result possible, as may seen from the following example.

Let $(\mathcal{S}, d)$ be the bi-infinite euclidean strip $\mathbf{R} \times [-1, 1] \subseteq \mathbf{R}^2$. For $i \in \mathbf{N}$ let X_i be the point $(4i, (-1)^i) \in \mathcal{S}$. Given $n \in \mathbf{N}$, let V_n be the set of points $\{X_i \mid 0 \leq i \leq n\}$. The (unique) Steiner tree joining these points consists of the piecewise geodesic arc $\alpha = \bigcup_{i=1}^n [X_i, X_{i-1}]$. Clearly, for $n \geq 3$, we have length $\alpha \geq d(X_0, X_n) + \epsilon n$ for a fixed $\epsilon > 0$. The construction of Section 3.3 relied on a choice of ordering of the points of V. If we order the points in the obvious way, $X_0, X_2, \ldots, X_n$, then this construction also gives $T_V = \alpha$.

It seems quite likely that if we were to judiciously choose the ordering of the points of V, then the construction of Section 3.3 would always give an $O(\log n)$-approximating tree.

However, I have been unable to prove this, and so Theorem 7.6.1 uses a slightly different construction (based on that of [G, Chapter 6]). This will in general give an immersed, rather than an embedded spanning tree. Perhaps the proof can be modified so as always to give an embedded tree. However, for general hyperbolic spaces, the assumption of embeddedness does not seem particularly natural (unless, for example, $\mathcal{S}$ happens to be a 2-manifold). Note that by a small perturbation, we can always arrange that a tree be embedded in $\mathcal{S} \times [0, \epsilon]$.

The main steps in the proof of Theorem 7.6.1. are as follows. We begin by giving a construction of abstract spanning trees which approximate a metric on a set V of $n+1$ elements to $O(\log n)$ (Proposition 7.3.1). This gives us some hint as to the combinatorial structure of our desired immersed spanning tree when $V \subseteq \mathcal{S}$. However, we need to decide how to partition our tree into arcs, destined to become geodesics in $\mathcal{S}$. This is the purpose of the combinatorial result, Lemma 7.4.1. The final construction of our immersed spanning tree is based on a variant of that of Section 3.3. To show that this works, we need the result, stated above, that the construction of Section 3.3 is always $O(n)$. This is the only real geometric input. It can be viewed as a corollary of a result about piecewise geodesic paths in $\mathcal{S}$. This can in turn be interpreted as a statement about finite metric spaces (Proposition 7.3.4).

7.2. A note on parameters.

Throughout this paper, we have used the notation $x \simeq y$ to mean that $|x - y| \leq K$, where K has been assumed to be some function only of the parameter of hyperbolicity. Similarly, we have used $x \preceq y$ for $x \leq y + K$. If we look at Chapter 3, we see that all such numbers K arising in this way are the result of applying a transitivity law a certain finite number of times, starting with the parameter of hyperbolicity—either k_1 or k_2. We see that we can reinterpret the notation $x \simeq y$ to mean that $|x - y|$ is bounded by some universal multiple of k_1 (or k_2). Similarly for $\preceq$. We see that all the functions of the hyperbolicity parameter arising in Chapter 3 can be written in the form λk_1 for some $\lambda \in [0, \infty)$. This will apply equally well to this chapter, though we shall not always state this explicitly.

Another way to view this is to note that, if $k_1 \neq 0$, then after rescaling the metric by a factor of $1/k_1$, any k_1-H1 space becomes 1-H1. Now, any 0-H1 geodesic space is metric tree (Proposition 3.4.2), so we could restrict attention to 1-H1 spaces, and thus not worry about the dependence on k_1.

7.3. Abstract spanning trees.

Let V be a finite set. A *pseudometric* on V is a symmetric function $\rho : V \times V \longrightarrow [0, \infty)$ satisfying the triangle inequalities. Note that the definition H1 makes equally good sense applied to an arbitrary pseudometric space, in particular to (V, ρ). We may thus define

$$\operatorname{hyp} \rho = \min\{k \in [0, \infty) \mid (V, \rho) \text{ is } k - \text{H1}\}.$$

Equivalently,

$$\operatorname{hyp}\rho = \max\{\rho(x,y)+\rho(z,w)-\max\big(\rho(x,w)+\rho(y,z),\rho(x,z)+\rho(y,w)\big) \mid x,y,z,w\in V\}.$$

The following result is copied from [G, Chapter 6], except that we have added more detail to the proof.

Proposition 7.3.1 : *Let V be a set of $n+1$ elements, and let ρ be a pseudometric on V. Then, there is an abstract spanning tree τ for V, such that*

$$\rho_\tau \leq \rho \leq \rho_\tau + (1+\log_2 n)\operatorname{hyp}\rho.$$

Proof : We construct $\tau = (T,\sigma,f)$ as follows.

First, choose any $v_0 \in V$, and write $V' = V\setminus\{v_0\}$. Let $\mathcal{I} = \{(t,v) \subseteq \mathbf{R}\times V' \mid 0 \leq t \leq \rho(v,v_0)\}$. Thus, we imagine $\mathcal{I}$ as a disjoint union $\mathcal{I} = \bigsqcup_{v\in V'} I_v$ of intervals I_v, where $I_v = [0,\rho(v,v_0)]\times\{v\}$. We topologise $\mathcal{I}$ accordingly.

Given $v,w \in V'$, write

$$\langle u,w\rangle = \frac{1}{2}\big(\rho(u,v_0)+\rho(w,v_0)-\rho(u,w)\big).$$

Thus, $\langle u,w\rangle \leq \min\big(\rho(u,v_0),\rho(w,v_0)\big)$. We define a relation

$$R_{u,w} \subseteq I_u \times I_w$$

by $\big((s,u),(t,w)\big) \in R_{u,w}$ if and only if $s = t \leq \langle u,w\rangle$. In other words, we identify the initial segments $[0,\langle u,w\rangle]\times\{u\}$ and $[0,\langle u,w\rangle]\times\{w\}$. Let $R \subseteq \mathcal{I}\times\mathcal{I}$ be the transitive closure of all the relations $R_{u,w}$ for $u,w\in V'$, and let $T = \mathcal{I}/R$ be the quotient.

One sees easily that T has the structure of a tree, with a path-metric σ induced from the parameterisation of the intervals I_v. Define $f: V \longrightarrow T$ as follows. For $v \in V'$, we let $f(v)$ be the projection of the point $(\rho(v,v_0),v)\in I_v$ to the quotient under R. We define $f(v_0)$ to be the projection of $(0,v)\in I_v$ for any $v\in V$, noting that all such points are identified under R. We see that $\operatorname{ext} T \subseteq f(V)$, so that $\tau = (T,\sigma,f)$ is a spanning tree for V.

By construction, we have

$$\rho_\tau(v,v_0) = \rho(v,v_0)$$

for all $v \in V'$.

Now, given $u,w\in V'$, write $\alpha(u,w)$ for the arc joining $f(u)$ to $f(w)$ in T. Let

$$\langle u,w\rangle_\tau = \frac{1}{2}\big(\rho_\tau(u,v_0)+\rho_\tau(w,v_0)-\rho_\tau(u,w)\big).$$

Thus, $\langle u,w\rangle_\tau$ is the distance, $\sigma(f(v_0),\alpha(u,w))$, from $f(v_0)$ to $\alpha(u,w)$.

From the construction, it is clear that

$$\langle u,w\rangle_\tau \geq \langle u,w\rangle$$

and thus,

$$\rho_\tau(u,w) \leq \rho(u,w).$$

To complete the proof, we need to show that

$$\langle u,w\rangle_\tau \leq \langle u,w\rangle + \frac{1}{2}k(1+\log_2 n),$$

where $k = \operatorname{hyp}\rho$.

Now, from the definition of $k = \operatorname{hyp}\rho$, we find that for any $u, v, w \in V'$, we have

$$\rho(v,v_0) + \rho(u,w) \leq \max\big(\rho(u,v_0)+\rho(v,w), \rho(w,v_0)+\rho(u,v)\big) + k.$$

Rearranging the terms, we deduce that

$$\langle u,w\rangle \geq \min\big(\langle u,v\rangle, \langle v,w\rangle\big) - k/2.$$

More generally, therefore, if $v_1, v_2, \ldots, v_p$ is any sequence of p points of V', we have

$$\langle v_1,v_p\rangle \geq \min\{\langle v_i,v_{i+1}\rangle \mid 1 \leq i \leq p-1\} - \frac{1}{2}k(1+\log_2 p).$$

Thus, given $u, w \in V'$, it is enough to find some sequence $u = v_1, v_2, \ldots, v_p = w$ of distinct points of V' so that

$$\langle u,w\rangle_\tau \leq \min\{\langle v_i,v_{i+1}\rangle \mid 1 \leq i \leq p-1\}.$$

This may be accomplished as follows.

Let $x \in \alpha(u,w)$ be the nearest point of $\alpha(u,w)$ to $f(v_0)$. Let $d = \sigma(f(v_0),x) = \sigma(f(v_0),\alpha(u,w)) = \langle u,w\rangle_\tau$.

Recall that $T = \mathcal{I}/R$, where $\mathcal{I} = \bigsqcup_{v\in V'} I_v$. We see that the points $(d,u) \in I_u$ and $(d,w) \in I_w$ project to the same point $x \in T$. In other words, (d,u) and (d,w) are identified under R. From the definition of R as a transitive closure, this means that there are points $u = v_1, v_2, \ldots, v_p = w$ in V' with $\big((d,v_i),(d,v_{i+1})\big) \in R_{v_i,v_{i+1}}$ for each $i \in \{1,2,\ldots,p-1\}$. Thus $d \leq \langle v_i,v_{i+1}\rangle$ for $i = 1,2,\ldots,p-1$. In other words,

$$\langle u,w\rangle_\tau \leq \min\{\langle v_i,v_{i+1}\rangle \mid 1 \leq i \leq p-1\},$$

as required. (In fact we have equality in this expression.)

◊

Note that f might not be injective, even if ρ is a metric. (For example, take $V = \{1,2,3,4\}$, and $\rho(1,2) = \rho(2,3) = \rho(3,4) = \rho(4,1) = 1$ and $\rho(2,4) = \rho(1,3) = 2$.)

Corollary 7.3.2 : *Suppose ρ is a pseudometric on a finite set V. Then, $\operatorname{hyp}\rho = 0$ if and only if there is some spanning tree τ for V such that $\rho = \rho_\tau$.*

Proof : It is a simple exercise that $\operatorname{hyp}\rho_\tau = 0$ for any spanning tree τ. The reverse implication is an immediate consequence of Proposition 7.3.1.

◇

By hypothesis, a spanning tree $\tau = (T, \sigma, f)$ for V satisfies $\operatorname{ext} T \subseteq f(V)$. It would be convenient if we could always take f to be a bijection to $\operatorname{ext} T$, so that we may identify V with $\operatorname{ext} T$. In fact, this can be arranged provided we allow σ to be a "path-pseudometric". If we think of a path-metric on T as assigning a positive length to each edge of T, then a *path-pseudometric* assigns a non-negative length to each edge. We do this as follows.

Given any spanning tree $\tau = (T, \sigma, f)$ for V, we define $\tau' = (T', \sigma', f')$ by attaching an arc of length 0 to $f(v)$ for each $v \in V$. More formally, define $T' = \big(([0,1] \times V) \sqcup T\big)/\sim$, where $(1, v) \sim f(v)$ for each $v \in V$. Let $f'(v)$ be the projection of $(0, v)$ to T'. Define σ' by $\sigma'|T = \sigma$ and $\sigma'([0,1] \times \{v\}) = 0$ for all v. Thus, f' is injective, $f'(V) = \operatorname{ext} T'$, and $\rho_{\tau'} = \rho_\tau$.

The remainder of this section is aimed at proving a result about finite metric spaces (Proposition 7.3.4), which may be interpreted, in the context of almost-hyperbolic geodesic spaces, as a statment about piecewise geodesic paths. The result is best motivated in that context (see the beginning of Section 7.5), though it fits more logically into this section.

Suppose that (V, ρ) is a pseudometric space. We introduce the following notation.

Given $x, y, z, w \in V$, write

$$xy \wedge_\rho zw = \frac{1}{2}\big(\rho(x,y) + \rho(z,w) - \rho(x,z) - \rho(y,w)\big).$$

We shall usually abbreviate $xy \wedge_\rho zw$ to $xy \wedge zw$.

We list the following properties of $\wedge$, though we shall only find explicit use for parts (1), (2) and (4).

Lemma 7.3.3 : *Suppose that x, y, z, w, u are any points in the pseudometric space (V, ρ), then*

(1) $xy \wedge zw = yx \wedge wz = -xz \wedge yw$,

(2) $xy \wedge zw + xw \wedge yz + xz \wedge wy = 0$,

(3) $xy \wedge zw \leq \min\big(\rho(x,y), \rho(z,w)\big)$,

(4) $xy \wedge zy \geq 0$,

(5) $xy \wedge xy = \rho(x,y)$,

(6) $uy \wedge wz + uw \wedge yx + wx \wedge zy = 0$.

Proof : Elementary.

◇

We remark that if V is finite, then we can define $\operatorname{hyp}\rho$ in terms of this notation, thus:

$$\operatorname{hyp}\rho = 2\max\{h(xy \wedge zw, xw \wedge yz, xz \wedge wy) \mid x, y, z, w \in V\},$$

where $h(a,b,c) \geq 0$ is defined for any $a, b, c \in \mathbf{R}$ with $a+b+c = 0$, as follows. If it happens that $a \geq \max(|b|, |c|)$, then we set $h(a,b,c) = -b$. From this, we may define $h(a,b,c)$ in general, by insisting that it have the symmetry $h(a,b,c) = h(b,c,a) = h(-a,-c,-b)$. It is readily checked that this agrees with the definition given earlier in this section.

Suppose that $\operatorname{hyp}\rho = 0$, so that $\rho = \rho_\tau$ for some spanning tree $\tau = (T, \sigma, f)$ (Corollary 7.3.2). Given $x, y \in V$, write $\alpha(x,y) = \operatorname{span}(f(x), f(y))$. In this case, if $xy \wedge zw \geq 0$, then $xy \wedge zw = \sigma(\alpha(x,y), \alpha(z,w)) = \sigma(\beta)$ (i.e. the σ-length of β), where $\beta = \alpha(x,y) \cap \alpha(z,w) = \alpha(x,w) \cap (y,z)$. Lemma 7.5.1 gives an interpretation of the quantity $xy \wedge zw$ in the context of almost-hyperbolic geodesic spaces.

Suppose now that $V = \{v_0, v_1, \ldots, v_n\}$ is a set of $n+1$ points with pseudometric ρ. We set

$$\Lambda(v_0, v_1, \ldots, v_n; \rho) = \max\{v_i v_{i+1} \wedge v_{j+1} v_j \mid 0 \leq i < j \leq n\}.$$

Thus, $\Lambda(v_0, v_1, \ldots, v_n; \rho) \geq 0$, by Lemma 7.3.3(4).

Proposition 7.3.4 : *Suppose that $V = \{v_0, v_1, \ldots, v_n\}$ is a set of $n+1$ points with $n \geq 2$, and that ρ is a pseudometric on V. Let $k = \operatorname{hyp}\rho$ and $l = \Lambda(v_0, v_1, \ldots, v_n; \rho)$. Then,*

$$\sum_{i=1}^{n} \rho(v_i, v_{i-1}) \leq \rho(v_0, v_n) + 2(2n-3)l + 2(3n-4)k.$$

The essential point is that the terms in k and l are both linear in n. In fact, the term in l is the best possible, though there is some room for improvement in the term in k.

