Available online at www.sciencedirect.com
SCIENCE DIRECT®
ELSEVIER
C. R. Acad. Sci. Paris, Ser. I 342 (2006) 595–600
COMPTES RENDUS MATHEMATIQUE
http://france.elsevier.com/direct/CRASS1/

Topology/Geometry

Cut loci in lens manifolds

Sergei Anisov [1]

Department of Mathematics, Utrecht University, P.O. Box 80.010, NL-3508 TA Utrecht, The Netherlands

Received 25 December 2005; accepted 25 January 2006

Available online 9 March 2006

Presented by Étienne Ghys

Abstract

Cut loci in geometric three-manifolds equipped with their natural metrics are an interesting source of spines with small number of vertices. An application of this principle to lens manifolds reveals an interplay between their geometry and topology, combinatorial types of convex hulls of group orbits, and estimates of rotation distance between certain triangulations. ***To cite this article: S. Anisov, C. R. Acad. Sci. Paris, Ser. I 342 (2006).***

Résumé

Cut loci dans les espaces lenticulaires. Les cut loci dans les variétés géométriques de dimension 3 par rapport à leurs métriques naturelles forment une classe remarquable d'épines. Par exemple, ces épines ont un petit nombre de sommets. En appliquant cette idée aux espaces lenticulaires, nous étudions des rapports entre leurs géométrie et topologie, les types combinatoires des enveloppes convexes des $\mathbb{Z}_p$-orbites, et des estimations de distance de rotation entre triangulations spécifiques d'un p-gone. ***Pour citer cet article : S. Anisov, C. R. Acad. Sci. Paris, Ser. I 342 (2006).***

Version française abrégée

Nous étudions les cut loci dans les espaces lenticulaires munis de leur métrique naturelle ; c'est-à-dire, nous considérons l'espace-quotient $L_{p,q} = S^3/\mathbb{Z}_p$, où la métrique riemannienne dans $S^3 \subset \mathbb{C}^2$ est de courbure constante $+1$, et le groupe $\mathbb{Z}_p = \langle g \mid g^p = 1\rangle$ est engendré par la transformation linéaire suivante de $\mathbb{C}^2$: $g(z,w) = (\xi z, \xi^q w)$, où $\operatorname{pgcd}(p,q) = 1$ et $\xi = \mathrm{e}^{2\pi \mathrm{i}/p}$. Rappelons [2,5], que le cut locus $C(x)$ d'un point $x \in M$ dans une variété riemanienne M est l'adhérence de l'ensemble des points $y \in M$ tels que la géodésique la plus courte entre x et y n'est pas unique.

On montre que le diamètre de $L_{p,q}$ est égal à $\pi/2$. Un point $x \in L_{p,q}$ est appelé *spécial* s'il existe un point $y \in L_{p,q}$ tel que $d(x,y) = \operatorname{diam} L_{p,q} = \pi/2$. Tous les autres points de $L_{p,q}$ sont appelés *ordinaires*. On montre que tous les points $x \in L_{p,q}$ sont spéciaux si $q = \pm 1 \bmod p$, mais si $q \neq \pm 1 \bmod p$, alors les points spéciaux forment deux cercles dans $L_{p,q}$ (on obtient donc que dans ce cas les points ordinaires forment un sous-ensemble dense de $L_{p,q}$).

E-mail address: anisov@math.uu.nl (S. Anisov).
[1] The author thanks Netherlands Organisation for Scientific Research (NWO) for financial support.

1631-073X/$ – see front matter
doi:10.1016/j.crma.2006.01.025

Théorème 0.1. (a) *Soit $x \in L_{p,q}$ un point spécial. Alors $C(x)$ est constitué d'un disque de dimension* 2 *et d'un cercle $S(x)$. Le bord du disque est collé à $S(x)$ de façon que l'application de bord est un revêtement de degré p. Si $q = 1$ ou $q = p - 1$, alors la monodromie de $C(x)$ le long de $S(x)$ soit une rotation d'angle $2\pi/p$. Si $1 < q < p - 1$, alors la monodromie de $C(x)$ le long de $S(x)$ est une rotation d'angle $2\pi q/p$ ou $2\pi r/p$ (où $r = q^{-1} \bmod p$) suivant que $z = 0$ ou $w = 0$ pour les préimages de x. Le cercle $S(x)$ est l'ensemble des points situés à distance $\pi/2$ de x.*

