Intersection numbers and the hyperbolicity of the curve complex.

Brian H. Bowditch
Faculty of Mathematical Studies, University of Southampton,
Highfield, Southampton SO17 1BJ, Great Britain.
http://www.maths.soton.ac.uk/staff/Bowditch

[First draft: July 2002. Revised: January 2003]

0. Introduction.

In [MM1], Masur and Minsky showed that the curve complex associated to a surface is hyperbolic in the sense of Gromov. In this paper we give another proof of this result. Our constructions are more combinatorial in nature, and allow for certain refinements and elaborations. The result applies to a compact surface, possibly with boundary (apart from a few trivial cases), or equivalently to a closed surface with a prefered finite subset, thought of as punctures. We use the latter formulation. For simplicity, we describe only the orientable case.

Let Σ be a closed orientable surface, and let $\Pi \subseteq \Sigma$ be a (possible empty) finite set. In [Harv], Harvey associated a "curve complex" to (Σ, Π) as follows. The vertex set, $X = X(\Sigma, \Pi)$, consists of the set homotopy classes of simple closed curves in $\Sigma \setminus \Pi$ (which we refer to simply as "curves"). A set of curves is deemed to span a simplex in the curve complex if they can be realised disjointly in $\Sigma \setminus \Pi$. There are a few "exceptional cases", namely if Σ is a 2-sphere and $|\Pi| \leq 3$, then $X = \emptyset$, and if Σ is either a 2-sphere with $|\Pi| = 4$, or a torus with $|\Pi| \leq 1$, then the associated complex is just a countable set points. On the other hand, if (Σ, Π) is non-exceptional, then it's not hard to see that the curve complex is connected, and has dimension $3\,\mathrm{genus}(\Sigma) + |\Pi| - 4$. We shall take $C(\Sigma, \Pi) = 3\,\mathrm{genus}(\Sigma) + |\Pi| - 4$ as a convenient measure of the "complexity" of (Σ, Π). Thus (Σ, Π) is non-exceptional if $C(\Sigma, \Pi) > 0$.

The main result of [MM1] can be stated as follows.

Theorem 0 : *If $C(\Sigma, \Pi) > 0$, then the curve complex is hyperbolic.*

The hyperbolicity constant may depend on (Σ, Π). In [MM1] part of the argument is non-constructive, and thus there is no explicit estimate of this constant (though in principle, their arguments can be adapted to give some computable bound). In this paper we show that the hyperbolicity constant is bounded by a logarithmic function of complexity — see Proposition 6.1.

Note that all that is relevant here is the 1-skeleton of the curve complex. We shall denote this graph by $\mathcal{G} = \mathcal{G}(\Sigma, \Pi)$. We write d for the induced combinatorial path-metric on X which assigns unit length to each edge of $\mathcal{G}$. We shall confine our discussion of hyperbolicity to graphs.

The notion of a hyperbolic metric space was introduced by Gromov [Gr1]. Other expositions are [Bo1,CDP,GhH,S]. Hyperbolicity is a quasiisometry invariant, from which one can deduce immediately that certain variations on the curve complex are also hyperbolic (see Section 1). The result is thus quite robust.

We remark that the curve complex has some nice combinatorial and topological properties, which have been used to illuminate the structure of the mapping class group. See, for example, [Hare,I,L] and further references therein.

The large scale geometry of these objects has only come the fore more recently, notably in [MM1]. In a sequel, [MM2], the same authors build on this work to investigate certain hierarchical finiteness properties of the curve complex. This is an important ingredient in the study of ends of hyperbolic 3-manifolds by Minsky and his collaborators, culminating in proof of the ending lamination conjecture [BrCM]. A description of the Gromov boundary of the curve complex is given in [K]. The geometry of the curve complex is also used in [BeF] to show that subgroups of the mapping class group that are not virtually abelian have infinite dimensional second bounded cohomology.

Of course, all this would be rather trivial if the curve complex had finite diameter. However, a simple argument given in [MM1], which the authors attribute to Luo, shows that any non-exceptional curve complex has infinite diameter. Indeed, it is also shown in [MM1] that any pseudoanosov mapping class is loxodromic in curve complex, i.e. has positive stable length.

The proofs of the paper are logically independent of those of [MM1], though the latter served as the main source of inspiration for the present paper. The two main components of [MM1] are a sophisticated study of nested train tracks, and an analysis of the geometry of Teichmüller geodesics. Some ideas from the latter are used in this paper, though they are mostly phrased more combinatorially in terms of intersection numbers. Indeed, one can give a description of geodesics and centres etc. purely in terms of intersection numbers. Some of this laid out explicitly in Section 6.

Much of the work for this paper was carried out at the Centre de Recerca Matemàtica, Barcelona and the Max-Planck Institut für Mathematik, Bonn and the Universität Bonn. I am grateful for the hospitality of these institutions. I thank Thomas Delzant and Ursula Hamenstädt for their interest and helpful discussions.

1. Intersection numbers.

In this section, we prove a couple of simple results relating intersection numbers to distances in the curve complex. We go on to briefly comment on a few variations of the curve complex, which can also be seen to be hyperbolic using Theorem 0.

Given curves, $\alpha, \beta \in X$, we write $i(\alpha, \beta)$ for the intersection number of α and β, i.e. the minimal cardinality of $\alpha \cap \beta$ among realisations of α and β in $\Sigma \setminus \Pi$. (Thus, $i(\alpha, \alpha) = 0$.) By definition, $d(\alpha, \beta) \leq 1$ if and only if $i(\alpha, \beta) = 0$.

We describe two inequalities relating $d(\alpha, \beta)$ and $i(\alpha, \beta)$. The first is simple and explicit, and the second, which we describe in more detail, has optimal asymptotics.

For the purposes of proving Theorem 0, any upper bound on $d(\alpha, \beta)$ in terms of $i(\alpha, \beta)$ will suffice. The logarithmic bound of Lemma 1.2 is required to obtain the bound in Proposition 6.1.

Lemma 1.1 : *If $C(\Sigma, \Pi) > 0$, then for all $\alpha, \beta \in X$ we have $d(\alpha, \beta) \leq i(\alpha, \beta) + 1$.*

Proof : Write $j(\alpha, \beta) = i(\alpha, \beta) + 1$. If $j(\alpha, \beta) \geq 2$ then except in one particular case, we can find a curve $\gamma \in X$ with $j(\alpha, \gamma) + j(\gamma, \beta) \leq j(\alpha, \beta)$. This can be proven by an argument similar to that of Lemma 1.3. We omit the details here. The exceptional case consists of two curves, α, β on a twice punctured torus, one obtained from the other by double Dehn twist, so that $i(\alpha, \beta) = 2$. In this case, one verifies directly that $d(\alpha, \beta) = 3$. In general, by continuing this subdivision, we arrive at a sequence of curve, $\alpha = \gamma_0, \ldots, \gamma_n = \beta$ with $j(\gamma_i, \gamma_{i+1}) = 1$ for all i. This gives a path of length n from α to β. Moreover, the above inequality implies that $n \leq j(\alpha, \beta)$, and so the result follows. ♢

The following variation is better for large distances:

Lemma 1.2 : *There is a function, $F : \mathbb{N} \longrightarrow \mathbb{N}$ with $F(n) = O(\log n)$ such that if $C(\Sigma, \Pi) > 0$ and $\alpha, \beta \in X$ then $d(\alpha, \beta) \leq F(i(\alpha, \beta))$.*

In general, one can not do better than logarithmic, as can be seen for example by considering the images of a fixed curve α under the iterates of a pseuoanosov, ψ, (see [MM1]). In this case $d(\alpha, \psi^n(\alpha))$ grows linearly in n, whereas $i(\alpha, \psi^n(\alpha))$ grows exponentially. Of course, there is no general lower bound, and in some sense, it would seem that in the "generic" situation one should expect something better than logarithmic.

Before continuing, we should note that, since we are assuming Σ to be orientable, the intersections of two curves, α, β, can be natural partitioned into two subsets according to the relation of the orientations of α, β and Σ.

To prove Lemma 1.2 we need:

Lemma 1.3 : *Suppose $C(\Sigma, \Pi) > 0$. If $\alpha, \beta \in X$ and $a, b \in \mathbb{N}$ with $ab \geq 2i(\alpha, \beta)$, then there is some $\gamma \in X$ with $i(\alpha, \gamma) \leq a$ and $i(\beta, \gamma) \leq b$.*

Proof : Let $n = i(\alpha, \beta)$. We can assume that $b \leq a$, and so $c(b + 1) \geq n$, where c is the integer part of $a/2$. We can thus write $\alpha = \alpha_1 \cup \cdots \cup \alpha_c$ where each α_i is a subarc of α containing at most $b + 1$ points of $\alpha \cap \beta$. We can also assume that $n \geq 2c + 1$, otherwise $n \leq a$, and we could simply take $\gamma = \beta$. Now let β_0 be any subarc of β containing exactly $2c + 1$ points of $\alpha \cap \beta$. There must be some i so that $|\alpha_i \cap \beta_0| \geq 3$. Thus, at least two intersections, say $x, y \in \alpha_i \cap \beta_0$, have the same orientation. Let α' and β' be the subarcs of α_i and β_0 respectively lying between x and y. By passing to smaller subarcs, if necessary, we can suppose that the interiors α' and β' meet, if at all, in a single point with the opposite orientation.

Now if $\alpha' \cup \beta'$ is simple then it represents an element, $\gamma \in X$ with $i(\alpha, \gamma) \leq 2c < a$ and $i(\beta, \gamma) \leq b$ (as can be seen by representing γ by a parallel curve running close on one side of $\alpha' \cup \beta'$). If α' and β' meet at an interior point, z, then z cuts $\alpha' = \alpha'_1 \cup \alpha'_2$ and $\beta' = \beta'_1 \cup \beta'_2$, where α'_1, β'_1 join x to z and α'_2, β'_2 join z to y. At least one of the curves $\alpha'_1 \cup (-\beta'_1)$, $\alpha'_2 \cup (-\beta'_2)$ or $\alpha'_1 \cup \beta'_2 \cup (-\alpha'_2) \cup (-\beta'_1)$ is essential and represents an element $\gamma \in X$. This satisfies $i(\alpha, \gamma) \leq 2c \leq a$ and $i(\beta, \gamma) \leq b$ as required. ♢

Proof of Lemma 1.2 : Write $h(n) = \log_3(n) - 1$. If $n \geq 6$ is an integer, then there are integers $a, b \geq 4$ with $2n \leq ab \leq 3n$. Thus $h(a) + h(b) \leq h(n)$. Given $\alpha, \beta \in X$, write $j(\alpha, \beta) = \max\{h(4), h(i(\alpha, \beta))\}$. Applying the previous observation and Lemma 1.3, we see that if $\alpha, \beta \in X$ with $i(\alpha, \beta) \geq 6$, then we can find $\gamma \in X$ with $j(\alpha, \gamma) + j(\gamma, \beta) \leq j(\alpha, \beta)$. Continuing to subdivide in this fashion, we eventually arrive at a sequence $\alpha = \gamma_0, \gamma_1, \ldots, \gamma_m = \beta$, with $\sum_{i=1}^{m} j(\gamma_{i-1}, \gamma_i) \leq j(\alpha, \beta)$, and with $i(\gamma_{i-1}, \gamma_i) \leq 5$ for all i. Thus, $mh(4) \leq j(\alpha, \beta)$. Moreover, applying Lemma 1.1, we have $d(\gamma_{i-1}, \gamma_i) \leq 5 + 1 = 6$ for all i. Thus $d(\alpha, \beta) \leq 6m/h(4)$. Setting $F(n) = 6 \max\{1, h(n)/h(4)\}$, the result follows. $\diamondsuit$

Before continuing, we note that if $d(\alpha, \beta) \geq 3$, then α, β *fill* the surface $\Sigma \setminus \Pi$. This means that for any realisation of α and β in $\Sigma \setminus \Pi$, every complementary component of $\Sigma \setminus \Pi$ is *trivial*, that is a topological disc containing at most one point of Π.

We can use intersection numbers to define a number of variations of the curve graph $\mathcal{G}(\Sigma, \Pi)$ which we briefly comment on. Given any $p \geq 0$, write $\mathcal{G}_p = \mathcal{G}_p(\Sigma, \Pi)$ for the graph with vertex set $X = X(\Sigma, \Pi)$, where $\alpha, \beta \in X$ are deemed to be adjacent if $i(\alpha, \beta) \leq p$ (so that $\mathcal{G}_0 = \mathcal{G}$). The embedding of $\mathcal{G}$ in $\mathcal{G}_p$ is a quasiisometry, and so $\mathcal{G}_p$ is hyperbolic if $C(\Sigma, \Pi) > 0$. We note that if $\mathrm{genus}(\Sigma) = 1$, $|\Pi| \leq 1$ and $p = 1$ or if $\mathrm{genus}(\Sigma) = 0$, $|\Pi| = 4$ and $p = 2$, then $\mathcal{G}_p(\Sigma, \Pi)$ is a Farey graph, and hence hyperbolic (indeed quasiisometric to a regular infinite valence tree). The graphs are thus also hyperbolic for any larger p.

