Abstract Let $\mathbf{G}$ be the reductive $\mathbb{Q}$-group $R_{F/\mathbb{Q}}\mathrm{GL}_n$, where $F/\mathbb{Q}$ is a number field. Let $\Gamma \subset \mathbf{G}$ be an arithmetic group. We discuss some techniques to compute explicitly the cohomology of Γ and the action of the Hecke operators on the cohomology. This is a writeup of a lecture course given at the summer school *Computations with modular forms*, Heidelberg, Germany, in August 2011.

Lectures on computing cohomology of arithmetic groups

Paul E. Gunnells

1 Introduction

This is a writeup of five lectures given at the summer school *Computations with modular forms*, Heidelberg, Germany, in August 2011. The course covered essentially all the material here, although I have made some corrections and modifications with the benefit of hindsight, and have taken the opportunity to elaborate the presentation. I've tried to preserve the informal nature of the lectures.

I thank the organizers for the opportunity to speak, and the participants of the summer school for a stimulating environment. I thank my collaborators Avner Ash, Mark McConnell, and Dan Yasaki, for many years of fun projects, and for all that they've taught me about this material. Thanks are also due to an anonymous referee, who carefully read the lectures and made many valuable suggestions. Finally, I thank the NSF for supporting the research described in these lectures.

2 Cohomology and holomorphic modular forms

The goal of our lectures is to explain how to explicitly compute some automorphic forms via cohomology of arithmetic groups. Thus we begin by reviewing modular symbols and how they can be used to compute with holomorphic modular forms. For more details we refer to [19, 61]. This material should be compared with that in Rob Pollack's lectures, which contains a different perspective on similar material.

Paul E. Gunnells
University of Massachusetts Amherst, Amherst, MA 01002, USA, e-mail: gunnells@math.umass.edu

Let $N \geq 1$ be an integer, and let $\Gamma_0(N) \subset \mathrm{SL}_2(\mathbb{Z})$ be the subgroup of matrices that are upper-triangular mod N. Let $\mathfrak{H} \subset \mathbb{C}$ be the upper halfplane of all z with positive imaginary part. The group $\Gamma_0(N)$ acts on $\mathfrak{H}$ by fractional linear transformations:

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix} \cdot z = \frac{az+b}{cz+d}. \tag{1}$$

We let $Y_0(N)$ be the quotient $\Gamma_0(N)\backslash\mathfrak{H}$. Then $Y_0(N)$ is a smooth algebraic curve defined over $\mathbb{Q}$, called an *(open) modular curve*.

The curve $Y_0(N)$ is not compact, and there is a standard way to compactify it. Let $\mathfrak{H}^* = \mathfrak{H} \cup \mathbb{P}^1(\mathbb{Q})$, where we think of $\mathbb{P}^1(\mathbb{Q})$ as being $\mathbb{Q} \cup \{\infty\}$ with $\mathbb{Q} \subset \mathbb{R} \subset \mathbb{C}$ and ∞ lying infinitely far up the imaginary axis. The points $\partial\mathfrak{H}^* = \mathfrak{H}^* \smallsetminus \mathfrak{H}$ are called *cusps*. The action of $\Gamma_0(N)$ extends to the cusps, and after endowing $\mathfrak{H}^*$ with an appropriate topology, the quotient $X_0(N) = \Gamma_0(N)\backslash\mathfrak{H}^*$ has the structure of a smooth projective curve over $\mathbb{Q}$. This is what most people call the *modular curve*.