The case $k = 0$ is not hard to deduce (set $k = 0$ in the argument presented below, and ignore most of the proof). Suppose we know this, and we are given a general pseudometric ρ on V. Then Proposition 7.3.1 gives us a ρ_τ, with $\operatorname{hyp}\rho_\tau = 0$, and

$$\rho_\tau \leq \rho \leq \rho_\tau + k(1 + \log_2 n).$$

We see that

$$l_\tau = \Lambda(v_0, v_1, \ldots, v_n; \rho_\tau) \leq l + k(1 + \log_2 n).$$

Thus,

$$\begin{aligned}\sum_{i=1}^{n} \rho(v_i, v_{i-1}) &\leq \sum_{i=1}^{n} \rho_\tau(v_i, v_{i-1}) + kn(1 + \log_2 n) \\ &\leq \rho_\tau(v_0, v_n) + 2(2n-3)l_\tau + kn(1 + \log_2 n) \\ &\leq \rho(v_0, v_n) + 2(2n-3)l + (5n-6)(1 + \log_2 n)k.\end{aligned}$$

So this gives the term in k to be $O(n \log n)$. If we want it $O(n)$, we need a more careful argument.

Proof of Proposition 7.3.4 : Let $V = \{v_0, v_1, \ldots, v_n\}, \rho, l, k$ be as in the hypothesis. We shall construct a spanning tree for V. We begin exactly as in the proof of Proposition 7.3.1. Let $V' = V \backslash \{v_0\}$, and let

$$\mathcal{I} = \{(t, v) \subseteq \mathbf{R} \times V' \mid 0 \leq t \leq \rho(v, v_0)\} = \bigsqcup_{v \in V'} I_v.$$

Write

$$\begin{aligned}\langle v,w\rangle &= vv_0 \wedge wv_0 \\ &= \frac{1}{2}\big(\rho(v,v_0)+\rho(w,v_0)-\rho(v,w)\big),\end{aligned}$$

so that, for $u,v,w \in V'$, we have

$$\langle u,w\rangle \geq \min(\langle u,v\rangle,\langle v,w\rangle) - k/2.$$

For $i \in \{1,2,\ldots,n\}$, let R_i be the relation $R_{v_{i-1},v_i} \subseteq I_{v_{i-1}} \times I_{v_i}$, i.e.

$$\big((s,v_{i-1}),(t,v_i)\big) \in R_i \Leftrightarrow s=t \leq \langle v_{i-1},v_i\rangle.$$

So far, there has been no difference from Proposition 7.3.1. This time, however, we set $R \subseteq \mathcal{I} \times \mathcal{I}$ to be the equivalence relation generated by the R_i for $i \in \{1,2,\ldots,n\}$ (rather than using all the $R_{v,w}$). Again, $T = \mathcal{I}/R$ is a tree with path-metric σ and a natural map $f : V \longrightarrow T$. Let $\tau = (T,\sigma,f)$. By attaching arcs of length 0, as described above, we can identify V with $\operatorname{ext} T$, and $\sigma|V$ with ρ_τ. We shall abbreviate $\rho(v_i,v_j)$, $\rho_\tau(v_i,v_j)$ and $\langle v_i,v_j\rangle$ respectively to $\rho(i,j)$, $\rho_\tau(i,j)$ and $\langle i,j\rangle$. We shall write $\alpha(i,j)$ for $\operatorname{span}(v_i,v_j)$. (Thus, $\rho_\tau(i,j)$ is the σ-length of $\alpha(i,j)$.)

From the construction of T, we have that $\rho_\tau(0,i)=\rho(0,i)$ and $\rho_\tau(i-1,i)=\rho(i-1,i)$ for each $i \in \{1,2,\ldots,n\}$.

Given $x,y \in T$, we shall write $x \leq y$ to mean that $x \in \operatorname{span}(v_0,y)$. Thus, $\leq$ is a partial ordering on the points of T. We write $x < y$ if $x \leq y$ and $x \neq y$. Given $x \in T$, we write $V(x) = \{v \in V \mid x \leq v\}$. From the construction of T, we see that $V(x)$ necessarily consists of a consecutive sequence of points $\{v_{p(x)}, v_{p(x)+1},\ldots,v_{q(x)}\}$ of points of V, where $0 \leq p(x) \leq q(x) \leq n$.

Suppose that $x \in T\backslash(\operatorname{ext} T \cup \operatorname{node} T)$. We see that $x \in \alpha(i-1,i)$ if and only if i equals either $p(x)$ or $q(x)+1$, using the convention that $n+1 \equiv 0$. Thus, the closed path $\bigcup_{i=1}^{n+1} \alpha(i-1,i)$ traverses each component of $T\backslash(\operatorname{ext} T \cup \operatorname{node} T)$ precisely twice. Thus,

$$\sum_{i=1}^{n} \rho_\tau(i-1,i) = \rho_\tau(0,n) + 2\sigma(T\backslash\alpha(0,n)).$$

Since $\rho_\tau(i-1,i)=\rho(i-1,i)$ and $\rho_\tau(0,n)=\rho(0,n)$, we see that we need to show that

$$\sigma(T\backslash\alpha(0,n)) \leq (2n-3)l + (3n-4)k.$$

We shall split this into two parts. To each node $x \in \operatorname{node} T$, we shall associate a subset $\beta(x) \subseteq T$ containing the point x, which is either a closed arc or a single point. Writing $\beta = \bigcup\{\beta(x) \mid x \in \operatorname{node} T\}$, we shall show that

$$\sigma(\beta) \leq (n-1)k,$$

and that

$$\sigma\big(T\backslash(\beta \cup \alpha(0,n))\big) \leq (2n-3)(l+k).$$

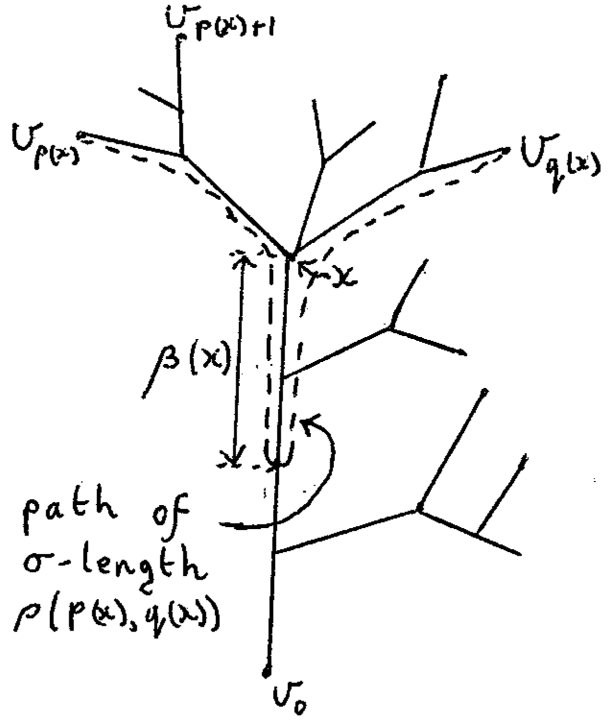

Figure 7a.

We shall begin with the first part. In fact, we shall show that for any $x \in \mathrm{node}\, T$, we have $\sigma(\gamma(x)) \leq (\deg x - 2)k$, where

$$\gamma(x) = \beta(x) \setminus \bigcup \{\beta(y) \mid y \in \mathrm{node}\, T \text{ and } x < y\}.$$

If $z \in \beta$, and we choose $x \in \mathrm{node}\, T$, maximal with respect to $<$ such that $z \in \beta(x)$, then $z \in \gamma(x)$. It follows that $\beta = \bigcup\{\gamma(x) \mid x \in \mathrm{node}\, T\}$. We may then deduce that

$$\sigma(\beta) \leq \sum_{x \in \mathrm{node}\, T} (\deg x - 2)k = (|\mathrm{ext}\, T| - 2)k = (n-1)k,$$

as required.

So, suppose that $x \in \mathrm{node}\, T$ so that $V(x) = \{v_{p(x)}, v_{p(x)+1}, \ldots, v_{q(x)}\}$. We define $\beta(x)$ as follows. If $\sigma(x, v_0) \leq \langle p(x), q(x) \rangle$, we set $\beta(x) = \{x\}$. If $\sigma(x, v_0) > \langle p(x), q(x) \rangle$, then we set $\beta = \mathrm{span}\,(x, y)$, where $y \in \sigma(x, v_0)$ is the point with $\sigma(v_0, y) = \langle p(x), q(x) \rangle$. (It is conceivable that there may be some ambiguity in the choice of y, arising from the fact that σ is only a pseudometric. Though this does not give us serious problem— all distances and lengths are well-defined.) We could alternatively define $\beta(x)$ as the closed arc (or point) in $\mathrm{span}\,(x, v_0)$, with one endpoint at x, and of σ-length equal to $\frac{1}{2}\max\big(0, \rho(p(x), q(x)) - \rho_\tau(p(x), q(x))\big)$. (See Figure 7a.) If $x \in \mathrm{ext}\, T$, we shall write $\beta(x) = \{x\}$. In each case, we have

$$\sigma(v_0, \beta(x)) = \min\big(\sigma(v_0, x), \langle p(x), q(x) \rangle\big).$$

For any $x \in \operatorname{node} T$, the set $V(x) = \{v_{p(x)}, v_{p(x)+1}, \dots, v_{q(x)}\}$ has a naturally partition as $V(x) = \bigsqcup_{i=1}^{r} V(x_i)$, where $r = \deg x - 1$ and $x_1, x_2, \dots, x_r \in \operatorname{node} T \cup \operatorname{ext} T$ are vertices of the tree T, adjacent to x. In fact, we can express the sequence $p(x), p(x)+1, \dots, q(x)$ as the concatenation of the consecutive sequences $p(x_i), p(x_i)+1, \dots, q(x_i)$ for $i = 1, 2, \dots, r$. Thus $p(x_1) = p(x)$, $q(x_r) = q(x)$ and $q(x_i) + 1 = p(x_{i+1})$ for $i \in \{1, 2, \dots, r-1\}$. (See Figure 7b.)

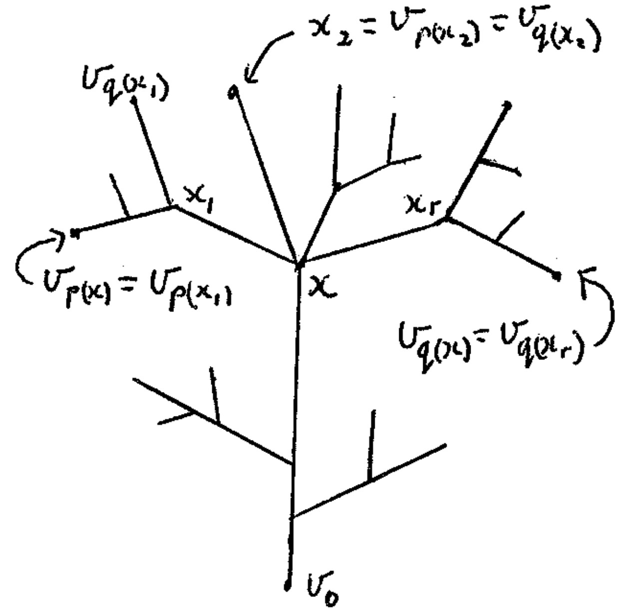

Figure 7b.

By the construction of T, we have that $\langle q(x_i), q(x_i)+1 \rangle = \sigma\big(v_0, \alpha(q(x_i), q(x_i)+1)\big) = \sigma(v_0, x)$ for each $i \in \{1, 2, \dots, r-1\}$.

Now, for any $i \in \{1, 2, \dots, r\}$, we see that

$$\begin{aligned} \sigma(\beta(x) \backslash \beta(x_i)) &= \max\big(0, \sigma(v_0, \beta(x_i)) - \sigma(v_0, \beta(x))\big) \\ &= \max\big(0, \langle p(x_i), q(x_i) \rangle - \langle p(x), q(x) \rangle\big). \end{aligned}$$

Also, $\sigma(\beta(x)) = \max\big(0, \sigma(v_0, x) - \langle p(x), q(x) \rangle\big)$. Since $\gamma(x) \subseteq \beta(x) \backslash \bigcup_{i=1}^{r} \beta(x_i)$, we have that

$$\sigma(\gamma(x)) \leq \max\big(0, \min(\{\langle p(x_i), q(x_i) \rangle \mid 1 \leq i \leq r\} \cup \{\sigma(v_0, x)\}) - \langle p(x), q(x) \rangle\big).$$

Now hyp $\rho = k$. So, after $(2r-2)$ applications of the inequality $\langle u, w \rangle \geq \min(\langle u, v \rangle, \langle v, w \rangle) - k/2$ (to the sequence $p(x) = p(x_i), q(x_1), q(x_1)+1 = p(x_2), q(x_2), \dots, p(x_r), q(x_r) = q(x)$), we find that

$$\langle p(x), q(x) \rangle \geq \min(\{\langle p(x_i), q(x_i) \rangle \mid 1 \leq i \leq r\} \cup \{\sigma(v_0, x)\}) - (2r-2)(k/2).$$

We conclude that

$$\sigma(\gamma(x)) \leq (2r-2)(k/2) = (r-1)k = (\deg x - 2)k,$$

as required.

We have shown that $\sigma(\beta) \leq (n-1)k$. It remains to show that $\sigma\big(T\backslash(\beta \cup \alpha(0,n))\big) \leq (2n-3)(l+k)$.

By an "edge" of T, we mean the closure of a component of $T \backslash \operatorname{node} T$. Now T has at most $2|\operatorname{ext} T| - 3 = 2n-1$ edges. Thus, there are at most $2n-3$ edges not lying in $\alpha(0,n)$. Let $e = \operatorname{span}(x,y)$ be such an edge, where $y \in \operatorname{node} T$, $x \in \operatorname{node} T \cup \operatorname{ext} T$ and $y \in \operatorname{span}(v_0, x)$. We claim that $\sigma(e \backslash \beta(x)) \leq l + k$. The result then follows by summing over all such edges.

Suppose first that $x \in \operatorname{node} T$. Write $p = p(x)$ and $q = q(x)$. If $e \subseteq \beta(x)$, then there is nothing to prove, so we can assume that

$$\sigma(v_0, y) = \sigma(v_0, e) < \sigma(v_0, \beta(x)) \leq \langle p, q \rangle.$$

Now, $V(x) = \{v_p, v_{p+1}, \ldots, v_q\}$ is, by definition, the set of points of V separated from v_0 by x. Since y is the adjacent node to x in the direction of v_0, we must have that both $\alpha(p-1,p)$ and $\alpha(q,q+1)$ pass through y. Thus,

$$\langle p-1, p \rangle = \sigma(v_0, \alpha(p-1,p)) \leq \sigma(v_0, y) \leq \langle p, q \rangle.$$

Similarly,

$$\langle q, q+1 \rangle \leq \langle p, q \rangle.$$

Writing

$$\langle i, j \rangle_\tau = \frac{1}{2}\big(\rho_\tau(0,i) + \rho_\tau(0,j) - \rho_\tau(i,j)\big) = \sigma(v_0, \alpha(i,j)),$$

(so that $\langle i, i+1 \rangle_\tau = \langle i, i+1 \rangle$ for each i) we see that

$$\begin{aligned}
\langle p-1, q+1 \rangle &\geq \min(\langle p-1,p \rangle, \langle p,q \rangle, \langle q,q+1 \rangle) - 2(k/2) \\
&= \min(\langle p-1,p \rangle_\tau, \langle q,q+1 \rangle_\tau) - k \\
&= \langle p-1, q+1 \rangle_\tau - k.
\end{aligned}$$

Thus,

$$\rho_\tau(p-1, q+1) \geq \rho(p-1, q+1) - 2k.$$

(Recall that $\rho_\tau(0,p-1) = \rho(0,p-1)$ and $\rho_\tau(0,q+1) = \rho(0,q+1)$.) From the definition of $\beta(x)$, we have that $\sigma(\beta(x)) = \frac{1}{2}\max\big(0, \rho(p,q) - \rho_\tau(p,q)\big)$, so that

$$\rho_\tau(p,q) + 2\sigma(\beta(x)) \geq \rho(p,q).$$

Also, we have $\rho_T(p-1,p) = \rho(p-1,p)$ and $\rho_T(q,q+1) = \rho(q,q+1)$. From the construction of T, we have that $e = \alpha(p-1,p) \cap \alpha(q,q+1)$, and so

$$\begin{aligned}\sigma(e\backslash\beta(x)) &= \sigma(e) - \sigma(\beta(x)) \\ &= \frac{1}{2}\big(\rho_T(p-1,p) + \rho_T(q,q+1) - \rho_T(p,q) - \rho_T(p-1,q+1)\big) - \sigma(\beta(x)) \\ &\leq \frac{1}{2}\big(\rho(p-1,p) + \rho(q,q+1) - \rho(p,q) - (\rho(p-1,q+1) - 2k)\big) \\ &= v_{p-1}v_p \wedge v_{q+1}v_q + k \\ &\leq l + k,\end{aligned}$$

as required.

