Minimality of the well-rounded retract

ALEXANDRA PETTET
JUAN SOUTO

We prove that the well-rounded retract of $\mathrm{SO}_n\backslash\mathrm{SL}_n\mathbb{R}$ is a minimal $\mathrm{SL}_n\mathbb{Z}$–invariant spine.

20F65; 12G10, 53C35, 11F75

1 Introduction

In this note we are interested in a certain $\mathrm{SL}_n\,\mathbb{Z}$–invariant deformation retract of the symmetric space $S_n = \mathrm{SO}_n \backslash \mathrm{SL}_n\,\mathbb{R}$. To every element $A \in \mathrm{SL}_n\,\mathbb{R}$ one can associate the lattice $A\mathbb{Z}^n$ in $\mathbb{R}^n$. The element A is *well-rounded* if the set of shortest nonzero vectors of the lattice $A\mathbb{Z}^n$ generate $\mathbb{R}^n$ as a real vector space. This property is invariant under the left action of SO_n and hence there is no ambiguity in saying that an element in S_n is well-rounded. The subset $\mathcal{X}$ of S_n consisting of well-rounded elements is homeomorphic to an $(n(n-1)/2)$–dimensional CW-complex and the right action of $\mathrm{SL}_n\,\mathbb{Z}$ on S_n induces a cocompact action on $\mathcal{X}$. Observe that if $n = 2$ then $\mathcal{X}$ is the dual to the Farey tesselation of $S_2 = \mathbb{H}^2$ and hence homeomorphic to the Bass–Serre tree of $\mathrm{SL}_2\,\mathbb{Z}$. For larger n, the set $\mathcal{X}$ does not have such a simple description, but Lannes and Soulé proved that $\mathcal{X}$ is a deformation retract of S_n and hence contractible (see Soulé [8] for the case of $n = 3$, and Ash [3] for all n, treated in a more general setting). This is why the subset $\mathcal{X}$ is known as the *well-rounded retract* of S_n. Our goal is to show that $\mathcal{X}$ is a minimal $\mathrm{SL}_n\,\mathbb{Z}$–invariant spine of S_n.

Definition 1.1 Let Γ be a group acting discretely on a contractible space S. We say that a closed subset X of S is a *minimal* Γ*–invariant spine* if it is Γ–invariant, contractible and does not properly contain any closed set with these properties.

We prove:

Theorem 1.2 *The well-rounded retract $\mathcal{X}$ is a minimal $\mathrm{SL}_n\,\mathbb{Z}$–invariant spine of the symmetric space $S_n = \mathrm{SO}_n \backslash \mathrm{SL}_n\,\mathbb{R}$.*

Published: 10 June 2008 DOI: 10.2140/gt.2008.12.1543

It has long been known that the well-rounded retract does not contain any smaller dimensional $\mathrm{SL}_n\,\mathbb{Z}$–invariant spines. This follows namely from the fact due to Borel–Serre [5] that the group $\mathrm{SL}_n\,\mathbb{Z}$ has virtual cohomological dimension

$$\mathrm{vcdim}(\mathrm{SL}_n\,\mathbb{Z}) = \frac{n(n-1)}{2} = \dim \mathcal{X}.$$

In order to appreciate the difference between this statement and the claim of Theorem 1.2 it should be observed that the well-rounded retract contains interesting $\mathrm{SL}_n\,\mathbb{Z}$–invariant subsets of dimension $n(n-1)/2$. For instance, recall that an element $A \in \mathrm{SL}_n\,\mathbb{R}$ is well-rounded if the set of shortest nonzero vectors of the lattice $A\mathbb{Z}^n$ generate $\mathbb{R}^n$ as a vector space; equivalently, they generate, as a group, a finite index lattice of $A\mathbb{Z}^n$. We will say that $A \in \mathrm{SL}_n\,\mathbb{R}$ is *extremely well-rounded* if the shortest nonzero vectors of $A\mathbb{Z}^n$ generate the whole lattice $A\mathbb{Z}^n$. The subset $\mathcal{X}'$ of S_n consisting of extremely well-rounded elements is $\mathrm{SL}_n\,\mathbb{Z}$–invariant and has dimension $n(n-1)/2$. While $\mathcal{X}' = \mathcal{X}$ for $n = 2, 3$ and 4, the set $\mathcal{X}'$ is a proper subset of the well-rounded retract for $n \geq 5$. In [7] we proved that $\mathcal{X}'$ is not contractible for $n \geq 5$. This result follows now directly from Theorem 1.2:

Corollary 1.3 [7] *The subset $\mathcal{X}' \subset S_n$ of extremely well-rounded elements is not contractible.* □

In order to prove Theorem 1.2 it suffices to show that whenever $\mathcal{Y}$ is a closed proper $\mathrm{SL}_n\,\mathbb{Z}$–invariant subset of $\mathcal{X}$, there is a torsion-free, finite index subgroup $\Gamma \subset \mathrm{SL}_n\,\mathbb{Z}$ such that the inclusion $\mathcal{Y}/\Gamma \hookrightarrow \mathcal{X}/\Gamma$ is not a homotopy equivalence. We proceed as follows: First we show that there is $A \in \mathcal{X} \setminus \mathcal{Y}$ with the property that there is a torsion-free, finite index subgroup Γ of $\mathrm{SL}_n\,\mathbb{Z}$ and a nontrivial homology class $[\alpha] \in H_{n-1}(\bar{M}_\Gamma, \partial \bar{M}_\Gamma)$ represented by a cycle α which intersects the well-rounded retract exactly at A. Here $\bar{M}_\Gamma$ is the Borel–Serre compactification of the locally symmetric space $M_\Gamma = S_n/\Gamma$ and the homology is with coefficients in the ring $\mathbb{Z}/2\mathbb{Z}$. The class $[\alpha]$ is dual to some class $[\beta] \in H_{n(n-1)/2}(M_\Gamma)$. The fact that the cycle α does not intersect $\mathcal{Y}$ implies that $[\beta]$ is not in the image of $H_*(\mathcal{Y}/\Gamma)$ in $H_*(\mathcal{X}/\Gamma)$. This shows that the inclusion $\mathcal{Y}/\Gamma$ in $\mathcal{X}/\Gamma$ is not a homotopy equivalence.