By work of Eichler, Haberland, and Shimura, the cohomology of the spaces $Y_0(N)$ and $X_0(N)$ has connections with modular forms. These are holomorphic functions $f\colon \mathfrak{H} \to \mathbb{C}$ satisfying the transformation law

$$f\Big(\frac{az+b}{cz+d}\Big) = (cz+d)^k f(z), \qquad \begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \Gamma_0(N),$$

where $k \geq 1$ is a fixed integer; f is also required to satisfy a growth condition as z approaches any cusp. The space of such functions $M_k(N)$ is a finite-dimensional complex vector space with a subspace $S_k(N)$ of *cusp forms*: these are the f that undergo exponential decay as z approaches any cusp. There is a natural complement $\mathrm{Eis}_k(N)$ to $S_k(N)$, called the space of *Eisenstein series*. Then we have

$$H^1(Y_0(N);\mathbb{C}) \xrightarrow{\sim} S_2(N) \oplus \overline{S}_2(N) \oplus \mathrm{Eis}_2(N), \tag{2}$$

$$H^1(X_0(N);\mathbb{C}) \xrightarrow{\sim} S_2(N) \oplus \overline{S}_2(N). \tag{3}$$

For example, let $N = 11$. Then it is known that $\dim M_2(11) = 2$ and $\dim S_2(11) = 1$. The curve $X_0(11)$ has genus 1, which is consistent with (3). The complement of $Y_0(11)$ in $X_0(11)$ consists of two points. Thus $Y_0(11)$ deformation retracts onto a graph with one vertex and three loops. This implies $H^1(Y_0(11);\mathbb{C}) \simeq \mathbb{C}^3$, again consistent with (2).

We can say even more about (2)–(3):

- We don't have to limit ourselves to quotients by $\Gamma_0(N)$. Indeed, we can use other finite-index subgroups, such as the subgroup $\Gamma_1(N)$ of matrices congruent to $\left(\begin{smallmatrix} 1 & * \\ 0 & 1 \end{smallmatrix}\right)$ modulo N, or the principal congruence subgroup $\Gamma(N)$ of matrices congruent to the identity modulo N.[1] We could also work with

[1] Throughout these lectures we only work with congruence subgroups. For $\mathrm{SL}_2(\mathbb{Z})$ this means any group containing $\Gamma(N)$ for some N.

cocompact subgroups of $\mathrm{SL}_2(\mathbb{R})$, such as arithmetic groups coming from orders in quaternion algebras, and (3) still holds (if we suitably modify our definitions).

- We can work with modular forms of higher weight $k > 2$ by taking cohomology with twisted coefficients [13,60,63]. More precisely, $\mathrm{SL}_2(\mathbb{Z})$ acts on the complex vector space of homogeneous polynomials of degree k by

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix} \cdot P(x, y) = P(ax + cy, bx + dy).$$

This induces a local system $\mathscr{M}_k$ on the quotients $X_0(N)$, $Y_0(N)$, and we have

$$H^1(Y_0(N); \mathscr{M}_{k-2}) \xrightarrow{\sim} S_k(N) \oplus \overline{S}_k(N) \oplus \mathrm{Eis}_k(N), \tag{4}$$

$$H^1(X_0(N); \mathscr{M}_{k-2}) \xrightarrow{\sim} S_k(N) \oplus \overline{S}_k(N). \tag{5}$$

- Let p be a prime. Then there are Hecke operators

$$\begin{aligned} &T_p, \qquad (p, N) = 1, \\ &U_p, \qquad (p, N) > 1 \end{aligned}$$

that generate an algebra of operators acting on $M_k(N)$. The action preserves the decomposition $M_k(N) = S_k(N) \oplus \mathrm{Eis}_k(N)$. There are corresponding operators acting on the cohomology spaces, and the isomorphisms (2)–(5) are isomorphisms of Hecke modules.

Together these facts imply that we can use topological tools to study modular forms and the action of the Hecke operators on them, and brings us to the main point of our lectures:

One can explicitly compute with certain automorphic forms of arithmetic interest by generalizing the left-hand sides of (2)–(5).

How this can be done will be explained in §4 onward. In the next section, we continue to discuss the classical case and modular symbols.

3 Modular symbols

Modular symbols provide an extremely convenient way to use topology to compute with modular forms. They form the main inspiration for the higher-dimensional computations we discuss later. We review modular symbols here. For simplicity we stick to weight $k = 2$, to avoid the notational complexity of twisted coefficients. For more details we refer to [61] or to R. Pollack's lectures.