The case where $x \in \operatorname{ext} T$ is similar, but simpler. We have, by definition, that $\beta(x) = \{x\}$. If we set $p = q$ so that $x = v_p = v_q$ then the argument goes through, more or less, as before.

◇

7.4. Pinnate Structures.

In this section, we describe a few combinatorial properties of trees.

Let T be a tree (i.e. a finite acyclic connected 1-complex). Let $V = \operatorname{ext} T$ be the set of extreme points of T. Suppose x, y, z, w are distinct points of V. We shall write $xy|zw$ to mean that $\operatorname{span}(x,y)$ meets $\operatorname{span}(z,w)$ in at most one point. (Figure 7c.)

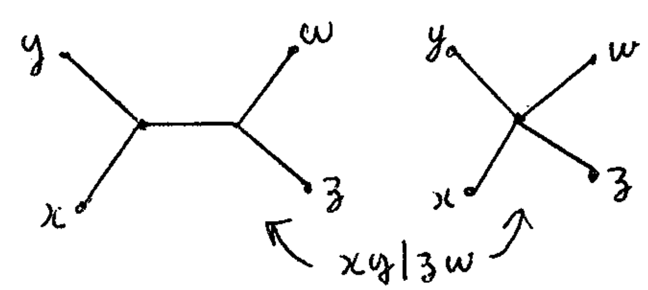

Figure 7c.

We have the following properties.

(1) $xy|zw \Leftrightarrow xy|wz \Leftrightarrow zw|xy$,

(2) $xy|zw$ and $xz|yw \Rightarrow xw|yz$.

Remark : Given a finite set V, we can regard the the quaternary relation $(..)|(..)$ as defining a structure on V equivalent to the notion of a combinatorial spanning tree for V, with V equal to the set of extreme points. As axioms for $(..)|(..)$, we can take properties

(1) and (2) stated above, as well as an enumeration of the possibilities for the relation restricted to any set of five distinct points $\{a,b,c,d,e\} \subseteq V$. That is, we need to assume that the relation on $\{a,b,c,d,e\}$ is derived from one of the possible combinatorial spanning trees for five points. (There are twenty-six such, in total.) We leave as an exercise the observation that one can reconstruct uniquely a spanning tree for V from such data.

In this section, we want to describe an additional structure on V, which we shall call a "pinnate structure".

Let T be a tree with $V = \operatorname{ext} T$.

Definition : A *pinnate decomposition* of T is a partition of T into disjoint subsets indexed by V, written $T = \bigsqcup_{x \in V} \gamma(x)$, such that $x \in \gamma(x)$ for all x, and $\gamma(x)$ is homeomorphic to a half-open interval for each $x \in V$, except one point $x_0 \in V$ for which $\gamma(x_0) = \{x_0\}$.

We refer to x_0 as the *root*. We call $\gamma(x)$ for $x \in V \setminus \{x_0\}$ a *branch* of the pinnate decomposition.

If $x \in V$, then there is a unique $\phi x \in V$ such that the closure $\bar{\gamma}(x)$ of $\gamma(x)$ meets $\gamma(\phi x)$. Thus, $\phi x_0 = x_0$, otherwise $\phi x \neq x$. Also, since x_0 is an extreme point of T, there is a unique point $x_1 \in V$ with $\phi x_1 = x_0$. We call $\operatorname{span}(x_0, x_1) = \{x_0\} \cup \gamma(x_1)$ the *stem* of the pinnate decomposition.

Thus, we may imagine a tree with a pinnate decomposition as resembling a fern frond with one main stem, and a sequence of successive branching. (Figure 7d.)

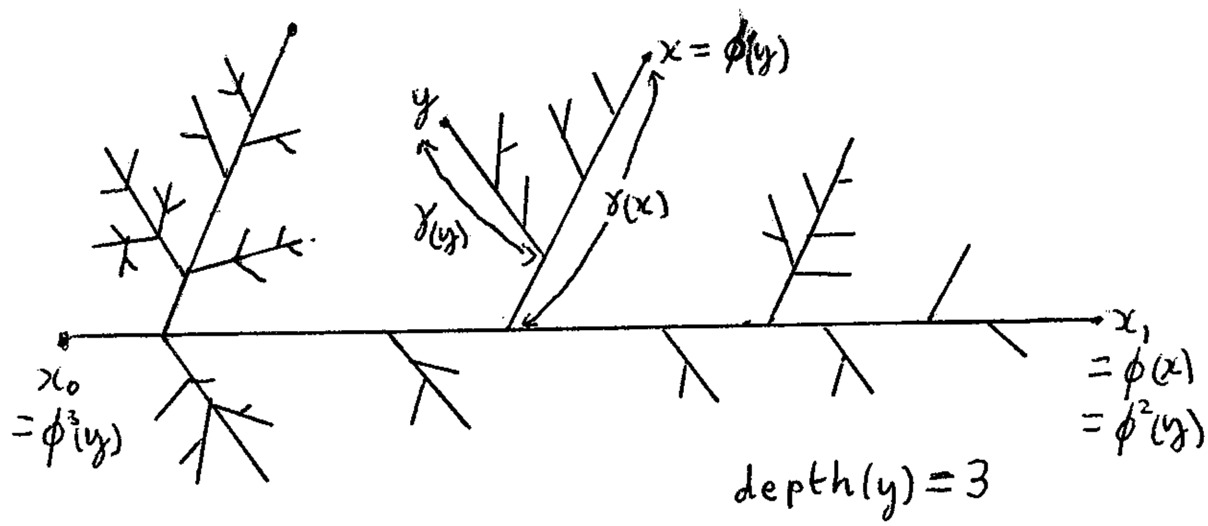

Figure 7d.

Given T, we may recover the pinnate decomposition from the map ϕ by taking

$$\gamma(x) = \operatorname{span}(x, \phi x, \phi^2 x) \setminus \operatorname{span}(\phi x, \phi^2 x)$$

for $x \neq x_0$, and

$$\gamma(x_0) = \{x_0\},$$

where x_0 is the unique fixed point of ϕ.

We call a map $\phi : V \longrightarrow V$ arising in this way a *pinnate structure* on T (or a pinnate structure on V compatible with T).

Remark : We may give a set of axioms for a pinnate structure $\phi : V \longrightarrow V$ compatible with $(..)|(..)$ as follows.

(1) There is some $x_0 \in V$ with $\phi x_0 = x_0$.

(2) If $x, y \in V \backslash \{x_0\}$ and $\phi x = \phi y = x_0$, then $x = y$.

(3) If $\phi^2 x \neq x_0$, then $x, \phi x, \phi^2 x, \phi^3 x$ are all distinct.

(4) If $x, y, \phi x, \phi y$ are all distinct, then $x(\phi x)|y(\phi y)$.

(5) If $x, y, \phi x, \phi^2 x$ are all distinct, and $\phi x = \phi y$, then either $xy|(\phi x)(\phi^2 x)$ or $yx|(\phi x)(\phi^2 x)$.

We leave as an exercise that this data suffices to reconstruct a pinnate decomposition $\{\gamma(x) \mid x \in V\}$ as described above.

Let $\phi : V \longrightarrow V$ be a pinnate structure compatible with T. If $\phi^t x \neq x_0$, then it follows that $x, \phi x, \phi^2 x, \ldots, \phi^t x$ are all distinct. Thus, there is some r such that $\phi^r x = x_0$. We write $\operatorname{depth}_\phi(x)$ (or just $\operatorname{depth}(x)$) for the smallest such r. We write

$$\operatorname{depth} \phi = \max_{x \in V}(\operatorname{depth}_\phi(x)).$$

Proposition 7.4.1 : *Any tree T admits a pinnate structure of depth at most $1 + \log_2(\frac{m+1}{3})$, where $m = |\operatorname{ext} T|$.*

The worst case is, in fact, a tree with 3.2^{p-2} extreme points, and each node of degree three, as shown in Figure 7e for $p = 5$. Any pinnate structure on such a tree must have depth at least p.

We shall first prove the following lemma (7.4.2). In fact, this lemma suffices to give a logrithmic bound on the depth of a pinnate structure, which is all we shall need for Theorem 7.6.1.

Lemma 7.4.2 : *Suppose T is a tree with at most 2^p extreme points, and that $x_0 \in \operatorname{ext} T$ is any extreme point. Then T admits a pinnate structure of depth at most p and with root x_0.*

Proof : Let $V = \operatorname{ext} T$. We construct a pinnate decomposition for T. We begin by constructing its stem α as follows. We imagine moving along T, starting at x_0. Each time we come to a node $y \in \operatorname{node} T$, we follow the edge of T which separates a maximal number of extreme points from x_0. In other words, we move into a component C of $T \backslash \{y\}$ which maximises $|C \cap V|$ for $x_0 \notin C$. We continue until we reach an extreme point $x_1 \in V$. Let $\alpha = \operatorname{span}(x_0, x_1)$ be the path we have followed. Let $\gamma(x_1) = \alpha \backslash \{x_0\}$.

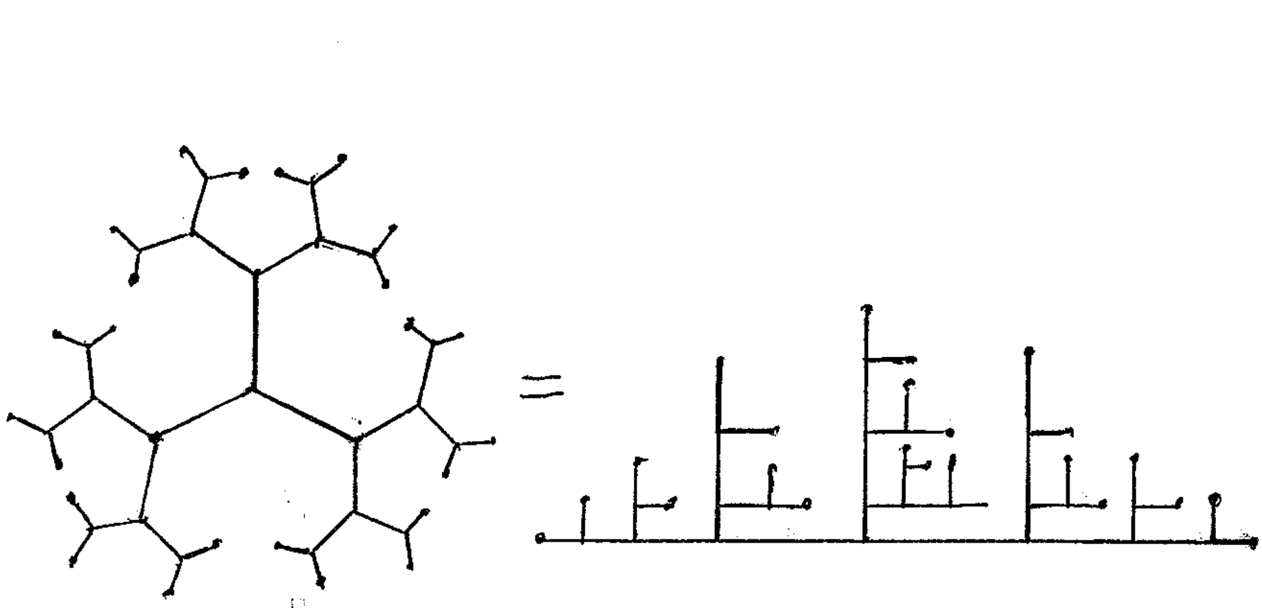

Figure 7e.

Now, if $D \subseteq T$ is a component of $T\backslash\alpha$, then by construction, $|D \cap V| \leq 2^{p-1} - 1$. The closure $\bar{D} \subseteq T$ is a subtree with $\operatorname{ext}\bar{D} = (D \cap V) \cup \{x_D\}$, where x_D is the point where $\bar{D}$ meets α. Thus, $|\operatorname{ext}\bar{D}| \leq 2^{p-1}$. By induction, $\bar{D}$ has a pinnate decomposition $\bar{D} = \bigsqcup_{x \in \operatorname{ext}\bar{D}} \gamma_D(x)$, with root x_D, and depth at most p–1. If $x \in D \cap V$, set $\gamma(x) = \gamma_D(x)$.

Doing this for each such component D, we arrive at a pinnate decomposition $\{\gamma(x) | x \in V\}$ of depth at most p, and root x_0.

$\diamondsuit$

Proof of Proposition 7.4.1 : Let $V = \operatorname{ext} T$, and suppose that $|V| \leq 3.2^{p-1} - 1$ with $p \geq 2$. We want to construct a pinnate decomposition of depth at most p. (The case $p = 1$ is trivial.)

Suppose that y is a node of T. Let $C_1, C_2, \ldots, C_r$ ($r \geq 2$) be the connected components of $T\backslash\{y\}$. Let $e_i = |C_i \cap V|$, so that $\sum_{i=1}^{r} e_i = |V|$. For each $j \in \{1, 2, \ldots, r\}$, we must have $e_j \leq \sum_{i \neq j} e_i$, otherwise we would do better choosing the node in C_j adjacent to y. It follows that

$$e_j \leq \frac{1}{2} \sum_{i=1}^{r} e_i \leq \frac{1}{2}(3.2^{p-1} - 1) \leq 2^p - 1.$$

Now, without loss of generality, we have $e_1 \geq e_2 \geq \cdots \geq e_r$. Since $\sum_{i=1}^{r} e_i \leq 3.2^{p-1} - 1$, we must have $e_i \leq 2^{p-1} - 1$ for $i \geq 3$.

Considering each closures $\bar{C}_i$ as a subtree of T, we have $\operatorname{ext}\bar{C}_i = (C_i \cap V) \cup \{y\}$. Thus, $|\operatorname{ext}\bar{C}_i| \leq 2^p$ for $i = 1, 2$, and $|\operatorname{ext}\bar{C}_i| \leq 2^{p-1}$ for $i \geq 3$. For each i, let $\{\gamma_i(x) \mid x \in \operatorname{ext}\bar{C}_i\}$ be the pinnate decomposition of $\bar{C}_i$, with root y, given by Lemma 7.4.2. Let $x_0 \in C_1 \cap V$ and $x_1 \in C_2 \cap V$ be, respectively, the unique points such that $\bar{\gamma}_1(x_0)$ and $\bar{\gamma}_2(x_1)$ meet y. (Thus span(y, x_0) and span(y, x_1) are, respectively, the stems for C_1 and C_2.)