Indeed we can generalise further. Given $q \geq 0$, let X_q be the set of curves with self-intersection number at most q (i.e. those that can be realised with at most q self-intersections). If $p \geq q$, we define $\mathcal{G}_{p,q}$ to be the graph with vertices $\alpha, \beta \in X_q$ adjacent if $d(\alpha, \beta) \leq p$. Clearly $\mathcal{G}_p$ embeds in $\mathcal{G}_{p,q}$ as a full subgraph. Unless $\mathrm{genus}(\Sigma) = 0$ and $|\Pi| \leq 3$, every point of X_q is adjacent to a point of X. Indeed if $\alpha \in X_q$ then there is some $\beta \in X$ which can be homotoped into the image of α such that there is at most one preimage of any point other that a self-intersection of α. We have $d(\alpha, \beta) \leq q \leq p$, so α, β are are adjacent. Moreover, the map that sends α to β sends cannot increase intersection numbers, and so sends adjacent points to adjacent points. It thus defines a retraction of $\mathcal{G}_{p,q}$ onto $\mathcal{G}_p$. It follows that the inclusion of $\mathcal{G}_p$ into $\mathcal{G}_{p,q}$ is a quasiisometry. In the special case where Σ is a sphere and $|\Pi| = 3$, then $\mathcal{G}_{p,q}$ is a finite connected graph whenever $p > q \geq 1$ or $p \geq q \geq 2$.

From all this, we may conclude:

Proposition 3.4 : *Suppose Σ is a closed orientable surface, $\Pi \subseteq \Sigma$ is finite, and $p, q \in \mathbb{N}$ with $p \geq q$. If $\mathcal{G}_{p,q}(\Sigma, \Pi)$ has positive dimension, then it is hyperbolic.* $\diamondsuit$

Finally another variation on the curve complex is used in [MM2]. Suppose $P \subseteq \Pi$. In addition to classes of simple closed curves in $\Sigma \setminus \Pi$, we allow classes of arcs with endpoints in P, but otherwise disjoint from Π. An arc might have its two endpoints identified to a single point of P. Two arcs are in the same class if one can be deformed to the other through such arcs. This defines the vertex set, $X(\Sigma, \Pi, P)$. Two vertices are deemed adjacent if they can be realised so as to be disjoint outside P. This defines a graph $\mathcal{G}(\Sigma, \Pi, P)$. Note that $\mathcal{G}(\Sigma, \Pi, \emptyset) = \mathcal{G}(\Sigma, \Pi)$. Moreover, the inclusion of $\mathcal{G}(\Sigma, \Pi)$ into $\mathcal{G}(\Sigma, \Pi, P)$ is a quasiisometry, and so the latter is hyperbolic if $C(\Sigma, \Pi) > 0$. Indeed, we can include the cases where $\mathrm{genus}(\Sigma) = 1$ and $|P| = |\Pi| = 1$, and where $\mathrm{genus}(\Sigma) = 0$, $|\Pi| = 4$, and $P \neq \emptyset$.

2. The idea of the proof.

In this section, we sketch the strategy for proving Theorem 0. We also introduce some terminology regarding multicurves that will also be used later.

Suppose $C(\Sigma,\Pi) > 0$ and write $X = X(\Sigma,\Pi)$ and $\mathcal{G} = \mathcal{G}(\Sigma,\Pi)$. The basic idea, as in [MM1], is to construct a preferred family of paths connecting any pair of vertices. Thus, if $\alpha,\beta \in X$, we have a path $\pi_{\alpha\beta}$ in $\mathcal{G}$ from α to β. In [MM1] the authors construct a kind of uniform retraction from X onto each $\pi_{\alpha\beta}$, which far from $\pi_{\alpha\beta}$ decreases distances by a definite factor. However the latter point is difficult to establish, and relies, among other things, on a certain combinatorial lemma resulting from a sophisticated analysis of nested train tracks. Here, we do not need such an estimate. Instead, we show directly that any triangle formed by three paths $\pi_{\alpha\beta}$, $\pi_{\beta\gamma}$ and $\pi_{\gamma\alpha}$ is "thin" in an appropriate sense. In particular, there is a "centre", $\phi(\alpha,\beta,\gamma) \in X$, which is a bounded distance from all three sides. A key point in the argument is to show that if $\gamma,\delta \in X$ are adjacent, then $d(\phi(\alpha,\beta,\gamma),\phi(\alpha,\beta,\delta))$ is bounded. (A similar statement is also proven in [MM1] by a different method.) Given this, one sees that the paths $\pi_{\alpha\beta}$ are uniformly quasigeodesic. From this, the hyperbolicity of $\mathcal{G}$ follows via a subquadratic isoperimetric inequality (Proposition 3.1).

In practice, we first construct a "line" from α to β. This is a subset $\Lambda_{\alpha\beta} \subseteq X$ which will be within a bounded Hausdorff distance of $\pi_{\alpha\beta}$. It also carries a "coarse order", $\leq_{\alpha\beta}$, which measures the approximate order of points along $\pi_{\alpha\beta}$.

To this end, we define a *weighted curve* formally as a pair, (λ,α) where $\lambda \in (0,\infty)$ and $\alpha \in X$. We denote this by $\lambda\alpha$, and think of λ as a *weight* assigned to α. We write WX for the set of all weighed curves. By identifying 1α with α, we can regard $X \subseteq WX$. Given $\lambda\alpha,\mu\beta \in WX$, we write $d(\lambda\alpha,\mu\beta) = d(\alpha,\beta)$ and $i(\lambda\alpha,\mu\beta) = \lambda\mu i(\alpha,\beta)$.

To define $\Lambda_{\alpha\beta}$, we can assume $i(\alpha,\beta) > 0$, and choose $\lambda,\mu > 0$ so that $i(\lambda\alpha,\mu\beta) = 1$. One can show (Lemma 4.3) that, for some fixed $R \geq 0$, depending only on $C(\Sigma,\Pi)$, there is some $\delta \in X$ with $i(\lambda\alpha,\delta) \leq R$ and $i(\mu\beta,\delta) \leq R$. Moreover, for given λ,μ, any two such curves are a bounded distance apart in X. We now let λ,μ vary, and let $\Lambda_{\alpha\beta}$ the set of all curves arising in this way.

The construction of a centre of $\alpha,\beta,\gamma \in X$ is based on a similar idea. We can assume that no two of these curves are equal or adjacent, and so we can renormalise them so as to give weighed curves, $\bar{\alpha},\bar{\beta},\bar{\gamma} \in WX$, with $i(\bar{\alpha},\bar{\beta}) = i(\bar{\beta},\bar{\gamma}) = i(\bar{\gamma},\bar{\alpha}) = 1$. In this case, we can take a centre, $\phi(\alpha,\beta,\gamma)$, to be a curve $\delta \in X$ with $i(\bar{\alpha},\delta) \leq R$, $i(\bar{\beta},\delta) \leq R$ and $i(\bar{\gamma},\delta) \leq R$ (Lemma 4.7).

It turns out that one can do a similar construction for multicurves — a fact that used to put a bound on $d(\phi(\alpha,\beta,\gamma),\phi(\alpha,\beta,\delta))$ for γ,δ adjacent (Propostion 4.11).

To this end, we will need to introduce some notation regarding multicurves. A *multicurve* is a non-empty subset, $\{\alpha_1,\ldots,\alpha_n\}$, of X that spans a simplex in the curve complex, i.e. the curves α_i can be realised disjointly. A *weighted multicurve* is a formal sum, $\alpha = \sum_{i=1}^n \lambda_i\alpha_i$ where $\lambda \in [0,\infty)$ and $\sum_{i=1}^n \lambda_i > 0$. If $\alpha' = \sum_{i=1}^n \lambda'_i\alpha_i$ and $\nu > 0$, we write $\alpha + \alpha' = \sum_{i=1}^n (\lambda + \lambda'_i)\alpha_i$ and $\nu\alpha = \sum_{i=1}^n (\nu\lambda_i)\alpha_i$. We say that α and $\nu\alpha$ are in the same *projective class*. We write MX and WMX for the sets of multicurves and weighed multicurves respectively. We view $X \subseteq MX \subseteq WMX$ and $X \subseteq WX \subseteq WMX$. (In practice, we only need to deal with multicurves with at most two elements, but this makes not difference to the discussion.)

If $\alpha \in \sum_i \lambda_i \alpha_i, \quad \beta = \sum_j \mu_j \beta_j \in WMX$, we write $i(\alpha, \beta) = \sum_{i,j} \lambda_i \mu_j i(\alpha_i, \beta_j)$ and $d(\alpha, \beta) = \min_{i,j} d(\alpha_i, \beta_j)$.

3. A characterisation of hyperbolicity.

In this section, we describe the characterisation of hyperbolicity we shall use. It can readily be interpreted for any path-metric space, though we state it only for a connected graph, $\mathcal{G}$, with vertex set, X, and distance function d, that assigns unit length to each edge.

We assume that X has associated to it the following structures. To each pair, $a, b \in X$, we have an associated subset $\Lambda_{ab} \subseteq X$, an well as a "coarse order", $\leq_{ab}$, on Λ_{ab}. By a *coarse order* we mean that $\leq_{ab}$ is reflexive and transitive and and satisfies the dichotomy rule (for all $x, y \in \Lambda_{\alpha\beta}$, either $x \leq_{ab} y$ or $y \leq_{ab} x$). However it need not be antisymmetric. We define minima and maxima in the usual way. Moreover, we assume we are given a ternary function, $\phi : X \times X \times X \longrightarrow X$. We refer to $(\Lambda_{ab}, \leq_{ab})$ as the *line* from a to b, and to $\phi(a, b, c)$ as the *centre* of a, b, c.

We assume that ϕ and $[(a, b) \mapsto (\Lambda_{ab}, \leq_{ab})]$ satisfy the following conditions. We have $\Lambda_{ab} = \Lambda_{ba}$ for all $a, b \in X$. Moreover, Λ_{ab} and Λ_{ba} have reverse order, i.e. $\leq_{ab} = \geq_{ba}$. Given $x, y \in \Lambda_{ab}$, with $x \leq_{ab} y$, we shall write

$$\Lambda_{ab}[x, y] = \Lambda_{ab}[y, x] = \{z \in \Lambda_{ab} | x \leq_{ab} z \leq_{ab} y\}.$$

We suppose that ϕ has the symmetry $\phi(a, b, c) = \phi(b, c, a) = \phi(b, a, c)$ and that $\phi(a, a, b) = a$, for all $a, b, c \in X$. We suppose that $\phi(a, b, c) \in \Lambda_{ab} \cap \Lambda_{bc} \cap \Lambda_{ca}$. Moreover, there is a constant, $K \geq 0$ with the following properties.

(1) If $a, b, c \in X$ then $\mathrm{HausDist}(\Lambda_{ab}[a, \phi(a, b, c)], \Lambda_{ac}[a, \phi(a, b, c)]) \leq K$.

(2) If $x, y \in X$ with $d(x, y) \leq 1$, then $\mathrm{diam}\, \Lambda_{ab}[\phi(a, b, x), \phi(a, b, y)] \leq K$.

(3) If $c \in \Lambda_{ab}$, then $\Lambda_{ab}[c, \phi(a, b, c)] \leq K$.

Here diam and HausDist denote respectively, diameter and Hausdorff distance with respect to the metric d on X.

We shall show:

Proposition 3.1 : *If X admits a system of lines and centres satisfying the above axioms then X is hyperbolic with hyperbolicity constant depending only on K. Moreover, for all $a, b \in X$, the line Λ_{ab} is a bounded Hausdorff distance (again depending only on K) from any geodesic connecting a to b.*

We remark that the converse in clear. Given $a, b \in X$, let $[a, b]$ be any geodesic connecting a to b. We can take Λ_{ab} to be a uniform neighbourhood of $[a, b]$. Thus $\phi(a, b, c)$ is a centre of the triangle $[a, b] \cup [b, c] \cup [c, a]$. The coarse order on Λ_{ab} can be obtained from the linear order of nearest points on $[a, b]$.

We now set about proving Proposition 3.1. We fix, for the moment, $a, b \in X$. Given $x \in X$, we write $\phi(x) = \phi_{ab}(x) = \phi(a, b, x)$. Thus $\phi : X \longrightarrow \Lambda_{ab}$ with $\phi(a) = a$ and $\phi(b) = b$.