In [7], we used this strategy to prove Corollary 1.3. In that particular case we faced much simpler technical problems since it was possible to explicitly find a rational maximal flat intersecting $\mathcal{X}$ exactly once, at a point outside of $\mathcal{X}'$. Even in the case $n = 2$, it is easy to see that for a generic point $A \in \mathcal{X}$, every maximal flat through A intersects $\mathcal{X}$ many times. To bypass this problem we give an elementary, though somewhat involved, construction of the cycle α.

The paper is organized as follows: In Section 2 we review some facts about the symmetric space $S_n = \mathrm{SO}_n \backslash \mathrm{SL}_n \mathbb{R}$ and its quotients. In Section 3 we discuss some properties of the well-rounded retract, proving that a generic well-rounded element in S_n has exactly $2n$ shortest vectors. In Section 4 we show that certain homology classes are nontrivial; all the results in this section are surely well known. In Section 5 we derive Theorem 1.2 from a result, Proposition 5.1, proved in Section 6. Proposition 5.1, the key point of this paper, yields nontrivial cycles in $C_{n-1}(\bar{M}_\Gamma, \partial \bar{M}_\Gamma)$ which intersect the well-rounded retract at a single point.

We are grateful to Mladen Bestvina for enduring us while we were working on this project and for suggesting the strategy behind the proof of Theorem 1.2. We thank the referee and Christophe Soulé for telling us about some of the history of the well-rounded retract. The first author is indebted to the University of Southampton and to Juan's mama for their hospitality during the realization of this paper. The second author thanks the Stanford University Department of Mathematics. The results of this paper were obtained while the second author was a member of the Department of Mathematics of the University of Chicago. Finally, both authors thank the scientific organizers of the Fall 2007 programs *Geometric group theory* and *Teichmüller theory and Kleinian groups* at MSRI.

Notation We denote by $\{e_1, \dots, e_n\}$ and $|\cdot|$ the standard basis and Euclidean norm of $\mathbb{R}^n$. Sometimes we will write elements in $\mathbb{R}^n$ as columns and sometimes as rows; we hope that this does not cause any confusion. If U is a linear subspace of $\mathbb{R}^n$, denote by $U^\perp$ its orthogonal complement with respect to the standard Euclidean product. We will use the same symbol to denote both an equivalence class and a representative of the equivalence class. For example, we use the same notation for an element in $\mathrm{SL}_n \mathbb{R}$ and for the corresponding element in the symmetric space $S_n = \mathrm{SO}_n \backslash \mathrm{SL}_n \mathbb{R}$, or in even smaller quotients such as $S_n / \mathrm{SL}_n \mathbb{Z}$. We will however consistently denote the homology class corresponding to a cycle α by $[\alpha]$. All the homology groups considered below have coefficients in the field $\mathbb{Z}/2\mathbb{Z}$ of two elements, although everything remains true with respect to any other commutative ring with unit.

2 The symmetric space $S_n = \mathrm{SO}_n \backslash \mathrm{SL}_n \mathbb{R}$

Up to scaling, the manifold $S_n = \mathrm{SO}_n \backslash \mathrm{SL}_n \mathbb{R}$ admits a unique symmetric metric invariant under the right action of $\mathrm{SL}_n \mathbb{R}$; we shall always assume S_n to be endowed with such a metric. The restriction of the right action of $\mathrm{SL}_n \mathbb{R}$ on S_n to $\mathrm{SL}_n \mathbb{Z}$ is discrete. Moreover, any torsion-free subgroup Γ of $\mathrm{SL}_n \mathbb{Z}$ acts freely and hence the quotient $M_\Gamma = S_n / \Gamma$ is a smooth locally symmetric manifold. It is well known that

$\mathrm{SL}_n\,\mathbb{Z}$ contains torsion-free finite index subgroups. If $\Gamma \subset \mathrm{SL}_n\,\mathbb{Z}$ is any such subgroup, then the manifold M_Γ is not compact, but is homeomorphic to the interior of a compact manifold $\bar{M}_\Gamma$, the so-called Borel–Serre compactification of M_Γ [5].

For every $v \in \mathbb{R}^n$, the *length function*

$$l_v\colon S_n \to \mathbb{R}, \quad l_v(A) = |Av|$$

is well-defined, analytic and convex. In particular we have

$$l_v(A'') \leq \max\{l_v(A), l_v(A')\} \tag{2–1}$$

for all $A, A' \in S_n$ and every A'' in the unique geodesic segment $[A, A']$ joining A and A' in S_n. It should be observed that for every $B \in \mathrm{SL}_n\,\mathbb{R}$ we have $l_v(AB) = l_{Bv}(A)$. Since $\mathrm{SL}_n\,\mathbb{Z}$ acts on the set $\mathbb{Z}^n \setminus \{0\}$, this implies that the function

$$\mathrm{syst}_1\colon S_n \to (0,\infty), \quad \mathrm{syst}_1(A) = \min_{v\in\mathbb{Z}^n, v\neq 0} l_v(A) \tag{2–2}$$

is $\mathrm{SL}_n\,\mathbb{Z}$–invariant. The quantity $\mathrm{syst}_1(A)$ is said to be the *systole*, or *first minimum*, of $A \in S_n$. The elements of the set

$$\mathcal{S}_1(A) = \{v \in \mathbb{Z}^n \mid l_v(A) = \mathrm{syst}_1(A)\} \tag{2–3}$$

are said to be the *systoles* or *shortest vectors* of A.