(b) *Soit $x \in L_{p,q}$ un point ordinaire. Alors $C(x)$ est un polyèdre simple avec $E(p,q) - 3$ sommets, où $E(p,q) = n_1 + \cdots + n_k$ est la somme des dénominateurs de la fraction continue $p/q = n_1 + 1/(n_2 + 1/(n_3 + \cdots + 1/n_k)\ldots)$. De plus $C(x)$ a le même type combinatoire que l'épine de $L_{p,q}$ construite dans* [6].

Soit $\widetilde{C}(x)$ la préimage de $C(x) \subset L_{p,q}$ dans le revêtement canonique $S^3 \to L_{p,q}$. La démonstration du Théorème 0.1 se fait en étudiant $\widetilde{C}(x) \subset S^3 \subset \mathbb{R}^4$. Notons que $\widetilde{C}(x) = V(x) \cap S^3$, où $V(x)$ est le diagramme de Voronoi [9] de la préimage de $x \in L_{p,q}$ dans S^3. De plus, la décomposition cellulaire de S^3 définie par $\widetilde{C}(x)$ et l'enveloppe convexe de $\mathbb{Z}_p\tilde{x}$ dans $\mathbb{R}^4$ sont duales. Grâce à cette dualité, il suffit d'étudier l'enveloppe convexe de la $\mathbb{Z}_p$-orbite d'un point ordinaire $\tilde{x} \in S^3$.

Les considérations géométriques précédentes ont des conséquences suivantes de nature topologique et combinatoire.

Théorème 0.2. *Les cut loci dans une variété géométrique M^3 (de dimension 3) par rapport à leur métrique naturelle (ou leurs petites perturbations, dans le cas dégénéré) sont des épines simples de M^3 ayant le nombre minimal de sommets parmi les épines connues à ce jour dans au moins les deux cas suivants* :

(a) *$c(M^3) \leqslant 6$, où $c(\cdot)$ désigne la complexité* [6] ;
(b) *M^3 est un espace lenticulaire $L_{p,q}$.*

Théorème 0.3. *Soit Δ une triangulation quelconque d'un p-gone régulier, et $R_{2\pi q/p}\Delta$ l'image de Δ par la rotation d'angle $2\pi q/p$.*

(a) *On a $\rho(\Delta, R_{2\pi q/p}\Delta) \geqslant E(p,q) - 3$ (où la distance $\rho(\Delta_1, \Delta_2)$ est définie dans* [10]) ;
(b) *on peut choisir une triangulation Δ de façon que l'inégalité précédente devienne une égalité.*

1. Introduction

Let M^n be a compact Riemannian manifold. Fix a point $x \in M$. For any unit tangent vector $v \in T_xM$, consider the geodesic $\gamma_v(t) = \exp tv$. Set $s(v) = \sup\{t \mid d(x, \gamma_v(t)) = t\}$, where d is the distance in M. The *cut locus* of x is $C(x) = \{\gamma_v(s(v)) \mid v \in T_xM, \|v\| = 1\}$, see [2], §6.5.4. Recall [2,5] that $M \setminus C(x)$ is an n-dimensional cell. If $n = 3$ and M^3 is equipped with a metric without conjugate points, then $C(x)$ is a two-dimensional polyhedron which is a simple (see below) spine of M^3 [3].