Lemma 3.2 : *There is a sequence $a = x_0 < x_1 < \cdots < x_n = b$ of points of Λ_{ab} with $\operatorname{diam} \Lambda_{ab}[x_i, x_{i+1}] \leq K$ for all i, with $n \leq d(a, b)$ and $|i - j| \leq d(x_i, x_j) + 2$ for all i, j.*

Proof : By a *chain* we mean a sequence $(x_i)_{i=0}^n$ in Λ_{ab} with $x_0 = a$, $x_n = b$ and with $\operatorname{diam} \Lambda_{ab}[x_i, x_{i+1}] \leq K$ for all i. Such a chain exists, since if we connect a to b by a geodesic, $a = y_0, y_1, \ldots, y_n = b$ in X, then $(\phi_i(y_i))_i$ is a chain. We now choose a chain, $(x_i)_{i=0}^n$, with n minimal. Thus $n \leq d(a, b)$. Moreover, if $i < j$, let $x_i = z_0, z_1, \ldots, z_m = x_j$ be a geodesic in X from x_i to x_j, so that $m = d(x_i, x_j)$. Let $w_i = \phi(z_i)$. Replacing $x_i, x_{i+1}, \ldots, x_j$ by $x_i, w_0, w_1, \ldots, w_m, x_j$, we get another chain, so by minimality of n, it follows that $j - i \leq m + 2$ as claimed. Finally suppose $x_{i+1} \leq x_i$. Then there is some $j > i$ with $x_j \leq x_i$ and $x_{j+1} > x_i$. Thus $\Lambda_{ab}[x_i, x_{j+1}] \subseteq \Lambda_{ab}[x_j, x_{i+1}]$. We could thus omit the points $x_{j+1}, \ldots, x_j$ to obtain a shorter chain. It follows that $x_0 < x_1 < \cdots < x_n$ as claimed. ◊

Note that $\Lambda_{ab} = \bigcup_{i=0}^{n-1} \Lambda_{ab}[x_i, x_{i+1}]$. In particular, $\{x_0, x_1, \ldots, x_n\}$ is K-dense in Λ_{ab}. We write π_{ab} for the concatenation of the paths $[x_i, x_{i+1}]$ where $[x, y]$ denotes any geodesic connecting x to y. Thus, $\operatorname{HausDist}(\pi_{ab}, \Lambda_{ab}) \leq K$. We refer to each x_i as a *breakpoint* of π_{ab}. We can assume that $\pi_{ab} = \pi_{ba}$. We write $T(a, b, c)$ for the triangle $\pi_{ab} \cup \pi_{bc} \cup \pi_{ca}$. Note that the total number of breakpoints in $T(a, b, c)$ is at most $d(a, b) + d(b, c) + d(c, a) \leq 3 \operatorname{diam}\{a, b, c\}$. Given $x, y \in \pi_{ab}$, write $\pi_{ab}[x, y]$ for the segment of π_{ab} between x and y. We note:

Lemma 3.3 : *For all $a, b \in X$, the paths π_{ab} are uniformly quasigeodesic.*

Proof : In other words, if $x, y \in \pi_{ab}$, then the length of π_{ab} is bounded above by a fixed linear function of $d(x, y)$. We can assume that $x = x_i$ and $y = x_j$ are breakpoints. But now, $\operatorname{length}(\pi_{ab}[x_i, x_j]) \leq K|i - j| \leq K(d(x_i, x_j) + 2)$ as required. ◊

Before continuing, we need to formulate the notion of an "isoperimetric inequality". We need to make sense of the notion of a curve γ in the graph $\mathcal{G}$ "spanning a disc" of area to most A. There are several ways to do this. One is as follows.

We can think of γ as a map, $\gamma : S^1 \longrightarrow \mathcal{G}$, where S^1 is the boundary of the unit disc D. A *cellulation* of D is a representation of D as a CW-complex. A "spanning disc" for γ is then an extension of γ to the 1-skeleton of some cellulation of D. Its *mesh* is the maximal length of the boundary of a 2-cell of the cellulation. Its *area* is the number of 2-cells.

Hyperbolicity can be characterised by a subquadratic isoperimetric inequality [Gr1, CDP,O,P,Bo2] as follows:

Theorem 3.4 : *Given $M \geq 0$ and a function $g : \mathbb{N} \longrightarrow \mathbb{N}$ with $g(n) = o(n^2)$ as $n \to \infty$, there is a constant $k \geq 0$ such that if $\mathcal{G}$ is a graph in which every loop in $\mathcal{G}$ bounds a spanning disc of mesh at most M and area at most $g(\operatorname{length}(\gamma))$, then X is k-hyperbolic.*

This is a standard form that will suffice for Theorem 0. To get better control on the hyperbolicity constant for Proposition 6.1, we use another formulation in terms of homology and where the mesh is defined in terms of diameter rather than length. This is discussed in Section 7.

We now return to our particular space X.

Lemma 3.5 : *There is a constant, M, depending only on K, such that for all $a, b, c \in X$, the triangle $T(a, b, c)$ bounds a spanning disc of mesh at most M and area at most $3 \operatorname{diam}\{a, b, c\}$.*

Proof : Let $w = \phi(a, b, c)$. Let p_{ab} and p_{ba} be equal or adjacent breakpoints of π_{ab} with $p_{ab} \leq_{ab} w \leq_{ab} p_{ba}$ in Λ_{ab}. We similarly define p_{bc}, p_{cb}, p_{ca} and p_{ac}. Let $\sigma_{ab} = \pi_{ab}[a, p_{ab}]$ etc. Thus, from the construction of π_{ab}, $\operatorname{HausDist}(\sigma_{ab}, \Lambda_{ab}[a, p_{ab}]) \leq K$. Also, since $\Lambda_{ab}[p_{ab}, w] \subseteq \Lambda_{ab}[p_{ab}, p_{ba}]$, it follows that $\operatorname{diam}(\Lambda_{ab}[p_{ab}, w]) \leq K$. One of our assumptions is that the Hausdorff distance between $\Lambda_{ab}[a, w]$ and $\Lambda_{ac}[a, c]$ is at most K. We conclude that $d(p_{ac}, p_{ab}) \leq 2K$ and that $\operatorname{HausDist}(\sigma_{ab}, \sigma_{ac}) \leq 3K$.

Now since σ_{ab} and σ_{ac} are uniformly quasigeodesic (depending on K), it is an easy exercise to show that $\sigma_{ab} \cup \sigma_{ac} \cup [p_{ab}, p_{ac}]$ bounds a spanning disc of bounded mesh, and with area bounded by the number of segments in σ_{ab}. Indeed this can be achieved by adding a set of geodesics between distinct breakpoints of σ_{ab} and distinct breakpoints of σ_{ac}. We perform similar constructions with respect to the vertices b and c. By including the hexagonal curve $p_{ab}p_{ba}p_{bc}p_{cb}p_{ca}p_{ba}p_{ab}$ as the boundary of a 2-cell, we obtain a spanning disc for $T(a, b, c)$, with at most $3 \operatorname{diam}\{a, b, c\}$ 2-cells, and with mesh bounded in terms of K as required. ♢

Proof of Proposition 3.1 : We first show that $\mathcal{G}$ is hyperbolic by deriving a subquadratic isoperimetric inequality, as described by Theorem 3.4.

To this end, let γ be a loop of length n in $\mathcal{G}$. Given a natural number p, we shall divide γ into 2^{p+1} segments, $\gamma = \gamma_1^p \cup \gamma_2^p \cup \cdots \gamma_{2^{p+1}}^p$ by binary subdivision as follows. We start by writing $\gamma = \gamma_1^0 \cup \gamma_2^0$, with $|\operatorname{length}(\gamma_1^0) - \operatorname{length}(\gamma_2^0)| \leq 1$. We now continue inductively. If we have constructed $\gamma_1^{p-1}, \ldots, \gamma_{2^p}^{p-1}$, we write each $\gamma_i^{p-1} = \gamma_{2i-1}^p \cup \gamma_{2i}^p$, with $|\operatorname{length}(\gamma_{2i-1}^p) - \operatorname{length}(\gamma_{2i}^p)| \leq 1$. After $q = [\log_2(n-1)] \leq \log_2 n$ steps, we find that $\operatorname{length}(\gamma_i^q) \leq 1$ for all i and we stop.

Now let β_i^p be a geodesic in $\mathcal{G}$ with the same endpoints as γ_i^p. We can suppose that $\beta_1^0 = \beta_2^0$ and that $\beta_i^q = \gamma_i^q$ for all i. We let $\beta^p = \beta_1^p \cup \beta_2^p \cup \cdots \cup \beta_{2^{p+1}}^p$. Thus β^p is a 2^{p+1}-gon with vertices in γ. Moreover $\beta^q = \gamma$, and the image of β^0 consists of a single geodesic.

Now the region between β^{p-1} and β^p consists of a cycle of 2^p triangles of the form $\beta_i^{p-1} \cup \beta_{2i-1}^p \cup \beta_{2i}^p$, for $i = 1, \ldots, 2^p$. By Lemma 3.5, each such triangle bounds a spanning disc of mesh at most M and area at most three times the diameter of its vertex set, and hence at most $3(\operatorname{length}(\gamma_{2i-1}^p) + \operatorname{length}(\gamma_{2i}^p))$. Thus, the region between β^{p-1} and β^p is spanned by a union of spanning discs of total area at most $3 \sum_{i=1}^{2^{p+1}} \operatorname{length}(\gamma_i^p) = 3 \operatorname{length}(\gamma) = 3n$. Assembling these for $p = 1, \ldots, q \leq \log_2 n$, we obtain a spanning disc for γ of mesh at most M, and area at most $3n \log_2 n$. This proves that $\mathcal{G}$ is hyperbolic.

Now, by Lemma 3.3, the path π_{ab} is uniformly quasigeodesic, for any $a, b \in X$. Since $\mathcal{G}$ is hyperbolic, it follows that π_{ab} is within a bounded Hausdorff distance of any geodesic connecting a to b. Thus, Λ_{ab} is also within a bounded Hausdorff distance of any such geodesic. This proves Proposition 3.1. $\diamondsuit$

4. Construction of lines and centres.

In this section, we describe how to construct a system of lines and centres in the curve complex satisfying the hypotheses of Proposition 3.1. This will prove Theorem 0 modulo one result, namely Lemma 4.1, whose proof we postpone until Section 5.

Let (Σ, Π) be a non-exceptional surface. (In other words, we insist that $|\Pi| \geq 5$ if Σ is a sphere, and that $|\Pi| \geq 2$ if Σ is a torus.) Recall that X, WX, MX, WMX are the sets of curves, weighted curves, multicurves and weighted multicurves respectively. We view $X \subseteq WX \subseteq WMX$ and $X \subseteq MX \subseteq WMX$. Given $\alpha \in X$, we write $N(\alpha) = \{\beta \in X \mid d(\alpha, \beta) \leq 1\}$.

We shall need the following notions. Suppose $\alpha, \beta \subseteq WMX$ with $i(\alpha, \beta) = 1$. Given $\delta \in WMX$, set

$$l(\delta) = l_{\alpha\beta}(\delta) = \max\{i(\alpha, \delta), i(\beta, \delta)\}.$$

If $\delta \in X$, write

$$m(\delta) = \sup(\{l(\delta)\} \cup \{\frac{i(\gamma, \delta)}{l(\gamma)} \mid \gamma \in X\}).$$

For the latter, we adopt the convention that $0/0 = 0$ and that $n/0 = \infty$ for $n > 0$. Thus $m(\delta) \in [0, \infty]$. Note that, for all $\gamma, \delta \in X$, we have $l(\delta) \leq m(\delta)$ and $i(\gamma, \delta) \leq l(\gamma) m(\delta)$. Indeed, by linearity of intersection number, the second inequality holds for all $\gamma \in WMX$. Moreover, if $m(\delta) < \infty$ and $l(\gamma) = 0$, then $i(\gamma, \delta) = 0$. (A more intuitive description of $m(\delta)$ in terms of the reciprocal of the width of an annular neighbourhood of δ in $\Sigma \setminus \Pi$ will be given in Section 5.)

Given $r \geq 0$, write

$$L(\alpha, \beta, r) = \{\delta \in X \mid l(\delta) \leq r\}$$

$$M(\alpha, \beta, r) = \{\delta \in X \mid m(\delta) \leq r\}.$$

Clearly, $M(\alpha, \beta, r) \subseteq L(\alpha, \beta, r)$.

In Section 5, we shall show:

Lemma 4.1 : *There is some $R \geq 0$, depending only on the topological type of (Σ, Π) such that if $\alpha, \beta \in WMX$ with $i(\alpha, \beta) = 1$ and $d(\alpha, \beta) \geq 2$, then $M(\alpha, \beta, R) \neq \emptyset$.*

We deduce:

Lemma 4.2 : *For all $r \geq 0$, $\operatorname{diam} L(\alpha, \beta, r) \leq 2Rr + 2$.*

Proof : Choose any $\delta \in M(\alpha,\beta)$. If $\gamma \in L(\alpha,\beta,r)$ by Lemma 1.1, $d(\gamma.\delta) \leq i(\gamma,\delta)+1 \leq l(\gamma)m(\delta)+1 \leq Rr+1$. The result follows. ♢

In particular, setting $D = 2R^2+2$, we see:

Lemma 4.3 : *$L(\alpha,\beta,R)$ is non-empty and has diameter at most D.* ♢

We remark that using Lemma 1.2 in place of Lemma 1.1, we can replace $2Rr+2$ in Lemma 4.2 by a function that is $O(\log(Rr))$. This gives $D = O(\log(R))$. This observation also applies to the proof of Lemma 4.12 below.

Now, given $\alpha,\beta \in WMX$ with $d(\alpha,\beta) \geq 2$, write

$$\kappa(\alpha,\beta) = \log i(\alpha,\beta).$$

Given $t \in \mathbb{R}$, write $\alpha_t = e^{-t}\alpha$. Setting $u = \kappa(\alpha,\beta)-t$ and $\beta_u = e^{-u}\beta$, we get $i(\alpha_t,\beta_u)=1$.

