Comment. Math. Helvetici 59 (1984) 600–634 0010–2571/84/040600–35$01.50 + 0.20/0

Reduction theory using semistability*

DANIEL R. GRAYSON

Introduction

The purpose of this paper is to develop some new techniques for proving theorems about arithmetic groups. Strong theorems about these infinite groups of matrices were proved by Borel and Serre [1973], including finite presentation of the group, finite generation of the cohomology and homology, and a form of generalized Poincaré duality. Their technique is to produce a compact aspherical manifold with boundary whose fundamental group is isomorphic to the arithmetic group; then they show the boundary of the universal cover is homotopy equivalent to the Tits building, which yields precise knowledge about its homology groups.

We will produce the required manifold in a slightly different way, still arriving at the same end results. Whereas Borel and Serre adjoin a boundary to an open manifold, we will instead delete an open neighborhood of infinity. Of course, for such a neighborhood we could simply choose a collar of Borel and Serre's boundary, but this choice is not canonical. Our approach appears to be as canonical and explicit as possible, and is independent of Borel and Serre's results. In fact, it amounts to a reworking of part of reduction theory for quadratic forms, as developed by Minkowski [1896, 1911], Hermite [1905], Siegel [1957], and Borel [1966]. We dispense with the "Siegel sets", and replace them with the study of "semistability" for lattices in Euclidean space. Roughly speaking, this involves the part of reduction theory which provides lower bounds on lengths of vectors in lattices. The other, more classical part of reduction theory is concerned with getting upper bounds on lengths of shortest vectors in lattices, and with combining upper and lower bounds to prove finiteness assertions.

The first advance along these lines was made by Stuhler [1976, 1977]. Serre [1977] and Quillen [see Grayson, 1982] used the notion of semistable vector bundle on an algebraic curve to study $Sl_n(\mathcal{O})$ when $\mathcal{O}$ is a Dedekind domain

* This research was supported by the National Science Foundation through the Institute for Advanced Study and the University of Illinois.

finitely generated over a finite field. Stuhler knew of the analogy between function fields and number fields formulated by Weil [1939], and was inspired to apply it here. He saw that some work of Harder and Narasimhan [1975] on stable vector bundles carries over, and leads to new facts about lattices in Euclidean space.

One way to reformulate Stuhler's result is this. Let L be a lattice in Euclidean space. The real span of any nonzero subgroup M of L is again an inner product space, so it makes sense to speak of the (nonzero) covolume $\operatorname{vol} M$ of M in its span, regardless of the dimension of M. Now plot (for all M) the points $(\dim M, \log \operatorname{vol} M)$ in the (x, y) plane. These points are bounded below, so their convex hull is bounded below by a certain convex polygon. The result is this: each vertex of the polygon is represented by a unique subgroup M; moreover, the subgroups representing the vertices form a chain. This chain is dubbed the Harder–Narasimhan canonical filtration of L because it is analogous to the filtration Harder and Narasimhan obtain for a vector bundle on a projective algebraic curve.

In this paper we use the canonical filtration to undertake a more detailed study of the structure of the space of lattices in a fixed Euclidean space. The idea is to check that more of the work of Serre and Quillen for function fields can be transferred, by analogy, to number fields. Imagine moving L in the space of lattices; motion towards infinity (towards a cusp) can be detected by a decrease in the angles at the vertices of the convex polygon of L, and the canonical filtration itself tells us in which direction we are moving off toward infinity. We make this precise: we get functions which measure the distance from infinity, and use them to determine the open neighborhoods of infinity, the deletion of which gives a manifold. These neighborhoods are small enough so that they do not change the homotopy type of the manifold, and they allow proving that the boundary of the universal cover is homotopy equivalent to the Tits building. This allows recovering all the results of Borel–Serre.

In the first part of the paper we consider $Gl_n\mathcal{O}$, where $\mathcal{O}$ is a ring of integers in a number field. This should be useful to K-theorists as a way of simplifying the proof that the higher K-groups $K_i\mathcal{O}$ are finitely generated. The reader may easily modify the arguments of this paper to apply to $Sl_n\mathcal{O}$.

In section 7 we consider orthogonal groups. This makes use of some ideas of Atiyah and Bott, who discuss semistability for principal G-bundles on a Riemann surface.

In the work of Atiyah and Bott, G is any semisimple Lie group. Thanks to Ofer Gabber, I know how to express the symmetric space of maximal compact subgroups of G as the space of those inner products on the Lie algebra of G which differ from the Killing form by a Cartan involution. I hope to be able to present this case in a future paper.

An interesting incidental consequence of this work is a more "intrinsic" construction of the Borel–Serre "geodesic action" (and their boundary, although we don't pursue that here) for $Gl_n\mathcal{O}$; here intrinsic means without reference to the chosen basis of $\mathcal{O}^n$.

Acknowledgements

I thank Armand Borel, Ofer Gabber, Jean-Pierre Serre, Spencer Bloch, Ranga Rao, Daniel Quillen, Hyman Bass, Steven Zucker, David Kazhdan, Avner Ash, Karen Vogtmann, Ulrich Stuhler, Friedhelm Waldhausen, Gerald Janusz, and Tom Goodwillie heartily for helpful discussions and ideas.

1. Lattices

Let $\mathcal{O}$ be the ring of integers in an algebraic number field F. Let ∞ be an archimedean place of F, and let F_∞ be the completion of F at ∞. The $\mathbb{R}$-algebra F_∞ is either equal to the real numbers $\mathbb{R}$ or is isomorphic to the complex numbers $\mathbb{C}$ (in one of two ways).

If V_∞ is a finite dimensional F_∞-vector space, then an inner product on V_∞ is a positive definite bilinear form $V_\infty \times V_\infty \to F_\infty$, which is required to be symmetric if ∞ is real and to be hermitian if ∞ is complex. When equipped with an inner product, V_∞ is called an inner product space.

We define an $\mathcal{O}$-lattice L to be a projective $\mathcal{O}$-module P of finite rank equipped with an inner product on each of the vector spaces $V_\infty = P \otimes_{\mathcal{O}} F_\infty$. We adopt the notations $V = L \otimes_{\mathbb{Z}} \mathbb{R}$ and $V_\infty = L \otimes_{\mathcal{O}} F_\infty$. We will maintain this notation later, and the addition of subscripts or superscripts will not interfere with the meanings of the letters L, P, and V. [The other logical (but too bulky) choice for the notation would be to mimic the notation for global sections of a sheaf over various open sets, for what we have resembles a sheaf on the topological space $\operatorname{Spec}(\mathcal{O}) \cup \operatorname{Spec}(\mathcal{O} \otimes_{\mathbb{Z}} R)$, the collection of all primes of $\mathcal{O}$, finite, infinite, and zero. Then L would be the "sheaf" on the whole space, P would be the sections over the finite primes, and V_∞ would appear as the "stalk" at ∞.]

If $V_{\mathbb{C}}$ is a hermitian (complex) inner product space, then there is a procedure called *restriction of scalars* which makes it into a symmetric (real) inner product space $V_{\mathbb{R}}$. One simply takes the real part of the inner product and forgets the scalar multiplication by complex, nonreal, numbers. If $\{v, w, \ldots\}$ is an orthonormal basis for $V_{\mathbb{C}}$, then $\{v, iv, w, iw, \ldots\}$ is an orthonormal basis for $V_{\mathbb{R}}$.