Let $\Gamma \subset \mathrm{SL}_2(\mathbb{Z})$ be a torsionfree subgroup. For instance, one could take $\Gamma = \Gamma(N)$ for $N \geq 3$. Put $Y_\Gamma = \Gamma\backslash\mathfrak{H}$ and $X_\Gamma = \Gamma\backslash\mathfrak{H}^*$ as before. We want to study the cohomology spaces $H^1(Y_\Gamma;\mathbb{C})$ and $H^1(X_\Gamma;\mathbb{C})$.

By Lefschetz duality [63, Chapter 6], we have an isomorphism

$$H^1(Y_\Gamma;\mathbb{C}) \xrightarrow{\sim} H_1(X_\Gamma,\partial X_\Gamma;\mathbb{C}), \tag{6}$$

where the right hand side is the homology of X_Γ relative to the cusps. This differs from the usual homology in that we allow not only 1-cycles, whose boundaries vanish, but also 1-chains that have boundary supported on the cusps.

According to basic algebraic topology, we can compute $H_1(X_\Gamma;\partial X_\Gamma;\mathbb{C})$ by taking a triangulation of X_Γ with vertices at the cusps. We then get a chain complex $C_*(X_\Gamma)$ with a subcomplex $C_*(\partial X_\Gamma)$, and the relative homology groups are by definition those of the quotient complex $C_*(X_\Gamma)/C_*(\partial X_\Gamma)$.

A quick way to construct the chain complexes $C_*(X_\Gamma)$, $C_*(\partial X_\Gamma)$ is via the *Farey tessellation* $\mathscr{T}$ of $\mathfrak{H}^*$. This is the ideal triangulation of $\mathfrak{H}$ given by the $\mathrm{SL}_2(\mathbb{Z})$-translates of the ideal triangle Δ with vertices at $\{0,1,\infty\}$ (Figure 1). It's easy to describe the edges of $\mathscr{T}$. Denote the cusps $\mathbb{P}^1(\mathbb{Q}) = \mathbb{Q}\cup\{\infty\}$ by column vectors of relatively prime integers, with ∞ corresponding to $(1,0)^t$. Thus the cusp $\alpha \in \mathbb{Q}$ corresponds to the column vector $(a,b)^t$ if $\alpha = a/b$; we think of ∞ as corresponding to the "fraction" $1/0$. Then two cusps are joined by an edge in the triangulation if and only if the corresponding column vectors form a matrix with determinant ± 1:

$$a/b \text{ joined to } c/d \text{ in } \mathscr{T} \quad\Longleftrightarrow\quad \det\begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \{\pm 1\}. \tag{7}$$

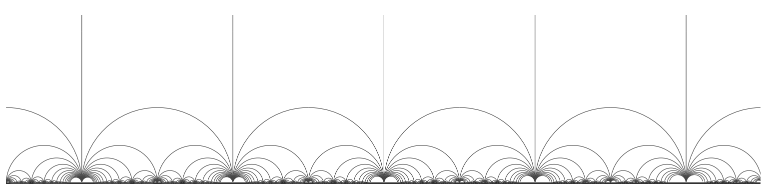

Fig. 1 The Farey tessellation.

Note that Δ is not a fundamental domain for $\mathrm{SL}_2(\mathbb{Z})$, but rather a union of three fundamental domains. This suffices for our purposes, since one can easily see that a fundamental domain for any torsionfree Γ can be assembled from finitely many copies of Δ. Thus $\mathscr{T}$ endows our quotient X_Γ with a finite triangulation, and by construction the vertices of this triangulation are exactly ∂X_Γ.