Given $\gamma \in WMX$ and $\delta \in X$, write $l_t(\gamma) = l_{\alpha_t\beta_u}(\gamma)$ and $m_t(\delta) = m_{\alpha_t\beta_u}(\delta)$. (We use $l_{t,\alpha\beta}(\gamma)$ and $m_{t,\alpha\beta}(\delta)$ if there is any ambiguity concerning α,β.) If $r \geq 0$, set:

$$L(t,r) = L_{\alpha\beta}(t,r) = L(\alpha_t,\beta_u,r) = \{\gamma \in X \mid l_t(\gamma) \leq r\}$$
$$M(t,r) = M_{\alpha\beta}(t,r) = M(\alpha_t,\beta_u,r) = \{\gamma \in X \mid m_t(\gamma) \leq r\}.$$

Note that $L_{\alpha\beta}(t,r)$ and $M_{\alpha\beta}(t,r)$ depend only on the projective class of β (since the weighting on α is fixed by t, and hence determines that on β). We abbreviate $L(t) = L(t,R)$ and $M(t) = M(t,R)$.

Suppose now that $\alpha,\beta,\gamma \in WMX$ with $\min\{d(\alpha,\beta),d(\beta,\gamma),d(\gamma,\alpha)\} \geq 2$. Write

$$\tau_\alpha = \tau_{\alpha\beta}(\gamma) = \tau_{\alpha\gamma}(\beta) = \frac{1}{2}(\kappa(\alpha,\beta)+\kappa(\alpha,\gamma)-\kappa(\beta,\gamma))$$

etc. permuting α,β,γ. Thus, $\tau_\alpha+\tau_\beta = \kappa(\alpha,\beta)$, $\tau_\beta+\tau_\gamma = \kappa(\beta,\gamma)$ and $\tau_\gamma+\tau_\alpha = \kappa(\gamma,\alpha)$. Setting $\bar\alpha = \alpha_{\tau_\alpha} = e^{-\tau_\alpha}\alpha$ etc. we obtain $i(\bar\alpha,\bar\beta) = i(\bar\beta,\bar\gamma) = i(\bar\gamma,\bar\alpha) = 1$. Note also that $L_{\alpha\beta}(\tau_\alpha) = L_{\beta\alpha}(\tau_\beta) = L(\bar\alpha,\bar\beta,R)$ and $M_{\alpha\beta}(\tau_\alpha) = M_{\beta\alpha}(\tau_\beta) = M(\bar\alpha,\bar\beta,R)$ etc.

We note:

Lemma 4.4 : *If $t \leq \tau_\alpha$, then $M_{\alpha\beta}(t) \subseteq L_{\alpha\gamma}(t)$.*

Proof : Let $u = \kappa(\alpha,\beta)-t$, $v = \kappa(\alpha,\gamma)-t$ and $\lambda = e^{\tau_\alpha - t} \geq 1$. Thus, $t+u = \kappa(\alpha,\beta) = \tau_\alpha+\tau_\beta$, so $\beta_u = e^{-u}\beta = e^{t-\tau_\alpha}e^{-\tau_\beta}\beta = \lambda^{-1}\bar\beta$. Similarly, $\gamma_v = \lambda^{-1}\bar\gamma$. We also have $\alpha_t = \lambda\bar\alpha$.

Suppose $\delta \in M_{\alpha\beta}(t)$. Now

$$\begin{aligned} i(\delta,\alpha_t) &\leq \max\{i(\delta,\alpha_t), i(\delta,\beta_u)\} \\ &= l_{t,\alpha\beta}(\gamma) \leq m_{t,\alpha\beta}(\gamma) \leq R. \end{aligned}$$

Also,

$$\begin{aligned} i(\delta,\gamma_v) &\leq m_{t,\alpha\beta}(\delta) l_{t,\alpha\beta}(\gamma_v) \\ &\leq R\max\{i(\alpha_t,\gamma_v), i(\beta_u,\gamma_v)\} \\ &= R\max\{i(\lambda\bar\alpha,\lambda^{-1}\bar\gamma), i(\lambda^{-1}\bar\beta,\lambda^{-1}\bar\gamma)\} \\ &= R\max\{i(\bar\alpha,\bar\gamma), \lambda^{-2} i(\bar\beta,\bar\gamma)\} \\ &= R\max\{1,\lambda^{-2}\} = R. \end{aligned}$$

Thus $l_{t,\alpha\gamma}(\delta) = \max\{i(\delta, \alpha_t), i(\delta, \gamma_v)\} \leq R$, and so $\delta \in L_{\alpha\gamma}(t)$ as required. ◊

Putting this together with Lemma 4.1, we see that $L_{\alpha\beta}(t) \cap L_{\alpha\gamma}(t) \neq \emptyset$, so by Lemma 4.3 we see:

Lemma 4.5 : *If* $t \leq \tau_\alpha$, *then* $\mathrm{diam}(L_{\alpha\beta}(t) \cup L_{\alpha\gamma}(t)) \leq 2D$. ◊

We are now in a position to define lines. Suppose $\alpha, \beta \in X$ and $d(\alpha, \beta) \geq 2$. Given $t \in \mathbb{R}$, we define $L(t)$ as above. Note as $t \to \infty$ the weights assigned to α and β tend to 0 and ∞ respectively. It follows that $L(t) = N(\beta)$ for all sufficiently large t. It is thus natural to define $L(\infty) = N(\beta)$. Similarly, set $L(-\infty) = N(\alpha)$.

Given $I \subseteq [-\infty, \infty]$, set $L(I) = L_{\alpha\beta}(I) = \bigcup_{t \in I} L(t)$. We define

$$\Lambda = \Lambda_{\alpha\beta} = L_{\alpha\beta}([-\infty, \infty]).$$

Suppose $\gamma \in X \setminus (L(-\infty) \cup L(\infty))$. Setting $A = i(\alpha, \gamma) > 0$ and $B = i(\beta, \gamma)/i(\alpha, \beta) > 0$, we see that $l_t(\gamma) = \max\{Ae^{-t}, Be^t\}$ attains a unique minimum when $t = \frac{1}{2}\log(A/B) = \tau_\gamma = \tau_{\alpha\beta}(\gamma)$. If $\gamma \in L(\infty)$, we set $\tau_{\alpha\beta}(\gamma) = -\infty$ and if $\gamma \in L(\infty)$ set $\tau_{\alpha\beta}(\gamma) = \infty$. In the "exceptional" case where $\gamma \in L(-\infty) \cap L(\infty)$ (so that $d(\alpha, \beta) = 2$), we leave $\tau_{\alpha\beta}(\gamma)$ undefined.

We write $T(\gamma) = T_{\alpha\beta}(\gamma) = \{t \in [-\infty, \infty] \mid \gamma \in L(t)\}$. Thus $T(\gamma) \neq \emptyset$ if and only if $\lambda \in \Lambda$. In this case $T(\gamma)$ is a closed interval centred on $\tau(\gamma)$ (where we deem $[-\infty, t]$ and $[t, \infty]$ to be centred on $-\infty$ and ∞ respectively, for all $t \in \mathbb{R}$).

Given $I \subseteq [-\infty, \infty]$ write $B(I) = B_{\alpha\beta}(I) = \{\gamma \in \Lambda \mid \tau(\gamma) \in I\}$. Thus $B(I) \subseteq L(I)$.

If $\gamma, \delta \in \Lambda \setminus (\Lambda(-\infty) \cup \Lambda(\infty))$, we write $\gamma \leq \delta$ or $\gamma \leq_{\alpha\beta} \delta$ to mean that $\tau(\gamma) \leq \tau(\delta)$. This defines a coarse order whose sets of minima and maxima are respectively $\Lambda(-\infty) \setminus \Lambda(\infty)$ and $\Lambda(\infty) \setminus \Lambda(-\infty)$. In the exceptional case where $d(\alpha, \beta) = 2$ we can (rather arbitrarily) deem that each point of $\Lambda(-\infty) \cap \Lambda(\infty)$ lies strictly between these sets of minima and maxima. If $\gamma \leq \delta$, we write:

$$\Lambda[\gamma, \delta] = \Lambda[\delta, \gamma] = \{\epsilon \in \Lambda \mid \gamma \leq \epsilon \leq \delta\}.$$

Note that by deinition, $\Lambda[\gamma, \delta] = B([\tau(\gamma), \tau(\delta)])$, so $\Lambda[\gamma, \delta] \subseteq L([\tau(\gamma), \tau(\delta)])$.

We note:

Lemma 4.6 : *For all* $\gamma \in X$, $\mathrm{diam}\, L(T(\gamma)) \leq 2D$.

Proof : If $\delta \in L(t)$ and $\epsilon \in L(u)$ with $t, u \in T(\gamma)$, then $\gamma \in L(t) \cap L(u)$, so by Lemma 4.3, $d(\delta, \epsilon) \leq d(\delta, \gamma) + d(\gamma, \epsilon) \leq 2D$. ◊

We next want to define centres. Let $\alpha, \beta, \gamma \in X$. Suppose first, that $\min\{d(\alpha, \beta), d(\beta, \gamma), d(\gamma, \alpha)\}$ 2. Let τ_α, τ_β and τ_γ be as defined above, and set $\bar{\alpha} = \alpha_{\tau_\alpha}$ etc. We set

$$\begin{aligned}\Phi_{\alpha\beta\gamma} &= L_{\alpha\beta}(\tau_\alpha) \cap L_{\beta\gamma}(\tau_\beta) \cap L_{\gamma\alpha}(\tau_\gamma) \\ &= \{\delta \in X \mid \max\{i(\bar{\alpha}, \delta), i(\bar{\beta}, \delta), i(\bar{\gamma}, \delta)\} \leq R\} \\ &\subseteq \Lambda_{\alpha\beta} \cap \Lambda_{\beta\gamma} \cap \Lambda_{\gamma\alpha}.\end{aligned}$$

We note:

Lemma 4.7 : $\Phi_{\alpha\beta\gamma} \neq \emptyset$ *and* $\operatorname{diam} \Phi_{\alpha\beta\gamma} \leq D$.

Proof : By Lemma 4.1, $M_{\alpha\beta}(\tau_\alpha) = M_{\beta\alpha}(\tau_\beta) \neq \emptyset$. But $M_{\alpha\beta}(\tau_\alpha) \subseteq L_{\alpha\beta}(\tau_\alpha)$ and by Lemma 4.4, $M_{\alpha\beta}(\tau_\beta) \subseteq L_{\alpha\gamma}(\tau_\beta)$ and $M_{\beta\alpha}(\tau_\beta) \subseteq L_{\beta\gamma}(\tau_\beta)$. Thus $\Phi_{\alpha\beta\gamma} \neq \emptyset$. Finally, since $\Phi_{\alpha\beta\gamma} \subseteq L_{\alpha\beta}(\tau_\alpha)$, it follows by Lemma 4.3 that $\operatorname{diam} \Phi_{\alpha\beta\gamma} \leq D$. ♢

We can thus define $\phi(\alpha, \beta, \gamma)$ to be an arbitrary element of $\Phi_{\alpha\beta\gamma}$, subject to the constraints that $\phi(\alpha, \beta, \gamma) = \phi(\beta, \gamma, \alpha) = \phi(\gamma, \beta, \alpha)$ for all $\alpha, \beta, \gamma \in X$. (Unlike our previous constructions, the above choice is arbitrary. However, choosing a particular point only really serves for notational convenience.)

It remains to define these notions in a few exceptional cases, where two or more curves are close. If $d(\alpha, \beta) \leq 1$, we define $\Lambda_{\alpha\beta} = N(\alpha) \cup N(\beta)$, and set $\gamma \leq_{\alpha\beta} \delta$ for all $\gamma, \delta \in L_{\alpha\beta}$. For all $\alpha, \beta \in X$, we set $\phi(\alpha, \alpha, \beta) = \alpha$. If $\alpha, \beta, \gamma \in X$ and $d(\alpha, \beta) = 1$, $d(\alpha, \gamma) \geq 2$ and $d(\beta, \gamma) \geq 2$ set $\phi(\alpha, \beta, \gamma)$ to be an arbitrary element of $N(\alpha) \cap N(\beta)$. If $\alpha, \beta, \gamma \in X$ and $d(\alpha, \beta) = 1$ and $d(\beta, \gamma) = 1$ and $\beta \neq \gamma$, then set $\phi(\alpha, \beta, \gamma)$ to be an arbitrary element of $N(\alpha) \cap N(\beta) \cap N(\gamma)$. Again, these choices are subject to symmetry.

Lemma 4.8 : *If* $\alpha, \beta \in X$ *and* $d(\alpha, \beta) \leq 2$, *then* $\operatorname{diam} \Lambda_{\alpha\beta} \leq 2D + 2$.