There is also a procedure called *restriction of scalars* which makes an $\mathcal{O}$-lattice into a $\mathbb{Z}$-lattice. Notice that $V=\prod V_\infty$. Define an inner product on the real vector space V by

$$\langle v, w\rangle = \sum_{\text{real}} \langle v_\infty, w_\infty\rangle + \sum_{\text{complex}} \operatorname{Re}\langle v_\infty, w_\infty\rangle.$$

Let f_*L denote the $\mathbb{Z}$-lattice obtained by equipping L, regarded as a $\mathbb{Z}$-module, with this inner product (at the unique infinite place of Q).

Let L be an $\mathcal{O}$-lattice, and let P denote the underlying $\mathcal{O}$-module. Any submodule $P_1 \subseteq P$ can be made into an $\mathcal{O}$-lattice by restricting the inner product on each V_∞ to $V_{1,\infty} = P_1 \otimes_0 F_\infty$; call the resulting $\mathcal{O}$-lattice L_1, and write $L_1 \subseteq L$. We will use the notation $L_1 = L \cap P_1$. Assume now that P/P_1 is projective; then we say that L_1 is a sublattice of L. The orthogonal projections $p_\infty: V_\infty \to V_{1,\infty}^\perp$ provide isomorphisms $(P/P_1)\otimes_{\mathcal{O}} F_\infty \to V_{1,\infty}^\perp$ which can be used to make P/P_1 into an $\mathcal{O}$-lattice which we will call L/L_1. We say that L/L_1 is a quotient lattice of L, and that $E: 0 \to L_1 \to L \to L/L_1 \to 0$ is an exact sequence of lattices.

We do not define a notion of "morphism" of $\mathcal{O}$-lattices. The arrows in our diagrams will be simply maps of the underlying $\mathcal{O}$-modules. An isomorphism of $\mathcal{O}$-lattices is an isomorphism of underlying $\mathcal{O}$-modules which preserves the inner products at the infinite places.

Recall that a finitely generated $\mathcal{O}$-module is projective if and only if it has no torsion, and it doesn't matter whether we look for torsion by elements of $\mathcal{O}$ or of $\mathbb{Z}$. From this observation comes the following collection of trivia.

LEMMA 1.1. *Suppose L is an $\mathcal{O}$-lattice.*

(a) *If L_1 is a sublattice of L_2, and L_2 is a sublattice of L, then L_1 is a sublattice of L.*

(b) *If $L_1 \subseteq L_2 \subseteq L$, and L_1 is a sublattice of L, then L_1 is a sublattice of L_2.*

(c) *If $L_1 \subseteq L \subseteq L_2$ are both sublattices of L, then $L_1 \cap L_2$ is a sublattice of L.*

(d) *If $L_1 \subseteq L$, then $L_2 \cap (L_1 \otimes_{\mathcal{O}} F)$ is a sublattice of L, and contains L_1 as a subgroup of finite index. If L_1 is, in addition, a sublattice, then $L_1 = L_2$.*

(e) *If $L_1 \subseteq L$ has volume minimal among volumes of submodules of L of the same rank, then L_1 is a sublattice of L.*

LEMMA 1.2. (a) *Suppose L, L' are $\mathcal{O}$-lattices with the same underlying module P. Then $f_*L = f_*L'$ iff $L = L'$.*

(b) *Suppose $E: 0 \to L_1 \to L \to L_2 \to 0$ is an exact sequence of $\mathcal{O}$-lattices. Then the inner products on L_1 and L_2 are uniquely determined by the inner products on L.*

(c) *Suppose E is a sequence of $\mathcal{O}$-lattices, as above. Then E is exact iff f_*E is.*

Proof. (a) This part is clear, because a hermitian inner product can be recovered from its real part, or indeed from the associated norm.

(b) This is clear from the definitions.

(c) Suppose E is exact. It is clear that f_*L_1 is a sublattice of f_*L. We must check that $f_*L \to f_*L_2$ comes from orthogonal projection. This has two ingredients: the first is that if p_∞ is an orthogonal projection at a complex place, then it remains an orthogonal projection after restriction of scalars to $\mathbb{R}$. The second is that $V = \prod V_\infty$ is an orthogonal sum, and similarly for V_2; the orthogonal sum of orthogonal projections is still an orthogonal projection.

Now suppose f_*E is exact. Then if E isn't exact, we may modify the inner products on L_1 and L_2 to make it exact, obtaining a new exact sequence $E': 0 \to L_1' \to L \to L_2' \to 0$, with the same underlying $\mathcal{O}$-modules. Now f_*E' is exact, so $f_*L_i' = f_*L_i$ for $i = 1, 2$. By part (a), $E' = E$, so is exact. QED

Not all of the usual isomorphism theorems for $\mathcal{O}$-modules go through for $\mathcal{O}$-lattices, so we must be careful at this stage.

LEMMA 1.3. *If L_1 is a sublattice of L, and L_2 is a sublattice of L_1, then L_1/L_2 is a sublattice of L/L_2, and $(L/L_2)/(L_1/L_2) = L/L_2$.*

Proof. By Lemma 1.2, we may as well apply restriction of scalars to everything in sight, achieving the case $\mathcal{O} = \mathbb{Z}$. The first assertion is clear, and the second amounts to the fact that the composite of two orthogonal projections $V \to V/V_2 \to (V/V_2)/(V_1/V_2)$ is again an orthogonal projection. QED

We let $rk\,(L)$ denote the $\mathcal{O}$-module rank of L, and let $\dim(L)$ denote the rank of L as $\mathbb{Z}$-module. Of course, $\dim(L) = rk\,(L)\dim(\mathcal{O})$.

We define the volume of L, $\mathrm{vol}(L)$, to be the covolume of the lattice f_*L inside its inner product space V. This may be computed as $|\det\langle l_i, e_j\rangle|$, where $\{l_i\}$ is a $\mathbb{Z}$-basis of f_*L, and $\{e_j\}$ is an orthonormal basis of V. Thus if $\dim L = 0$, then $\mathrm{vol}\, L = 1$. If $\dim L = 1$, then $\mathrm{vol}\, L$ is the length of a generator of L. If $\dim L = 2$, then $\mathrm{vol}\, L$ is the area of a fundamental parallelogram, and so on. It is worth reiterating that the volumes are not measured with respect to a fixed dimension, and they are always nonzero.

FACT 1.4. If L' is a submodule of finite index in L. then $\mathrm{vol}\, L' = [L : L']\,\mathrm{vol}\, L$.

EXAMPLE 1.5. Take $L = \mathcal{O}$, and for each place ∞ declare $\{1\}$ to be an orthonormal basis of $V_\infty = F_\infty$. This makes $\mathcal{O}$ into an $\mathcal{O}$-lattice in a natural way, and it turns out that $\mathrm{vol}\,\mathcal{O} = 2^{-r_2}\sqrt{|d|}$, where r_2 is the number of complex places of F, and d is the discriminant of $\mathcal{O}$. See Lang [1970, p. 115] for a proof.

EXAMPLE 1.6. Take $L = \mathcal{O}^n$, and give it inner products at each place. Let z_∞ be the matrix of the inner product at ∞; then

$$\operatorname{vol} L = \prod (\det z_\infty)^{e(\infty)/2} (\operatorname{vol} \mathcal{O})^n$$

where $e(\infty) = [F_\infty : R]$.

Proof. If each $z_\infty = 1$, then the direct sum $L = \mathcal{O}^n$ is orthogonal, and the result is immediate.