Ash proved in [2] that the systole function is a topological Morse function (see also Bavard [4] and Akrout [1]). Moreover, the induced function on $S_n/\,\mathrm{SL}_n\,\mathbb{Z}$ is proper by the following theorem:

Mahler's compactness theorem *A closed subset $K \subset S_n/\,\mathrm{SL}_n\,\mathbb{Z}$ is compact if and only if there is $\epsilon > 0$ with $\mathrm{syst}_1(A) \geq \epsilon$ for all $A \in K$.*

We deduce from (2–1) and Mahler's compactness theorem the following important observation:

Lemma 2.1 *Let Γ be a torsion-free subgroup of $\mathrm{SL}_n\,\mathbb{Z}$, N a manifold, and $f, g\colon N \to S_n$ two continuous maps such that for all $\epsilon > 0$ there is a compact set $K_\epsilon \subset N$ with the following property:*

(∗) *For all $x \notin K_\epsilon$ there is $v \in \mathbb{Z}^n \setminus \{0\}$ with $l_v(f(x)), l_v(g(x)) < \epsilon$.*

Then the compositions of f and g with the projection $\pi\colon S_n \to M_\Gamma$ are properly homotopic.

Proof Let $H\colon N\times[0,1]\to S_n$ be the geodesic homotopy from f to g, ie $t\to H_t(x)$ traverses with constant velocity the geodesic segment $[f(x),g(x)]$. We claim that $h=\pi\circ H$ is proper. Let C be a compact subset of $M_\Gamma = S_n/\Gamma$. By Mahler's compactness theorem there is some ϵ positive with $\mathrm{syst}_1(A)\geq\epsilon$ for all $A\in C$. For such an ϵ, let $K_\epsilon\subset N$ be the compact subset provided by $(*)$. Then for $x\notin K_\epsilon$ there is some $v_x\in\mathbb{Z}$, $v_x\neq 0$, with $l_{v_x}(f(x)), l_{v_x}(g(x))<\epsilon$. By (2–1) we have then $l_{v_x}(H_t(x))<\epsilon$ for all $t\in[0,1]$. This implies that $h^{-1}(C)\subset K_\epsilon\times[0,1]$, proving that it is proper. □

We will use Lemma 2.1 several times in the following situation.

Corollary 2.2 *Assume that Γ is a finite index subgroup of $\mathrm{SL}_n\mathbb{Z}$, and that $N\subset\mathrm{SL}_n\mathbb{R}$ projects properly to $M_\Gamma=\mathrm{SO}_n\backslash\mathrm{SL}_n\mathbb{R}/\Gamma$. Then for every $B\in\mathrm{SL}_n\mathbb{R}$ the projections of N and of $BN=\{Bx, x\in N\}$ to M_Γ are properly homotopic.* □

3 The well-rounded retract

In this section we discuss briefly some of the properties of the well-rounded retract. Recall the definition of the systole (2–2) and of the set of systoles (2–3) of a point $A\in S_n$. Let also

$$\Lambda_1(A)=\mathrm{Span}_{\mathbb{R}}(\mathcal{S}_1(A)) \tag{3–1}$$

be the linear subspace of $\mathbb{R}^n$ generated by the set of systoles of A.

Definition 3.1 An element $A\in S_n$ is *well-rounded* if $\Lambda_1(A)=\mathbb{R}^n$. The subset $\mathcal{X}$ of S_n consisting of all well-rounded elements is called the *well-rounded retract*.

As mentioned in the introduction, Soulé [8] and Ash [3] proved that $\mathcal{X}$ is an $\mathrm{SL}_n\mathbb{Z}$–invariant deformation retract. The idea behind this result is simple and beautiful, and so we explain it briefly here:

Theorem 3.2 (Soulé, Ash) *The well-rounded retract $\mathcal{X}$ is a deformation retract of S_n.*

For $k=1,\dots,n$ let $\mathcal{X}_k$ be the set of those $A\in S_n$ for which we have $\dim\Lambda_1(A)\geq k$. We have the following chain of nested $\mathrm{SL}_n\mathbb{Z}$–invariant subspaces:

$$\mathcal{X}=\mathcal{X}_n\subset\mathcal{X}_{n-1}\subset\cdots\subset\mathcal{X}_1=S_n$$

In order to prove Theorem 3.2 it suffices to show that for $k = 1, \dots, n-1$ the space $\mathcal{X}_{k+1}$ is an $\mathrm{SL}_n\,\mathbb{Z}$–equivariant spine of $\mathcal{X}_k$; we construct a retraction. Given $A \in \mathcal{X}_k$ and $\lambda \in \mathbb{R}$, consider the one-parameter family of linear maps:

$$T_A^\lambda \in \mathrm{SL}_n\,\mathbb{R}, \quad T_A^\lambda(v) = \begin{cases} e^{(n-k)\lambda} v & \text{for } v \in A\Lambda_1(A) \\ e^{-k\lambda} v & \text{for } v \in (A\Lambda_1(A))^\perp \end{cases}$$

In other words, for positive λ the map T_A^λ expands the subspace generated by the image of the shortest vectors of A, while contracting the orthogonal complement. Observe that for $U \in \mathrm{SO}_n$ we have $T_{UA}^\lambda UA = UT_A^\lambda A$; hence the point $T_A^\lambda A \in S_n$ depends only on A and not on the choice of representative.