Recall [4] that a 2-polyhedron is called *simple* if the link of any of its points is homeomorphic to a circle or to a circle with a diameter or to a circle with three radii. The points whose links are circles with diameters form *triple lines* of $C(x)$, and the points whose link is a circle with three radii are called *vertices* of $C(x)$. Let P be a simple polyhedron with at least one vertex. By *SP* denote the union of all triple lines and vertices of P. Then P is called a *special* polyhedron if it contains no closed triple lines (without vertices) and every connected component of $P \setminus SP$ is a cell. A spine $P \subset M^3$ is *simple*, respectively, *special* if it is a simple (respectively, special) polyhedron. Every 3-manifold has infinitely many special spines, and can be uniquely (up to homeomorphism) reconstructed from any of its special spines [6,7]. The *complexity* $c(M)$ of a 3-manifold M is the minimal number of vertices of an almost simple spine of M (which is called a *minimal* spine in this case); for the manifolds considered below, minimal spines are always special [6].

Now suppose that M^3 admits a geometric structure, e.g., is elliptic or flat. Though cut loci are defined with respect to arbitrary Riemannian metric, the cut loci corresponding to canonical metrics on M clearly deserve special attention (in particular, as a natural class of spines of M). In this paper, we study geometry, topology, and combinatorial properties of cut loci in lens manifolds equipped with the canonical metric. In Section 3 we compare these cut loci with the spines constructed in [6].

2. Geometry

Let p and q be coprime positive integers, $p > q$. The cyclic group with p elements, $\mathbb{Z}_p$, acts freely on the unit sphere $S^3 \subset \mathbb{C}^2$: the generator of the group takes $(z, w) \in \mathbb{C}^2$ to $(\xi z, \xi^q w)$, where $\xi = e^{2\pi i/p}$. The quotient space is the lens manifold $L_{p,q}$ endowed with the natural metric of constant curvature $+1$. In what follows, we always consider this standard metric on $L_{p,q}$ and this action of $\mathbb{Z}_p$.

Lemma 2.1. *The diameter of $L_{p,q}$ equals $\pi/2$.*

Proof. The distance in S^3 between the $\mathbb{Z}_p$ orbits of $(1, 0)$ and $(0, 1)$ is $\pi/2$, therefore $\operatorname{diam}(L_{p,q}) \geqslant \pi/2$. Now suppose that $\operatorname{diam}(L_{p,q}) = d > \pi/2$. Then there are two points $x, y \in S^3$ such that the d-neighborhood of x contains no points of the orbit $\mathbb{Z}_p y$ of y. Then, if x is the North pole, the whole orbit $\mathbb{Z}_p y$ lies in the Southern hemisphere, and the baricenter of $\mathbb{Z}_p y$ differs from the origin. □

Definition 2.2. A point $x \in L_{p,q}$ is a *special point* if there exists a point $y \in L_{p,q}$ such that the distance $d(x, y)$ equals $\pi/2$; otherwise x is called an *ordinary point*.

Lemma 2.3. *If $q = \pm 1 \bmod p$, then all points $x \in L_{p,q}$ are special points. If $q \neq \pm 1 \bmod p$, then the special points form two circles in $L_{p,q}$ (thus, in this case ordinary points form a dense subset of $L_{p,q}$).*

Proof. The proof of Lemma 2.1 implies that $d(x, y) = \pi/2$ for points $x, y \in L_{p,q}$ if and only if the unit vectors Ox_i and Oy_j are orthogonal for any $i, j \in \{1, \ldots, p\}$, where $x_1, \ldots, x_p$ and $y_1, \ldots, y_p$ are the pre-images of x and y under the natural covering $S^3 \to L_{p,q}$. Consider the spans (in $\mathbb{R}^4$) of $\{Ox_1, \ldots, Ox_p\}$ and of $\{Oy_1, \ldots, Oy_p\}$. These subspaces of $\mathbb{R}^4$ are at least two-dimensional (as $p > 2$) and orthogonal, which means that they are two-dimensional. It can easily be shown that the span of the $\mathbb{Z}_p$ orbit of a point $(z, w) \in \mathbb{C}^2$ has (real) dimension 2 only in the following cases: (1) $z = 0$ or $w = 0$; (2) $q = 1$ or $q = p - 1$.