Choose an orthonormal basis $\{e_{i,\infty}\}$ for $V_\infty = F_\infty^n$, and let y_∞ be the F_∞-automorphism of V_∞ such that the standard basis vectors b_i are given by $b_i = y \cdot e_{i,\infty}$. Let y be the direct product of the y_∞'s it is an R-automorphism of $V = \prod V_\infty$. If we let $L' = y^{-1}L$, then by the first line of the proof, we know $\operatorname{vol}(L') = \operatorname{vol}(\mathcal{O})^n$. Thus $\operatorname{vol} L = |\det y| (\operatorname{vol} \mathcal{O})^n = \prod |\det y_\infty|^{e(\infty)} (\operatorname{vol} \mathcal{O})^n = \prod (\det z_\infty)^{e(\infty)/2} (\operatorname{vol} \mathcal{O})^n$. The last equality comes from $z_\infty = {}^tY_\infty \cdot Y_\infty$, where Y_∞ denotes the matrix of y_∞ with respect to the basis $\{e_{i,\infty}\}$. Notice also that $\det z_\infty > 0$. The middle equality makes use of the formula

$$\det_{\mathbb{R}} h = |\det h|^2$$

which holds when h is an endomorphism of a complex vector space, and $\det_{\mathbb{R}} h$ denotes its determinant when considered as an endomorphism of the underlying real vector space; to prove it, one reduces to the case where the complex dimension is 1 by row and column reduction.

Remark 1.7. For any $\mathcal{O}$-lattice L, we can find its volume as follows. There is a sublattice $L' \subseteq L$ with the same rank as L, and which is free. Fix a basis for it, and for each ∞ let z_∞ be the matrix of the inner product with respect to it. Then

$$\operatorname{vol} L = \prod (\det z_\infty)^{e(\infty)/2} (\operatorname{vol} \mathcal{O})^n / [L : L'].$$

LEMMA 1.8. *If L is an $\mathcal{O}$-lattice and L' is a sublattice of L, then* $\operatorname{vol}(L) = \operatorname{vol}(L') \cdot \operatorname{vol}(L/L')$.

Proof. By restriction of scalars, we may assume that $\mathcal{O} = \mathbb{Z}$. Choose bases $\{l_i\}$ for L' and $\{l_i\} \cup \{m_j\}$ for L. Choose orthonormal bases $\{e_i\}$ for V' and $\{e_i\} \cup \{f_j\}$ for V, and let p be the orthogonal projection onto $V'^{\perp}$. Then, omitting all subscripts for clarity, we have

$$\begin{aligned}\mathrm{vol}\,(L) &= \left|\det\begin{pmatrix}\langle l, e\rangle & \langle l, f\rangle\\ \langle m, e\rangle & \langle m, f\rangle\end{pmatrix}\right|\\ &= \left|\det\begin{pmatrix}\langle l, e\rangle & 0\\ * & \langle pm, f\rangle\end{pmatrix}\right|\\ &= \mathrm{vol}\,(L')\cdot\mathrm{vol}\,(L/L'). \quad \text{QED}\end{aligned}$$

DEFINITION 1.9 [Stuhler, 1976, Definition 1]. The slope of a nonzero lattice L is the number $(\log \mathrm{vol}\,(L))/\dim L$, and can be thought of as the log of an average length. The log is thrown in solely to convert the multiplicativity of the volumes (provided by Lemma 1.8) into additivity. The slope is undefined when $L=0$.

DEFINITION 1.10. Suppose we plot all submodules of a nonzero lattice L as points in the plane, where the horizontal axis is the dimension, and the vertical axis is log vol. Call this plot the *canonical plot* of L. The slope of L appears in this plot as the slope of the line segment joining 0 and L. The import of Lemma 1.8 is that slope (L/L') appears in this plot as the slope of the line segment joining L' to L (see Figure 1.11). In fact, the canonical plot for L/L' appears (translated) in the canonical plot for L as those points represented by sublattices of L containing L'.

If A and B are subgroups of an abelian group C, then a basic fact is that $A/A\cap B = A+B/B$. For lattices this is false, as can be seen in easy examples. Nevertheless, we can make do with a certain inequality for the volumes, which we now derive.

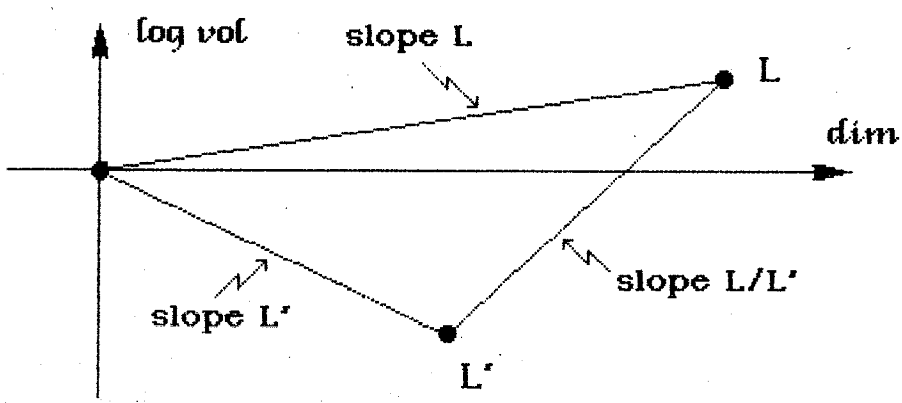

Figure 1.11

In the notation of the proof of Lemma 1.8, there is a formula

$$p\left(\sum a_i e_i + \sum b_j f_j\right) = \sum b_j f_j$$

for the orthogonal projection p. It follows that orthogonal projection is length decreasing, i.e. for all v, $|v| \geq |pv|$; but it also is volume decreasing in the following sense.

THEOREM 1.12 [Stuhler 1976, Proposition 2]. *Suppose L is an $\mathcal{O}$-lattice, and L_1 and L_2 are sublattices. Then*

(i) $\mathrm{vol}\,(L_2/L_2 \cap L_1) \geq \mathrm{vol}\,(L_1 + L_2/L_1)$

and

(ii) $\mathrm{vol}\,(L_1 \cap L_2)\,\mathrm{vol}\,(L_1 + L_2) \leq \mathrm{vol}\,(L_1)\,\mathrm{vol}\,(L_2)$.

Proof. Part (ii) follows from part (i) together with Lemma 1.8. Let's show part (i). First, we may assume $\mathcal{O} = \mathbb{Z}$, by restriction of scalars. Let P_1 and P_2 denote the underlying modules. Now choose a filtration $P_1 \cap P_2 = Q_0 \subseteq Q_1 \subseteq \cdots \subseteq Q_n = P_2$ in which each subquotient Q_i/Q_{i-1} is free with rank 1. By Lemma 1.8, we may replace P_2 by Q_i and P_1 by $Q_{i-1} + P_1$, thereby achieving $\dim\,(L_2/L_1 \cap L_2) = 1$. We may also achieve $L_1 \cap L_2 = 0$ by replacing L_i by $L_i/L_1 \cap L_2$ for $i = 1, 2$; this reduction uses Lemma 1.3. Now let m be a generator for L_2, and let p be the orthogonal projection fo V_1. Then $\mathrm{vol}\,(L_2) = |m| \geq |pm| = \mathrm{vol}\,(L_1 + L_2/L_1)$. QED

Remark: One can use 1.4 and 1.1(d) to extend 1.12 to any pair of submodules, as Stuhler does.