For example, if $\Gamma = \Gamma(N)$ then the quotient $X(N) := X_\Gamma$ equipped with this triangulation is beautifully symmetric: it has an action of $\mathrm{PSL}_2(\mathbb{Z}/N\mathbb{Z})$ induced by the isomorphism $\mathrm{SL}_2(\mathbb{Z})/\Gamma(N) \simeq \mathrm{SL}_2(\mathbb{Z}/N\mathbb{Z})$. (Don't forget that the center of $\mathrm{SL}_2(\mathbb{Z})$ acts trivially on $\mathfrak{H}$.) This finite group acts transitively on the cells in the triangulation. For $N = 3, 4, 5$ the surfaces $X(N)$ have genus 0, and the induced triangulations are familiar to anyone who inhabits three dimensions (cf. [52]). For $N = 6$ the quotient is a torus with a triangulation consisting of 24 triangles, 36 edges, and 12 vertices. For $N = 7$ we have $|\mathrm{PSL}_2(\mathbb{Z}/7\mathbb{Z})| = 168$, and the Riemann surface $X(7)$ realizes Hurwitz's upper bound for the size of the automorphism group of a surface of genus three. It is a pleasant exercise to draw the triangulations for $N \leq 7$. The tessellated Riemann surfaces $X(N), N > 1$ are called *Platonic surfaces* [12].

Thus we can compute the right hand side of (6) if we can understand (1) a generating set for the relative homology group, and (2) all the relations between our generating set. We will only sketch what happens, since we don't need more precision for our discussion.

The first is easy. The images of the Farey edges become edges in the triangulation, and so their classes will span $H_1(X_\Gamma, \partial X_\Gamma; \mathbb{C})$. Each such edge corresponds to a pair of cusps of determinant ± 1, as in (7). We only need to work with representatives of these pairs modulo Γ, since these will give all edges in the triangulation.

The relations are also not hard to understand. They come from the finite subgroups of $\mathrm{SL}_2(\mathbb{Z})$. For instance, the subgroup generated by $\left(\begin{smallmatrix} 0 & 1 \\ -1 & 0 \end{smallmatrix}\right)$ stabilizes the edge in $\mathscr{T}$ from 0 to ∞, and this tells us how to find the boundary of this edge. The subgroup generated by $\left(\begin{smallmatrix} 0 & 1 \\ -1 & 1 \end{smallmatrix}\right)$ stabilizes Δ. This tells us how to find the boundary of Δ, and thus to compute a relation between three elements of $H_1(X_\Gamma, \partial X_\Gamma; \mathbb{C})$.

The upshot: we can compute the relative homology $H_1(X_\Gamma, \partial X_\Gamma; \mathbb{C})$ as the $\mathbb{C}$-vector space generated by certain pairs of cusps modulo Γ, divided out by certain relations imposed by the finite subgroups of $\mathrm{SL}_2(\mathbb{Z})$. For full details, including the extension to Γ with torsion, we refer to [48, 61]. Later (§9) we will see how to define an action of the Hecke operators on this model.

4 Algebraic groups and symmetric spaces

The first step in generalizing (2)–(3) is understanding exactly how the spaces arise from group theory. In fact they are examples of *locally symmetric spaces*.

Let G be the Lie group $\mathrm{SL}_2(\mathbb{R})$. The subgroup $K = \mathrm{SO}(2)$ of matrices satisfying $gg^t = \mathrm{Id}$ is maximal compact, and is the unique subgroup with this property up to G-conjugacy. The group G acts on $\mathfrak{H}$, again by fractional linear transformations (1), and the action is transitive. Indeed, the subgroup of upper-triangular matrices already acts transitively, since

$$\begin{pmatrix} \sqrt{y} & x/\sqrt{y} \\ 0 & 1/\sqrt{y} \end{pmatrix} \cdot i = x + iy.$$

The stabilizer in G of i is K, and so we have a diffeomorphism

$$G/K \xrightarrow{\sim} \mathfrak{H}.$$

This exhibits $\mathfrak{H}$ as a *Riemannian globally symmetric space* [40]. We recall that such a space is an analytic Riemannian manifold D with a family of involutive isometries $\sigma_p \colon D \to D$, one for each $p \in D$, such that p is the unique fixed point of σ_p. It is known that any such space D can be written as a quotient G/K, where G is the connected component of the group of isometries of D and $K \subset G$ is a compact subgroup stabilizing a chosen point $p_0 \in D$. If Γ is any finite index subgroup of $\mathrm{SL}_2(\mathbb{Z})$, then the quotient $\Gamma\backslash\mathfrak{H} = \Gamma\backslash G/K$ inherits this structure locally. Such double quotients are known as *locally symmetric spaces*. These spaces, and their compactifications, will be the replacements for $Y_0(N)$, $X_0(N)$.