Now $T_A^0 A = A$, and if $A \in \mathcal{X}_k \setminus \mathcal{X}_{k+1}$, there is some λ positive with $T_A^\lambda A \in \mathcal{X}_{k+1}$. For $A \in \mathcal{X}_k$, let $\tau(A) \geq 0$ be maximal such that

$$T_A^\lambda A \in \mathcal{X}_k \setminus \mathcal{X}_{k+1} \quad \text{for all } \lambda \in [0, \tau(A)).$$

By definition $\tau(A) = 0$ for $A \in \mathcal{X}_{k+1}$. The function $A \mapsto \tau(A)$ is continuous on $\mathcal{X}_k$, which implies that

$$[0,1] \times \mathcal{X}_k \to \mathcal{X}_k, \quad (t, A) \mapsto T_A^{t\tau(A)} A$$

is continuous as well. By definition, this homotopy is $\mathrm{SL}_n\,\mathbb{Z}$–equivariant, starts with the identity, and ends with a projection of $\mathcal{X}_k$ to $\mathcal{X}_{k+1}$. This proves that $\mathcal{X}_{k+1}$ is an $\mathrm{SL}_n\,\mathbb{Z}$–equivariant spine of $\mathcal{X}_k$ for $k = 1, \dots, n-1$, concluding the sketch of the proof of Theorem 3.2. □

It is not difficult to prove that $\mathcal{X}_k$ is a codimension $k-1$ semi-algebraic set, ie, that it is given by a locally finite collection of inequalities and (quadratic) algebraic equations. Hence $\mathcal{X}$ is homeomorphic to a CW-complex of dimension $\dim(\mathcal{X}) = \dim S_n - (n-1) = n(n-1)/2$. It is also easy to see that $\mathcal{X}/\Gamma$ is compact. We prove now that a generic point in $\mathcal{X}$ has exactly $2n$ shortest vectors:

Proposition 3.3 *The set of those $A \in \mathcal{X}$ for which there are $v_1, \dots, v_n \in \mathbb{Z}^n$ linearly independent with $\mathcal{S}_1(A) = \{\pm v_1, \dots, \pm v_n\}$ is dense in $\mathcal{X}$.*

In order to prove Proposition 3.3 we will use the following not very surprising but also not completely obvious geometric lemma.

Lemma 3.4 *Assume that $\mathcal{S}$ is a finite subset of the sphere $\mathbb{S}^{n-1}$ in $\mathbb{R}^n$ with the property that $\mathbb{R}^n = \mathrm{Span}_{\mathbb{R}}\,\mathcal{S}$ and assume that if $v \in \mathcal{S}$ then $-v \in \mathcal{S}$ as well. Then there is basis $\mathcal{B}$ of $\mathbb{R}^n$ contained in $\mathcal{S}$ and a linear map $F\colon \mathbb{R}^n \to \mathbb{R}^n$ close to the identity such that for $v \in \mathcal{S}$ we have $|Fv| = |v|$ if $\pm v \in \mathcal{B}$ and $|Fv| > |v|$ otherwise.*

Assuming Lemma 3.4, we prove Proposition 3.3. Given $A \in \mathcal{X}$ choose a representative in $\mathrm{SL}_n \mathbb{R}$, again denoted by A. By definition, the image $A\mathcal{S}_1(A)$ of the set of systoles of A generates $\mathbb{R}^n$ and is contained in the round sphere $\mathbb{S}^{n-1}_{\mathrm{syst}_1(A)}$ of radius $\mathrm{syst}_1(A)$. Let $\mathcal{B} \subset A\mathcal{S}_1(A)$ and $F\colon \mathbb{R}^n \to \mathbb{R}^n$ be the basis and the linear map provided by Lemma 3.4. We set $A^{-1}\mathcal{B} = \{v_1, \dots, v_n\}$ and $A' = (1/\sqrt[n]{\det(F)})FA$. Since we may assume that F is very close to the identity, we have that A' is very close to A, and hence $\mathcal{S}_1(A') \subset \mathcal{S}_1(A)$. It follows now from Lemma 3.4 that $\mathcal{S}_1(A') = \{\pm v_1, \dots, \pm v_n\}$. This concludes the proof of Proposition 3.3. □

We prove now Lemma 3.4:

Proof of Lemma 3.4 We use induction on the number of elements in $\mathcal{S}$. There is nothing to show if $\mathcal{S}$ has $2n$ elements, so assume that we have proved the lemma for all sets with at most $2k \geq 2n$ elements, and that $\mathcal{S}$ has $2(k+1)$ elements. Observe that there is a codimension one linear subspace $U \subset \mathbb{R}^n$ generated by $U \cap \mathcal{S}$ such that there are at least four elements in $\mathcal{S}$ which don't belong to U (recalling that if $v \in \mathcal{S}$, then $-v \in \mathcal{S}$ as well). We first describe a map $F_1\colon \mathbb{R}^n \to \mathbb{R}^n$ which will allow us to apply our inductive hypothesis.

We choose $v \in \mathcal{S}$, $v \notin U$ with minimal angle $\angle(U, v) = \theta \in (0, \pi/2)$. Let V be the codimension one linear subspace containing v and the intersection $(\mathbb{R}v)^{\perp} \cap U$ of the orthogonal complement of $\mathbb{R}v$ and U. The planes U and V have angle θ and divide $\mathbb{R}^n$ into two open sectors, C_1 and C_2 with angle θ, and two also open sectors, C_3 and C_4 with angle $\pi - \theta$. By the minimality of θ, any vector in S which is not in $U \cup \{\pm v\}$ has angle at least θ with U and so is not contained in V. Moreover, for the same reason, we have $\mathcal{S} \cap (C_1 \cup C_2) = \varnothing$, but $\mathcal{S} \cap (C_3 \cup C_4) \neq \varnothing$.

For $\eta > \theta$ with $\eta - \theta$ small we can consider the linear map $F_1\colon \mathbb{R}^n \to \mathbb{R}^n$ which is the identity on U, an isometry when restricted to V, and which opens C_1 and C_2 to angle η. The map F_1 preserves the length of vectors in $U \cup V$, reduces the length of vectors in $C_1 \cup C_2$ and increases the length of vectors in $C_3 \cup C_4$. In particular, F_1 maps $(\mathcal{S} \cap U) \cup \{\pm v\}$ to the subset $(\mathcal{S} \cap U) \cup \{\pm F_1(v)\}$ of $\mathbb{S}^n$ which still generates $\mathbb{R}^n$, and increases the length of the (at least two) remaining vectors in $\mathcal{S}$.