DISCUSSION 1.13. Theorem 1.12 is fundamental – it mimics the equality for the dimensions:

$$\dim\,(L_1 \cap L_2) + \dim\,(L_1 + L_2) = \dim\,(L_1) + \dim\,(L_2).$$

We can interpret this in terms of the canonical plot of L, from Definition 1.10. Consider just the four points obtained from $L_1 \cap L_2$, $L_1 + L_2$, L_1, and L_2; any three of them determine a parallelogram, and then the theorem can be visualized as an assertion about the relationship of the fourth vertex of that parallelogram to the fourth point. If the fourth point comes from L_1 or L_2, then it lies at or above the corresponding vertex of the parallelogram (and on the same vertical line, by the equality for the ranks). If the fourth point is from $L_1 \cap L_2$ or $L_1 + L_2$, then it lies at or below the corresponding vertex of the parallelogram (and on the same vertical line). This situation is easily visualized: see Figure 1.14. We will call it the "parallelogram constraint."

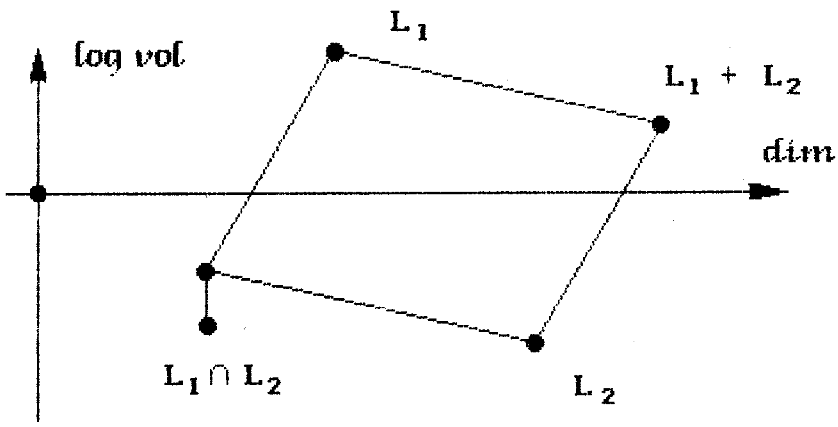

Figure 1.14

LEMMA 1.15. *Given a lattice L and a number c, there exist only a finite number of submodules $L_1 \subseteq L$ with* $\operatorname{vol}(L_1) < c$.

Proof [compare Stuhler 1976, Proposition 1]. By restriction of scalars, we may assume $\mathcal{O} = \mathbb{Z}$. Choose an integer r and require also that $\dim L_1 = r$. In case $r = 1$, finiteness follows from the fact that L_1 is discrete and the sphere of radius c is compact. For $r > 1$ we may replace L by $\Lambda^r L$ and each L_1 by $\Lambda^r L_1$. (As inner product on $\Lambda^r V$ we take the one satisfying $\langle (l_1 \cdots l_r), (m_1 \cdots m_r) \rangle = \det \langle l_i, m_j \rangle$.) The assignment $L_1 \mapsto \Lambda^r L_1$ is finite-to-one because $\Lambda^r L_1$ determines $L \cap (L_1 \otimes \mathbb{Q})$ and the index in it of L_1. We have $\dim \Lambda^r L_1 = 1$ and $\operatorname{vol}(L_1) = \operatorname{vol}(\Lambda^r L_1)$, so the finiteness for $r > 1$ follows from the finiteness for $r = 1$. QED

DISCUSSION 1.16. Lemma 1.15 tells us that the canonical plot of L is bounded below. Thus the convex hull of the canonical plot of L will be bounded on left and right by two vertical lines at 0 and $\dim L$; it is unbounded above unless $L = 0$ (because L has submodules of arbitrarily large finite index). Its lower boundary is a convex polygon stretching from the origin to the point corresponding to L: we call it the canonical polygon of L. The interesting thing for us will be to study its vertices, each of which, according to 1.15, is represented by a sublattice of L. By "vertex", we mean either an endpoint of the polygon, or a point on the polygon where the slope actually changes; points on the polygon represented by submodules may not be at vertices.

Suppose now that L_1 and L_2, submodules of L, are chosen to lie on the canonical polygon of L in such a way that they do not both lie in the interior of the same straight line segment of the boundary (this will happen, for example, if either of them lies on a vertex). Since they have minimal volume for their ranks,

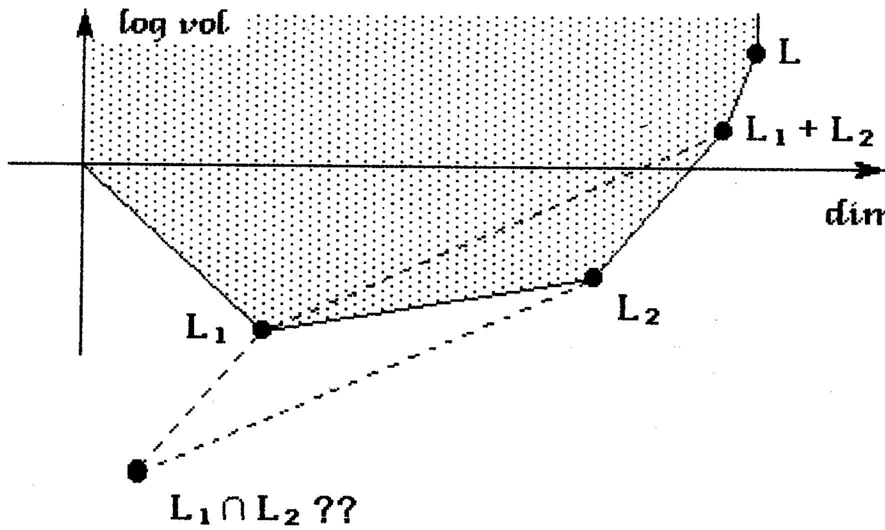

Figure 1.17

both L_1 and L_2 are actually sublattices of L, according to Lemma 1.1(e). Assume for the sake of definiteness, that $\dim L_1 \leqslant \dim L_2$. Then the slope of the segment of the polygon just to the right of L_2 is strictly steeper than the slope just to the left of L_1. This means that it is not possible that $\dim L_1 + L_2 > \dim L_2$, without violating the parallelogram constraint from Discussion 1.13, so therefore $\dim L_1 + L_2 = \dim L_2$ (see Figure 1.17). It follows that $L_1 + L_2 = L_2$, for else its volume would be strictly smaller than the volume of L_2. Thus we've shown that $L_1 \subseteq L_2$.

Now suppose L_1 and L_2 represent the same vertex. The preceding argument shows both $L_1 \subseteq L_2$ and $L_2 \subseteq L_1$, so $L_1 = L_2$.

We have proved the following theorem.

THEOREM 1.18. *The vertices of the canonical polygon of L are represented by unique sublattices of L, and they form a chain.*

DEFINITION 1.19. The filtration of L consisting of those sublattices of L which represent vertices of the canonical polygon of L, is called the *canonical filtration* of L. By convention, it always includes 0 and L. The canonical filtration is called canonical because it depends only on L, and not on any choices.

Theorem 1.18 and Definition 1.19 are roughly equivalent to [Stuhler 1976, Satz 1, Folgerung aus Satz 1].

DEFINITION 1.20. We say that L is *semistable* if its canonical filtration contains only 0 and L (i.e. its canonical polygon is a single line segment). In all other cases we say L is *unstable*.

If $\mathrm{rk}(L)=1$, then L is semistable. The successive subquotients of the canonical filtration of L are all semistable, and their slopes are (strictly) increasing.

L is semistable if and only if it satisfies the inequalities slope $M \geqslant$ slope L for every submodule M.

Remark 1.21. It follows immediately from the definition that if $h: L \xrightarrow{\sim} M$ is an isomorphism of lattices (i.e. an isometry), then h carries the canonical filtration of L into the canonical filtration of M. It is this fundamental fact that enables equivariant constructions in the symmetric space in chapter 2.