Define $\gamma(x_0) = \{x_0\}$, $\gamma(x_1) = \operatorname{span}(x_0, x_1)\backslash\{x_0\}$, and $\gamma(x) = \gamma_i(x)$ for $x \in (C_i \cap V)\backslash \{x_0, x_1\}$. We check that $\{\gamma(x) \mid x \in V\}$ is a pinnate decomposition of T of depth at most p.

◇

Suppose that $\phi : V \longrightarrow V$ is a pinnate structure compatible with the tree T. Let $\{\gamma(x) \mid x \in V\}$ be the corresponding decomposition of T. For any $x \in V$, we have that $\gamma(z)$ meets $\operatorname{span}(x_0, x)$ if and only if $z = \phi^t x$ for some t with $0 \le t \le r = \operatorname{depth} x$. We call $x_0 = \phi^r x, \phi^{r-1} x, \ldots, \phi x, x$ the *path sequence* for x. (Figure 7f.)

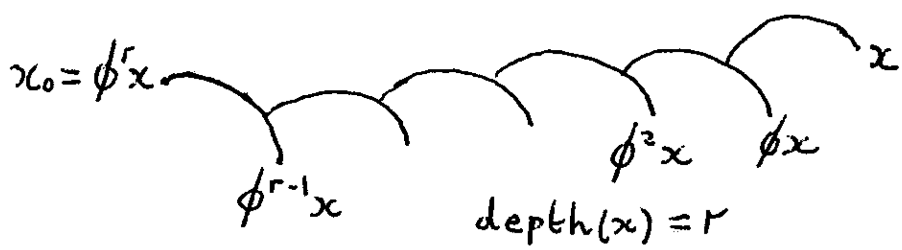

Figure 7f.

More generally, suppose that $x, y \in V$ with $r = \operatorname{depth} x$ and $s = \operatorname{depth} y$. Let $t \ge 0$ be the largest integer such that $\phi^{r-t} x = \phi^{s-t} y$. Then, $\gamma(z)$ meets $\operatorname{span}(x, y)$ if and only if z belongs to the sequence

$$x, \phi x, \ldots, \phi^{r-t} x = \phi^{s-t} y, \ldots, \phi y, y.$$

We write $w = w(x, y) = \phi^{r-t} x = \phi^{s-t} y$. The path sequences for x and y agree on the first $t+1$ terms, namely, $x_0 = \phi^t w, \ldots, \phi w, w$. Note that this is precisely the path sequence for w. We must have either $x x_0 | y w$ or $y x_0 | x w$. (Figure 7g.)

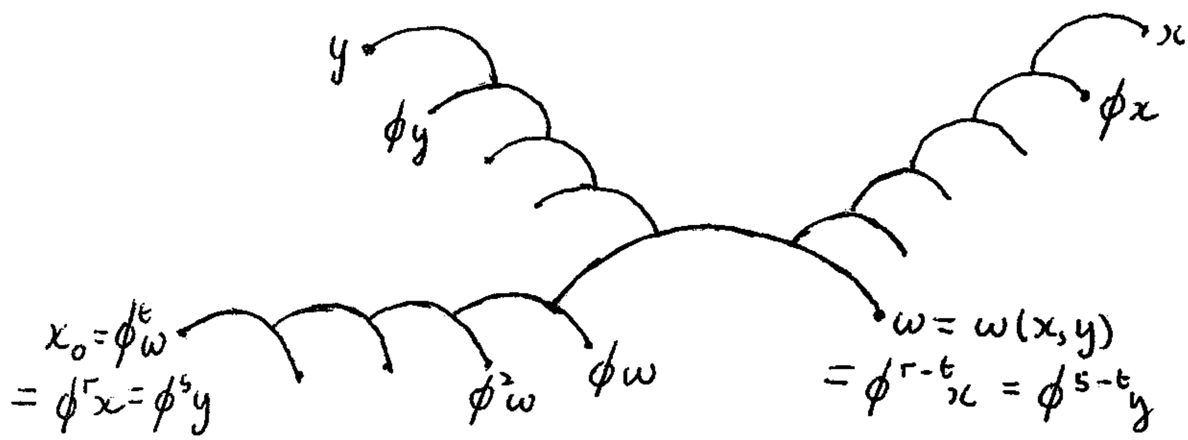

Figure 7g.

This dicussion will be relevant to Section 7.6.

7.5. Piecewise geodesic paths.

In this section, we return to the context of almost-hyperbolic geodesic spaces.

Suppose that V is a set of $n+1$ points in such a space. We have already stated (in Section 7.1) that both Steiner trees, and the construction of Section 3.3, give trees which approximate distances up to an additive term $kO(n)$. The property both trees have in common which gives rise to this linear estimate may be summarised as follows. Suppose $X, Y \in V$, and that β is the path in the tree joining X to Y. Then β consists of the union of at most n geodesic segments, no two of which run "almost parallel" over a long distance. The result can thus be rephrased as a property of such paths, namely that $\text{length}\,\beta \leq d(X,Y) + kO(n)$. We shall see that this is a simple consequence of Proposition 7.3.4. In fact, we are able to weaken the hypothesis, taking account of the direction of geodesic segments. Thus, we need only assume that the path β should not double back on itself over a large distance.

We introduce the following notation. Suppose that $(\mathcal{S}, d)$ is a (pseudo)-metric space. If $X_0, X_1, \ldots, X_p \in \mathcal{S}$, we shall write $\langle X_0 X_1 \cdots X_p \rangle$ to mean that

$$d(X_0, X_p) = \sum_{i=1}^{p} d(X_i, X_{i-1}).$$

If $(\mathcal{S}, d)$ happens to be a geodesic space, then we shall always choose geodesics so that $[X_i, X_j] \subseteq [X_0, X_p]$ for any $i, j \in \{0, 1, \ldots, p\}$. Thus we may imagine the sequence of points $X_0, X_1, \ldots, X_p$ occurring in order along $[X_0, X_p]$. If X, Y, Z, W are any points in $(\mathcal{S}, d)$, we have already defined (Section 7.3)

$$XY \wedge ZW = \frac{1}{2}\big(d(X,Y) + d(Z,W) - d(X,Z) - d(Y,W)\big).$$

To the properties (1)–(6) of Lemma 7.3.3, we may add that if $X, Y, Z, W, X', Y', Z', W' \in \mathcal{S}$ with $\langle XX'Y'Y \rangle$ and $\langle ZZ'W'W \rangle$, then $X'Y' \wedge Z'W' \leq XY \wedge ZW$. We have already observed that we can express property H1 in terms of this notation. In fact, we can write $XY : ZW$ as

$$XZ \wedge YW \simeq XW \wedge YZ \succeq 0.$$

If we already know that $(\mathcal{S}, d)$ is H1, then $XY : ZW$ becomes equivalent to the statement $XZ \wedge YW \simeq XW \wedge YZ$, or to the statement $XZ \wedge YW \simeq 0$.

Suppose that $(\mathcal{S}, d)$ is an almost hyperbolic space. If $XZ : YW$, and $(XZ)AB(YW)$ is a spanning tree for X, Y, Z, W (Section 3.1), then it is readily checked that

$$XY \wedge ZW \simeq XW \wedge ZY \simeq d(A, B).$$

Thus, the quantity $XY \wedge ZW$ may be thought of as measuring the distance over which the geodesics $[X, Y]$ and $[Z, W]$ run almost parallel.

The following lemma gives a more careful formulation of this idea.

Lemma 7.5.1 : *$\forall k\ \exists h$ such that the following holds.*

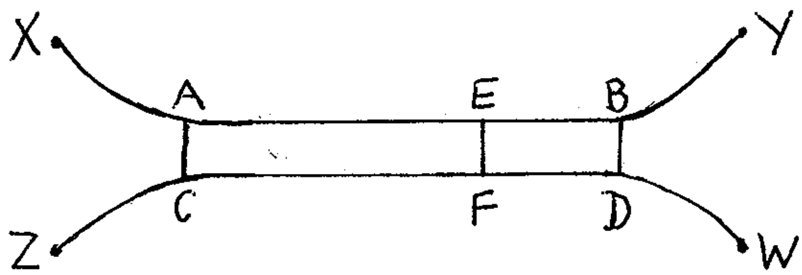

Figure 7h.

Suppose $(\mathcal{S}, d)$ is a k-H1 geodesic space, and that $X, Y, Z, W \in \mathcal{S}$ satisfy $XY \wedge ZW \geq 0$. Then, there exist $A, B, C, D \in \mathcal{S}$ with $\langle XABY \rangle$ and $\langle ZCDW \rangle$ and $d(A,B) = d(C,D) = XY \wedge ZW$. Also, if $E \in [A,B]$ and $F \in [C,D]$ satisfy $d(A,E) = d(C,F)$, then $d(E,F) \leq h$. (In particular, $d(A,C) \leq h$ and $d(B,D) \leq h$.) See Figure 7h.

Proof :

Case (1) : $XW : YZ$ (or equivalently, $XY \wedge ZW \simeq 0$).

We have $d([X,Y],[Z,W]) \simeq 0$. Since $XY \wedge ZW \leq \min(d(X,Y), d(Z,W))$ (Lemma 7.3.3(3)), we can find $A \sim B \sim C \sim D$ with $\langle XABY \rangle$, $\langle ZCDW \rangle$ and $d(A,B) = d(C,D) = XY \wedge ZW$.

Case (2) : $XY : ZW$.

We have $XY \wedge ZW = -XZ \wedge YW \preceq 0$. So, $XY \wedge ZW \simeq 0$, and we are back in Case (1).

Case (3) : $XZ : YW$.

Let $(XZ)PQ(YW)$ be a spanning tree for X, Y, Z, W (Section 3.1). Thus, $d(P,Q) \simeq XY \wedge ZW$.

Now, there exist $A', B' \in [X,Y]$ with $A' \sim P$ and $B' \sim Q$, so that $d(A',B') \simeq d(P,Q) \simeq XY \wedge ZW$. Since $d(X,Y) \geq XY \wedge ZW$, we can find $A, B \in [X,Y]$ with $A \sim A'$ and $B \sim B'$ and $d(A,B) = XY \wedge ZW$. Now,

$$\begin{aligned} d(X,Y) &\simeq d(X,P) + d(P,Q) + d(Q,Y) \\ &\simeq d(X,A) + d(A,B) + d(B,Y). \end{aligned}$$

Thus if $\langle XBAY \rangle$, then $d(A,B) \simeq 0$, and so $XY \wedge ZW \simeq d(A,B) \simeq 0$, and we are back in Case (1). Therefore, we can assume that $\langle XABY \rangle$.

Similarly, we can find $C, D \in [Z,W]$ with $C \sim P$, $D \sim Q$ and $\langle ZCDW \rangle$ and $d(C,D) = XY \wedge ZW$.

Now, suppose that $E \in [A,B]$ and $F \in [C,D]$ satisfy $d(A,E) = d(C,F)$. Applying Lemma 3.1.2 twice, we find $F' \in [C,D]$ with $F' \sim E$. Now, $d(C,F') \simeq d(C,E) \simeq d(A,E) \simeq d(C,F)$, and so $F \sim F'$. Thus $E \sim F$.

$\diamondsuit$

Suppose $X_0, X_1, \ldots, X_n$ is any sequence of points in the geodesic space $(\mathcal{S}, d)$. We shall write

$$[X_0, X_1, \ldots, X_n] = \bigcup_{i=1}^{n} [X_{i-1}, X_i]$$

for the piecewise geodesic path joining these points. If $\gamma = [X_0, X_1, \ldots, X_n]$, then, as in Section 7.3, we shall write

$$\Lambda(\gamma) = \Lambda(X_0, X_1, \ldots, X_n) = \max\{X_i X_{i+1} \wedge X_{j+1} X_j \mid 0 \leq i < j \leq n-1\}.$$

(Formally we should think of γ as consisting just of the sequence of points $(X_0, X_1, \ldots, X_n)$, though we are imagining it as a piecewise geodesic path, or as the image of this path in $\mathcal{S}$.) Note that if γ' is obtained by subdividing γ, i.e. by adding additional points along the geodesic segments of γ, then we have $\Lambda(\gamma') \leq \Lambda(\gamma)$.

If $(\mathcal{S}, d)$ is k-H1, then the quantity $\Lambda(\gamma)$ measures the maximum distance along which γ doubles back on itself (along a pair of geodesic segments). Proposition 7.3.4 tells us that the total length of γ is at most $d(X_0, X_n) + 2(2n-3)\Lambda(\gamma) + 2(3n-4)k$.

We shall apply this to the tree construction given in Section 3.3.

Let $V \subseteq \mathcal{S}$ be a set with $(n+1)$ elements. A linear ordering on V can be thought of as a bijection $\underline{X} : \{0, 1, \ldots, n\} \longrightarrow V$. We shall write $X_i = \underline{X}(i)$ and $\underline{X} = (X_0, X_1, \ldots, X_n)$. We define $T_V = T_{\underline{X}} \subseteq \mathcal{S}$ inductively as follows. We take $T_{(X_0, X_1)} = [X_0, X_1]$ and $T_{(X_0, X_1, \ldots, X_n)} = T_{(X_0, X_1, \ldots, X_{n-1})} \cup [X_n, Y_n]$, where Y_n is a nearest point to X_n on the tree $T_{(X_0, X_1, \ldots, X_{n-1})}$. We see that $T_{\underline{X}}$ an embedded tree in $\mathcal{S}$, and that $\operatorname{ext} T_{\underline{X}} \subseteq V$. Thus, $T_{\underline{X}}$ is an embedded spanning tree for V in the sense described in Section 7.1.

Suppose that $i, j \in \{0, 1, \ldots, n\}$, with $i > j$. Let α be the arc in $T_{\underline{X}}$ joining X_i to X_j. Then α is piecewise geodesic with at most n segments. In fact, α has one of the following forms (Figure 7i).

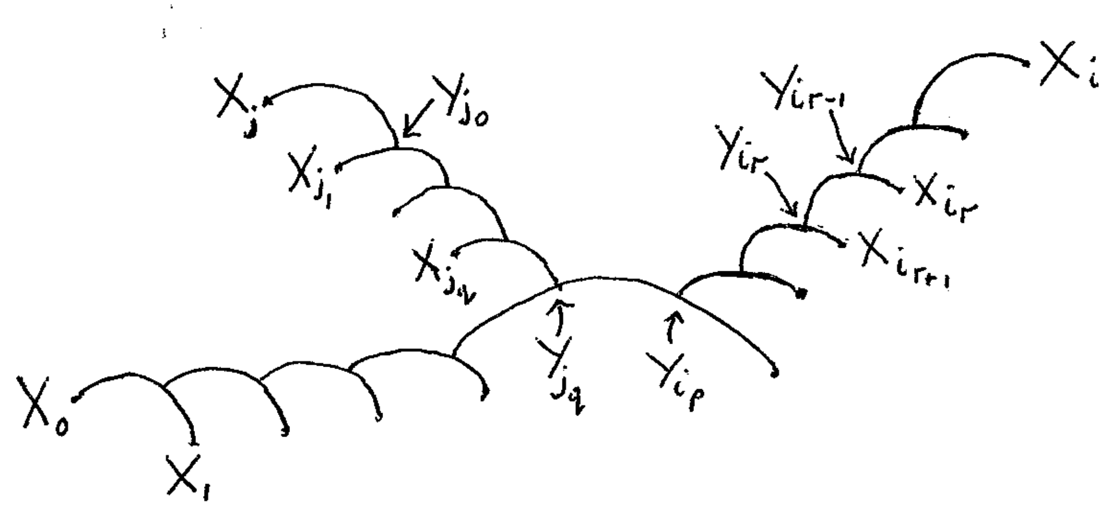

Figure 7i.