The induction hypothesis now applies to the set $(\mathcal{S} \cap U) \cup \{\pm F_1(v)\}$ of cardinality at most $2k$: there is a basis $\mathcal{B}_1$ of $\mathbb{R}^n$ contained in $(\mathcal{S} \cap U) \cup \{\pm F_1(v)\}$, and a map $F_2\colon \mathbb{R}^n \to \mathbb{R}^n$ which preserves the lengths of the elements of $\mathcal{B}_1$ (and their negatives) and increases the lengths of all other vectors in $(\mathcal{S} \cap U) \cup \{\pm F_1(v)\}$. We require that F_2 be close enough to the identity that the vectors in $F_1(\mathcal{S})$ of length greater than one remain so after applying F_2. Now the basis $\mathcal{B} = F_1^{-1}(\mathcal{B}_1)$ and the map $F = F_2 \circ F_1\colon \mathbb{R}^n \to \mathbb{R}^n$ satisfy the requirements of the lemma for the set $\mathcal{S}$. □

4 A bit of homology

In this section we give elementary proofs of some homological results which are probably well known to experts and nonexperts alike.

As mentioned above, $\mathrm{SL}_n \mathbb{Z}$ contains torsion-free subgroups of finite index, and any such subgroup acts freely and discretely on S_n; as always, we denote the quotient manifold by $M_\Gamma = S_n/\Gamma$ and its Borel–Serre compactification by $\bar{M}_\Gamma$. If $U \subset \bar{M}_\Gamma$ is a regular neighborhood of $\partial \bar{M}_\Gamma$, we have $H_*(\bar{M}_\Gamma, U) \simeq H_*(\bar{M}_\Gamma, \partial \bar{M}_\Gamma)$. In particular, we can consider every properly immersed submanifold of M_Γ as a cycle in $C_*(\bar{M}_\Gamma, \partial \bar{M}_\Gamma)$. Recall that we always consider homology with coefficients in $\mathbb{Z}/2\mathbb{Z}$.

Before stating the main result of this section, we recall that by Lefschetz duality there is a nondegenerate pairing

$$\iota\colon H_{n-1}(\bar{M}_\Gamma, \partial \bar{M}_\Gamma) \times H_{n(n-1)/2}(M_\Gamma) \to \mathbb{Z}/2\mathbb{Z}$$

which can be computed as follows. Given homology classes $[\alpha] \in H_{n-1}(\bar{M}_\Gamma, \partial \bar{M}_\Gamma)$ and $[\beta] \in H_{n(n-1)/2}(M_\Gamma)$, represent them by cycles α and β in general position. Then $\iota([\alpha],[\beta])$ is just the parity of the cardinality of the set $\alpha \cap \beta$. Observe that in order to prove that a cycle $\beta \in C_{n(n-1)/2}(M_\Gamma)$ represents a nontrivial homology class, it suffices to find a cycle $\alpha \in C_{n-1}(\bar{M}_\Gamma, \partial \bar{M}_\Gamma)$ which intersects β transversally at a single point; if this is the case we will say that the two classes $[\alpha]$ and $[\beta]$ are dual to each other. This is the argument used in [7] to prove:

Proposition 4.1 *Let Γ be a finite index torsion-free subgroup of $\mathrm{SL}_n \mathbb{Z}$, Δ the connected component of the identity in the diagonal subgroup of $\mathrm{SL}_n \mathbb{R}$ and Nil the subgroup of $\mathrm{SL}_n \mathbb{R}$ consisting of upper triangular matrices with units in the diagonal. Then the projection of Δ and Nil to M_Γ represent dual, and hence nontrivial, homology classes in $H_{n-1}(\bar{M}_\Gamma, \partial \bar{M}_\Gamma)$ and $H_{n(n-1)/2}(M_\Gamma)$, respectively.* □

Proposition 4.1 is surely well known, as is the following slightly more general version.

Corollary 4.2 *Given $B \in \mathrm{GL}_n \mathbb{Q}$ assume that $\Gamma \subset \mathrm{SL}_n \mathbb{Z}$ is a finite index torsion-free subgroup with $B^{-1}\Gamma B \subset \mathrm{SL}_n \mathbb{Z}$, and that Δ and Nil are as in Proposition 4.1. Then the projections of $B\Delta B^{-1}$ and $B\,\mathrm{Nil}\,B^{-1}$ to M_Γ represent dual, and hence nontrivial, homology classes in $H_{n-1}(\bar{M}_\Gamma, \partial \bar{M}_\Gamma)$ and $H_{n(n-1)/2}(M_\Gamma)$, respectively.*

Proof The map $\phi\colon S_n \to S_n$ given by $\phi(X) = XB^{-1}$ induces a diffeomorphism $\Phi\colon M_{B^{-1}\Gamma B} \to M_\Gamma$. By Proposition 4.1 the projections of Δ and Nil represent dual homology classes in $M_{B^{-1}\Gamma B}$. Pushing forward with Φ, we obtain dual cycles ΔB^{-1} and $\mathrm{Nil}\,B^{-1}$. By Corollary 2.2, these cycles are properly homotopic, and hence homologous, to the cycles $B\Delta B^{-1}$ and $B\,\mathrm{Nil}\,B^{-1}$. The claim follows. □

5 Proof of Theorem 1.2

In the next section we will show:

Proposition 5.1 *Assume that $A \in \mathcal{X}$ is such that there are $v_1, \dots, v_n \in \mathbb{Z}^n$ linearly independent with $\mathcal{S}_1(A) = \{\pm v_1, \dots, \pm v_n\}$. Let $B \in \mathrm{GL}_n \mathbb{Q}$ be the matrix with columns $v_1, \dots, v_n$, and let Γ be a finite index torsion-free subgroup of $\mathrm{SL}_n \mathbb{Z} \cap B \,\mathrm{SL}_n \mathbb{Z} B^{-1}$. Then the nontrivial homology class $[B \Delta B^{-1}]$ is represented by a cycle $\alpha \in C_{n-1}(\bar{M}_\Gamma, \partial \bar{M}_\Gamma)$ whose support intersects the well-rounded retract $\mathcal{X}$ only in A.*

Assuming Proposition 5.1, we prove the main theorem:

Theorem 1.2 *The well-rounded retract $\mathcal{X}$ is a minimal $\mathrm{SL}_n \mathbb{Z}$–invariant spine of the symmetric space $S_n = \mathrm{SO}_n \backslash \mathrm{SL}_n \mathbb{R}$.*

Proof Assume that $\mathcal{Y} \subset \mathcal{X}$ is a proper, closed, $\mathrm{SL}_n \mathbb{Z}$–invariant subset of $\mathcal{X}$. As mentioned in the introduction, in order to show that $\mathcal{Y}$ is not contractible, it suffices to prove that for some $\Gamma \subset \mathrm{SL}_n \mathbb{Z}$ the induced map $\mathcal{Y}/\Gamma \to \mathcal{X}/\Gamma$ is not a homotopy equivalence.

By Proposition 3.3 there is $A \in \mathcal{X} \setminus \mathcal{Y}$ and a linearly independent subset $\{v_1, \dots, v_n\} \subset \mathbb{Z}$ with $\mathcal{S}_1(A) = \{\pm v_1, \dots, \pm v_n\}$. Let $B \in \mathrm{GL}_n \mathbb{Q}$ be the matrix with columns $v_1, \dots, v_n$. The subgroups $\mathrm{SL}_n \mathbb{Z}$ and $B \,\mathrm{SL}_n \mathbb{Z} B^{-1}$ are commensurable and hence there is a torsion-free finite index subgroup $\Gamma \subset \mathrm{SL}_n \mathbb{Z} \cap B \,\mathrm{SL}_n \mathbb{Z} B^{-1}$. By Proposition 5.1, the homology class $[B \Delta B^{-1}] \in H_{n-1}(\bar{M}_\Gamma, \partial \bar{M}_\Gamma)$ is represented by a cycle α with $\alpha \cap \mathcal{X} = \{A\}$. On the other hand, the class $[B \Delta B^{-1}]$ is dual to some class $[\beta] \in H_{n(n-1)/2}(M_\Gamma)$ by Corollary 4.2. Since α represents $[B \Delta B^{-1}]$ and intersects $\mathcal{X}$ only at A, we deduce that every cycle contained in $\mathcal{X}/\Gamma$ and representing $[\beta]$ has to contain A in its support. In particular, the map

$$H_{n(n-1)/2}(\mathcal{Y}/\Gamma) \to H_{n(n-1)/2}(\mathcal{X}/\Gamma)$$

is not surjective. This implies that the map $\mathcal{Y}/\Gamma \to \mathcal{X}/\Gamma$ is not a homotopy equivalence. □

6 Flags of systoles

In this section we prove Proposition 5.1. The first step is to construct a certain continuous map

$$\Phi \colon S_n \times [0, \infty) \to S_n \tag{6–1}$$

which essentially pushes points in $S_n \setminus \mathcal{X}$ away from $\mathcal{X}$.

To begin with, recall the definition of the systole $\mathrm{syst}_1(A)$ of $A \in S_n$. We can extend this definition as follows: for $i = 1, \dots, n$, the *i–th systole* of A is given by

$$\mathrm{syst}_i(A) = \inf\{r \mid \dim_{\mathbb{R}}(\mathrm{Span}_{\mathbb{R}}\{v \in \mathbb{Z} \text{ with } |Av| < r\}) \geq i\}. \tag{6–2}$$

In other words, $\mathrm{syst}_i(A)$ is the infimum of those r for which the set of vectors v in $\mathbb{Z}^n$ whose image Av has length less than r generates a subspace of $\mathbb{R}^n$ with dimension at least i. Equivalently,

$$\mathrm{syst}_i(A) = \sup\{r \mid \dim_{\mathbb{R}}(\mathrm{Span}_{\mathbb{R}}\{v \in \mathbb{Z} \text{ with } |Av| < r\}) < i\}. \tag{6–3}$$

The i–th systole coincides with Minkowski's i–th successive minimum of the lattice $A\mathbb{Z}^n$ with respect to the ball B_1 of radius 1 in $\mathbb{R}^n$. See Martinet [6] for more about successive minima.

For $i = 1, \dots, n$, the i–th systole function

$$\mathrm{syst}_i \colon S_n \to (0, \infty)$$

is well-defined and $\mathrm{SL}_n \mathbb{Z}$–equivariant. We claim that it is continuous. In fact, if (A_k) is a sequence in S_n converging to some $A \in S_n$ then for all r the finite sets $\{v \in \mathbb{Z}^n, |A_k v| < r\}$ converge in the Gromov-Hausdorff topology to the (again finite) set $\{v \in \mathbb{Z}^n, |Av| < r\}$. Since $\mathbb{Z}^n$ is discrete, we have that for all sufficiently large k

$$\{v \in \mathbb{Z}^n, |A_k v| < r\} = \{v \in \mathbb{Z}^n, |Av| < r\}.$$

Together with (6–2), this implies that syst_i is lower semi-continuous. Likewise (6–3) and the same argument yield upper semi-continuity.