Proof : We can suppose that $d(\alpha, \beta) = 2$. Choose any $\epsilon \in N(\alpha) \cap N(\beta)$. Suppose $\gamma \in \Lambda_{\alpha\beta}$. Thus $\gamma \in L_{\alpha\beta}(t)$ for some $t \in [-\infty, \infty]$. If $t \in \{-\infty, \infty\}$, then $d(\gamma, \{\alpha, \beta\}) \leq 1$. If not, choose any $\delta \in M_{\alpha\beta}(t)$. Now $m_t(\delta) < \infty$, and $l_t(\gamma) = 0$ (since γ is disjoint from both α and β). Thus, $i(\delta, \epsilon) = 0$. In other words, $d(\delta, \epsilon) \leq 1$. By Lemma 4.3, $d(\gamma, \delta) \leq D$ and so $d(\gamma, \epsilon) \leq D + 1$. The result follows. ♢

Now suppose again that $\alpha, \beta \in X$ with $d(\alpha, \beta) \geq 2$. Given $\gamma \in X$, write $\phi(\gamma) = \phi_{\alpha\beta}(\gamma) = \phi(\alpha, \beta, \gamma)$. We note:

Lemma 4.9 : *If* $\gamma \in \Lambda_{\alpha\beta}$, *then* $\operatorname{diam} \Lambda[\gamma, \phi(\gamma)] \leq 2D$.

Proof : By definition, $\phi(\gamma) \in \Lambda_{\alpha\beta}(\tau(\gamma))$, so $\tau(\gamma) \in T(\phi(\gamma))$. Also, since $\phi(\gamma) \in \Lambda$, we have $\tau(\phi(\gamma)) \in T(\phi(\gamma))$. Thus, $\Lambda[\gamma, \phi(\gamma)] \subseteq L([\tau(\gamma), \tau(\phi(\gamma))]) \subseteq L(T(\phi(\gamma)))$. By Lemma 4.6, $\operatorname{diam} L(T(\phi(\gamma))) \leq 2D$. ♢

(Note that Lemma 4.9 also holds trivially in the case where $d(\alpha, \beta) \leq 1$.)

For the next proof we need to note that if $\gamma, \delta \in X$, with $T(\gamma) \cap T(\delta) \neq \emptyset$, then $L(T(\gamma)) \cap L(T(\delta)) \neq \emptyset$, so Lemma 4.6 shows that $\operatorname{diam} L(T(\gamma) \cup T(\delta)) \leq 4D$. (Indeed the same argument shows that $\operatorname{diam} L(T(\gamma) \cup T(\delta)) \leq 3D$.)

Lemma 4.10 : *If* $\alpha, \beta, \gamma \in X$ *with* $\min\{d(\alpha, \beta), d(\beta, \gamma), d(\gamma, \alpha)\} \leq 2D$ *and* $\delta = \phi(\alpha, \beta, \gamma)$, *then* $\operatorname{HausDist}(\Lambda_{\alpha\beta}[\alpha, \delta], \Lambda_{\alpha\gamma}[\alpha, \delta]) \leq 4D$.

Proof : Let $\epsilon \in \Lambda_{\alpha\beta}[\alpha, \delta]$, and let $t = \tau_{\alpha\beta}(\epsilon)$. Thus $t \leq \tau_{\alpha\beta}(\delta)$ and $t \in T(\epsilon)$. Writing $\tau_\alpha = \tau_{\alpha\beta}(\gamma)$ as above, we have, by definition, $\delta \in L_{\alpha\beta}(\tau_\alpha)$, and so $\tau_\alpha \in T_{\alpha\beta}(\delta)$.
Case (1) : $t \geq \tau_\alpha$.

Now $t \in [\tau_\alpha, \tau_{\alpha\beta}(\delta)] \subseteq \tau_{\alpha\beta}(\delta)$. Thus $T_{\alpha\beta}(\epsilon) \cap T_{\alpha\beta}(\delta) \neq \emptyset$, so by the above observation (Lemma 4.6) $d(\epsilon, \delta) \leq 4D$. But $\delta \in \Lambda_{\alpha\gamma}[\alpha, \delta]$ and so we're done.

Case (2) : $t \leq \tau_\alpha$.

By Lemma 4.1, $M_{\alpha\beta}(t) \neq \emptyset$, so we can choose some $\zeta \in M_{\alpha\beta}(t)$. Thus, $\epsilon, \zeta \in L_{\alpha\beta}(t)$, so by Lemma 4.3, $d(\epsilon, \zeta) \leq D$. By Lemma 4.4, $M_{\alpha\beta}(t) \subseteq L_{\alpha\gamma}(t)$, and so $\zeta \in L_{\alpha\gamma}(t)$. Let $u = \tau_{\alpha\gamma}(\zeta)$. We have $t, u \in T_{\alpha\gamma}(\zeta)$.

Case (2a) : $u \leq \tau_{\alpha\gamma}(\delta)$.

In other words, $\tau_{\alpha\gamma}(\zeta) \subseteq \tau_{\alpha\gamma}(\delta)$, so by definition, $\zeta \in \Lambda_{\alpha\beta}[\alpha, \delta]$ and we're done.

Case (2b) : $u \geq \tau_{\alpha\gamma}(\delta)$.

Recall that $t \leq \tau_\alpha$, that $t, u \in T_{\alpha\gamma}(\zeta)$ and that $\tau_\alpha, \tau_{\alpha\gamma}(\delta) \in T_{\alpha\gamma}(\delta)$. It follows that $T_{\alpha\gamma}(\zeta) \cap T_{\alpha\gamma}(\delta) \neq \emptyset$. Thus, by Lemma 4.2, $d(\zeta, \delta) \leq D$, and so $d(\epsilon, \delta) \leq d(\epsilon, \zeta) + d(\zeta, \delta) \leq 2D$. But $\delta \in \Lambda_{\alpha\gamma}[\alpha, \delta]$ and we're done. ◊

In fact, Lemma 4.10 remains true if one or more of the distances is at most 1. In view of Lemma 4.8, the only non-trivial case that need to be checked is where $d(\beta, \gamma) = 1$, $d(\alpha, \beta) \geq 3$ and $d(\alpha, \gamma) \geq 2$. In this case, from our definition of centre, $\tau_{\alpha\beta}(\delta) = \tau_{\alpha\gamma}(\delta) = \infty$ and so $\Lambda_{\alpha\beta}[\alpha, \delta] = \Lambda_{\alpha\beta}$ and $\Lambda_{\alpha\gamma}[\alpha, \delta] = \Lambda_{\alpha\gamma}$. We thus need to check that $\mathrm{HausDist}(\Lambda_{\alpha\beta}, \Lambda_{\alpha\gamma}) \geq 4D$. This is not hard to see this directly, though we omit the argument as it follows also from Lemma 4.12 below.

In order to verify the hypotheses of Proposition 3.1, it remains to check the following:

Proposition 4.11 : *Suppose that $\alpha, \beta, \gamma^0, \gamma^1 \in X$ with $d(\gamma^0, \gamma^1) \leq 1$, then*

$$\mathrm{diam}\, \Lambda_{\alpha\beta}[\phi_{\alpha\beta}(\gamma^0), \phi_{\alpha\beta}(\gamma^1)] \leq 18D.$$

We first consider the case where $d(\{\alpha, \beta\}, \{\gamma^1, \gamma^0\}) \geq 2$, and discuss the remaining cases at the end.

We begin with the following observation:

Lemma 4.12 : *Suppose that $\gamma^0, \gamma^1 \in WMX$ are both supported on a multicurve $\gamma \in MX$. Suppose that $\alpha \in WMX$ with $d(\alpha, \gamma) \geq 2$. Then, for all $t \in \mathbb{R}$, $\mathrm{diam}(L_{\alpha\gamma^0}(t) \cup L_{\alpha\gamma^1}(t)) \leq 4D$.*

Proof : We claim that $L_{\alpha\gamma^0}(t, 2R) \cap L_{\alpha\gamma^1}(t, 2R) \neq \emptyset$. By Lemma 4.2, the diameter of each of these sets is at most $2R(2R) + 2 = 4R^2 + 2 < 2D$, from which the result will follow.

We have observed that the set $L_{\alpha\gamma^0}(t, 2R)$ depends only on the projective class of γ^0. We can thus assume that $i(\alpha_t, \gamma^0) = 1$ (where $\alpha_t = e^{-t}\alpha$, as in the definition of $L_{\alpha\gamma^0}(t, 2R)$). Similarly, we can assume that $i(\alpha_t, \gamma^1) = 1$. Let $\beta = \frac{1}{2}(\gamma^0 + \gamma^1)$. Thus, $i(\alpha_t, \beta) = 1$ (so that $t = \kappa(\alpha, \beta)$). By Lemma 4.1, there is some $\delta \in M(\alpha_t, \beta, R) \subseteq L(\alpha_t, \beta, R)$. In other words, $l_{\alpha_t\beta}(\delta) \leq R$ and so $i(\alpha_t, \delta) \leq R$ and $i(\beta, \delta) \leq R$. Since $\gamma^0 + \gamma^1 = 2\beta$, we have $i(\gamma^0, \delta) \leq 2i(\beta, \delta) \leq 2R$, and $i(\gamma^1, \delta) \leq 2R$. Thus, $\delta \in L(\alpha_t, \gamma^0, 2R) \cap L(\alpha_t, \gamma^1, 2R) = L_{\alpha\gamma^0}(2R) \cap L_{\alpha\gamma^1}(2R)$ as claimed. ◊

Note that it follows that $\mathrm{HausDist}(\Lambda_{\alpha\gamma^0}, \Lambda_{\alpha\gamma^1}) \leq 4D$. Applying Lemma 4.8 and then 4.12 inductively to a geodesic sequence in X connecting α to β, it follows that

$$\operatorname{diam} \Lambda_{\alpha\beta} \leq (n-2)4D + (2D+2) \leq 4nD$$

where $n = d(\alpha, \beta)$. (However, this argument only a-priori gives an exponential bound on any sensible notion of "length" of $\Lambda_{\alpha\beta}$ — defined taking account of the coarse order of points on $\Lambda_{\alpha\beta}$.)

Now let $\alpha, \beta, \gamma^0, \gamma^1 \in X$ with $d(\alpha, \beta) \geq 2$ and $d(\{\alpha, \beta\}, \{\gamma^0, \gamma^1\}) \geq 2$. Let $\tau_0 = \tau_{\alpha\beta}(\gamma^0)$, $\tau_1 = \tau_{\alpha\beta}(\gamma^1)$, $\tau'_0 = \tau_{\alpha\beta}(\phi_{\alpha\beta}(\gamma^0))$ and $\tau'_1 = \tau_{\alpha\beta}(\phi_{\alpha\beta}(\gamma^1))$. By definition, $\Lambda_{\alpha\beta}[\phi_{\alpha\beta}(\gamma^0), \phi_{\alpha\beta}(\gamma^1)$ $B_{\alpha\beta}([\tau'_0, \tau_1]) \subseteq L_{\alpha\beta}([\tau'_0, \tau'_1])$. But $\tau_0, \tau'_0 \in T(\phi(\gamma^0))$ and so by Lemma 4.6, $\operatorname{diam} L_{\alpha\beta}([\tau_0, \tau'_0]) \leq 2D$. Similarly $\operatorname{diam} L_{\alpha\beta}([\tau_1, \tau'_1]) \leq 2D$. To prove Proposition 4.11 in this case, it is thus enough to show:

Lemma 4.13 : $\operatorname{diam} L_{\alpha\beta}([\tau_0, \tau_1]) \leq 14D$.

Proof : We can assume that $\tau_0 \leq \tau_1$. Suppose that $t_0, t_1 \in [\tau_0, \tau_1]$, and that $\delta^0 \in L_{\alpha\beta}(t_0)$ and $\delta^1 \in L_{\alpha\beta}(t_1)$. We want to show that $d(\delta^0, \delta^1) \leq 14D$. We can assume that $t_0 \leq t_1$. Let $t = \frac{1}{2}(t_0 + t_1)$. Thus $\tau_0 \leq t_0 \leq t \leq t_1 \leq \tau_1$. Given $\mu \in [0,1]$, let $\gamma^\mu = (1-\mu)\gamma^0 + \mu\gamma^1 \in WMX$. Note that

$$\tau_{\alpha\beta}(\gamma^\mu) = \frac{1}{2} \log \left(i(\alpha, \beta) \frac{(1-\mu)i(\alpha, \gamma^0) + \mu i(\alpha, \gamma^1)}{(1-\mu)i(\beta, \gamma^0) + \mu i(\beta, \gamma^1)} \right)$$

varies continuously between τ_0 and τ_1. In particular, we can find $\gamma = \gamma^\mu$ so that $\tau_{\alpha\beta}(\gamma) = t$. Let $u = \tau_{\beta\gamma}(\alpha)$ and $v = \tau_{\gamma\alpha}(\beta)$. Let $u_0 = \kappa(\alpha, \beta) - t_0$, $u_1 = \kappa(\alpha, \beta) - t_1$ and $v' = v - t + t_0 \leq v$. Thus

$$t + u = t_0 + u_0 = t_1 + u_1 = \kappa(\alpha, \beta)$$

$$t + v = t_1 + v' = \kappa(\alpha, \gamma)$$

$$u + v = u_0 + v' = \kappa(\beta, \gamma).$$

By assumption, $\delta^0 \in L_{\alpha\beta}(t_0)$ which equals $L_{\beta\alpha}(u_0)$ since $t_0 + u_0 = \kappa(\alpha, \beta)$. Now $u_0 = \kappa(\alpha, \beta) - t_0 \leq \kappa(\alpha, \beta) - \tau_0 = \tau_{\beta\alpha}(\gamma^0)$, and so by Lemma 4.5, $\operatorname{diam}(L_{\beta\alpha}(u_0) \cup L_{\beta\gamma^0}(u_0)) \leq 2D$. Thus there is some $\epsilon^0 \in L_{\beta\gamma^0}(u_0)$ with $d(\delta^0, \epsilon^0) \leq 2D$.