(1) $\alpha = [X_i, Y_{i_0}, Y_{i_1}, Y_{i_2}, \ldots, Y_{i_p}, Y_{j_q}, Y_{j_{q-1}}, \ldots, Y_{j_1}, Y_{j_0}, X_j]$ where $p, q \in \{0, 1, \ldots, n\}$ and $i = i_0 > i_1 > \cdots > i_p$ and $j = j_0 > j_1 > \cdots > j_q$. Also, the i_r and j_s are all distinct. Moreover, we have $Y_{i_r} \in [X_{i_{r+1}}, Y_{i_{r+1}}]$ for $0 \leq r \leq p-1$, and $Y_{j_s} \in [X_{j_{s+1}}, Y_{j_{s+1}}]$ for $0 \leq s \leq q-1$.
(2) $\alpha = [X_i, Y_{i_0}, Y_{i_1}, Y_{i_2}, \ldots, Y_{i_p}, X_j]$ where $p \in \{0, 1, \ldots, n\}$ and $i = i_0 > i_1 > \cdots > i_p$. Also $Y_{i_r} \in [X_{i_{r+1}}, Y_{i_{r+1}}]$ for $0 \leq r \leq p-1$.
(3) $\alpha = [X_0, X_1]$.

Suppose that $[A, B]$ and $[C, D]$ are two distinct geodesic segments of α, so that A, B, C, D occur in this order along α. (Possibly $B = C$.)

By inspection of the form of α given above, we see that we must have either $d(A, B) \leq d(A, [C, D])$ or $d(D, C) \leq d(D, [A, B])$. We claim that this implies that $AB \wedge DC \leq h$, where h is the constant of Lemma 7.5.1.

We can assume, without loss of generality, that $d(A, B) \leq d(A, [C, D])$, and (since $h \geq 0$) that $AB \wedge CD \geq 0$. Lemma 7.5.1 gives us points $A', B' \in [A, B]$ and $D' \in [C, D]$ with $\langle AA'B'B \rangle$, $d(A', B') = AB \wedge DC$ and $d(A', D') \leq h$. Now,

$$\begin{aligned} d(A, A') + AB \wedge DC &= d(A, A') + d(A', B') \leq d(A, B) \\ &\leq d(A, [C, D]) \leq d(A, A') + d(A', C') \\ &\leq d(A, A') + h. \end{aligned}$$

Thus, $AB \wedge DC \leq h$ as required.

Since this applies to any pair of geodesic segments of α, we conclude that

$$\Lambda(\alpha) \leq h.$$

As remarked in Section 7.2, the constant h must have the form λk for some universal $\lambda \in \mathbf{R}$. We may now apply Proposition 7.3.4 to find that

$$\begin{aligned} \operatorname{length} \alpha &\leq d(X_i, X_j) + 2(2n-3)(\lambda k) + 2(3n-4)k \\ &= d(X_i, X_j) + kH(n), \end{aligned}$$

where $H(n) = 2\lambda(2n-3) + 2(3n-4)$.

We have shown:

Proposition 7.5.2 : *There is a linear function $H : \mathbf{N} \longrightarrow \mathbf{R}$ such that the following holds.*

Suppose that $(\mathcal{S}, d)$ is a k-H1 geodesic space, and that $X_0, X_1, \ldots, X_n \in \mathcal{S}$. Let $T_{\underline{X}}$ be the embedded spanning tree as defined above. If $i, j \in \{0, 1, \ldots, n\}$ and α is the arc in $T_{\underline{X}}$ joining X_i to X_j, then

$$\operatorname{length} \alpha \leq d(X_i, X_j) + kH(n).$$

This proposition is used in the proof of Theorem 7.6.1.

We conclude this section with a brief discussion of Steiner trees.

Suppose that $(\mathcal{S}, d)$ is a geodesic space, and $V \subseteq \mathcal{S}$ a set of $(n+1)$ points. Among all spanning trees (T, σ, f, g) for V, we choose one of minimal length $\sigma(T)$. It is an exercise to show that this minimum is always attained in geodesic spaces. Such a tree must be embedded, so we may identify T with its image, S_V, in $\mathcal{S}$. Clearly, each edge of S_V must be geodesic. We call S_V a *Steiner tree* for V.

Suppose that $(\mathcal{S}, d)$ is k-H1 and that α an arc in S_V joining two points of V. A simple application of Lemma 7.5.1 shows that, as in the previous case, $\Lambda(\alpha) \leq h$ (otherwise we could construct a shorter tree). Applying Proposition 7.3.4 (as in the case of Proposition 7.3.2) we deduce:

Proposition 7.5.3 : *There is a linear function $J : \mathbf{N} \longrightarrow \mathbf{R}$ such that the following holds.*

Suppose that $(\mathcal{S}, d)$ is a k-H1 geodesic space, and that $V \subseteq \mathcal{S}$ is a set of $(n+1)$ points. Let S_V be a Steiner tree for V. If $X, Y \in V$, and α is the arc in S_V joining X to Y, then

$$\operatorname{length} \alpha \leq d(X, Y) + kJ(n).$$

7.6. A logrithmic spanning tree.

In this section we collect the various pieces together to show how to construct an $O(\log n)$-approximating tree.

Suppose $(\mathcal{S}, d)$ is a k-H1 geodesic space, and $V \subseteq \mathcal{S}$ is a set of $(n+1)$ points. We shall write ρ for the metric d restricted to V. Thus, $\operatorname{hyp} \rho \leq k$ (Section 7.3). Proposition 7.3.1 gives us an abstract spanning tree $\tau' = (T', \sigma', f')$ for V with

$$\rho_{\tau'} \leq \rho \leq \rho_{\tau'} + (1 + \log_2 n)k.$$

As remarked after the proof of Proposition 7.3.1, we can identify V with $\operatorname{ext} T$.

Now, Lemma 7.4.1 gives us a pinnate structure $\phi : V \longrightarrow V$ for T of depth p, where $p \leq 1 + \log_2 \left(\frac{n+2}{3}\right)$. Let X_0 be the root of this pinnate structure.

Suppose $X \in V$ has depth r. We defined, at the end of Section 7.4, the "path sequence" $\underline{X} = (X_0, X_1, \ldots, X_r)$ for X, where $X_i = \phi^{r-i}X$, so that $X = X_r$. If X' is some other point of V, of depth s, then the path sequence $\underline{X'} = (X_0, X'_1, \ldots, X'_s)$ for X' will agree with that for X precisely on some initial segment $(X_0, X_1, \ldots, X_t) = (X_0, X'_1, \ldots, X'_t)$. This initial segment is, itself, the path sequence, $\underline{W}$, for some point $W = W(X, X') \in V$ of depth t. (Possibly W is the same as X or X'.) In the tree T', we must have either $XX_0|X'W$ or $X'X_0|XW$ (see Section 7.4). We see that

$$\rho_{\tau'}(X, X') + \rho_{\tau'}(W, X_0) = \max\bigl(\rho_{\tau'}(X, W) + \rho_{\tau'}(X', X_0), \rho_{\tau'}(X', W) + \rho_{\tau'}(X, X_0)\bigr).$$

Thus,

$$\begin{aligned}\rho(X, X') &+ 2k(1 + \log_2 n) \\ &\geq \max\bigl(\rho(X, W) + \rho(X', X_0) - \rho(W, X_0), \rho(X', W) + \rho(X, X_0) - \rho(W, X_0)\bigr).\end{aligned}$$

We shall now use these path sequences to construct our immersed $\tau = (T, \sigma, f, g)$ as follows.

For each $X \in V$, we construct the embedded tree $T(X) = T_{\underline{X}}$, as defined in Section 7.5, where $\underline{X}$ is the path sequence for X. Since the definition of $T_{\underline{X}}$ was inductive, we can do this consistently over initial segments, so that for any $X, X' \in V$, the trees $T(X)$ and $T(X')$ agree on the common subtree $T(W(X, X'))$. Thus the trees $T(X)$, as X varies in V, form a directed set under inclusion. This allows us to define T as the direct limit of the $T(X)$.

More explicitly, let $\tilde{T} = \{(X, Y) \in V \times \mathcal{S} \mid Y \in T(X)\}$. We write $(X, Y) \sim (X', Y')$ if $Y = Y' \in T(W(X, X'))$. We check that $\sim$ is an equivalence relation, and set $T = \tilde{T}/\sim$. Let ι_X be the natural inclusion of $T(X)$ in T. We define $f : V \longrightarrow T$ by $f(X) = \iota_X(X)$. We define $g : T \longrightarrow \mathcal{S}$ by demanding, for each $X \in V$, that $g \circ \iota_X$ be the inclusion of $T(X)$ in $\mathcal{S}$. Thus, $g \circ f$ is the identity on V. Finally, the path-metrics on the $T(X)$, induced from $\mathcal{S}$, themselves induce a path-metric σ on T. Clearly, f is distance non-increasing from (T, σ) to $(\mathcal{S}, d)$. This defines our spanning tree $\tau = (T, \sigma, f, g)$.

Note that if $Y, Y' \in V$ both lie in the path sequence for X, then $f(Y), f(Y') \in \iota_X T(X) \subseteq T$. Applying Proposition 7.5.2, we see that

$$\rho(Y, Y') \leq \rho_\tau(Y, Y') = \sigma(f(Y), f(Y')) \leq \rho(Y, Y') + kH(p),$$

where $p = \operatorname{depth} \phi$.

Now, suppose that $X, X' \in V$ are arbitrary. Let $W = W(X, X')$. Now, X_0 and W each lie in the path sequences of both X and X'. Thus,

$$\begin{aligned}
&\rho(X, X') + 2k(1 + \log_2 n) \\
&\quad \geq \max\big(\rho(X, W) + \rho(X', X_0) - \rho(W, X_0), \rho(X', W) + \rho(X, X_0) - \rho(W, X_0)\big) \\
&\quad \geq \max\big(\rho_\tau(X, W) + \rho_\tau(X', X_0) - \rho_\tau(W, X_0), \rho_\tau(X', W) + \rho_\tau(X, X_0) - \rho_\tau(W, X_0)\big) \\
&\qquad\quad - 2kH(p).
\end{aligned}$$

But $\operatorname{hyp} \rho_\tau = 0$, and so

$$\rho_\tau(X, X') + \rho_\tau(W, X_0) \leq \max\big(\rho_\tau(X, W) + \rho_\tau(X', X_0), \rho_\tau(X', W) + \rho_\tau(X, X_0)\big).$$

We conclude that

$$\rho_\tau(X, X') \leq \rho(X, X') + 2k(1 + \log_2 n + H(p)).$$

Now $p \leq 1 + \log_2\left(\frac{n+2}{3}\right)$, and H is linear, so

$$\rho_\tau(X, X') \leq \rho(X, X') + kF(n),$$

where

$$\begin{aligned}
F(n) &= 2\big(1 + \log_2 n + H(1 + \log_2 \left(\frac{n+2}{3}\right))\big) \\
&= O(\log n).
\end{aligned}$$

We have shown:

Theorem 7.6.1 : *There is a function $F : \mathbf{N} \longrightarrow \mathbf{R}$ such that the following holds.*

Suppose $(\mathcal{S}, d)$ is a k-H1 geodesic space, and $V \subseteq \mathcal{S}$ is a set of $(n+1)$ points. Then, there is an immersed spanning tree τ for V in $\mathcal{S}$, such that for any $X, Y \in V$, we have

$$\rho_\tau(X,Y) \leq d(X,Y) + kF(n).$$

Moreover, $F(n) = O(\log n)$.

CHAPTER 8 : Propagation of hyperbolicity.

8.1. Summary.

The aim of this chapter is to show that the property of almost-hyperbolicity "propagates". Thus, we shall show that a path-metric space is almost-hyperbolic if and only if it is "locally almost-hyperbolic" and "almost simply-connected", provided, of course, that we properly quantify the various parameters involved.

Let $(\mathcal{S}, d)$ be a path-metric space. Recall, from Section 2.3, the notion of a cellular net, (P, ρ, f), and of its mesh, $m(P, \rho)$, as well as the notion of a loop, (γ, σ). Given $X \in \mathcal{S}$ and $r \in [0, \infty)$, write $N_r(X)$ for the uniform ball $\{Y \in \mathcal{S} \mid d(X, Y) \leq r\}$.

Definition : The space $(\mathcal{S}, d)$ is *m-simply-connected* if every loop in $\mathcal{S}$ bounds a cellular net of mesh at most m.

Definition : Given $i \in \{1,2,3,4,5\}$, and $\underline{k} \in [0, \infty) \sqcup [0, \infty)^2$, and $r \in (0, \infty)$, we shall say that $(\mathcal{S}, d)$ is *r-locally* $\underline{k}$*-*H(i) if, for each $X \in \mathcal{S}$, the uniform ball $N_r(X)$ about X is $\underline{k}$-H(i) in the induced path metric.

Given any $Q \subseteq \mathcal{S}$, we shall write d_Q for the induced path-metric on Q. Applying the equivalence of definitions H(i) to the spaces (N, d_N) for $N = N_r(X)$ as X varies in $\mathcal{S}$, we see that if $(\mathcal{S}, d)$ is r-locally $\underline{k}$-H(i), then it is r-locally $\underline{k}'$-H(j) for $\underline{k}' = \underline{k}'(k, i, j)$. This is the primary reason for making the definitions in this way. We shall use mostly the hypothesis of locally H1. Clearly any space is r-locally k-H1 if $r \leq k/2$, so the hypothesis is only useful if r is large in relation to k.

Proposition 8.1.1 :
(1) $\forall k\ \exists m$ such that any k-H1 geodesic space is m-simply-connected.
(2) $\forall k\ \exists k'\ \forall r$ any k-H1 geodesic space is r-locally k'-H1.

Part (1) is an immediate consequence of property H3ca (Section 6.1).

Part (2) is in not quite immediate, since we have defined local hyperbolicity in terms of the induced path metric on balls. However, it is a simple consequence of the uniform convexity of balls in an almost hyperbolic space, as we shall explain in Section 8.2.

The main result of this chapter is the following converse to Proposition 8.1.1.

Theorem 8.1.2 : *$\forall k\ \exists k'\ \forall m\ \exists R$ such that any m-simply-connected R-locally k-H1 geodesic space is k'-H1.*

As usual, we can take k' to be a certain universal multiple of k (see Section 8.7).

The necessity of some sort of simple-connectedness is easy to see. Consider, for example, the subset $W_R = \big([0, \infty) \times \{0\}\big) \cup \bigcup_{n=R}^{\infty} C_n$ of the euclidean plane $\mathbf{R}^2$, where $R \in \mathbf{N}$, and for each $n \in \mathbf{N}$, C_n is the circle of radius n about the origin $(0,0)$. If we take the induced path-metric on W_R, then it is R-locally 0-H1, but not k-H1 for any k.

The fact that simple-connectedness is central to this result suggests that we should make use of the linear isoperimetric inequality (property H3). The idea is roughly as follows.

Suppose that $(\mathcal{S}, d)$ is almost simply-connected, and locally almost-hyperbolic, for suitable parameters. Let γ be any loop in $\mathcal{S}$. We span γ by a "disc" of "minimal area". We may imagine this disc as having some kind of metric structure induced from $\mathcal{S}$. In this induced structure, the disc is, itself, locally almost-hyperbolic. The reason for this is as follows. Any loop in the disc of small diameter satisfies the linear isoperimetric inequality in $\mathcal{S}$ since $\mathcal{S}$ is locally H3, and thus intrinsically in the disc since it is area-minimising. We see that the disc is intrinsically locally H3.

We have essentially essentially reduced the problem to a 2-dimensional situation. In this context, we can imagine that the "curvature" of our disc is always negative when averaged on a certain scale. From this we might hope to deduce that the global average of the curvature of the disc is negative. This should then imply the global isoperimetric inequality, namely that the area of the disc is at most a certain linear function of the length of its boundary. This would show that $(\mathcal{S}, d)$ were H3.