Lemma 6.1 *The function* $\mathrm{syst}_i \colon S_n \to (0, \infty)$ *is continuous and* $\mathrm{SL}_n \mathbb{Z}$*–equivariant for* $i = 1, \dots, n$. □

Recall now the definition of $\Lambda_1(A)$ given in (3–1). We extend this definition, setting for $i = 1, \dots, n$

$$\Lambda_i(A) = \mathrm{Span}_{\mathbb{R}}(\{v \in \mathbb{Z}^n, |Av| \leq \mathrm{syst}_i(A)\}).$$

In order to avoid treating special cases we set $\Lambda_0(A) = 0$ for all $A \in S_n$. By definition

$$0 \subsetneq \Lambda_1(A) \subset \cdots \subset \Lambda_n = \mathbb{R}^n \tag{6–4}$$

and $\dim_{\mathbb{R}}(\Lambda_i(A)) \geq i$. Observe that for $i < n$ this last inequality is strict if A is well-rounded. In particular, we cannot expect that the subspaces $\Lambda_i(A)$ depend continuously on A. However we have the following weak continuity, which can be proved with essentially the same argument as Lemma 6.1:

Lemma 6.2 *Assume that (A_k) is a sequence in S_n converging to some $A \in S_n$. Then there is k_0 such that for all $k \geq k_0$ and $i \in \{1,\dots,n\}$ there is a unique $\kappa(k,i) \in \{1,\dots,n\}$ with*

- $\Lambda_{\kappa(k,i)}(A_k) = \Lambda_i(A)$, *and*
- *if $\kappa(k,i) \neq n$ then $\Lambda_{\kappa(k,i)+1}(A_k) \neq \Lambda_i(A)$.*

If moreover i' is minimal with $\mathrm{syst}_{i'}(A) = \mathrm{syst}_i(A)$ then

$$\lim_{k\to\infty} \mathrm{syst}_{j_k}(A_k) = \mathrm{syst}_i(A)$$

for all choices of j_k with $\kappa(k,i'-1) < j_k \leq \kappa(k,i)$. □

We use the flag (6–4) to construct the continuous map (6–1). To begin with we consider for $i = 1,\dots,n$ the subspace

$$\Theta_i(A) = (A\Lambda_{i-1}(A))^{\perp} \cap (A\Lambda_i(A)).$$

In more plain language, $\Theta_i(A)$ is the orthogonal complement of the image of $\Lambda_{i-1}(A)$ under A within the image of $\Lambda_i(A)$. We have thus the orthogonal decomposition

$$\mathbb{R}^n = \Theta_1(A) \oplus \cdots \oplus \Theta_n(A)$$

together with the associated orthogonal projections

$$\pi_{\Theta_i(A)}\colon \mathbb{R}^n \to \Theta_i(A).$$

We define now for $x \in \mathbb{R}^n$

$$\Phi_t(A)x = \frac{1}{\sqrt[n]{\prod_{i=1}^n \mathrm{syst}_i(A)^{t \dim_{\mathbb{R}} \Theta_i(A)}}} \sum_{i=1}^n \mathrm{syst}_i(A)^t \pi_{\Theta_i(A)}(Ax). \tag{6–5}$$

The multiplicative factor in (6–5) ensures that $\Phi_t(A) \in \mathrm{SL}_n\mathbb{R}$ for all $A \in \mathrm{SL}_n\mathbb{R}$. Moreover, for all $U \in \mathrm{SO}_n$ we have $\Phi_t(UA) = U\Phi_t(A)$. In particular, we have a well-defined map

$$\Phi_t\colon S_n \times [1,\infty) \to S_n \tag{6–6}$$

It is easy to check that the map (6–6) is $\mathrm{SL}_n\mathbb{Z}$–equivariant, and its continuity follows from Lemma 6.2. Moreover, since $\mathrm{syst}_1(A) \le \mathrm{syst}_i(A)$ for all i, we have for all $x \in \mathbb{R}^n$

$$|\Phi_t(A)x| \ge \left(\frac{\mathrm{syst}_1(A)}{\sqrt[n]{\prod_{i=1}^n \mathrm{syst}_i(A)^{\dim_{\mathbb{R}} \Theta_i(A)}}} \right)^t |Ax| \tag{6–7}$$

with equality if and only if $x \in \Lambda_1(A)$. In particular we see that $\Lambda_1(\Phi_t(A)) = \Lambda_1(A)$ for all $t \ge 0$. Moreover, if $\Lambda_1(A) \ne \mathbb{R}^n$ then the exponentiated quantity in (6–7) is less than 1 and hence

$$\lim_{t\to\infty} \mathrm{syst}_1(\Phi_t(A)) = 0$$

On the other hand, if $\Lambda_1(A) = \mathbb{R}^n$ then $\Phi_t(A) = A$ for all t.

Summing up, we have:

Proposition 6.3 *There is a continuous map $\Phi\colon S_n \times [0,\infty) \to S_n$, $\Phi(A,t) = \Phi_t(A)$, with the following properties:*

- $\Phi_0(\cdot) = \mathrm{Id}$,
- $\Phi_t(A) \in \mathcal{X}$ *if and only if* $A \in \mathcal{X}$, *and*
- *if* $A \notin \mathcal{X}$ *then* $\lim_{t\to\infty} |\Phi_t(A)v| = 0$ *for all* $v \in \Lambda_1(A)$. □

We are now ready to prove Proposition 5.1:

Proposition 5.1 *Assume that $A \in \mathcal{X}$ is such that there are $v_1,\dots,v_n \in \mathbb{Z}^n$ linearly independent with $\mathcal{S}_1(A) = \{\pm v_1,\dots,\pm v_n\}$, let $B \in \mathrm{GL}_n \mathbb{Q}$ be the matrix with columns $v_1,\dots,v_n$ and Γ a finite index torsion-free subgroup in $\mathrm{SL}_n \mathbb{Z} \cap B\, \mathrm{SL}_n \mathbb{Z} B^{-1}$. Then the nontrivial homology class $[B\Delta B^{-1}]$ is represented by a cycle $\alpha \in C_{n-1}(\bar{M}_\Gamma, \partial \bar{M}_\Gamma)$ whose support intersects the well-rounded retract $\mathcal{X}$ only at A.*

Recall that Δ is the connected component of the identity in the diagonal subgroup of $\mathrm{SL}_n \mathbb{R}$.