In the first case, special points form in $L_{p,q}$ two circles that are projections of the two circles $(0, e^{i\varphi})$ and $(e^{i\varphi}, 0)$, $\varphi \in \mathbb{R}/2\pi\mathbb{Z}$, in $S^3 \subset \mathbb{C}^2$, and they are $\pi/2$-equidistants of each other. In the second case, any $\mathbb{Z}_p$ orbit is contained in a Hopf circle $z/w = \text{const}$ (or $\bar{z}/w = \text{const}$) whose $\pi/2$-equidistant is another Hopf circle containing some $\mathbb{Z}_p$ orbits, which implies that $L_{p,1}$ and $L_{p,p-1}$ consist of special points only. □

The distance from any point of $L_{p,q}$ to its first conjugate point along any geodesic equals π (consider the unit sphere S^3, which covers $L_{p,q}$); by Lemma 2.1 this exceeds the diameter of $L_{p,q}$. Thus, for any $x \in L_{p,q}$, the cut locus $C(x) \subset L_{p,q}$ is the set of points $y \in L_{p,q}$ such that the shortest geodesic between x and y is not unique (see [5]).

Theorem 2.4. (a) *Let $x \in L_{p,q}$ be a special point. Then $C(x)$ consists of an open 2-disk and a circle $S(x)$. The boundary of the disk is glued to $S(x)$ so that the gluing map is a p-fold covering. Thus a small neighborhood of $S(x)$ in $C(x)$ is a fibration with fiber a wedge of p intervals. Its monodromy is a rotation by $2\pi/p$ if $q = 1$ or $q = p - 1$, and a rotation by $2\pi q/p$ or $2\pi r/p$ with $r = q^{-1} \bmod p$, depending on whether $z = 0$ or $w = 0$ for the pre-images of x in S^3, if $1 < q < p - 1$. The circle $S(x)$ is the set of points lying at the distance $\pi/2$ from x.*

(b) *Let $x \in L_{p,q}$ be an ordinary point. Then $C(x)$ is a simple polyhedron with $E(p, q) - 3$ vertices, where $E(p, q) = n_1 + \cdots + n_k$ is the sum of the entries of the continued fraction $p/q = n_1 + 1/(n_2 + 1/(n_3 + \cdots + 1/n_k)\ldots)$. Moreover, $C(x)$ has the same combinatorial type as the spine of $L_{p,q}$ constructed in* [6].

Let $\widetilde{C}(x)$ be the pre-image of $C(x)$ under the standard p-fold covering $S^3 \to L_{p,q}$. Then $\widetilde{C}(x)$ is the set of points $y \in S^3$ such that the element of the pre-image $\mathbb{Z}_p\tilde{x}$ of $x \in L_{p,q}$ nearest (in S^3) to y is not unique. Consider also the Voronoi diagram $V(x) \subset \mathbb{R}^4$ of the p elements of the orbit $\mathbb{Z}_p\tilde{x}$. By definition (see [9]), $V(x)$ is the set of points $y \in \mathbb{R}^4$ such that the point of $\mathbb{Z}_p\tilde{x}$ nearest to y is not unique.

Lemma 2.5. *We have $\widetilde{C}(x) = V(x) \cap S^3$.*

Proof. Let $u, v, w \in S^3$. By $d(u, v)$ (respectively, $|uv|$) denote the distance between u and v in S^3 (respectively, in ambient $\mathbb{R}^4$). Then $d(u, w) < d(v, w)$ if and only if $|uw| < |vw|$, and the claim follows. □

It follows from the proof of Lemma 2.3 that the orbit $\mathbb{Z}_p\tilde{x}$ is the set of vertices of a flat regular p-gon. Then $V(x) = \mathbb{R}^2 \times Y_p$, where the $\mathbb{R}^2$ passes through the origin O orthogonally to the plane containing $\mathbb{Z}_p\tilde{x}$, and Y_p is the wedge of p rays Oy_i, $i = 1, \ldots, p$, which emanate from the center of the regular p-gon $\mathbb{Z}_p\tilde{x}$ and pass through the midpoints of its sides. In other words, $V(x)$ consists of p copies of $\mathbb{R}^2 \times \mathbb{R}_{\geqslant 0}$ glued together along the plane $\mathbb{R}^2 \times \{0\}$. Therefore, $V(x) \cap S^3$ consists of p disks glued together along the circle $(\mathbb{R}^2 \times \{0\}) \cap S^3$. By Lemma 2.5, $C(x)$ is the projection of $V(x) \cap S^3$ under the natural covering $S^3 \to L_{p,q}$, and the statements made in part (a) of Theorem 2.4 follow.