There are many ways one might try to make sense of this argument. It seems that all lead into technical complications at some point or other. We shall try to keep these to a minimum by adopting a notion of spanning disc best suited to our purposes, namely what we shall call a "vertex net". We can thus contain the most unpleasant technicalities to relating this to our previous notion of cellular net (Lemma 8.5.1). The essential passage from local to global will take the form of a combinatorial lemma (8.4.1).

We shall give a more detailed outline of the proof in Section 8.6. For the moment, we return to a discussion of local hyperbolicity.

8.2. Local hyperbolicity.

Suppose $(\mathcal{S}, d)$ is a geodesic space. If Q is a closed subset of $\mathcal{S}$, and d_Q is the induced path-metric on Q, then (Q, d_Q) is also a geodesic space. (In general, we need to allow d_Q to take the value ∞ when two points are not joined by a rectifiable path in Q.)

Recall the discussion of convexity from Chapter 4.

Lemma 8.2.1 : *Suppose that $(\mathcal{S}, d)$ is a geodesic space, and that $Q \subseteq \mathcal{S}$ is closed and λ-convex. Let $N = N_\lambda(Q)$. Then, for all $X, Y \in N$, we have*

$$d_N(X,Y) \leq d(X,Y) + 4\lambda.$$

Proof: Let $Z \in \text{proj}_Q X$ and $W \in \text{proj}_Q Y$. Then $[Z,W] \subseteq N$, and the path $[X,Z,W,Y] = [X,Z] \cup [Z,W] \cup [W,Y] \subseteq N$ has length at most $d(X,Y) + 4\lambda$.
◊

We deduce Proposition 8.1.1(2) as follows.

Corollary 8.2.2 : *Suppose $(\mathcal{S}, d)$ is a k-H1 geodesic space, and $\lambda_0 = \lambda_0(k)$ is the constant of Lemma 4.2. For any $r \in [0, \infty)$ and $X \in \mathcal{S}$, we have that (N, d_N) is $(k + 8\lambda_0)$-H1, where $N = N_r(X)$.*

Proof : If $r \leq \lambda_0$, then clearly (N, d_N) is $(4\lambda_0)$-H1.

If $r > \lambda_0$, then $Q = N_{(r-\lambda_0)}(X)$ is λ_0-convex, by Lemma 4.2. Thus, $d \leq d_N \leq d + 4\lambda_0$ by Lemma 8.2.1. The result follows easily.

◊

Remark : There is an alternative way one might define local hyperbolicity. Given a geodesic space $(\mathcal{S}, d)$, we weaken the hypothesis of $\underline{k}$-H(i), for $i \in \{1, 2, 3, 4, 5\}$, by demanding that the set of points under consideration should have diameter at most some fixed number, which we call the *range*. Thus, for example, to say that $(\mathcal{S}, d)$ is locally k-H1 in this sense means that $XY : ZW$ or $XZ : WY$ or $XW : YZ$ for any $X, Y, Z, W \in \mathcal{S}$ satisfying $\operatorname{diam}\{X, Y, Z, W\} \leq R$ for some fixed range R. The equivalence of these five definitions for large range follows from the arguments of previous chapters, although we cannot apply these results directly. Relating these notions to the definitions of Section 8.1 is easiest using property H1. Suppose we are given $k, r \in [0, \infty)$. Then if the range, $R = R(k, r)$, is large enough, and $(\mathcal{S}, d)$ is a locally k-H1 geodesic space, in the sense just defined, then, using the arguments of previous chapters, we see that any uniform r-ball in $\mathcal{S}$ is almost convex. It follows (as in 8.2.1 an 8.2.2) that $(\mathcal{S}, d)$ is r-locally k'-H1 in the original sense (i.e. that of Section 8.1), where k' depends only on k. The converse of this statement (similarly quantified) is a consequence of Lemma 8.2.3.

We conclude this section with the following trivial observation.

Lemma 8.2.3 : *Suppose that $(\mathcal{S}, d)$ is a geodesic space, and $N = N_r(X)$, where $X \in \mathcal{S}$ and $r \geq 0$. Then for any $Y, Z \in N_{r/2}(X)$, we have $d_N(Y, Z) = d(Y, Z)$.*

Proof : $[Y, Z] \subseteq N$.

◊

This means that if we know that $(\mathcal{S}, d)$ is r-locally almost hyperbolic, we can apply any result obtained for globally almost hyperbolic spaces, provided that we ensure that our constructions do not take us outside a set of diameter r.

8.3. A geometric lemma.

Suppose $(\mathcal{S}, d)$ is a geodesic space, and that $X, Y, Z, W, X_0, X_1, \ldots, X_n \in \mathcal{S}$. Recall, from Section 7.5, the notation $XY \wedge ZW$, $\Lambda(X_0, X_1, \ldots, X_n)$, $[X_0, X_1, \ldots, X_n]$ and

$\langle X_0 X_1 \ldots X_n \rangle$. Note that if $X_0 = X_n$, then

$$\Lambda(X_0, X_1, \ldots, X_n) = \Lambda(X_1, X_2, \ldots, X_{n-1}, X_0, X_1).$$

Thus, if $\gamma = [X_0, X_1, \ldots, X_n]$ is a closed piecewise geodesic path, then $\Lambda(\gamma)$ is defined independently of the choice of basepoint. We shall write $\text{length}\,\gamma = \sum_{i=1}^{n} d(X_i, X_{i-1})$.

Lemma 8.3.1 : *There is a map $F : \mathbf{N} \longrightarrow \mathbf{R}$ and a fixed $\lambda \geq 0$ such that the following holds.*

Suppose that $(\mathcal{S}, d)$ is a k-H1 geodesic space, and that γ is a closed path in $\mathcal{S}$ consisting on n geodesic segments. Then,

$$\text{length}\,\gamma \leq (\Lambda(\gamma) + \lambda k) F(n).$$

Proof : This is just a special case of Proposition 7.3.4 (which gives F linear in n).

Alternatively, one may give a short proof of Lemma 8.3.1 as follows.

Let $L = \text{length}\,\gamma$. Without loss of generality, we can suppose that $d(X_{n-1}, X_0) \geq d(X_i, X_{i+1})$ for each $i \in \{0, 1, \ldots, n-1\}$. Thus, $d(X_{n-1}, X_0) \geq \frac{1}{n} L$. For $i = 0, 1, \ldots, n-1$, let Y_i be a nearest point in $[X_0, X_{n-1}]$ to X_i. Thus, $Y_0 = X_0$ and $Y_{n-1} = X_{n-1}$. Clearly there is some $i \in \{0, 1, \ldots, n-1\}$ with $\langle X_0 Y_i Y_{i+1} X_{n-1} \rangle$ and

$$d(Y_i, Y_{i+1}) \geq \frac{1}{(n-1)} d(X_{n-1}, X_0) \geq \frac{1}{n(n-1)} L.$$

Then, either $\frac{1}{n(n-1)} L \simeq 0$, or else one can show without difficulty that

$$\Lambda(\gamma) \geq X_0 X_{n-1} \wedge Y_i Y_{i+1} \succeq \frac{1}{n(n-1)} L.$$

(This argument gives F quadratic in n.)
$\diamondsuit$

Lemma 8.3.2 : *$\forall k \, \exists a \, \forall H \, \exists L$ such that the following holds. Suppose $(\mathcal{S}, d)$ is a k-H1 geodesic space, and that $\gamma = [X_0, X_1, \ldots, X_n]$ is a closed piecewise geodesic path with at most 13 segments. (Thus, $X_0 = X_n$ and $n \leq 13$.) Then, either $\text{length}\,\gamma \leq L$, or else we can find $0 \leq i < j \leq n$, and $A, B \in [X_j, X_{j+1}]$, such that $\langle X_j A B X_{j+1} \rangle$, $d(A, B) \geq H$ and*

$$d(X_i, B) + d(B, A) + d(A, X_{i+1}) \leq d(X_i, X_{i+1}) + a.$$

Proof : Let $L = (H + \lambda k) \max\{F(r) \mid 1 \leq r \leq 13\}$.

If $\Lambda(\gamma)$, then by Lemma 8.3.1, $\text{length}\,\gamma \leq L$. So, suppose $\Lambda(\gamma) \geq H$. This means there exist i, j with $0 \leq i < j \leq n$ and with $X_i X_{i+1} \wedge X_{j+1} X_j \geq H$. Applying Lemma 7.5.1, we find $A, B \in [X_j, X_{j+1}]$ as required, with $a = 2h$ (h being the constant of Lemma 7.5.1).

$\diamondsuit$

8.4. A combinatorial lemma.

Let P be a presentation of the disc D as a CW-complex. Write $C_i(P)$ for the set of i-cells of P, and $\Sigma(P)$ for the 1-skeleton of P. We say that Σ is *j-edge-connected*, for $j \in \mathbf{N}$, if no set of $j-1$ edges disconnects Σ. We say that Σ is *trivalent* if every vertex of Σ has degree 3.

Note that if Σ is trivalent and 2-edge-connected, then the boundary, ∂c, of every 2-cell $c \in C_2(P)$ is an embedded topological circle. Moreover, the two endpoints of any edge are distinct. Thus, in this case, P, is a cellulation in the sense of Section 2.3. If Σ is trivalent and 3-edge-connected, then any two distinct 2-cells of P meet either along a single common edge or not at all. Moreover, any 2-cell of P meets ∂D either in a single edge or not at all.

For $i = 0, 1$, let $C_i^\partial(P)$ be the set of i-cells lying in ∂D. Write $C_i^I(P) = C_i(P) \backslash C_i^\partial(P)$. Clearly, $|C_0^\partial(P)| = |C_1^\partial(P)|$. Given $c \in C_2(P)$, write $\nu(c) = |C_0(P) \cap \partial c|$ for the number of vertices in the boundary of c (or equivalently, the number of edges of c).

We call $c \in C_2(P)$ an *interior* 2-cell if $c \cap \partial D = \emptyset$. Write $C_2^I(P)$ for the set of interior 2-cells of P, and $C_2^\partial(P) = C_2(P) \backslash C_2^I(P)$.

Lemma 8.4.1 : *Suppose that P is a cellulation of the disc D, with $\Sigma(P)$ trivalent and 3-edge-connected. Suppose that for any two distinct interior 2-cells, $c_1, c_2 \in C_2^I(P)$, satisfying $\nu(c_1) \leq 13$ and $\nu(c_2) \leq 13$, we have that $c_1 \cap c_2 = \emptyset$. Then, it follows that*

$$|C_2(P)| \leq 15|C_1^\partial(P)|.$$

Proof : We construct a new cellulation P' of the disc by contracting each interior 2-cell with fewer than 14 edges to a single point. This is a well-defined process, since, by hypothesis, no two such cells intersect. The set of interior 2-cells of P' corresponds bijectively to the the set of interior 2 cells of P with at least 14 edges. Now each such cell has lost at most half its edges through this process of contraction, and so each $c \in C_2^I(P')$ has at least 7 edges. Clearly, there is a bijective correspondence between $C_1^\partial(P)$ and $C_1^\partial(P')$, since ∂D is unchanged. Note also that P' is 3-edge-connected, and so $C_2^\partial(P')$ is also in bijective correspondence with $C_1^\partial(P)$. Also, each vertex of P' has degree at least 3.

For $i = 0, 1, 2$, write $\Lambda_i = |C_i(P)|$, $\Lambda_i^\partial = |C_i^\partial(P)|$, $\lambda_i = |C_i(P')|$, $\lambda_i^\partial = |C_i^\partial(P')|$ and $\lambda_i^I = |C_i^I(P')|$.

The following relations are easily verified.

(1) $\lambda_0 + \lambda_2 - \lambda_1 = 1$,

(2) $3\lambda_0 \leq 2\lambda_1$,

(3) $\lambda_1 = \lambda_1^\partial + \lambda_1^I$,

(4) $\lambda_2 = \lambda_2^\partial + \lambda_2^I$,

(5) $7\lambda_2^I \leq 2\lambda_1^I$,

(6) $\Lambda_1^\partial = \lambda_1^\partial = \lambda_2^\partial$,

(7) $\Lambda_2 \leq \lambda_0 + \lambda_2$.

Now, from (1) and (2), we obtain the following inequality (∗),

$$\lambda_1 \leq 3(\lambda_2 - 1).$$

Applying (3), (4) and (6) to this, we get

$$\lambda_1^I \leq 3\lambda_2^I + 2\lambda_1^\partial - 3.$$

From (5), we get

$$7\lambda_2^I \leq 2(3\lambda_2^I + 2\lambda_1^\partial - 3),$$

and so

$$\lambda_2^I \leq 4\lambda_1^\partial - 6.$$

From (4) and (6),

$$\lambda_2 = \lambda_2^I + \lambda_1^\partial \leq (4\lambda_1^\partial - 6) + \lambda_1^\partial = 5\lambda_1^\partial - 6.$$

Using (∗) again, we get

$$\lambda_1 \leq 3(\lambda_2 - 1) \leq 15\lambda_1^\partial - 21.$$

By (7) and (1),

$$\Lambda_2 \leq \lambda_0 + \lambda_2 = 1 + \lambda_1 \leq 15\lambda_1^\partial - 20.$$

Using (6),

$$\Lambda_2 \leq 15\Lambda_1^\partial - 20.$$

So certainly,

$$|C_2(P)| \leq 15|C_1^\partial(P)|.$$

◊

8.5. Vertex nets.

This is a technical section. We define the notion of a "vertex net", which is another formulation of the idea of a spanning disc. We relate this to the idea of a cellular net from Section 2.3.

Let G be a finite graph. We write $C_0(G)$ for the set of vertices of G, and $C_1(G)$ for the set of edges of G. By a *subgraph*, G', of G, we mean a subset $C_1(G')$ of the edges of G, together with all their endpoints.

Suppose that $(\mathcal{S}, d)$ is metric space, and that $g : C_0(G) \longrightarrow \mathcal{S}$ is any map. If $e \in C_1(G)$, we shall write

$$\text{length}(e, g) = d(g(v), g(w)),$$

where $v, w \in C_0(G)$ are the endpoints of e. If G' is a subgraph of G, then we write

$$\text{length}(G', g) = \sum_{e \in C_1(G')} \text{length}(e, g)$$

and

$$\text{coarseness}(G', g) = \max_{e \in C_1(G')} \text{length}(e, g).$$

As an example, consider the circle S^1 as the unit circle in the complex plane. For $n \in \mathbf{N}$ and $j \in \mathbf{Z}$, let v_j^n be the point $e^{2\pi i(j/n)} \in S^1$. Let $V^n = \{v_1^n, v_2^n, \ldots, v_n^n\} \subseteq S^1$. For any n, we can represent S^1 as a graph with vertex set $C_1(S^1) = V^n$. Suppose $(\mathcal{S}, d)$ is a geodesic space, and $g : V^n \longrightarrow \mathcal{S}$ any map. Then, we can identify (S^1, g) with the closed piecewise geodesic path $\gamma = [X_0, X_1, \ldots, X_n]$. Thus, $\text{length}(S^1, g) = \text{length}\gamma = \sum_{i=1}^n d(X_i, X_{i-1})$. We shall call such an object a *cycle* of n points in $\mathcal{S}$.

Let P be a cellulation of the disc D. We call such a cellulation *good* if the 1-skeleton, $\Sigma = \Sigma(P)$, is trivalent and 2-edge-connected (see Section 8.4), and if each 2-cell of P meets ∂D either in a single edge or not at all.

We identify $S^1 \equiv \partial D$.

Definition : Let $(\mathcal{S}, d)$ be a geodesic space. A *vertex net*, (P, g), consists of a good cellulation P of the disc D, together with a map $g : C_0(P) \longrightarrow \mathcal{S}$.