OBSERVATION 1.22. The (finite) orthogonal group G of L leaves invariant the canonical filtration of L; the same applies if we tensor L with the rationals. If G acts irreducibly on L, then L must be semistable. This happens, for example, if we take $L=f_*\mathcal{O}$ where $\mathcal{O}$ is the ring of integers in a cyclotomic field, because the roots of unity of $\mathcal{O}$ are in G. This gives an interesting explicit lower bound on volumes of subgroups of $f_*\mathcal{O}$.

DEFINITION 1.23. Let $\max L$ denote the largest slope of a segment of the canonical polygon of L, and let $\min L$ denote the smallest.

DIVERSION 1.24. If r is a positive real number, then from L we may produce a new $\mathcal{O}$-lattice called $L[r]$ by multiplying each of the norms on L by r (or equivalently, by multiplying the inner products by r^2). Clearly $f_*(L[r])=(f_*L)[r]$, thus the following formulas hold.

$$\mathrm{vol}\, L[r] = r^{\dim L}\, \mathrm{vol}\, L$$
$$\mathrm{slope}\, L[r] = \log r + \mathrm{slope}\, L$$

The canonical plot for $L[r]$ can be obtained from the canonical plot for L by applying the affine transformation $(x, y) \mapsto (x, y + x \log r)$. This transformation preserves straight lines, and thus transforms the canonical polygon for L into the canonical polygon for $L[r]$. It follows that L and $L[r]$ have the same canonical filtration (as far as the underlying $\mathcal{O}$-modules are concerned), and the following formulas hold.

$$\max L[r] = \log r + \max L$$
$$\min L[r] = \log r + \min L$$

Another thing to notice is that the rescaling $L \mapsto L[r]$ preserves exact sequences of lattices; if it were a functor, we could call it an exact functor. It is analogous to tensoring a vector bundle with a power of a fixed line bundle (on an algebriac curve).

It is also possible to rescale the norms by independent factors at the infinite places; were we studying $Sl_n\mathcal{O}$ we would do this.

EXAMPLE 1.25. Consider $\mathbb{C} = \mathbb{R}^2$ as the Euclidean plane, and let $\mathcal{H}$ be the upper half plane. For any $t \in \mathcal{H}$ we may form the lattice $L = L(t) = \mathbb{Z} \cdot 1 + \mathbb{Z} \cdot t$. Let $\mathcal{D} = \{z : |z| \geq 1 \text{ and } |\mathrm{Re}\, z| \leq \frac{1}{2}\}$ be the usual fundamental domain for the action of $Sl_2\mathbb{Z}$ on $\mathcal{H}$. Assume that $t \in \mathcal{D}$; then it is clear that 1 is a vector of minimal length in L. Since $\mathrm{vol}\, L = \mathrm{Im}\, t$, it follows that L is semistable if and only if $\mathrm{Im}\, t \leq 1$. The set B of all $t \in \mathcal{H}$ such that L is semistable is invariant under $\Gamma = Sl_2\mathbb{Z}$; for, given $g \in \Gamma$, we see easily that $L(gt) = zL(t)$ for some complex number z. Write $z = ru$, where r is real and $|u| = 1$. Then $uL(t)$ and $L(t)$ are isomorphic lattices (the isomorphism is multiplication by u), and $ruL(t)$ has the same canonical filtration as $uL(t)$ by diversion 1.24.

We know now that B is Γ-invariant, and we know its intersection with the fundamental domain $\mathcal{D}$. This allows us to determine B – it is the complement of countably many disjoint open disks, namely all the translates of the half-plane $C = \{t \in \mathcal{H} : \mathrm{Im}\, t > 1\}$. See Figure 1.26: this is the same picture which appears in Rademacher's work on partitions [1973]. Many of these disks are tangent (at points corresponding to those lattices with two independent vectors of minimal length), so clearly B is not a manifold (with boundary). If, however, we shrink C slightly by strengthening the inequality in its definition to $\mathrm{Im}\, t > 1 + \varepsilon$, letting the

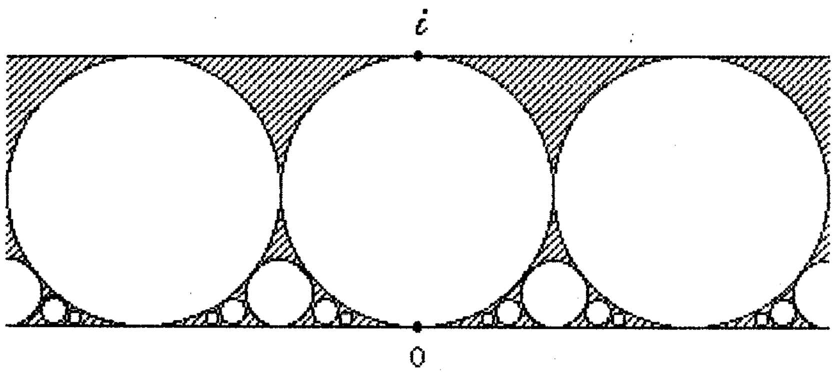

Figure 1.26

other disks shrink the same way, then the tangencies disappear, and this enlargement of B is a manifold with boundary. This was explained by Serre [1979], and it is this which I generalize to Gl_n in the sequel.

DISCUSSION 1.27. Suppose $\mathscr{F}: 0 = L_0 \subseteq L_1 \subseteq \cdots \subseteq L_s = L$ is a filtration by sublattices of an $\mathcal{O}$-lattice L. Consider the plot formed by plotting log vol and dim (as in Definition 1.10) for only those submodules L' of L such that $L_i \subseteq L' \subseteq L_{i+1}$ for some i; call it the *canonical plot of L subordinate to the filtration* $\mathscr{F}$. Consider also, as before, the convex hull of this plot, and the corresponding convex polygon C bonding it below. Suppose now that each L_i happens to sit on C: I claim then that C actually is the canonical polygon. It is equivalent to show that every sublattice L' of L lies on or above C, and this we can do by induction on s (the case $s = 1$ being obvious). Consider $L' + L_{s-1}$ and $L' \cap L_{s-1}$: the former clearly lies on or above C, and the latter, by induction, does, too. Now L_{s-1} is on C, which is convex, so the parallelogram constraint of discussion 1.13 tells us that L' must be on or above C. See Figure 1.28.

The canonical filtration of L will include those L_i which sit at vertices of C. Notice that we didn't assume that each L_i occurs at a vertex of C; thus this argument might easily lead to the conclusion that L is semistable. Indeed, it follows that $\mathcal{O}^n$ is a semistable $\mathcal{O}$-lattice for any n, where $\mathcal{O}^n$ denotes the n-fold orthogonal sum of the $\mathcal{O}$-lattice $\mathcal{O}$ from example 1.5.