Proof In order to construct the cycle α we start with the map

$$g_1\colon \Delta \to M_\Gamma, \quad g_1(X) = BXB^{-1}$$

The cycle $g_1(\Delta)$ represents a nontrivial homology class in $H_{n-1}(\bar{M}_\Gamma, \partial \bar{M}_\Gamma)$ by Corollary 4.2. The point A may not belong to the image of $g_1(\Delta)$, but this can be easily corrected by considering the map

$$g_2\colon \Delta \to M_\Gamma, \quad g_2(X) = ABXB^{-1}$$

Corollary 2.2 implies that $g_1(\Delta)$ and $g_2(\Delta)$ are properly homotopic and hence homologous.

Now we have $g_2(\mathrm{Id}) = A$, but it is not clear at all how many other times $g_2(\Delta)$ may intersect $\mathcal{X}$. We correct this problem by constructing a third map g_3 properly homotopic to g_2. Before going further we identify Δ with $\mathbb{R}^{n-1}$ via the following map

$$(a_1, \dots, a_{n-1}) \mapsto \begin{pmatrix} e^{a_1} & 0 & \dots & 0 & 0 \\ 0 & e^{a_2} & \dots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \dots & e^{a_{n-1}} & 0 \\ 0 & 0 & \dots & 0 & e^{-a_1-a_2-\dots-a_{n-1}} \end{pmatrix}$$

A simple computation shows:

Lemma 6.4 *There is some $\epsilon > 0$ such that for all $x \in B_\epsilon \subset \mathbb{R}^{n-1} = \Delta$, $g_2(x) \in \mathcal{X}$ if and only if $x = 0$. If moreover $x \in B_\epsilon$, $x \neq 0$ and $v \in \mathcal{S}_1(g_2(x))$ then we have*

$$\lim_{t\to\infty} l_v(g_2(tx)) = 0. \tag{6–8}$$

Here B_ϵ is the ball of radius ϵ centered at 0 in $\mathbb{R}^{n-1} \simeq \Delta$. □

We can now define the map $g_3\colon \mathbb{R}^{n-1} \to M_\Gamma$. With ϵ as in Lemma 6.4 and Φ the map provided by Proposition 6.3, we set

$$g_3(x) = \begin{cases} g_2(x) & |x| \le \epsilon \\ \Phi_{|x|-\epsilon}(g_2(\epsilon \frac{x}{|x|})) & |x| \ge \epsilon. \end{cases}$$

In other words we extend radially, using the map Φ and the restriction of g_2 to B_ϵ. Since $g_2(x) \notin \mathcal{X}$ for x with $|x| = \epsilon$, we deduce from Proposition 6.3 that $g_3(x) \notin \mathcal{X}$ for all x with $|x| \ge \epsilon$. On the other hand, for $|x| \le \epsilon$ we have $g_3(x) = g_2(x)$. Hence

$$g_3(\mathbb{R}^{n-1}) \cap \mathcal{X} = \{A\}.$$

If $v \in \mathbb{Z}^n$ is a systole for $g_2(x)$ with $|x| = \epsilon$, then we have by (6–8)

$$\lim_{t\to\infty} l_v(g_2(tx)) = 0$$

and by Proposition 6.3

$$\lim_{t\to\infty} l_v(g_3(tx)) = \lim_{t\to\infty} l_v(\Phi_{t-1}(g_2(x))) = 0.$$

Lemma 2.1 implies now that the maps g_2 and g_3 are properly homotopic to each other. Hence the cycle $\alpha = g_3(\Delta)$ represents the nontrivial homology class $[B\Delta B^{-1}] \in H_{n-1}(\bar{M}_\Gamma, \partial \bar{M}_\Gamma)$ and $\alpha \cap \mathcal{X} = \{A\}$. □

References

[1] **H Akrout**, *Singularités topologiques des systoles généralisées*, Topology 42 (2003) 291–308 MR1941437

[2] **A Ash**, *On eutactic forms*, Canad. J. Math. 29 (1977) 1040–1054 MR0491523

[3] **A Ash**, *Small-dimensional classifying spaces for arithmetic subgroups of general linear groups*, Duke Math. J. 51 (1984) 459–468 MR747876

[4] **C Bavard**, *Systole et invariant d'Hermite*, J. Reine Angew. Math. 482 (1997) 93–120 MR1427658

[5] **A Borel**, **J-P Serre**, *Corners and arithmetic groups*, Comment. Math. Helv. 48 (1973) 436–491 MR0387495Avec un appendice: Arrondissement des variétés à coins, par A Douady et L Hérault

[6] **J Martinet**, *Perfect lattices in Euclidean spaces*, Grundlehren series 327, Springer, Berlin (2003) MR1957723

[7] **A Pettet**, **J Souto**, *The spine that was no spine*, to appear in Enseign. Math.

[8] **C Soulé**, *The cohomology of* $\mathrm{SL}_3(\mathbb{Z})$, Topology 17 (1978) 1–22 MR0470141

Department of Mathematics, Stanford University
Stanford, California 94305, USA

Department of Mathematics, University of Michigan
Ann Arbor, Michigan 48109, USA

apettet@math.stanford.edu, jsouto@umich.edu

Proposed: Benson Farb Received: 1 December 2007
Seconded: Walter Neumann, Martin Bridson Revised: 18 April 2008