Now γ^0 and γ^1 are both supported on the multicurve $\{\gamma^0, \gamma^1\}$ with $d(\beta, \{\gamma^0, \gamma^1\}) \geq 2$, and so by Lemma 4.12, $\operatorname{diam}(L_{\beta\gamma^0}(u_0) \cup L_{\beta\gamma}(u_0)) \leq 4D$. Thus there is some $\zeta^0 \in L_{\beta\gamma}(u_0)$ with $d(\epsilon^0, \zeta^0) \leq 4D$.

Since $u_0 + v' = \kappa(\beta, \gamma)$, we have $L_{\beta\gamma}(u_0) = L_{\gamma\beta}(v')$. Moreover since $v' \leq v$, by Lemma 4.5, $\operatorname{diam}(L_{\gamma\beta}(v') \cup L_{\gamma\alpha}(v')) \leq 2D$. Since $t_1 + v' = \kappa(\alpha, \gamma)$, we have $L_{\gamma\alpha}(v') = L_{\alpha\gamma}(t_1)$. Thus, there is some $\zeta^1 \in L_{\alpha\gamma}(t_1)$ with $d(\zeta^0, \zeta^1) \leq 2D$.

We now invert the process, swapping the roles of α and β and of t and u_0. Thus applying Lemma 4.12, we find $\epsilon^1 \in L_{\alpha\gamma^1}(t_1)$ with $d(\zeta^1, \epsilon^1) \leq 4D$. Applying Lemma 4.5 (since $t_1 \leq \tau_1$ and $\delta^1 \in L_{\alpha\beta}(t_1)$) we see that $d(\epsilon^1, \delta^1) \leq 2D$.

We conclude that $d(\delta^0, \delta^1) \leq 2D + 4D + 2D + 4D + 2D = 14D$ as required. ◊

To complete the proof of Proposition 4.11, it remains to deal with a few exceptional cases. First note that if $d(\alpha,\beta) \leq 3$, then by the remark following Lemma 4.12, we have $\operatorname{diam} \Lambda_{\alpha\beta} \leq 6D + 2$ and so there is nothing to prove in this case. We can thus assume that $d(\alpha,\beta) \geq 4$. Note that if $\gamma \in X$ with $d(\alpha,\gamma) \leq 1$ then $\delta = \phi_{\alpha\beta}(\gamma) \in N(\alpha)$. Thus, $\tau_{\alpha\beta}(\delta) = -\infty$ and so $\Lambda_{\alpha\beta}[\alpha,\delta] = L(-\infty) = N(\alpha)$. The only case that needs to be considered is where $d(\alpha,\gamma^0) = d(\gamma^0,\gamma^1) = 1$ and $d(\alpha,\gamma^1) = 2$. Let $\delta^1 = \phi_{\alpha\beta}(\gamma^1)$. By Lemma 4.10, $\operatorname{HausDist}(\Lambda_{\alpha\beta}[\alpha,\delta^1], \Lambda_{\alpha\gamma^1}[\alpha,\delta^1]) \leq 4D$. But $\Lambda_{\alpha\gamma^1}[\alpha,\delta^1] \subseteq \Lambda_{\alpha\gamma^1}$ which has diameter at most $2D+2$ by Lemma 4.8, and $\Lambda_{\alpha\beta}[\phi_{\alpha\beta}(\gamma^0), \phi_{\alpha\beta}(\gamma^1)] = \Lambda_{\alpha\beta}[\alpha,\delta^1]$. It follows that $\operatorname{diam} \Lambda_{\alpha\beta}[\phi_{\alpha\beta}(\gamma^0), \phi_{\alpha\beta}(\gamma^1)] \leq (2D + 2) + 2(2D) = 10D + 2 < 18D$.

This completes the proof of Propositon 4.11. We have thus proven Theorem 0 modulo Lemma 4.1, which we prove in Section 5.

5. Singular euclidean structures.

In this section, we give a proof of Lemma 4.1. We give a geometric argument using singular euclidean structures. There is a similar argument in [MM1], phrased in terms of quadratic differentials. However, their argument involves a limiting process, and so does not a-priori give explicit constants. For this reason we offer an alternative proof below that will give the bound required for Proposition 6.1.

Let S be a compact surface, and let $Q \subseteq S$ be a finite set. By a *singular euclidean* structure on S, we mean a metric locally modelled on the euclidean plane $\mathbb{E}^2 \equiv \mathbb{R}^2$ away from Q, and with cone singularities at points of Q. We shall assume in addition that one can restrict the holonomy to translations of $\mathbb{R}^2$ possibly composed with rotations through π. This implies, in particular, that all cone angles are integral multiples of π. Locally we have preferred "horizontal" and "vertical" coordinates, (ζ,ξ) corresponding the axes of $\mathbb{R}^2$. The euclidean metric, ρ^E, is given infinitesimally by $\sqrt{d\zeta^2 + d\xi^2}$. We also have "horizontal" and "vertical" pseudometrics given by $|d\zeta|$ and $|d\xi|$, as well as an L^∞ metric given by $\max\{d\zeta, d\xi\}$. By integrating, we can define the euclidean, horizontal, vertical and L^∞ lengths of a piecewise smooth curve, γ, in S, which we denote respectively by $l^E(\gamma)$, $l^H(\gamma)$, $l^V(\gamma)$ and $l^I(\gamma)$. Clearly $l^I(\gamma) \leq l^E(\gamma) \leq \sqrt{2} l^I(\gamma)$. We say that γ is "horizontal" ("vertical") if $l^V(\gamma) = 0$ ($l^H(\gamma) = 0$). We say γ is *inefficient* if there is some subpath $\alpha \subseteq \gamma$ and a horizontal or vertical path β such that $\alpha \cup \beta$ is a closed curve bounding a disc containing no point of Q. Otherwise, we say that γ is *efficient*. Efficiency is really a local property. Any (local) geodesic with respect to the euclidean metric ρ^E is efficient. An efficient closed curve will minimise the L^∞ length in its homotopy class.

Now suppose that (Σ,Π) is a non-exceptional surface (i.e. if Σ is a 2-sphere we insist that $|\Pi| \geq 5$ and if Σ is a torus, we insist that $|\Pi| \geq 2$, so that $C(\Sigma,\Pi) > 0$. Suppose that $\alpha,\beta \in MX$ with $d(\alpha,\beta) \geq 2$. We realise α and β so that $|\alpha \cap \beta| = i(\alpha,\beta)$. (For example, take closed geodesics on a complete hyperbolic metric on $S \setminus P$.) Now $\alpha \cup \beta$ has the structure of a connected 4-valent graph, Υ, with vertex set $\alpha \cap \beta$. We partition the edge set into "horizontal" and "vertical" edges lying in α and β respectively. We construct a complex, $S = S(\alpha,\beta)$ by taking a rectangle for each vertex of Υ and gluing them together so that Υ embeds in S as a dual to its 1-skeleton. More precisely, each rectangle has two "horizontal" and two "vertical" sides. Two rectangles associated to the endpoints of a horizontal/vertical edge of Γ are glued together along a common vertical/horizontal edge. We see that S is a closed surface. We write P^+ set of vertices (i.e. 0-cells) of S of degree 2.

If $\alpha \cup \beta$ fills (Σ, Π) (for example, if $d(\alpha, \beta) \geq 3$), then there is a homeomorphism from Σ to S carrying Π to a subset, P, of the vertex set of S, and carrying Υ to the dual graph to the 1-skeleton. Note that $P^+ \subseteq P$.

The special case where α, β do not fill (Σ, Π) (so that $d(\alpha, \beta) = 2$) is a bit more complicated. In this case, Σ and S have a common quotient, which is a "nodal surface" N. More precisely, the quotient map $\pi_S : S \longrightarrow N$ identifies certain 0-cells of S. There is a (possibly disconnected) subsurface, $\Sigma' \subseteq \Sigma$, disjoint from $\alpha \cup \beta$, so that the quotient map $\pi_\Sigma : \Sigma \longrightarrow N$ is obtained by collapsing each component of Σ' to a point, in such a way that $\pi_\Sigma(\Pi \cup \Sigma')$ lies in the image under π_S of the 0-skeleton of S. We write $P = \pi_S^{-1} \pi_\Sigma(\Pi \cup \Sigma')$. Again, $P^+ \subseteq P$. The case of the preceding paragraph corresponds to $\Sigma' = \emptyset$ and $\Sigma = N = S$. In all cases, $|P|$ is linearly bounded in terms of $C(\Sigma, \Pi)$.

Now suppose $\bar{\alpha} = \sum_{i=1}^m \lambda_i \alpha_i \in WMX$ and $\bar{\beta} = \sum_{i=1}^n \mu_i \beta_i \in WMX$, where $\alpha = \{\alpha_1, \ldots, \alpha_m\}$ and $\beta = \{\beta_1, \ldots, \beta_n\}$, and $\lambda_i, \mu_j > 0$ for all i, j. We put a singular euclidean structure on S as follows. If R is a rectangle which corresponds to an intersection of α_i with β_j, then we give R the structure of a euclidean rectangle with vertical side length λ_i and horizontal side length μ_j. We write $S(\bar{\alpha}, \bar{\beta})$ for the resulting surface. Note that $\operatorname{area} S(\bar{\alpha}, \bar{\beta}) = i(\bar{\alpha}, \bar{\beta})$. The cone points of positive and negative curvature are precisely the 0-cells of degree 2 and at least 6 respectively. Clearly, $S(\bar{\alpha}, \bar{\beta})$ comes equipped with local horizontal and vertical coordinates, and we can define the length functions l^E, l^H, l^V and l^I as above.

We assume henceforth that $i(\bar{\alpha}, \bar{\beta}) = 1$ so that $\operatorname{area} S(\bar{\alpha}, \bar{\beta}) = 1$.

Suppose that $\alpha \cup \beta$ fills (Σ, Π), so we can identify Σ with $S(\bar{\alpha}, \bar{\beta})$ and Π with P. Suppose $\gamma \in X$. By definition, γ represents a free homotopy class in $\Sigma \setminus \Pi$. If we allow ourselves the freedom to homotop γ so that it touches Π (but not to homotop across points of Π), then we can realise γ as an efficient path in $S(\bar{\alpha}, \bar{\beta})$, for example a euclidean geodesic. (More presicely, we can take the lift of γ to the completion of the universal cover of $\Sigma \setminus \Pi$, realise it as a biinfinite geodesic, and then project back to Σ.) In this case, it follows from the definition that $l^H(\gamma) = i(\bar{\beta}, \gamma)$ and $l^V(\gamma) = i(\bar{\alpha}, \gamma)$. Thus, $l^I(\gamma) = l(\gamma)$. Note also, that any representative of γ will satisfy $l(\gamma) \leq l^I(\gamma)$.

In the case where $\alpha \cup \beta$ does not fill (Σ, Π), we can realise γ in Σ so that its premage in S of its image in N is a union of efficient paths, say $\gamma_1, \ldots, \gamma_n$. (Either this is a single closed curve, or else each γ_i is a path connecting two points of Π.) In this case, $l(\gamma) = \sum_{i=1}^n l^I(\gamma_i)$.

In order to prove Lemma 4.1, we need to find in $S(\bar{\alpha}, \bar{\beta})$ an essential annulus whose "width" is bounded below in terms of $(\operatorname{genus}(S), |P|)$ and hence also $(\operatorname{genus}(\Sigma), |\Pi|)$. In fact, this is an immediate consequence of Lemma 5.1 of [MM1]. However, their proof uses a limiting argument, and so does not give an explicit constant. We therefore offer another proof below.

The argument will work for a large class of metric, though to avoid technicalities, we assume that S is a singular riemannian surface, i.e. it has a metric ρ which is riemannian away from a finite set of cone points. Let $P \subseteq S$ be a finite set (which need have no relation to the set of cone points). Suppose $A \subseteq S \setminus P$ is a closed embedded annulus, with boundary curves γ_1 and γ_2. We define $\mathrm{width}(A) = \rho(\gamma_1, \gamma_2)$. (Intrinsic distance in A would serve equally well for our purposes.) We say that A is *essential* if it is not null-homotopic in $S \setminus P$.

By a *trivial region* in S we mean an open disc containing at most one point of P.