COROLLARY 1.29. *Suppose L has filtration $0 = L_0 \subseteq L_1 \subseteq \cdots \subseteq L_s = L$ by*

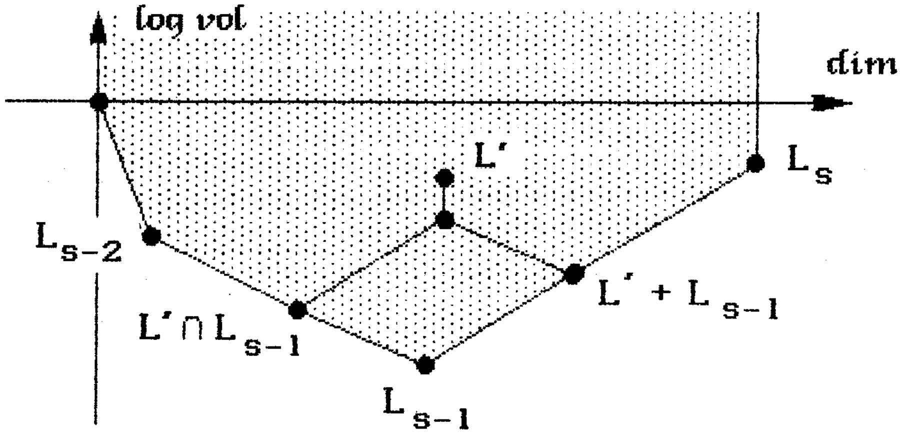

Figure 1.28

sublattices, so that $\min L_{i+1}/L_i \geqslant \max L_i/L_{i-1}$. *Then*

(a) *The canonical polygon of L is formed by laying the canonical polygons of the subquotients* L_i/L_{i-1} *end to end.*

(b) *Each* L_i *lies on the canonical polygon of L.*

(c) $\min L/L_i = \min L_{i+1}/L_i$.

(d) $\max L_i = \max L_i/L_{i-1}$

(e) *If* $\min L_{i+1}/L_i > \max L_i/L_{i-1}$, *then* L_i *is the canonical filtration of L.*

(f) *If* $L_i \subsetneqq L' \subsetneqq L_{i+1}$ *and* L' *is in the canonical filtration of* L_{i+1}/L_i, *then* L' *is in the canonical filtration of L.*

(g) *The canonical filtration of L consists solely of sublattices arising as in* (e) *or* (f).

COROLLARY 1.30. *Suppose L has a filtration* $0 = L_0 \subseteq L_1 \subseteq \cdots \subseteq L_s = L$, *whose subquotients* L_i/L_{i-1} *are semistable, with strictly increasing slopes. Then this filtration is the canonical filtration.*

COROLLARY 1.31. *Suppose* L' *is a sublattice of L. Then* L' *is in the canonical filtration of L is and only if* $\max L' < \min L/L'$.

2. Spaces of lattices

In this section we investigate the way the canonical filtration behaves when the lattice moves.

Let P be a finitely generated projective $\mathcal{O}$-module of rank n, let $\Gamma = Gl(P)$. Let $\tilde{X} = \tilde{X}(P)$ be the space of lattices L whose underlying $\mathcal{O}$-module is P. Let $\tilde{X}_\infty$ be the space of inner products on V_∞; if a basis is chosen for V_∞, then $\tilde{X}_\infty$ is seen to be an open subspace of a real or complex vector space. We have $\tilde{X} = \prod \tilde{X}_\infty$, and this provides us with a topology for $\tilde{X}$.

We consider Γ to act on P on the left. If a basis is chosen, our vectors will be thought of as column vectors, and matrices of linear maps will be written on the left, as usual.

Given $L \in \tilde{X}$ and $v, w \in V_\infty$, let $\langle v, w; L\rangle_\infty$ denote the value of the inner product on the vectors v and w. If $g \in \Gamma$, we define a new lattice gL in $\tilde{X}$ by the formula $\langle v, w; gL\rangle_\infty = \langle g^{-1}v, g^{-1}w; L\rangle_\infty$. This defines an action of Γ on $\tilde{X}$ on the left.

It happens that there is an isomorphism $L \xrightarrow{\sim} gL$ of lattices defined by $v \mapsto gv$ which we may as well call g, also.

Conversely, suppose $g: L_1 \xrightarrow{\sim} L_2$ is an isomorphism of $\mathcal{O}$-lattices, each of which is in $\tilde{X}$. Since P is the underlying module for both of them, g gives rise to an element of Γ, which we may also call g. We see clearly that $L_1 = gL_2$.

Thus the orbit set $\Gamma\backslash\tilde{X}$ can be regarded as the set of isomorphism classes of $\mathcal{O}$-lattices whose underlying $\mathcal{O}$-module is isomorphic to P. (It is the analogue of the moduli space for vector bundles on an algebraic curve, and will become compact once we throw out the unstable points.)

Scaling the norms commutes with changing the basis, i.e. if $r>0$ and $g\in\Gamma$, then $g(L[r])=(gL)[r]$. Let X be the quotient of $\tilde{X}$ by the equivalence relation $L\sim L[r]$; it is clear that X is a manifold. The difference between $\tilde{X}$ and X becomes important only for assertions about compactness.

DEFINITION 2.1. By an F-subspace of V, we will mean either an F-subspace of $P\otimes_{\mathcal{O}}F$, or the real span of such a subspace in V: no confusion should result from this blurring of the distinction between an F-subspace and its real span, for each can be recovered from the other. If $L\in\tilde{X}$, then the sublattices M of L are in one-to-one correspondence with the F-subspaces $W\subseteq V$. We will use the notation $M=L\cap W$. If $0\subseteq L\cap W_1\subseteq\cdots\subseteq L\cap W_s=L$ is the canonical filtration of L, then we will refer to $0\subseteq W_1\subseteq\cdots\subseteq W_s=V$ also as the canonical filtration of L. For F-subspaces W of V we may define a real function d_W on $\tilde{X}$ by the formula

$$d_W(L)=d(W,L)=\exp\left((\min L/L\cap W)-(\max L\cap W)\right)$$

This function is concocted so that, by corollary 1.31, W occurs in the canonical filtration is and only if $d(W,L)>1$ (and in that case, the canonical filtration for L is obtained by splicing the canonical filtrations for $L\cap W$ and $L/L\cap W$). In terms of the polygon, $d(W,L)>1$ iff W is at a vertex, $d(W,L)=1$ iff W is in the interior of an edge, and $d(W,L)<1$ iff W is not on the polygon. A larger value of $d(W,L)$ corresponds to a more acute slope change at the vertex W. Moreover, $d(W,L[r])=d(W,L)$, for any $r>0$, so d_W descends to a function on X. We may imagine that the larger $d(W,L)$ is, the further L is out toward infinity; alternatively, $d(W,L)^{1}$ measures the distance to the cusp corresponding to W.

For any $t\geq 1$, define $\tilde{X}_W(t)=\tilde{X}(W,t)=\{x\in\tilde{X}:d(W,x)>t\}$, and let $\tilde{X}_W=\tilde{X}(W,1)$. Define $X(W,t)$ and X_W similarly. We let $X_{ss}(t)=X-\bigcup_W X(W,t)$, and $X_{ss}=X_{ss}(1)$. We call $\tilde{X}_{ss}$ the semistable part of $\tilde{X}$, for its points are those L which are semistable. For $t=1$ and $\mathcal{O}=\mathbb{Z}$, these sets agree with those defined in [Stuhler, 1976]; what we prove in this case was already obtained by Stuhler.

We state the following easy lemma:

LEMMA 2.1.

(a) $d(gW, gL) = d(W, L)$, *for* $g \in \Gamma$.

(b) $X(gW, t) = gX(W, t)$, *for* $g \in \Gamma$.

(c) $X_{ss}(t)$ *is closed, and is stable under* Γ.