The proof of part (b) of Theorem 2.4 begins with the following statement.

Lemma 2.6. *The cell decomposition of S^3 defined by $\widetilde{C}(x)$ is dual to the convex hull of $\mathbb{Z}_p\tilde{x}$.*

Proof. By the proof of Lemma 2.3, the origin O is an interior point of $\operatorname{Conv}\mathbb{Z}_p\tilde{x}$ whenever $\tilde{x} \in S^3$ represents an ordinary point of $L_{p,q}$; thus, support hyperplanes of $\operatorname{Conv}\mathbb{Z}_p\tilde{x}$ do not pass through O.

For any $A \in S^3$, a support hyperplane $\alpha(A)$ of $\operatorname{Conv}\mathbb{Z}_p\tilde{x}$ is given by the following construction: move the plane tangent to S^3 at A parallelly to itself towards the origin until it meets some vertices $A_1, \ldots, A_i$ of $\operatorname{Conv}\mathbb{Z}_p\tilde{x}$. Similarly, a point $A(\alpha) \in S^3$ is assigned to a support hyperplane α; here we use that $O \notin \alpha$.

All points of $S^3 \cap \alpha$ (including $A_1, \ldots, A_i$) are equally distant from $A(\alpha)$, while all other points of $\mathbb{Z}_p\tilde{x}$ are more distant from A than $A_1, \ldots, A_i$. This means that a point $A \in S^3$ belongs to the closure of the cell of $V(x)$ defined by $A_1, \ldots, A_i$ (or of the Voronoi domain of A_1 if $i = 1$) if and only if the support hyperplane $\alpha(A)$ contains the face $A_1 \cdots A_i$ of $\operatorname{Conv}\mathbb{Z}_p\tilde{x}$, and the duality follows. □

By Lemma 2.6, the combinatorial type of the polytope $\operatorname{Conv}\mathbb{Z}_p\tilde{x}$ governs the structure of $\widetilde{C}(x) \subset S^3$ and of $C(x) \subset L_{p,q}$. Note that $\operatorname{Conv}\mathbb{Z}_p u$ is taken to $\operatorname{Conv}\mathbb{Z}_p v$ by the diagonal linear transformation of $\mathbb{C}^2$ that takes u to v; such a transformation always exists provided that u and v correspond to ordinary points of $L_{p,q}$. Consequently, the combinatorial type of $\operatorname{Conv}\mathbb{Z}_p\tilde{x}$ is independent of the choice of ordinary point x.

The proof of part (b) of Theorem 2.4 is reduced to the study of the combinatorial type of $\operatorname{Conv}\mathbb{Z}_p\tilde{x}$; it is quite involved and will be presented in detail elsewhere.

3. Topology

In this section we compare cut loci in lens manifolds with the spines of $L_{p,q}$ constructed in [6], and cut loci in manifolds of complexity $c(M) \leqslant 6$ with their minimal spines (listed in [7]).

The complexity $c(L_{p,q})$ of an arbitrary lens manifold $L_{p,q}$ is not known. However, special spines of $L_{p,q}$ with exactly $E(p,q) - 3$ vertices do exist, and it is conjectured that $c(L_{p,q}) = E(p,q) - 3$, see [6]. No spines of $L_{p,q}$ with less than $E(p,q) - 3$ vertices are known.

Consider a lens manifold $L_{p,q}$, $q \neq \pm 1 \bmod p$. Most of its points are ordinary and, by Theorem 2.4(b), their cut loci are the special spines constructed in [6] topologically; now we get a geometrical construction.

If $q = 1$, then any point $x \in L_{p,q}$ is special, so the cut loci are not simple spines whenever $p > 3$. Nevertheless, there exist small perturbations of the cut loci that are special spines of $L_{p,q}$ with exactly $p - 3$ vertices (note that $E(p,1) = p$).