So how do we build locally symmetric spaces? The first step is to build globally symmetric spaces, and that is the focus of this section. We begin with a *linear algebraic group* $\mathbf{G}$. This is a group that also has the structure of an affine algebraic variety, with the group operations being morphisms. For instance, the group GL_n of $n \times n$ invertible matrices can be realized as a closed subgroup of affine $n^2 + 1$-space. We take the ring

$$\mathbb{C}[x_{11}, x_{12}, \ldots, x_{1n}, x_{21}, \ldots, x_{nn}, \delta] \tag{8}$$

with variables x_{ij} corresponding to the entries of an indeterminate $n \times n$ matrix. The group GL_n is then the zero set of the polynomial $\delta \det(x_{ij}) = 1$. The group operations can be written as polynomials in these variables, so GL_n is a linear algebraic group.

More generally, $\mathbf{G}$ is a linear algebraic group if it is a subgroup of GL_n defined by polynomial equations. If the coordinate ring of $\mathbf{G}$ can be defined by an ideal generated by polynomials in a subfield $F \subset \mathbb{C}$, then we say $\mathbf{G}$ is *defined over* F.

Basic examples are the classical groups: SL_n, the subgroup of GL_n of matrices of determinant 1; SO_n, the subgroup of SL_n preserving a fixed nondegenerate symmetric bilinear form; and Sp_n, the subgroup of SL_{2n} preserving a nondegenerate alternating bilinear form.

Other examples are provided by tori. By definition $\mathbf{G}$ is a torus if $\mathbf{G} \simeq (\mathrm{GL}_1)^d$ for some d, called the *rank* of G. We note that this isomorphism need not be defined over F, even if $\mathbf{G}$ is defined over F. If it is, we say that $\mathbf{G}$ is *F-split*. The integer d is then called the *F-rank* of $\mathbf{G}$ and is denoted $r_F(\mathbf{G})$. More generally, the F-rank $r_F(\mathbf{G})$ of any algebraic group $\mathbf{G}$ is defined to be the F-rank of the maximal F-split torus in $\mathbf{G}$.

The most important linear algebraic groups for us, which are also the most familiar, are the reductive and semisimple groups. By definition, the *radical*

$R(\mathbf{G})$ is the maximal connected solvable normal subgroup of $\mathbf{G}$, where *connected* means irreducible as an algebraic variety. The *unipotent radical* $R_u(\mathbf{G})$ is the maximal connected unipotent normal subgroup of $\mathbf{G}$, where *unipotent* means all eigenvalues are 1. A group is called *reductive* if $R_u(\mathbf{G})$ is trivial, and *semisimple* if $R(\mathbf{G})$ is trivial. We have $R_u(\mathbf{G}) \subset R(\mathbf{G})$, so semisimple is a special case of reductive. Any connected group contains a reductive and semisimple quotient: if $\mathbf{G}$ is connected, then $\mathbf{G}/R_u(\mathbf{G})$ is reductive and $\mathbf{G}/R(\mathbf{G})$ is semisimple.

For example, the classical groups $\mathrm{SL}_n, \mathrm{SO}_n, \mathrm{Sp}_n$ are semisimple. The group GL_n is reductive and is not semisimple. For an example of a group that is neither reductive nor semisimple, one can take the Borel subgroup $\mathbf{B} \subset \mathrm{GL}_2$ of upper-triangular matrices. The unipotent radical $R_u(\mathbf{B})$ is the subgroup of $\mathbf{B}$ with 1s on the diagonal. This example generalizes to the subgroups $\mathbf{P} \subset \mathrm{GL}_n$ of block upper-triangular matrices; these are examples of *parabolic* subgroups. In general, even if one is ultimately interested in phenomena involving reductive groups, one must consider non-reductive groups, since they often provide an inductive tool to understand structures on reductive groups (cf. (16)).[2]