(d) *If* $X(W_1, t) \cap \cdots \cap X(W_n, t) \neq \emptyset$, *then* $\{W_1, \ldots, W_n\}$; *ordered by inclusion, is a chain.*

2.3. Examples

EXAMPLE 2.3.1. Take $\mathcal{O} = \mathbb{Z}$ and $n = 2$. Let $\mathcal{H} \subseteq \mathbb{C}$ be the upper half plane. Given $z \in \mathcal{H}$ we may embed $P = \mathbb{Z}^2 = \mathbb{Z}e_1 \oplus \mathbb{Z}e_2$ into $\mathbb{C}$ by sending e_1 to z and e_2 to 1. The plane $\mathbb{C} = \mathbb{R}^2$, with its standard (real) inner product, makes P into a lattice which we will call $L = L(z)$. The number z can be recovered from $L(z)$ (forgetting the embedding into $\mathbb{C}$) because $\langle e_1, e_1 \rangle = |z|^2$ and $\langle e_1, e_2 \rangle = \operatorname{Re} z$. We can also extract $L(z)$ from its equivalence class up to scaling because of the normalization condition $\langle e_2, e_2 \rangle = 1$. This shows that the resulting map $\mathcal{H} \tilde{\to} X$ is a diffeomorphism. It turns out that $gL(z)$ and $L(gz)$ differ only by scaling, provided we define $gz = (dz - c)/(-bz + a)$. If we replace the usual action of $Sl_2\mathbb{Z}$ on $\mathcal{H}$ by its composite with transpose inverse, we may say that the map $\mathcal{H} \tilde{\to} X$ is equivariant.

Now suppose a $\mathbb{Q}$-subspace $W = \operatorname{span}(re_1 + se_2) \subseteq V$ is given, where r and s are relatively prime integers. Then $L \cap W$ corresponds to the subgroup $\mathbb{Z}(rz + s)$ of $\mathbb{C}$, so

$$\begin{aligned}
&\operatorname{vol} L \cap W = |rz + s| \\
&\operatorname{vol} L = \operatorname{Im} z \\
&d_W(z) = d(W, L(z)) = (\operatorname{vol} L/L \cap W)/(\operatorname{vol} L \cap W) = (\operatorname{vol} L)/(\operatorname{vol} L \cap W)^2 \\
&\qquad = (\operatorname{Im} z)/|rz + s|^2
\end{aligned}$$

If $r = 0$, then $|s| = 1$, so $d_W(z) = \operatorname{Im} z$ and $X(W, t) = \{z \in \mathcal{H} : \operatorname{Im} z > t\}$. If $r \neq 0$, then $X(W, t)$ is the open disk of diameter $1/(tr^2)$ tangent to the real line at the rational number s/r. For $t = 1$, we recover the situation in example 1.25; for $t > 1$, the closed set $X_{ss}(t)$ is a manifold with boundary, and was described by Serre [1979].

The next two examples use Remark 1.7 implicitly for computing volumes.

EXAMPLE 2.3.2. This time take $\mathcal{O}$ quadratic imaginary, $n = 2$, and let $\mathcal{H}^3 = \{(z, w) \in \mathbb{C}^2 : \operatorname{Im} w = 0, \operatorname{Re} w > 0\}$ be the hyperbolic 3-space, sitting in $\mathbb{C}^2$ endowed with the standard hermitian inner product. Choose an embedding $\mathcal{O} \subseteq \mathbb{C}$.

Given $h=(z,w)\in\mathcal{H}^3$ we may embed $P=\mathcal{O}^2=\mathcal{O}e_1\oplus\mathcal{O}e_2$ into $\mathbb{C}^2$ by sending e_1 to h and e_2 to $(1,0)$. This makes P into an $\mathcal{O}$-lattice we call $L(h)$, and we get a diffeomorphism $\mathcal{H}^3 \xrightarrow{\sim} X$, equivariant for $Sl_2\mathbb{C}$. If we identify $\mathbb{C}^2$ with the quaternions $\mathbb{C}\oplus\mathbb{C}j$, and take

$$g=\begin{pmatrix} a & b \\ c & d \end{pmatrix}$$

in $Sl_2\mathcal{O}$, the formula $g\cdot h=(dh-c)(-bh+a)^{-1}$ gives the action of $Sl_2\mathcal{O}$ on $\mathcal{H}^3$. We can repeat the discussion from Example 2.3.1 up to a point. We have

$$\begin{aligned}
&\operatorname{vol} L = w^2(\operatorname{vol}\mathcal{O})^2\\
&\operatorname{vol} L\cap W = (\operatorname{vol} L\cap \operatorname{span}(re_1+se_2))/I = |rh+s|^2(\operatorname{vol}\mathcal{O})/I\\
&d(W,L(h)) = (\operatorname{vol} L/L\cap W)^{1/2}/(\operatorname{vol} L\cap W)^{1/2}\\
&\qquad = (\operatorname{vol} L)^{1/2}/(\operatorname{vol} L\cap W)\\
&\qquad = wI/|rh+s|^2
\end{aligned}$$

Here r and s are chosen from $\mathcal{O}$ so that $L\cap W\supseteq\mathcal{O}(re_1+se_2)\neq 0$, we identify $s\in\mathcal{O}$ with $(s,0)\in\mathcal{H}^3$ so the expression $rh+s$ makes sense, and the integer I is defined by $I=[\mathcal{O}:\mathcal{O}r+\mathcal{O}s]$. If $\mathcal{O}$ happens to be a principal ideal domain, we can always take $I=1$ by choosing r and s properly. This is the same function used by Mendoza [1980], except that he gives a slightly different (but equivalent) definition for I.

EXAMPLE 2.3.3. Take $\mathcal{O}$ real quadratic (with its real places labeled 1, 2), $n=2$, and let $K=\mathcal{H}\times\mathcal{H}$, where H is the upper half plane. A point $(z,z')\in K$ gives two inner products on $\mathbb{R}^2$, each one defined as in Example 2.3.1; this is just what's required to make $P=\mathcal{O}^2$ into a lattice. For $W=\operatorname{span}(e_2)$, we find that

$$\begin{aligned}
&\operatorname{vol} L = (\operatorname{vol}\mathcal{O})^2(\operatorname{Im} z)(\operatorname{Im} z')\\
&\operatorname{vol} L\cap W = \operatorname{vol}\mathcal{O}\\
&d(W;z,z') = (\operatorname{vol} L)^{1/2}/(\operatorname{vol} L\cap W) = ((\operatorname{Im} z)(\operatorname{Im} z'))^{1/2}
\end{aligned}$$

The regions $X_W(t)$ turn out to be the same as those used in [Ash, et al., 1975, p. 41–42], where the function d_W^{-2} was called "distance to the cusp".

DISCUSSION 2.4. Now we describe our interpretation of the "geodesic action" of Borel–Serre. Suppose we are given $L\in\tilde{X}$, an F-subspace $W\subseteq V$, and

$r>0$. We construct a new lattice $L[W,r]$ in $\tilde{X}$ by changing the norms of L. Writing $V_\infty = L\otimes_{\mathcal{O}} F_\infty$ and $W_\infty = W\otimes_{\mathcal{O}} F_\infty$, we can use the inner product on V_∞ provided by L to write $V_\infty = W_\infty \oplus W_\infty^\perp$, an orthogonal sum. Now multiply the norm on $W_\infty^\perp$ by r, but leave the norm on W_∞ unchanged; assemble these by orthogonal sum to form a new inner product for V_∞. Doing this for each ∞ defines a new lattice $L[W,r]\in\tilde{X}$, there is an obvious exact sequence

$$0\to L\cap W\to L[W,r]\to (L/L\cap W)[r]\to 0$$

and

$$d(W, L[W,r]) = r\cdot d(W,L).$$

This procedure is easy to visualize as a dilation of V in the directions perpendicular to W. Since it commutes with scaling (i.e. $L[W,r_1][r_2] = L[r_2][W,r_1]$), we can also regard it as acting on X. For $g\in\Gamma$ we have $g(L[W;r]) = (gL)[gW;r]$.