Theorem 3.1. *Cut loci in geometrical 3-manifolds M^3 with respect to their natural metrics or small perturbations of those cut loci are simple spines of M^3 (a perturbation of $C(x)$, by shifting x or perturbing the metric, is only necessary if $C(x)$ is not a simple spine) that minimize the number of vertices among all currently known spines of M^3 in (at least) the following two cases*:

(a) $c(M^3) \leqslant 6$, *where $c(\cdot)$ denotes the complexity*;
(b) *M^3 is an arbitrary lens manifold $L_{p,q}$.*

Sketch of the proof. According to [7], there are 129 elliptic and 6 flat manifolds of complexity at most 6. Let us start with one of the flat manifolds, the torus $T^3 = \mathbb{R}^3/\mathbb{Z}^3$, which has complexity 6. In the most natural flat metric $\mathrm{d}s^2 = \mathrm{d}\varphi^2 + \mathrm{d}\psi^2 + \mathrm{d}\theta^2$, the cut locus of $(1/2, 1/2, 1/2) \in T^3$ is the union of three 'coordinate' 2-tori given by the equations $\varphi = 0$, $\psi = 0$, and $\theta = 0$. This is not a simple spine of T^3 because the intersections of the tori are not triple lines and the link of the most singular point $(0, 0, 0)$ is the 1-skeleton of the octahedron, not of the tetrahedron. However, in a slightly perturbed flat metric $\mathrm{d}s^2 = \mathrm{d}\varphi^2 + \mathrm{d}\psi^2 + \mathrm{d}\theta^2 + \varepsilon\,\mathrm{d}\varphi\,\mathrm{d}\psi + \varepsilon\,\mathrm{d}\varphi\,\mathrm{d}\theta + \varepsilon\,\mathrm{d}\psi\,\mathrm{d}\theta$, the cut locus of any point x becomes a simple spine of T^3 with 6 vertices, which is a small perturbation of $C(x)$ with respect to the standard flat metric on T^3. The other five flat manifolds can be examined similarly.

The other 129 manifolds of small complexity are covered by S^3 (note that 51 of them are not lenses, see [7], Ch. 9). Lemma 2.5 relates a cut locus in such a manifold to the Voronoi diagram of a $\pi_1(M^3)$-orbit in $\mathbb{R}^4$, and Lemma 2.6 enables us to extract the information that we need from the combinatorial type of the convex hull of that orbit. The computations are facilitated by the package [8].

If $M^3 = L_{p,q}$ is a lens manifold, there are two cases: $q = \pm 1 \bmod p$ and $q \neq \pm 1 \bmod p$. In the first case, a simple spine of $L_{p,q}$ with $p - 3$ vertices can be obtained as a small perturbation of the cut locus described in part (a) of Theorem 2.4. In the second case the result follows from part (b) of Theorem 2.4. □

4. Combinatorics

Theorem 3.1 suggests the following construction. Consider the cut locus $C(x) \subset L_{p,q}$ of a special point $x \in L_{p,q}$. The structure of $C(x)$ is described in Theorem 2.4(a): its 'non-generic' part is S^1 with p 'leaves' attached to it. A small generic perturbation (if $q \neq \pm 1 \bmod p$, it suffices to move x to an ordinary point) would break this S^1 into a gasket of triple lines crossing each other at a number of vertices.

Consider the sections of the perturbed cut locus by small disks transversal to the multiple circle S^1 of the unperturbed $C(x)$. Most of these sections (except those passing through the vertices) are trivalent graphs with $p-2$ internal vertices and p 'external' legs; the legs correspond to the p leaves of $C(x)$ that gather around S^1, and the trivalent vertices are cross-sections of the triple lines. Duality provides a natural bijection between the isotopy classes of these graphs (with indexed legs) and the triangulations of the regular p-gon (with indexed sides).

By genericity, we can assume that each nongeneric section contains at most one vertex of the perturbed cut locus. Therefore, as a transversal disk moves along the gasket, the corresponding triangulation of the p-gon undergoes a sequence of *flips*—simplest transformations of a triangulation, where two triangles ABC and ACD with a common side AC get replaced by the other pair of triangles, ABD and BCD, whose union is the same quadrilateral $ABCD$.