As we said above, semisimple is a special case of reductive. In fact being reductive is not that far from being semisimple. Let $\mathbf{G}$ be reductive and let $\mathbf{S}$ be the connected component of the center of $\mathbf{G}$. Then $\mathbf{S}$ is a torus. If we put $\mathbf{H} = [\mathbf{G}, \mathbf{G}]$ (the derived subgroup, which is semisimple), then

$$\mathbf{G} = \mathbf{H} \cdot \mathbf{S},$$

an *almost direct product*; this means $\mathbf{H} \cap \mathbf{S}$ is finite, not necessarily $\{1\}$.

Hence a reductive group looks like a semisimple one, up to a central torus factor. For instance, for GL_n the group $\mathbf{S}$ is the subgroup of scalar matrices $a\,\mathrm{Id}$ and the derived subgroup $\mathbf{H}$ is SL_n. Certainly $\mathrm{GL}_n = \mathbf{S} \cdot \mathrm{SL}_n$: the intersection $\mathbf{S} \cap \mathrm{SL}_n$ is the group of nth roots of unity.

Now that we've identified our groups of interest, let's explain how to find our spaces. Let $G = \mathbf{G}(\mathbb{R})$, the group of real points of $\mathbf{G}$. This has the structure of a Lie group, although not all Lie groups arise this way. Let $K \subset G$ be a maximal compact subgroup. We now have exactly the objects we needed to define $\mathfrak{H}$, and indeed we're done if $\mathbf{G}$ is semisimple: the relevant symmetric space is G/K. But if $\mathbf{G}$ is reductive we need to go further and divide by a bit more. Thus we introduce the *split component* A_G of $\mathbf{G}$. By definition A_G is the connected component of the identity of the group of real points of the maximal $\mathbb{Q}$-split torus in the center of $\mathbf{G}$. That's quite a mouthful, but it's easy to understand in examples, as we shall see. In any event, we define our symmetric space to be

$$D = G/A_G K. \tag{9}$$

[2] This is Harish-Chandra's "Philosophy of cusp forms" [39]; see also [15, Chapter 49].

One should think of this quotient as being the Lie group (G/A_G) divided by its maximal compact subgroup K.

For a first example let $\mathbf{G} = \mathrm{SL}_n$. Then $G = \mathbf{G}(\mathbb{R}) = \mathrm{SL}_n(\mathbb{R})$ and $K = \mathrm{SO}(n)$. The split component A_G is trivial, and our symmetric space is

$$D = \mathrm{SL}_n(\mathbb{R})/\mathrm{SO}(n). \tag{10}$$

If $n = 2$, then (10) becomes $\mathfrak{H}$. The dimension of D is $n(n+1)/2 - 1$. In particular note that $\dim D$ can be odd, and thus in general D does not have a complex structure. This is quite different from $\mathfrak{H}$.

Now let $\mathbf{G} = \mathrm{GL}_n$, so that $G = \mathrm{GL}_n(\mathbb{R})$. The maximal compact subgroup is $K = \mathrm{O}(n)$, which is only a little bigger than $\mathrm{SO}(n)$ (it has the same dimension as $\mathrm{SO}(n)$, just an extra component). The split component A_G consists of the real scalar matrices $\{a \,\mathrm{Id} \mid a \in \mathbb{R}_{>0}\}$. Our symmetric space is $G/A_G K$, and in fact is isomorphic to (10). Hence by dividing out by the split component we kill exactly the extra dimension we introduced by using GL_n instead of SL_n.

We now come to examples that will be our main focus, namely general linear groups over number fields. Before we can explain how they work, we need the important notion of restriction of scalars. This is a useful construction that allows us to focus our attention on groups defined over $\mathbb{Q}$, even if our group is most naturally written in terms of a bigger number field.