We can thus define $\text{length}(P, g)$ and $\text{coarseness}(P, g)$ by putting $G' = G = \Sigma(P)$ in the formulae given above. We set

$$\text{mesh}(P) = \max_{c \in C_2(P)} \text{length}(\partial c, g).$$

Definition : Suppose that $\gamma \equiv (S^1, g_0)$ is a cycle of n points in $\mathcal{S}$. We say that γ *bounds* a vertex net (P, g), if $C_0^\partial(P) = V^n$, and $g | C_0^\partial(P) = g_0$. (Recall that $C_0^\partial(P) = C_0(P) \cap \partial D$.)

Clearly, $\text{coarseness}\,\gamma \le \text{coarseness}(P, g) \le \text{mesh}(P, g)$.

We now relate all this to the definitions of Section 2.3.

Lemma 8.5.1 : *Let σ be a path-metric on the circle S^1. Suppose that the points $x_1, x_2, \ldots, x_n \in S^1$ are cyclically ordered, and cut S^1 into segments of σ-length at most m_0. Suppose that $(\mathcal{S}, d)$ is a geodesic space, and that $\beta : (S^1, \sigma) \longrightarrow (\mathcal{S}, d)$ is distance non-increasing (i.e. β is a "loop" in the sense of Section 2.3). Let γ be the cycle $[X_0, X_1, \ldots, X_n]$, where $X_i = \beta(x_i)$ for $i \ge 1$, and $X_0 = X_n$. (Thus* $\text{coarseness}\,\gamma \le m_0$.*) Then, we have the following.*

(1) Suppose that γ bounds a vertex net (P, g), then for any $\epsilon > 0$, β bounds a cellular net (P, ρ, f) with

$$m(P, \rho) \le m_0 + \text{mesh}(P, g) + \epsilon.$$

(2) Suppose that β bounds a cellular net (P, ρ, f), then γ bounds a vertex net (P', g) with

$$\text{mesh}(P', g) \le m_0 + m(P, \rho).$$

Proof :

(1) Write $C_1^\partial(P) \subseteq C_1(P)$ for the set of 1-cells lying in ∂D. Write $C_1^I(P) = C_1(P) \backslash C_1^\partial(P)$. By hypothesis, each 2-cell of P meets ∂D either in single edge, or not at all. Let $\eta = \epsilon / |C_1(P)|$.

We define a path-metric ρ on $\Sigma(P)$ as follows. If $e \in C_1^\partial(P)$, set $\rho(e) = m_0$. If $e \in C_1^I(P)$, set $\rho(e) = \max(\eta, \text{length}(e, g))$.

Now (by definition), $C_0^\partial(P) = C_0(P) \cap \partial D = V^n = \{v_1^n, v_2^n, \ldots, v_n^n\}$, and $g(v_i^n) = X_i$. We define $\partial f : (\partial D, \rho) \longrightarrow (S^1, \sigma)$ by setting $\partial f(v_i^n) = x_i$ for $i = 1, 2, \ldots, n$, and mapping each component of $\partial D \backslash C_1^\partial(P)$ onto the corresponding component of $S^1 \backslash \{x_1, x_2, \ldots, x_n\}$, linearly with respect to the parameterisations induced by ρ and σ respectively. Note that $\beta \circ \partial f | C_0^\partial(P) = g | C_1^\partial(P)$, and that ∂f is homotopic to the identification $\partial D \equiv S^1$.

We define $f : \Sigma(P) \longrightarrow (\mathcal{S}, d)$ as follows. We take $f | C_0(P) = g$ and $f | \partial D = \beta \circ \partial f$. (Note that this is consistent on $C_0^\partial(P)$.) Finally, if $e \in C_1^I(P)$, with endpoints $v, w \in C_0(P)$, then we map e (with parameterisation induced by ρ) linearly onto the geodesic $[g(v), g(w)]$ in $\mathcal{S}$.

We have defined (P, ρ, f). It is readily checked that $m(P, \rho) \leq m_0 + \text{mesh}(P, g) + \epsilon$. This completes the proof of part (1).

(2) We have that β bounds a cellular net (P, ρ, f). We can suppose that $m_0 > 0$. We shall construct (P', g) in a series of stages.

We identify $\partial D \equiv S^1$. Let D_0 be the topological disc $D_0 = D \cup_{\partial D} (S^1 \times [0,2])$, where we identify $S^1 = \partial D$ with $S^1 \times \{0\}$ via $x \leftrightarrow (x, 0)$.

Let Q be any triangulation of the annulus $S^1 \times [0,1]$. Let $h : S^1 \times [0,1] \longrightarrow S^1$ be a homotopy from $h(.,0) = \partial f : (\partial D, \rho_{\partial D}) \longrightarrow (S^1, \sigma)$ to the identity $h(.,1) = 1_{S^1} : S^1 \longrightarrow S^1$. By subdividing Q enough times, we obtain a new triangulation Q_1 of $S^1 \times [0,1]$ with the property that $\sigma(h(p), h(q)) \leq m/3$ for the endpoints, p and q, of any edge of Q_1. We define a cellulation, Q_2, of $S^1 \times [1,2]$ (i.e. a presentation of $S^1 \times [1,2]$ as a CW-complex) as follows. Set $x_0 = x_n$, and for each $i \in \{1, 2, \ldots, n\}$, write $[x_{i-1}, x_i] \subseteq S^1$ for the closed interval joining x_{i-1} to x_i in S^1 (in the positive direction). We take

$$C_0(Q_2) = \{x_1, x_2, \ldots, x_n\} \times \{1, 2\},$$

$$C_1(Q_2) = \{[x_{i-1}, x_i] \times \{j\} \mid 1 \leq i \leq n,\ j = 1, 2\} \cup \{\{x_i\} \times [1,2] \mid 1 \leq i \leq n\},$$

and

$$C_2(Q_2) = \{[x_{i-1}, x_i] \times [1,2] \mid 1 \leq i \leq n\}.$$

The cellulations P, Q_1 and Q_2, together give us a cellulation, P_0, of the disc D_0, after taking a common subdivision on the circles $\partial D \equiv S^1 \times \{0\}$ and $S^1 \times \{1\}$. By taking the maximal such subdivision, we will have that each 0-cell of P_0 has degree at least 3 in $\Sigma(P_0)$. Moreover each 0-cell in $\partial D_0 = S^1 \times \{2\}$ has degree exactly 3. We define $g_0 : C_0(P_0) \longrightarrow \mathcal{S}$ as follows.

$$g_0 | C_0(P) = f | C_0(P),$$

$$g_0 | C_0(Q_1) = \beta \circ h | C_0(Q_1),$$

and

$$g_0(x_i, j) = \beta(x_i)$$

for $i = 1, 2, \ldots, n$ and $j = 1, 2$. One can check that g_0 is well-defined on any 0-cells P, Q_1 or Q_2 may have in common. Moreover, one may check that $\text{mesh}(P_0, g_0) \leq m_0 + m(P, \rho)$.

Now let $\phi : D \longrightarrow D_0$ be any homeomorphism with $\phi(v_i^n) = (x_i, 2)$ for $i = 1, 2, \ldots, n$. Let P_1 be the pull-back of the cellulation P_0 to D. Let $g_1 = g_0 \circ \phi : C_0(P_1) \longrightarrow \mathcal{S}$. Thus $g_0(v_i^n) = X_i = \beta(x_i)$. Clearly, $\text{mesh}(P_1, g_1) = \text{mesh}(P_0, g_0)$. Note that every vertex of $C_0^\theta(P_1)$ has degree 3, and that no 2-cell meets ∂D in more than one edge.

Finally, we define P' by arbitrarily splitting apart each vertex of $C_0(P_1)$, with degree at least 4, to give a tree in $\Sigma(P_1)$. Thus, we can arrange that each vertex of P' has degree 3. There is a natural map $\theta : C_0(P') \longrightarrow C_0(P_1)$ given by contracting all these trees back to points. We define $g = g_1 \circ \theta : C_0(P') \longrightarrow \mathcal{S}$. Clearly, $\text{mesh}(P', g) = \text{mesh}(P_1, g_1) \leq m_0 + m(P, \rho)$.

◊

Corollary 8.5.2 : *In an m-simply-connected geodesic space, any cycle of coarseness m_0 bounds a vertex net of mesh at most $m + m_0$.*

Proof : Apply Lemma 8.5.1(2) to the closed piecewise geodesic path given by the cycle (i.e. by joining the points of the cycle by geodesic segments).

◊

Corollary 8.5.3 : $\forall m, m_0, \lambda, \mu\ \exists k_3, h_3$ *such that if $(\mathcal{S}, d)$ is a geodesic space in which every cycle of n points, with coarseness at most m_0, bounds a vertex net of mesh at most m and with at most $\lambda n + \mu$ 2-cells, then $(\mathcal{S}, d)$ is (k_3, h_3)-H3.*

Proof : Apply Lemma 8.5.1(1).

◊

8.6. Main Proof.

In this section, we give the proof that almost hyperbolicity propagates. It will remain in the final section (8.7) to refine this statement with regard to the parameters involved.

The idea of the proof is as follows. Suppose that $(\mathcal{S}, d)$ is an m-simply-connected r-locally k-H1 geodesic space, with r much larger than both m and k. Choose any cycle, $\gamma = [X_0, X_1, \ldots, X_n]$, in $\mathcal{S}$. Let $m_0 \geq \text{coarseness}\, \gamma = \max\{d(X_i, X_{i-1}) \mid 1 \leq i \leq n\}$. Corollary 8.5.2 tells us that γ bounds a vertex net of mesh at most $m + m_0$. We choose $M \geq m + m_0$ very large (but smaller than r). Among all vertex nets bounded by γ, and of mesh at most m, we select those of (close to) minimum length. Then, among these, we choose one, (P, g), with a minimal number of 2-cells. We aim to show that P satisfies the hypotheses of Lemma 8.4.1.

First, we note that P must be 3-edge-connected. This is a simple argument.

Next, we show that there is some upper bound, L, on length$(\partial c, g)$ for any interior 2-cell c with fewer than 14 edges. The reason is that, if length$(\partial c, g)$ were very large, we could use Lemma 8.3.2 to reduce length(P, g) by partially collapsing the 2-cell c. In doing this, however, we disturb one of the neighbouring cells, so it is conceivable that we may increase the mesh of the net above M. In this case, however, Lemma 6.1.4 allows us to cut the offending 2-cell in half, without increasing the total length of the net too much.

Finally, suppose that c_1 and c_2 are adjacent interior 2-cells, each with at most 13 edges. We have shown that length$(\partial c_i) \leq L$, and we can assume that $L \leq M/2$. Thus, if we delete the common edge of c_1 and c_2, we will not increase the mesh beyond M. However, we reduce the number of 2-cells. This contradicts the choice of (P, g).

We conclude that P satisfies the hypothesis of Lemma 8.4.1. Thus P has at most $15n$ 2-cells. We now apply Corollary 8.5.1 to deduce that $(\mathcal{S}, d)$ is almost hyperbolic.

To make this argument precise, we need to describe, more carefully, the modifications of vertex nets we use, and to quantify all the parameters involved.

Suppose that (P, g) is a vertex net. As before, write $C_1^\partial(P)$ for the set of 1-cells of P which lie in ∂D. Write $C_1^I = C_1(P) \setminus C_1^\partial(P)$. We describe the following ways to modify (P, g) to give a new net (P', g').

(1) Addition of an edge.

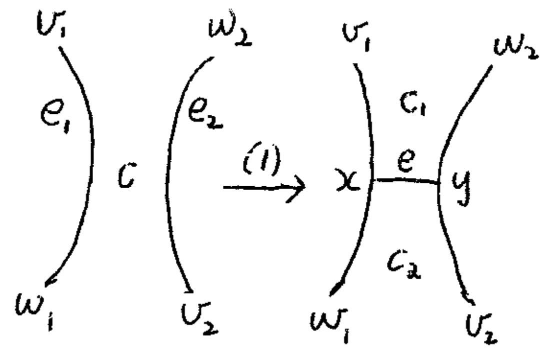

Figure 8a.

Suppose $c \in C_2(P)$. Suppose $e_1, e_2 \in C_1^I(P)$ are distinct edges of c, with endpoints v_1, w_1 and v_2, w_2 respectively. We assume that v_1, w_1, v_2, w_2 (not necessarily distinct) are cyclically ordered in this way around ∂c. (Figure 8a.) Suppose that $C \in \mathcal{S}$ lies in the geodesic segment $[g(v_1), g(w_1)]$ and that $D \in [g(v_2), g(w_2)]$. We form P' by adding a new edge, e, between the points $x \in e_1$ and $y \in e_2$. Thus, $C_0(P') = C_0(P) \cup \{x, y\}$. We define g' by $g'|C_0(P) = g$ and $g'(x) = C$ and $g'(y) = D$. We have

$$\begin{aligned} \text{length}(P', g') &\leq \text{length}(P, g) + d(C, D) \\ \text{mesh}(P', g') &\leq \max\big(\text{mesh}(P, g), \text{length}(\partial c_1, g'), \text{length}(\partial c_2, g')\big), \end{aligned}$$

where $c_1, c_2 \in C_2(P')$ are the 2-cells meeting e.

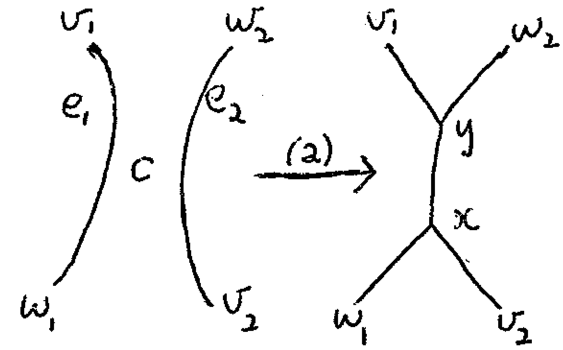

Figure 8b.

(2) Partial contraction of a 2-cell.

Let $c, e_1, e_2, v_1, w_1, v_2, w_2$ be as in modification (1). Suppose that $A, B \in \mathcal{S}$ with $\langle g(v_2)ABg(w_2)\rangle$. We form P' by contracting a strip joining e_1 to e_2 as shown in the Figure 8b. Let x and y be the new vertices introduced. Thus, $C_0(P') = C_0(P) \cup \{x, y\}$. Define g' by $g'|C_0(P) = g$ and $g'(x) = A$ and $g'(y) = B$. We have

$$\text{length}(P', g') \leq \text{length}(P, g) - d(A, B) + \delta$$
$$\text{mesh}(P', g') \leq \text{mesh}(P, g) + \delta,$$

where $\delta = d(g(v_1), B) + d(B, A) + d(A, g(w_1)) - d(g(v_1), g(w_1))$.

(3) Removal of an edge.

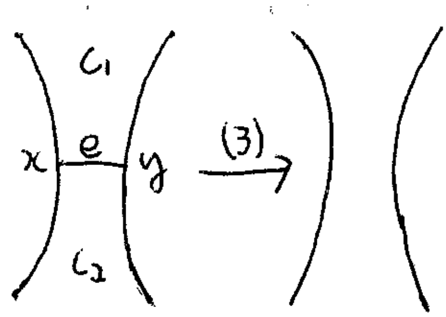

Figure 8c.

Suppose $e \in C_1(P)$ with endpoints $x, y \in C_0(P)$. Suppose that $x, y \notin \partial D$. Thus $C_0(P') = C_0 \setminus \{x, y\}$. Let $g' = g|C_0(P)$. We have

$$\text{length}(P', g') \leq \text{length}(P, g)$$
$$\text{mesh}(P', g') \leq \max\big(\text{mesh}(P, g), \text{length}(\partial c_1, g) + \text{length}(\partial c_2, g)\big),$$

where $c_1, c_2 \in C_2(P)$ are the 2-cells of P meeting e.