Lemma 5.1 : *Suppose that ρ is a singular riemannian metric on an orientable closed surface, S, with $\mathrm{area}(S) = 1$. Let $P \subseteq S$ be finite. If S is a 2-sphere, we suppose that $|P| \geq 5$. Suppose there is a homeomorphism $f : [0, \infty) \longrightarrow [0, \infty)$ such that $\mathrm{area}\, D \leq f(\mathrm{length}(\partial D))$ for any trivial region D. Then there is an essential annulus in $S \setminus P$, whose width is bounded below by a positive constant $\eta > 0$, which depends only on $\mathrm{genus}(S)$, $|P|$ and f.*

Proof : First note that it is enough to show that there are two essential curves α, β in $S \setminus P$ with $\rho(\alpha, \beta) > \eta_0$, where $\eta_0 > 0$ depends only on $\mathrm{genus}(S)$, $|P|$ and f. In this case, let $m = 3\,\mathrm{genus}(S) + |P| - 2$, and set $\eta = \eta_0 / m$. Take a distance decreasing map, $g : S \longrightarrow [0, \eta_0]$ with $g(\alpha) = 0$ and $g(\beta) = \eta_0$. For each $i = 0, \ldots, m$, we can suppose, after small perturbation, that $P \cap g^{-1}(i\eta) = \emptyset$. We can thus find an essential simple closed curve, γ_i, in $g^{-1}(i\eta) \subseteq S \setminus P$. Since these curves are disjoint, at least two, say γ_i and γ_j are homotopic, and hence bound an essential annulus, $A \subseteq S \setminus P$. Now $\mathrm{width}(A) = \rho(\gamma_i, \gamma_j) \geq \eta$, and so the result follows.

We need to find the curves α and β. To this end, we define a *spine* to be an embedded graph $\sigma \subseteq S \setminus P$ such that each component of $S \setminus \sigma$ is a trivial region. In other words, $\pi_1(S \setminus P)$ is supported on σ. Note that any spine has length at least $\eta_1 = f^{-1}(1/\max\{1, |P|\})$. This follows since, after deleting certain edges if necessary, we can assume that either each component of $S \setminus P$ meets P, or else $P = \emptyset$ and $S \setminus \sigma$ is connected. Thus, at least one such component has area at least $1/\max\{1, |P|\}$.

Now, let $\eta_0 > 0$ be chosen as described below, and let $\alpha \subseteq S \setminus P$ be an essential simple closed curve whose length is within η_0 of the infimum of all possible such lengths. Let N be the closed η_0-neighbourhood of α. We can find a set, $\{\beta_1, \ldots, \beta_n\}$, of disjoint arcs in $S \setminus P$, each of length at most $2\eta_0$ and each meeting α precisely in its endpoints, such that the image of $\pi_1(N \setminus P)$ in $\pi_1(S \setminus P)$ is supported on $\sigma = \alpha \cup \beta_1 \cup \cdots \cup \beta_n$. Moreover we can suppose that $n \leq \mathrm{genus}(S) + |P|$. Now if all the components of $S \setminus N$ were trivial regions, then $\pi_1(S \setminus P)$ would be supported on σ. In other words, σ is a spine, and so $\mathrm{length}(\alpha) \geq \mathrm{length}(\sigma) - 2n\eta_0 \geq \eta_1 - 2n\eta_0 \geq 100\eta_0$, provided we choose $\eta_0 \leq \eta_1/(100 + 2n)$. If α does not separate S, then at least one of the β_i connects opposite sides of α. If α does separate, then (given our clause regarding $|P|$ in the hypotheses) we must have at least one β_i on one side of α and at least two on the other. In either case, it is easy to find an essential curve (whose image lies in σ) of length at most $\mathrm{length}(\alpha) - 2\eta_0$. This contradicts the choice of α. We deduce that at least one component, R, of $S \setminus N$ is non-trivial. We take β to be any essential curve in $R \setminus P$. Thus $d(\alpha, \beta) \geq \eta_0$ as required. ◇

Corollary 5.2 : *Suppose that S is a singular euclidean surface of unit area, and that $P \subseteq S$ contains all singular points of cone angle strictly less than 2π. Suppose that each* *cone angle is at least* π. *Then* $S \setminus P$ *contains an essential annulus of width at least* η, *where* $\eta > 0$ *depends only on* $\mathrm{genus}(S)$ *and* $|P|$.

Proof : The case where S is a 2-sphere and $|P| = 4$ is easily dealt with explicitly. Otherwise apply Lemma 5.1, taking $f(t) = \frac{1}{2}\pi t^2$, for example. $\diamondsuit$

(Note that any trivial region is a euclidean disc possibly quotiented by a rotation through π. The form $[t \mapsto \frac{1}{2}\pi t^2]$ thus follows from the isoperimetric inequality $[t \mapsto \pi t^2]$ in the plane. Of course, it is fairly easy to obtain some quadratic bound, which would suffice for our purposes.)

To prove Lemma 4.1, we will need the following observation. By a *core curve* of an annulus, we mean any curve in A homotopic to a boundary curve.

Lemma 5.3 : *If* A *is an annulus in a singular riemannian surface, then we can find a core curve* γ *in* A *satisfying* $\mathrm{length}(\gamma)\,\mathrm{width}(A) \leq 1$.

Proof : This is an immediate consequence of the Besicovitch Lemma (or coarea formula) see for example [Gr2]. $\diamondsuit$

In fact it is fairly easy to give some upper bound on $\mathrm{length}(\gamma)\,\mathrm{width}(A)$, which will suffice for our purposes.

Proof of Lemma 4.1 : Let $\alpha, \beta \in WMX$ with $d(\alpha, \beta) \geq 2$. Let $S = S(\alpha, \beta)$ be the singular euclidean surface, and $P \subseteq S$, be as constructed above. Let $A \subseteq S \setminus P$ be an essential annulus of width at least $\eta > 0$ as given by Corollary 5.2. Let $R = \sqrt{2}/\eta$.

First consider the case where $\alpha \cup \beta$ fills (Σ, Π), so that $\Sigma \equiv S$ and $\Pi \equiv P$. Let δ be the homotopy class of a core curve of A. We claim that $m(\delta) \leq R$.

To see this, first note that by Lemma 5.3, we can realise δ so that $l^E(\delta)\eta \leq l^E(\delta)\,\mathrm{width}(A) \leq 1$, and so $l(\delta) \leq l^I(\delta) \leq 1/\eta < R$. Moreover, suppose that $\gamma \in X$. Allowing ourselves to homotop γ so that it touches Π, we can assume that γ is realised as an efficient curve. Thus, $l(\gamma) = l^I(\gamma)$. Now each essential intersection of γ and δ contributes at least $\mathrm{width}(A) \geq \eta$ to $l^E(\gamma)$. Thus $l^E(\gamma) \geq \eta i(\gamma, \delta)$, and so $l(\gamma) = l^I(\gamma) \geq l^E(\gamma)/\sqrt{2} \geq i(\gamma, \delta)/R$. Thus, $\frac{i(\gamma,\delta)}{l(\delta)} \leq R$ (noting that if $l(\delta) = 0$, then $i(\gamma, \delta) = 0$, and we have declared that $0/0 = 0$). Since this applies to any $\gamma \in X$, it follows that $m(\delta) \leq R$ as claimed.

It remains to comment on the case where $\alpha \cup \beta$ does not fill (Σ, Π). Since $A \subseteq S \setminus P$, $\pi_S(A)$ is an annulus in N, which is the injective image of an annulus, $B = \pi_\Sigma^{-1}\pi_S(A)$, in Σ. Let ϵ be a core curve for A, so that $\delta = \pi_\Sigma^{-1}\pi_S(\epsilon)$ is a core curve for B. Again, $m(\delta) \leq R$. To see this, note that any $\gamma \in X$ can be realised in Σ so that π_Σ is the image under π_S of a set $\{\gamma_1, \ldots, \gamma_n\}$ of efficient paths, with $l(\gamma) = \sum_{i=1}^n l^I(\gamma_i)$. Moreover, each essential intersection of γ with δ gives us a crossing of some γ_i with B. The proof thus proceeds as in the previous case. $\diamondsuit$

6. Some refinements.

In this section, we comment on a few further consequences of our arguments.

Firstly, we consider the matter of bounding the hyperbolicity constant. Recall that we have defined the *complexity* of the surface, (Σ, Π), as $C = C(\Sigma, \Pi) = 3\,\mathrm{genus}(\Sigma) + |\Pi| - 4$. (This is also the dimension of the curve complex.) We claim:

Proposition 6.1 : *There is a function $k : \mathbb{N} \longrightarrow \mathbb{N}$ with $k(n) = O(\log n)$, so that the graph of curves, $\mathcal{G}((\Sigma, \Pi))$, is $k(C(\Sigma, \Pi))$-hyperbolic, provided $C(\Sigma, \Pi) > 0$.*

It is not clear what is the best estimate. Indeed, it is conceivable that one may be able to take k constant.

To prove Proposition 6.1, we need to examine various critical points in the argument when we derive new constants from old.

We first note that the constant η of Lemma 5.1 satisfies $1/\eta = O(C_0^{5/2})$, provided that $f(t) = O(t^2)$, where $C_0 = C(S, P)$. To see this, note that $1/f^{-1}(1/x) = O(\sqrt{x})$, and so the constant η_1 of the proof satisfies $1/\eta_1 = 0(|P|^{1/2}) = O(C_0^{1/2})$. Thus $1/\eta_0 = (100 + 2\,\mathrm{genus}(S) + 2|P|)/\eta_1 = O(C_0^{3/2})$, and so $1/\eta = (3\,\mathrm{genus}(S) + |P| - 2)/\eta_0 = O(C_0^{5/2})$ as claimed. In Lemma 5.2, we have f quadratic in t, and so the constant η here is also $O(C_0^{5/2})$. (One can probably improve on $O(C_0^{5/2})$, but one can certainly not do better than $O(C_0)$, so after taking logarithms, this will not be critical to our argument.) In the "generic" case of Lemma 4.1, we have $(\Sigma, \Pi) \equiv (S, P)$. In general we have $C(S, P)$ bounded above by a linear function of $C(\Sigma, \Pi)$. We conclude that the constant, R, of Lemma 4.1 is $O(C^{5/2})$.

As described after Lemma 4.3, we can replace the bound $2Rr + 2$ of Lemma 4.2 by one which is $O(\log(Rr))$. This bound is used in the proofs of Lemmas 4.3 and 4.12, where $r = R$ and $r = 2R$ respectively. We thus get $D = O(\log R) = O(\log C)$. We have thus verified the hypotheses of Proposition 3.1 with $K = O(D) = O(\log C)$.

Now the hyperbolicity constant, k, obtained by Proposition 3.1 is linear in K. To see this, first note that the function $[n \mapsto 3n \log_2 n]$ that controls the area is fixed independently of (Σ, Π). The multiplicative constant of quasigeodecity in Lemma 3.3 is linear in K. Also the proof of Lemma 3.5, the Hausdorff distances between the relevant paths is also $O(K)$. Thus, if we use the homological characterisation of the isoperimetric inequality where mesh is defined in terms of diameters rather than lengths, as in Section 7, it is not hard to see that we get $M = O(K)$. Thus, Propostion 7.1 gives $k = O(M) = O(K) = O(\log C)$ as required.

This proves Proposition 6.1.

Next we give a description of geodesics in $\mathcal{G}$ up to bounded Hausdorff distance in terms of intersection numbers. Since we know that $\mathcal{G}$ is hyperbolic, any two geodesics between $\alpha, \beta \in X$ remain a bounded distance apart. Thus, the choice of geodesic, $[\alpha, \beta]$ will not matter.

Given $\alpha, \beta \in X$ with $d(\alpha, \beta) \geq 2$ and $Q \geq 0$, write

$$G_Q(\alpha, \beta) = \{\gamma \in X \mid i(\alpha, \gamma) i(\beta, \gamma) \leq Q i(\alpha, \beta)\}.$$

Proposition 6.2 : *There is a constant $q \in \mathbb{N}$, depending only on $C(\Sigma, \Pi)$ such that if $Q \geq q$ there is some $h \geq 0$ such that if $\alpha, \beta \in X$ with $d(\alpha, \beta) \geq 2$, then* $\mathrm{HausDist}([\alpha, \beta], G_Q(\alpha, \beta)) \leq h$.

As with Proposition 6.1, the argument will show that, in fact, $q = O(C^{5/2})$ and $k = O(\log C + \log Q)$.

Proof : Set $q = R^2$ where R is constant of Lemma 4.1, which defines $\Lambda_{\alpha\beta}$. By Proposition 3.1, $\mathrm{HausDist}([\alpha, \beta], \Lambda_{\alpha\beta})$ is bounded. We thus need to bound $\mathrm{HausDist}(\Lambda_{\alpha\beta}, G_Q(\alpha, \beta))$.

Suppose $\gamma \in \Lambda_{\alpha\beta}$. Thus $\gamma \in L_{\alpha\beta}(t)$ for some $t \in [-\infty, \infty]$. If $t = -\infty$ or $t = \infty$, then $i(\alpha, \gamma) = 0$ or $i(\beta, \gamma) = 0$, and so $\gamma \in G_Q(\alpha, \beta)$. If $t \in \mathbb{R}$, let $u = \kappa(\alpha, \beta) - t$ and let $\alpha_t = e^{-t}\alpha$ and $\beta_u = e^{-u}\beta$. Thus, $i(\alpha_t, \gamma) \leq R$ and $i(\beta_u, \gamma) \leq R$, and so $i(\alpha, \gamma) \leq Re^t$ and $i(\beta, \gamma) \leq Re^u$. Thus $i(\alpha, \gamma)i(\beta, \gamma) \leq R^2 e^{t+u} = R^2 i(\alpha, \beta) \leq Qi(\alpha, \beta)$ and so $\gamma \in G_Q(\alpha, \beta)$.