So suppose $\mathbf{G}$ is a linear algebraic group defined over a number field F. Then there is a group $R_{F/\mathbb{Q}}\mathbf{G}$, called the *restriction of scalars* of $\mathbf{G}$ from F to $\mathbb{Q}$, such that (i) $R_{F/\mathbb{Q}}\mathbf{G}$ is defined over $\mathbb{Q}$ and (ii)

$$(R_{F/\mathbb{Q}}\mathbf{G})(\mathbb{Q}) = \mathbf{G}(F).$$

This is something that is already familiar to you, even if you didn't realize it.[3] Consider the standard representation of the complex numbers as 2×2 real matrices:

$$a + bi \longmapsto \begin{pmatrix} a & -b \\ b & a \end{pmatrix}. \tag{11}$$

This is an example of restriction of scalars, after we pass to nonzero elements. Indeed, in that case the left hand side of (11) is the group of complex points of $\mathbf{G} = \mathrm{GL}_1$, thought of as a group defined over $\mathbb{C}$; the right hand side is the group of real points of $R_{\mathbb{C}/\mathbb{R}}\mathbf{G}$, a group clearly defined over $\mathbb{R}$. Note that $(R_{\mathbb{C}/\mathbb{R}}\mathbf{G})(\mathbb{R}) = \mathbf{G}(\mathbb{C})$.

For a more complicated example, suppose we take the group $\mathbf{G} = \mathrm{SL}_2$, again defined over $\mathbb{C}$, and build $R_{\mathbb{C}/\mathbb{R}}\mathrm{SL}_2$. We can do this using $\mathrm{GL}_4/\mathbb{R}$: simply use (11) to take the four complex entries of $\mathrm{SL}_2(\mathbb{C})$ to four 2×2 blocks A, B, C, D of a matrix in $\mathrm{GL}_4(\mathbb{R})$. The image will be determined by

[3] Like Molière's Monsieur Jourdain: "Par ma foi! il y a plus de quarante ans que je dis de la prose sans que j'en susse rien, et je vous suis le plus obligé du monde de m'avoir appris cela."

certain polynomial equations with real coefficients. Some of the equations encode the fact that the blocks arose via (11), whereas other equations come from the condition $AD - BC = \mathrm{Id}$ that defines SL_2. The general definition of restriction of scalars is no harder than this, although it's a bit messy to write out in these terms (cf. [53, §2.1.2]).

Let's build some symmetric spaces starting from number fields. As a first example, we take $F/\mathbb{Q}$ to be real quadratic and $\mathbf{G} = R_{F/\mathbb{Q}}\mathrm{SL}_2$. Then $G = \mathbf{G}(\mathbb{R}) \simeq \mathrm{SL}_2(\mathbb{R}) \times \mathrm{SL}_2(\mathbb{R})$, since $F \otimes_{\mathbb{Q}} \mathbb{R} \simeq \mathbb{R} \times \mathbb{R}$, corresponding to the two distinct embeddings of F into $\mathbb{R}$. One should think of $\mathbf{G}(\mathbb{Q}) = \mathrm{SL}_2(F)$ as mapping into $\mathbf{G}(\mathbb{R})$ via these embeddings, where we use a different one for each factor. The maximal compact subgroup K is $\mathrm{SO}(2) \times \mathrm{SO}(2)$, the split component is trivial ($\mathbf{G}$ is semisimple), and the symmetric space is

$$G/K = \mathfrak{H} \times \mathfrak{H}, \tag{12}$$

the product of upper halfplanes familiar from Hilbert modular forms [29].