Now suppose that we have a chain of F-subspaces $W_1\subseteq\cdots\subseteq W_m$ of V. Then for $r_1,\ldots,r_m>0$, let $L' = L[W_1,r_1]\cdots[W_m,r_m]$. Since $L'\cap W_i/L'\cap W_{i-1} = (L\cap W_i/L\cap W_{i-1})[r_1\cdots r_i]$, 1.29 implies that choosing r_i large enough ensures that W_i is in the canonical filtration of L', or even that $L'\in X(W_i,t)$. This proves the following converse to Lemma 2.2(d).

LEMMA 2.5. *Given F-subspaces* $0\neq W_1\subseteq\cdots\subseteq W_s\subseteq V$, the set $X(W_1,t)\cap\cdots\cap X(W_s,t)$ is nonempty.

Remark 2.6. We can also prove that any P can be given norms which make it into a semistable $\mathcal{O}$-lattice. To see this, choose the norms arbitrarily at first, producing an $\mathcal{O}$-lattice L. Then let $W_i\subseteq\cdots\subseteq W_m$ be its canonical filtration. Then for suitable choice of numbers r_i, the lattice $L' = L[W_1,r_1]\cdots[W_m,r_m]$ will have $L'\cap W_i/L'\cap W_{i-1}$ all of the same slope, and will be semistable by 1.27.

3. Continuity

In this section we prove that the functions d_W are continuous.

Suppose M is a topological space. We say that a family $\{f_n\}$ of real functions on M is locally equicontinuous if, for any ε, there is a covering of M by open sets U, such that for all n and each $x,y\in U$, $|f_n(x)-f_n(y)|<\varepsilon$. Equivalently, each x in M has a neighborhood U such that for all n and each $y\in U$, $|f_n(x)-f_n(y)|<\varepsilon$. It is clear that the supremum of a locally equicontinuous family, if finite, is itself a continuous function. The union of a finite number of locally equicontinuous families is locally equicontinuous.

LEMMA 3.1. (a) *For each nonzero* F*-subspace* $W \subseteq V$ *consider the real function* $\tilde{X}$ *defined by* $L \mapsto \log \operatorname{vol}(L \cap W)$; *this family of functions is locally equicontinuous*

(b) $L \mapsto \min L$ *is a continuous function on* $\tilde{X}$.
(c) $L \mapsto \max L$ *is a continuous function on* $\tilde{X}$.
(d) d_W *is a continuous function on* X.
(e) $X(W, t)$ *is an open subset of* X.

Proof. It is enough to prove (a). For example, to see that (a) $\rightarrow$ (c) simply observe that $\max L = \sup\{(\log \operatorname{vol} L - \log \operatorname{vol} L \cap W)/(\operatorname{rk} P - \operatorname{rk} L \cap W)\}$.

We may as well restrict scalars, achieving $\mathcal{O} = \mathbb{Z}$, for this only increases the family of functions being considered.

We may choose a number m, and restrict attention to F-subspaces W of V of dimension m. As in the proof of Lemma 1.15, we may apply the m^{th} exterior power to everything, achieving $m = 1$ (and possibly enlarging the family once again). Now let $n = \dim P$.

Choose a basis for V, and identify each L in $\tilde{X}$ with its (positive definite symmetric) matrix z. Let w be a generator for $L \cap W$, so $\log \operatorname{vol}(L \cap W) = \log |w| = (1/2) \log (w^t z w)$. Enlarge the family of functions once again by dropping the requirement that w be in W, and forget W; for any nonzero w in V we will consider the function $z \mapsto \log (w^t z w)$ on $\tilde{X}$, (forgetting the factor 1/2).

Fix a point $z \in \tilde{X}$ and a number $\varepsilon > 0$. We seek a small neighborhood of z, but we may as well first change the basis of V to make $z = 1$, the identity matrix. Let $\delta > 0$ be a small number (to be determined later) and consider arbitrary symmetric matrices Δz with $|(\Delta z)_{ij}| < \delta$ for each i, j. Then

$$\begin{aligned} |(w^t \cdot \Delta z \cdot w)/(w^t \cdot w)| &\leqslant \delta\Big(\sum\sum |w_i|\,|w_j|\Big)/\Big(\sum w_i^2\Big) \\ &\leqslant \delta\Big(\sum\sum (1/2)(w_i^2 + w_j^2)\Big)/\Big(\sum w_i^2\Big) \\ &= n\delta \end{aligned}$$

This leads to

$$\begin{aligned} |\log (w^t(z + \Delta z)w) - \log (w^t \cdot z \cdot w)| &= |\log (1 + (w^t \cdot \Delta z \cdot w)/(w^t \cdot w))| \\ &\leqslant |\log (1 - n\delta)|, \end{aligned}$$

which is smaller than ε if δ is chosen small enough. QED

COROLLARY 3.2. *Given a point $L \in \tilde{X}$ and a F-subspace $W \subseteq V$, there is a neighborhood of L on which d_W is the infinum of a finite set of smooth functions.*

Proof. We can write $\log d(W, L)$ as the infinum of all functions

$$\text{slope}\,(L \cap W_2/L \cap W) - \text{slope}\,(L \cap W/L \cap W_1)$$

where W_2 runs over all F-subspeaces of V containing W properly, and W_1 runs over all those contained properly in W. Each one of these functions is smooth, so it is enough to show that in some neighborhood of L, only a finite number of them are needed. In fact, only the ones which already achieve the minimum are needed, because by 1.15 and 3.1(a), the others stay far enough away on some small enough neighborhood U of L. As for the ones which do achieve the minimum, there are only a finite number of them (by 1.15, again). QED.

COROLLARY 3.3. *Given $t \geq 1$, the family of open subsets $X(W, t)$ of X (one for each F-subspace W of V) is locally finite. Moreover, if $t > 1$, each L in X has a neighborhood U so small that $\{W \mid X(W, t) \cap \neq \varnothing\}$, in addition to being finite, is a chain.*

Proof. Suppose $L \in X$. Consider first those $W \subseteq V$ for which $L \cap W$ lies above the canonical polygon of L. They are all further above it than a certain minimum distance, according to 1.15. By 3.1(a), we can find a neighborhood U of L so that whenever $L' \in U$, the points corresponding to the various $L' \cap W$ are still above the canonical polygon of L'; thus the only candidates for members of the canonical filtration of such an L' will be those W for which $L \cap W$ was on the canonical polygon of L; of these there are only a finite number, by 1.15 again.

For the second statement, the same argument shows that U can be chosen so that for all W, $d(W, L') < t$ if $d(W, L) < t$ for all L' in U. Q.E.D.

THEOREM 3.4. *Given $t > 1$, the spaces $\tilde{X}_{ss}(t)$ and $X_{ss}(t)$ are manifolds with boundary, and the boundary consists of those points x with $\sup_W d(W, x) = t$.*

Remark. Example 1.25 shows that the bound on t is sharp, for when $t = 1$, these spaces are not manifolds.

Proof. The following proof works equally well for either $\tilde{X}_{ss}(t)$ or $X_{ss}(t)$.

Define $h(x) = \max\,(1, \sup_W d(W, x))$, so that $X_{ss}(t)$ is $h^{-1}([1, t])$. It follows from 3.3 that h is continuous. Choose a point x_0 in X with $h(x_0) = t$. Since $t > 1$, in a small enough neighborhood U of x_0 we have $h(x) = \sup_i d(W_i, x)$, where $W_1 \subseteq \cdots \subseteq W_s$; this follows from 3.3. If U is small enough, then any x in U has