Conversely, suppose $\gamma \in G_Q(\alpha, \beta)$. Again, if $i(\alpha, \gamma) = 0$ or $i(\beta, \gamma) = 0$, then $\gamma \in \Lambda_{\alpha\beta}$. Otherwise, let $t = \tau_{\alpha\beta}(\gamma) \in \mathbb{R}$, and let $u = \tau_{\beta\alpha}(\gamma) \in \mathbb{R}$. Thus $t + u = \kappa(\alpha, \beta)$. Now $i(\alpha_t, \gamma) = i(\beta_u, \gamma)$. Thus, $i(\alpha_t, \gamma)^2 = i(\alpha_t, \gamma)i(\beta_u, \gamma) = e^{-(t+u)}i(\alpha, \gamma)i(\beta, \gamma) = i(\alpha, \gamma)i(\beta, \gamma)/i(\alpha, \beta) \leq Q$. Thus, $l_t(\gamma) = \max\{i(\alpha_t, \gamma), i(\beta_u, \gamma)\} \leq \sqrt{Q}$, and so $\gamma \in L_{\alpha\beta}(t, \sqrt{Q})$. Applying Lemma 4.2, we see that γ is a bounded distance from an element of $L_{\alpha\beta}(t) \subseteq \Lambda_{\alpha\beta}$ as required. $\diamondsuit$

It is apparent from the proof of Proposition 6.2 that the "approximate order" of points near $[\alpha, \beta]$ is determined by the ratio $i(\beta, \gamma)/i(\alpha, \gamma)$. This can be made more precise using Proposition 6.3 below.

First note that if $t, t' \in \mathbb{R}$, then $L_{\alpha\beta}(t') \subseteq L_{\alpha\beta}(t, Re^{|t-t'|})$. In particular, by Lemma 4.2, $\mathrm{diam}(L_{\alpha\beta}(t) \cup L_{\alpha\beta}(t'))$ is bounded in terms of $|t - t'|$.

Proposition 6.3 : *Given $Q > 0$, there is some $h \geq 0$ such that if $\alpha, \beta, \gamma, \delta \in X$ with $i(\alpha, \delta)i(\beta, \gamma) \leq Qi(\beta, \delta)i(\alpha, \gamma)$, then $d([\alpha, \gamma], [\beta, \delta]) \leq h$.*

Proof : We deal with the case where no two of $\alpha, \beta, \gamma, \delta$ are equal or adjacent, and leave the remaining trivial cases as an exercise.

Write $\tau(\gamma) = \tau_{\alpha\beta}(\gamma)$ and $\tau(\delta) = \tau_{\alpha\beta}(\delta)$. Thus $\tau(\delta) - \tau(\gamma) = \frac{1}{2}(\kappa(\alpha, \delta) + \kappa(\beta, \gamma) - \kappa(\beta, \delta) - \kappa(\alpha, \gamma)) \leq \log Q$. Let $\epsilon = \phi(\alpha, \beta, \delta) \in L_{\alpha\beta}(\tau(\delta))$. Now ϵ is a centre for α, β, δ, and so $d(\epsilon, [\beta, \delta])$ is bounded. Also $\mathrm{HausDist}([\alpha, \gamma], \Lambda_{\alpha\gamma})$ is bounded. We therefore need to bound $d(\epsilon, \Lambda_{\alpha\gamma})$.

If $\tau(\delta) \leq \tau(\gamma)$, then by Lemma 4.5, ϵ is a bounded distance from $\Lambda_{\alpha\gamma}(\tau(\delta)) \subseteq \Lambda_{\alpha\gamma}$ as required.

If $\tau(\delta) \geq \tau(\gamma)$, then $|\tau(\delta) - \tau(\gamma)| \leq \log Q$. Thus, from our earlier observation, $\mathrm{diam}(L_{\alpha\beta}(\tau(\gamma)) \cup L_{\alpha\beta}(\tau(\delta)))$ is bounded. But $\phi(\alpha, \beta, \gamma) \in \Lambda_{\alpha\beta}(\tau(\gamma)) \subseteq \Lambda_{\alpha\gamma}$ and so again the result follows. $\diamondsuit$

Note that since $\mathcal{G}$ is hyperbolic, if $\alpha, \beta, \gamma, \delta \in X$, then the largest two of the three quantities $d(\alpha, \beta) + d(\gamma, \delta)$, $d(\alpha, \gamma) + d(\beta, \delta)$ and $d(\alpha, \delta) + d(\beta, \gamma)$ have difference bounded above in absolute value. Moreover, if $d(\alpha, \delta) + d(\beta, \gamma) \leq d(\beta, \delta) + d(\alpha, \gamma)$, then $d([\alpha, \gamma], [\beta, \delta])$ is bounded. Thus Proposition 6.3 can be thought of as an analogous statement with d replaced by κ.

7. The isoperimetric inequality.

We give a formulation of the isoperimetric inequality for a graph in terms of homology, that is easier to apply than the usual one given in Section 3. This can be readily deduced from results in the literature, though I know of no explicit reference.

Let $\mathcal{G}$ be a connected graph with a positive length assigned to each edge. We write d for the resulting path metric. By a *loop*, γ, in $\mathcal{G}$ we mean a closed path. We write $[\gamma]$ for its homology class in $H_1(\mathcal{G};\mathbb{Z}_2)$. We write $\mathrm{length}(\gamma)$ and $\mathrm{diam}(\gamma)$ for its length and diameter respectively.

Proposition 7.1 : *Suppose there is some $M > 0$ and a function $g : \mathbb{N} \longrightarrow \mathbb{N}$, with $g(n) = o(n^2)$ as $n \to \infty$ such that if γ is any loop in $\mathcal{G}$ then there are loops $\gamma_1, \ldots, \gamma_n$ in $\mathcal{G}$ with $[\gamma] = [\gamma_1] + \cdots + [\gamma_n]$ in $H_1(\mathcal{G};\mathbb{Z}_2)$, with $\mathrm{diam}(\gamma_i) \leq M$ for all i, and with $n \leq g(\mathrm{length}(\gamma))$. Then $\mathcal{G}$ is hyperbolic with hyperbolicity constant depending only on M and g.*

Note that, in fact, the function g determines a constant $\lambda > 0$ such that for any $M > 0$, the hyperbolicity constant is bounded by λM. This can be seen immediately, by scaling the metric by a factor of $1/M$ and applying the result with $M = 1$, and rescaling back.

We remark that there are number of refinements of Proposition 7.1 (with the same proof). For example it is enough that $\limsup_{n\to\infty}(g(n)/n^2)$ be less than some fixed sufficiently small positive number (which we shall not determined here). Moreover, instead of bounding the diameters of each γ_i and controlling n, it would be enough to demand that $\sum_{i=1}^{n}(\mathrm{diam}(\gamma_i))^2$ is bounded above by a subquadratic function of $\mathrm{length}(\gamma)$.

Various proofs that a subquadratic inequality implies hyperbolicity can be found in [Gr1,CDP,O,P,Bo2]. To check that our formulation works, we need to verify that it satisfies the following "rectangle" or "coarea" inequality (cf. [Bo2]). This is a version of the Besicovitch Lemma, and we reinterpret the proof given for riemannian metrics in [Gr2].

Lemma 7.2 : *Suppose that $\gamma = \alpha_1 \cup \alpha_2 \cup \alpha_3 \cup \alpha_4$ is a loop in a graph $\mathcal{G}$ expressed as a concatenation of four paths $\alpha_1, \alpha_2, \alpha_3, \alpha_4$. Let $d_1 = d(\alpha_2, \alpha_4)$ and $d_2 = d(\alpha_1, \alpha_3)$. Suppose we have loops $\gamma_1, \ldots, \gamma_n$ with $[\gamma] = [\gamma_1] + \cdots + [\gamma_n]$. Then $d_1 d_2 \leq 2\sum_{i=1}^{n}(\mathrm{diam}(\gamma_i))^2$.*

Proof : Let R be the rectangle $R = [0, d_1] \times [0, d_2]$ with euclidean metric, ρ. Thus $\partial R = s_1 \cup s_2 \cup s_3 \cup s_4$, where $s_1 = [0, d_1] \times \{0\}$, $s_2 = \{0\} \times [0, d_2]$, $s_3 = [0, d_1] \times \{d_2\}$, $s_4 = \{d_1\} \times [0, d_2]$. Given a vertex x of $\mathcal{G}$, write $\pi_1(x) = \min\{d_1, d(x, \alpha_2)\}$ and $\pi_2(x) = \min\{d_2, d(x, \alpha_1)\}$. Let $\pi(x) = (\pi_1(x), \pi_2(x)) \in R$. We extend π to a continuous map $\pi : \mathcal{G} \longrightarrow R$ by sending each edge to a geodesic segment. Thus, $\rho(\pi(x), \pi(y)) \leq \sqrt{2}d(x, y)$ for all $x, y \in \mathcal{G}$. Also, π maps each α_i into s_i. It follows that $\pi \circ \gamma$ represents the non-trivial element of $H_1(\partial R;\mathbb{Z}_2)$. For each i, $\mathrm{diam}(\pi \circ \gamma_i) \leq \sqrt{2}\,\mathrm{diam}\,\gamma_i$, and so it bounds a singular 2-cell, σ_i, of diameter at most $\sqrt{2}\,\mathrm{diam}\,\gamma_i$. Now the formal sum of the σ_i is a singular 2-cycle representing an element of $H_2(R, \partial R;\mathbb{Z}_2)$. Under the boundary isomorphism to $H_1(\partial R;\mathbb{Z}_2)$, this gets mapped to $[\pi \circ \gamma]$ and is thus non-trivial. It follows that the images of the σ_i cover R. We deduce that $d_1 d_2 = \mathrm{area}(R) \leq 2\sum_{i=1}^{n}(\mathrm{diam}(\gamma_i))^2$, as required. $\diamondsuit$

In the application to Lemma 7.1, the right-hand side is bounded by $2M^2n$, so the result follows.

References.

[BeF] M.Bestvina, K.Fujiwara, *Bounded cohomology of subgroups of the mapping class groups* : Geometry & Topology **6** (2002) 69–89.

[Bo1] B.H.Bowditch, *Notes on Gromov's hyperbolicity criterion for path-metric spaces* : in "Group theory from a geometrical viewpoint" (ed. E.Ghys, A.Haefliger, A.Verjovsky), World Scientific (1991) 64–167.

[Bo2] B.H.Bowditch, *A short proof that a subquadratic isoperimetric inequality impies a linear one* : Michigan Math. J. **42** (1995) 103–107.

[BrCM] J.F.Brock, R.D.Canary, Y.N.Minsky : in preparation.

[CDP] M.Coornaert, T.Delzant, A.Papadopoulos, *Les groupes hyperboliques de Gromov* : Lecture Notes in Mathematics No. 1441, Springer Verlag (1990).

[GhH] E.Ghys, P.de la Harpe (eds.), *Sur les groupes hyperboliques d'après Mikhael Gromov* : Progress in Mathematics No. 83, Birkhäuser (1990).

[Gr1] M.Gromov, *Hyperbolic groups* : in "Essays in group theory", (ed. S.M.Gersten) M.S.R.I. Publications No. 8, Springer (1988) 75–263.

[Gr2] M.Gromov (edited by J.LaFontaine and P.Pansu), *Metric structures for Reimannian and non-Riemannian spaces* : Progress in Mathematics No. 152, Birkhäuser (1998).

[Hare] J.L.Harer, *The virtual cohomological dimension of the mapping class groups of orientable surfaces* : Invent. Math. **84** (1986) 157–176.

[Harv] W.J.Harvey, *Boundary structure of the modular group* : in "Riemann surfaces and related topics: Proceedings of the 1978 Stony Brook Conference" (ed. I.Kra, B.Maskit), Ann. of Math. Stud. No. 97, Princeton University Press (1981) 245–251.

[I] N.V.Ivanov, *Automorphism of complexes of curves and of Teichmüller spaces* : Internat. Math. Res. Notices (1997) 651–666.

[K] E.Klarreich, *The boundary at infinity of the curve complex and the relative Teichmüller space* : preprint, Michigan (1999).

[L] F.Luo, *Automorphisms of the complex of curves* : Topology **39** (2000) 283–298.

[MM1] H.A.Masur, Y.N.Minsky, *Geometry of the complex of curves I: hyperbolicity* : Invent. Math. **138** (1999) 103-149.

[MM2] H.A.Masur, Y.N.Minsky, *Geometry of the complex of curves II: hierarchical structure* : Geom. Funct. Anal. **10** (2000) 902–974.

[O] A.Yu.Ol'shanskii, *Hyperbolicity of groups with subquadratic isoperimetric inequality* : Internatonal J. Alg. and Comp. **1** (1991) 281–289.

[P] P.Papasoglu, *On the sub-quadratic isoperimetric inequality* : in "Geometric group theory (Columbus, OH, 1992)" (ed. R.Charney, M.Davis, M.Shapiro) Ohio State Univ. Math. Res. Inst. Publ. No. 3, de Gruyter (1995) 149–157.

[S] H.Short et al. *Notes on word hyperbolic groups* : in "Group theory from a geometrical viewpoint" (ed. E.Ghys, A.Haefliger, A.Verjovsky), World Scientific (1991) 3–63.