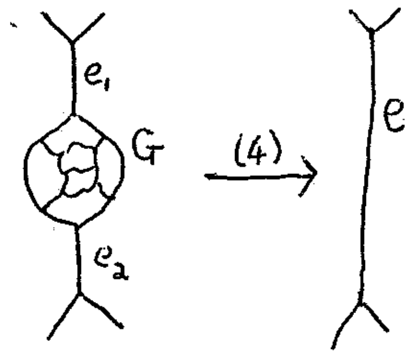

Figure 8d.

(4) Removal of a subgraph.

Suppose the edges e_1 and e_2 of P separate $\Sigma(P)$ into two components. Let G be the component not meeting ∂D. (Figure 8d.) We form P' by removing G and joining e_1 and e_2 to form a single edge e. Thus $C_0(P') = C_0(P) \setminus C_0(G)$. Let $g' = g|C_0(P')$. We have

$$\text{length}(P', g') \leq \text{length}(P, g)$$
$$\text{mesh}(P', g') \leq \text{mesh}(P, g).$$

We now give the proof in detail. We make no essential use of the completeness or local compactness assumptions on $\mathcal{S}$.

Proposition 8.6.1 : *$\forall k, m\, \exists r, k'$ such that if $(\mathcal{S}, d)$ is an m-simply-connected r-locally k-H1 geodesic space, then it is k'-H1.*

Proof : We are given k, m.

Choose $\epsilon > 0$ and $m_0 > 0$ arbitrarily (for example take $\epsilon = m_0 = 1$).

Let a be the constant arising from Lemma 8.3.2, given k.

Let $b = a + 2m_0$.

Let l be the constant arising from Lemma 6.1.4, given $k_1 = k$ and b.

Let $H = a + l + m_0 + 2\epsilon$.

Let L be the constant arising from Lemma 8.3.2, given k and H.

Let $M = \max(2L, m + m_0, 2(l + b))$.

Let $r = M + a$.

Suppose that $(\mathcal{S}, d)$ is an m-simply-connected r-locally k-H1 geodesic space. Suppose $\gamma = [X_0, X_1, \ldots, X_n]$ is some cycle in $\mathcal{S}$, with coarseness $\gamma \leq m_0$.

Let $\mathcal{P}$ be the set of all vertex nets in $\mathcal{S}$ spanning γ, and with mesh at most M. Since $m + m_0 \leq M$, applying Corollary 8.5.2, we have that $\mathcal{P} \neq \emptyset$.

Let

$$\lambda = \inf\{\text{length}(P, g) \mid (P, g) \in \mathcal{P}\},$$

and

$$\mathcal{P}_0 = \{(P,g) \in \mathcal{P} \mid \text{length}(P,g) \leq \lambda + \epsilon\}.$$

Choose a fixed $(P,g) \in \mathcal{P}_0$ with minimal number of 2-cells. We shall show that (P,g) satisfies the hypothesis of Lemma 8.4.1. We can deduce from this that $|C_2(P)| \leq 15n$. The result then follows by applying Corollary 8.5.3 and the equivalence of H1 and H3.

First note that P is 3-edge-connected — otherwise, we could use modification (4) to strictly reduce the number of 2-cells, without increasing the length or mesh.

Now, suppose (for contradiction) that there is some $c \in C_2^I(P)$ with length$(\partial c, g) > L$, and $\nu(c) \leq 13$. Let $y_1, y_2, \ldots, y_s$ be the set of vertices $\partial c \cap C_0(P)$, ordered cyclically around ∂c. Thus $s \leq 13$. Set $y_0 = y_s$, and let $Y_i = g(y_i) \in \mathcal{S}$ for each i. Let α be the closed piecewise geodesic path $[Y_0, Y_1, \ldots, Y_s]$. Now, diam$[Y_0, Y_1, \ldots, Y_s] \leq$ mesh$(P,g) \leq m < r$. Thus, we apply Lemmas 8.3.2 and 8.2.3 to find $0 \leq i < j \leq n$ and $A, B \in [Y_j, Y_{j+1}]$ with $\langle Y_j A B Y_{j+1} \rangle$, $d(A,B) \geq H$ and $d(Y_i, B) + d(B,A) + d(A, Y_{i+1}) \leq d(Y_i, Y_{i+1}) + a$. We perform modification (2), with $v_1 = y_i$, $w_1 = y_{i+1}$, $v_2 = y_j$ and $w_2 = y_{j+1}$ (Figure 8e).

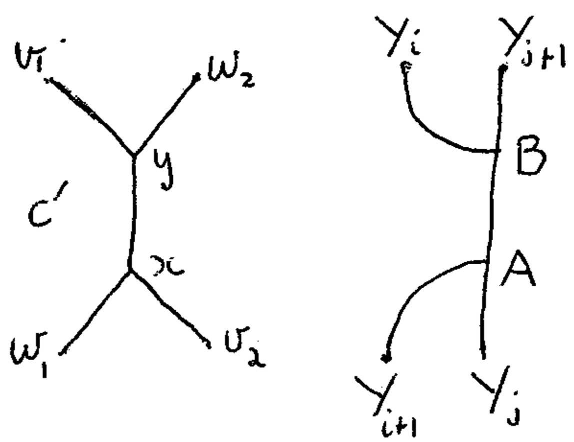

Figure 8e.

(Possibly $w_1 = v_2$ or $v_1 = w_2$.) This gives us a new net (P', g'). Let c_0 be the 2-cell of P, other than c, which meets the edge, e, between v_1 and w_1. Let c' be the corresponding cell of P'.

We have

$$\begin{aligned}\text{length}(P', g') &\leq \text{length}(P,g) - d(A,B) + a \\ &\leq \lambda + \epsilon - H + a.\end{aligned}$$

However we may have increased the mesh, so that, perhaps, (P', g') does not belong to $\mathcal{P}$. Since $d(A,B) \geq H > a$, the only cell of P' which may have boundary length greater than M is c'. We know, at least, that length$(\partial c', g) \leq M + a$.

If length$(\partial c', g') \leq M$, then mesh$(P', g') \leq M$. Since $H - a + \epsilon < 0$, we have have length$(P', g') < \lambda$, contradicting the definition of λ. Thus, we can assume that

$$M \leq \text{length}(\partial c', g') \leq M + a.$$

Let $z_1, z_2, \ldots, z_t$ be the vertices of $\partial c' \cap C_0(P')$ cyclically ordered around c. Let $z_0 = z_t$, and let $Z_i = g'(z_i)$ for $0 \leq i \leq t$. Let β be the closed piecewise geodesic path $[Z_0, Z_1, \ldots, Z_t]$. Now $\operatorname{diam}\beta \leq M + a \leq r$ and $\operatorname{length}\beta \geq M \geq 2(l+b)$. So, applying Lemmas 6.1.4 and 8.2.3, we find $C', D' \in \beta$ with $d(C', D') \leq l$, but with the distance between C' and D', measured along β, at least $l+b$. If $c' \in C_2^I$ (i.e. $c' \cap \partial D = \emptyset$), then set $C = C'$ and $D = D'$. If, on the other hand, $c' \cap \partial D \neq \emptyset$, then c' meets ∂D in precisely one edge, which, without loss of generality, we can take to have endpoints z_0 and z_1. Thus Z_0 and Z_1 must be consecutive points in our original cycle γ. It follows that $d(Z_0, Z_1) \leq m_0$. It is impossible that C' and D' lie in the same geodesic segment of β, so we can assume, without loss of generality, that $D' \neq [Z_0, Z_1]$. We set $D = D'$. If $C' \in [Z_0, Z_1]$, we set $C = Z_0$, otherwise, we set $C = C'$. Thus, in any case, $d(C, C') \leq m_0$. Now, $d(C, D) \leq l + m_0$, and the distance between C and D, measured along β is at least $l + b - m_0$. We have $C \in [Z_p, Z_{p+1}]$ and $D \in [Z_q, Z_{q+1}]$, where we can suppose that $p \neq 0$ and $q \neq 0$. Since $b - m_0 > m_0$, we must have $p \neq q$. Let e'_1 be the edge in $\partial c'$ bounded by z_p and z_{p+1}, and let e'_2 be the edge of $\partial c'$ bounded by z_{q+1}. We now perform modification (1), to give a new net (P'', g''). If $c''_1, c''_2 \in C_2(P'')$ are the 2-cells into which c' is split, then it is easily seen that, for $i = 1, 2$,

$$\operatorname{length}(\partial c''_i, g'') \leq \operatorname{length}(\partial c', g') - (b - 2m_0) \leq (M + a) - (b - 2m_0) = M.$$

Thus,

$$\operatorname{mesh}(P'', g'') \leq M,$$

and so $(P'', g'') \in \mathcal{P}$. However, we have

$$\begin{aligned} \operatorname{length}(P'', g'') &\leq \operatorname{length}(P', g') + (l + m_0) \\ &\leq (\lambda + \epsilon - H + a) + (l + m_0) \\ &= \lambda - \epsilon, \end{aligned}$$

which contradicts the definition of λ, and thus the existence of the original cell c.

We have shown that, for any 2-cell $c \in C_2^I(P)$ with $\nu(c) \leq 13$, we have $\operatorname{length}(\partial c, g) \leq L$.

Finally, suppose (for contradiction) that there exist $c_1, c_2 \in C_2^I(P)$ satisfy $\nu(c_1) \leq 13$, $\nu(c_2) \leq 13$ and meet along some edge $e \in C_1(P)$. Thus, $e \cap \partial D = \emptyset$. We know that $\operatorname{length}(\partial c_i, g) \leq L \leq M/2$. Thus, if we perform modification (3) to produce (P', g'), then we get

$$\operatorname{mesh}(P', g') \leq M$$

and

$$\operatorname{length}(P', g') \leq \operatorname{length}(P, g) \leq \lambda + \epsilon.$$

Thus $(P', g') \in \mathcal{P}_0$. However, $|C_0(P')| \leq |C_0(P)| - 1$. This contradicts the assumption that P has a minimal number of 2-cells, and so such c_1 and c_2 cannot exist.

We have shown (P, g) satisfies the hypotheses of Lemma 8.4.1, as required.

◊

8.7. Refinement of parameters.

The argument of the last section does not give any sensible way of relating the final hyperbolicity parameter k' to k. In fact, however, one can always take k' to be a certain universal multiple, ξk of k (independent of the mesh m). The main theorem of this chapter (Theorem 8.1.2) follows from Proposition 8.6.1, and the following proposition.

Proposition 8.7.1 : *There is some universal $\xi \in [0,\infty)$, and a function $R : [0,\infty) \longrightarrow [0,\infty)$ such that the following holds.*
Suppose that $(\mathcal{S}, d)$ is a κ-H1 and $R(\kappa)$-locally k-H1 geodesic space. Then it is ξk-H1.

We shall need two lemmas.

Lemma 8.7.2 : *$\forall k\, \exists h$ such that $(\mathcal{S}, d)$ if a k-H1 geodesic space, and if $X, Y, Z, W \in \mathcal{S}$ satisfy*

$$\min\big(d(X,Y), d(Z,W)\big) \geq \max\big(d(X,Z), d(Y,W)\big),$$

then $d([X,Y],[Z,W]) \leq h$.

Proof : $XY \wedge ZW \geq 0$. Apply Lemma 7.5.1.
◊

Lemma 8.7.3 : *$\forall k, h\, \exists K$ such that if $(\mathcal{S}, d)$ is a k-H1 geodesic space, and if the points $Y_1, Y_2, Y_3, Z_1, Z_2, Z_3 \in \mathcal{S}$ satisfy $d(Y_i, Z_i) \leq h$ for $i = 1,2,3$, then there is some $C \in \mathcal{S}$ with $d(C, [Y_{i+1}, Z_i]) \leq K$ for $i = 1,2,3$.*

Proof : Let C be a centre of $Y_1Y_2Y_3$. Apply Lemma 3.1.2.
◊

Proof of Proposition 8.7.1 : Suppose $(\mathcal{S}, d)$ is a κ-H1 geodesic space. Let X_1, X_2, X_3 be any three points of $\mathcal{S}$. Let A be a centre for $X_1X_2X_3$ (with respect to the parameter κ). For $i = 1,2,3$, let W_i be a nearest point on $[X_i, X_{i+1}]$ to A. (We take subscripts mod 3.) Applying Lemma 3.1.2, we can find $J = J(\kappa)$ such that $d(W_i, W_{i+1}) \leq J$ and $[X_i, W_i] \subseteq N_J[X_i, W_{i+1}]$ for $i = 1,2,3$. Let $R = 12J$. Thus R depends only on κ.

Suppose now that $(\mathcal{S}, d)$ is also R-locally k-H1. We want to find some point $C \in \mathcal{S}$ whose distance from each of the geodesics $[X_i, X_{i+1}]$ has a bound depending only on k.

Suppose $d(X_1, W_1) > 3J$. Then, let $E \in [W_1, X_1]$ be the point with $d(E, W_1) = 3J$. There is some $F \in [W_3, X_1]$ with $d(E,F) \leq J$. We must have $J \leq d(F, W_3) \leq 5J$. Thus, $\min\big(d(E, W_1), d(F, W_3)\big) \leq \max\big(d(E,F), d(W_1, W_3)\big)$. Also, $\operatorname{diam}[W_1, E, F, W_3, W_1] \leq 5J < R$. Thus, applying Lemmas 8.7.2 and 8.2.3, we can find $Y_1 \in [W_3, F]$ and $Z_1 \in [W_1, E]$ with $d(Y_1, W_3) \leq 5J$, $d(Z_1, W_1) \leq 5J$ and $d(Y_1, Z_1) \leq \min(h, J)$, where $h = h(k)$ comes from Lemma 8.7.2.

In the case when $d(X_1, W_1) \leq 3J$, we take $Y_1 = Z_1 = X_1$.

We perform similar constructions with respect to X_2 and X_3. This gives points $Y_i \in [X_i, W_{i-1}]$ and $Z_i \in [X_i, W_i]$ with $d(Y_i, W_{i-1}) \leq 5J$, $d(Z_i, W_i) \leq 5J$ and $d(Y_i, Z_i) \leq \min(h, J)$.

We thus have $\operatorname{diam}[Y_1, Z_1, Y_2, Z_2, Y_3, Z_3, Y_1] \leq 12J \leq R$. Thus, applying Lemmas 8.7.2 and 8.2.3, we find $C \in \mathcal{S}$ with $d(C, [Y_{i+1}, Z_i]) \leq K$ for $i = 1, 2, 3$, where $K = K(k, h(k))$ depends only on k.

We have shown that $(\mathcal{S}, d)$ is K-H2. It is thus k'-H1 for some k' depending only on k. As discussed in Section 7.2, we can assume that k' has the form ξk for some universal $\xi > 0$.

$\diamondsuit$

References.

[**ABCDFLMSS**] J.M.Alonso, T.Brady, D.Cooper, T.Delzant, V.Ferlini, M.Lustig, M. Mihalik, M.Shapiro, H.Short, *Notes on negatively curved groups* : preprint, M.S.R.I., Berkeley (1989).

[**BGHHSST**] W.Ballmann, E.Ghys, A.Haefliger, P.de la Harpe, E.Salem, R.Strebel, M.Troyanov, *Sur les groupes hyperboliques d'après Mikhael Gromov* : preprint, Geneva (1989).

[**CDP**] M.Coornaert, T.Delzant, A.Papadopoulos, *Notes sur les groupes hyperboliques de Gromov* : preprint, Université Louis Pasteur, Strasbourg (1989).

[**G**] M.Gromov, *Hyperbolic groups,* in "Essays in group theory" : ed. S.M.Gersten, M.S.R.I. Publications No. 8 (1988).