It follows from part (a) of Theorem 2.4 that the monodromy along the gasket takes the triangulation Δ_1 of the p-gon to its triangulation Δ_2 by rotating it by angle $2\pi q/p$: $\Delta_2 = R_{2\pi q/p}\Delta_1$. Recall [10] that the minimal number of flips required to convert Δ_1 to Δ_2 is called the *rotation distance* $\rho(\Delta_1, \Delta_2)$. By the construction described above, $\rho(\Delta, R_{2\pi q/p}\Delta)$ is bounded from below by the minimal possible number of vertices in a simple spine of $L_{p,q}$ that is a small perturbation of the cut locus $C(x)$ of a special point x. Moreover, this lower bound is sharp.

Theorem 4.1.

(a) *Let Δ be a triangulation of a regular p-gon. Then $\rho(\Delta, R_{2\pi q/p}\Delta) \geqslant E(p, q) - 3$;*
(b) *one can choose a triangulation Δ so that the inequality above becomes an equality.*

Sketch of the proof. For the proof of statement (a), see [1].

To prove statement (b), first assume p and q to be coprime. There exist simple spines of $L_{p,q}$ with $E(p, q) - 3$ vertices, see [7]. By part (b) of Theorem 3.1, such spines can be found among small perturbations of cut loci $C(x)$, where x is an arbitrary point of $L_{p,q}$, in particular, a special point. Therefore, $\rho(\Delta, R_{2\pi q/p}\Delta) \leqslant E(p, q) - 3$, where Δ is the p-gon triangulation dual to a 'cross-section of the gasket' as described above. Combined with the inequality of part (a), this gives the equality.

Now suppose that $p = p_1 d$ and $q = q_1 d$, where $d > 1$, and p_1 and q_1 are coprime. Note that $E(p, q) = E(p_1, q_1)$. Furthermore, there exists a triangulation Δ' of the regular p_1-gon such that $\rho(\Delta', R_{2\pi q/p}\Delta') = E(p_1, q_1) = E(p, q)$. Now inscribe the regular p_1-gon in the regular p-gon by selecting every d-th vertex. The rest of the p-gon consists of p_1 congruent $(d + 1)$-gons. Fix an arbitrary triangulation for one of them, repeat it in all other $(d + 1)$-gons (by rotations by $2\pi q/p$ around the center of the p-gon), and triangulate the p_1-gon by Δ'. This gives the required triangulation Δ in the case $\gcd(p, q) > 1$. □

Acknowledgements

I thank E. Ghys, V. Turchin, and the Referee for many valuable remarks.

References

[1] S. Anisov, Geometrical spines of lens manifolds, math.GT/0502326.
[2] M. Berger, A Panoramic View of Riemannian Geometry, Springer-Verlag, Berlin, 2003.
[3] M. Buchner, The structure of the cut locus in dimension less than or equal to six, Compositio Math. 37 (1978) 103–119.
[4] B. Casler, An embedding theorem for connected 3-manifolds with boundary, Proc. Amer. Math. Soc. 16 (1965) 559–566.
[5] Sh. Kobayasi, K. Nomizu, Foundations of Differential Geometry, vol. II, John Wiley & Sons, Inc., New York, 1969.
[6] S. Matveev, Complexity theory of three-dimensional manifolds, Acta Appl. Math. 19 (1990) 101–130.
[7] S. Matveev, Algorithmic Topology and Classification of 3-Manifolds, Springer-Verlag, Berlin, 2003.
[8] Polymake: A tool for the algorithmic treatment of polytopes and polyhedra, www.math.tu-berlin.de/polymake.
[9] F. Preparata, M. Shamos, Computational Geometry. An Introduction, Springer-Verlag, New York, 1985.
[10] D. Sleator, R. Tarjan, W. Thurston, Rotation distance, triangulations, and hyperbolic geometry, J. Amer. Math. Soc. 1 (3) (1988) 647–681.