Now let's try the reductive version. Put $\mathbf{G} = R_{F/\mathbb{Q}}\mathrm{GL}_2$. We have $G \simeq \mathrm{GL}_2(\mathbb{R}) \times \mathrm{GL}_2(\mathbb{R})$ and $K \simeq \mathrm{SO}(2) \times \mathrm{SO}(2)$. This time the split component A_G isn't trivial: the maximal $\mathbb{Q}$-split torus in the center of $\mathbf{G}$ has $\mathbb{Q}$-points $\{\left(\begin{smallmatrix} a & 0 \\ 0 & a \end{smallmatrix}\right) \mid a \in \mathbb{Q}^\times\}$. In $\mathbf{G}(\mathbb{R})$ this subgroup embeds the same in each factor, and after taking the connected component of the $\mathbb{R}$-points we find

$$A_G = \left\{ \left(\begin{pmatrix} a & 0 \\ 0 & a \end{pmatrix}, \begin{pmatrix} a & 0 \\ 0 & a \end{pmatrix} \right) \;\middle|\; a \in \mathbb{R}_{>0} \right\} \simeq \mathbb{R}_{>0},$$

which is one-dimensional. Counting dimensions, we see that $G/A_G K$ is *five*-dimensional. In fact, as Riemannian manifolds we have

$$G/A_G K = \mathfrak{H} \times \mathfrak{H} \times \mathbb{R}, \tag{13}$$

where $\mathbb{R}$ has the flat metric. The space (13) looks unnatural, especially to someone interested in the geometry of Hilbert modular surfaces, but as we will see later the "flat factor" is very convenient.

Now let F be imaginary quadratic and let $\mathbf{G} = R_{F/\mathbb{Q}}\mathrm{SL}_2$. Since $F \otimes \mathbb{R} \simeq \mathbb{C}$, we have $G = \mathrm{SL}_2(\mathbb{C})$. The maximal compact subgroup K is $\mathrm{SU}(2)$ and the split component A_G is trivial. The symmetric space is now $\mathfrak{H}_3$, three-dimensional hyperbolic space. If we take instead $\mathbf{G} = R_{F/\mathbb{Q}}\mathrm{GL}_2$, we find $G = \mathrm{GL}_2(\mathbb{C})$ and $K = \mathrm{U}(2)$. The maximal $\mathbb{Q}$-split torus in the center of $\mathbf{G}$ is again $\{\left(\begin{smallmatrix} a & 0 \\ 0 & a \end{smallmatrix}\right) \mid a \in \mathbb{Q}^\times\}$, so $A_G = \{\left(\begin{smallmatrix} a & 0 \\ 0 & a \end{smallmatrix}\right) \mid a \in \mathbb{R}_{>0}\}$. Thus the symmetric space is again $\mathfrak{H}_3$. Again, the situation is similar to the case of $F = \mathbb{Q}$: there is no flat factor, and replacing semisimple with reductive doesn't change the symmetric space.

Comparing the cases of F real/imaginary quadratic and the rationals suggests that the flat factors have something to do with the rank of $\mathscr{O}_F^\times$, the units of the integers of F. This is in fact true. Suppose $F \otimes \mathbb{R} \simeq \mathbb{R}^r \times \mathbb{C}^s$.

- If $\mathbf{G} = R_{F/\mathbb{Q}}\mathrm{SL}_2$, then

$$\begin{aligned} G &\simeq \mathrm{SL}_2(\mathbb{R})^r \times \mathrm{SL}_2(\mathbb{C})^s, \\ K &\simeq \mathrm{SO}(2)^r \times \mathrm{SU}(2)^s, \\ A_G &\simeq \{1\}, \\ G/K &\simeq \mathfrak{H}^r \times \mathfrak{H}_3^s. \end{aligned}$$

- If $\mathbf{G} = R_{F/\mathbb{Q}}\mathrm{GL}_2$, then

$$\begin{aligned} G &\simeq \mathrm{GL}_2(\mathbb{R})^r \times \mathrm{GL}_2(\mathbb{C})^s, \\ K &\simeq \mathrm{O}(2)^r \times \mathrm{U}(2)^s, \\ A_G &\simeq \mathbb{R}_{>0}, \\ G/A_G K &\simeq \mathfrak{H}^r \times \mathfrak{H}_3^s \times \mathbb{R}^{r+s-1}. \end{aligned}$$

  We can see that the dimension of the flat factor is the same as the rank of $\mathscr{O}_F^\times$.

5 Arithmetic groups, locally symmetric spaces, and cohomology

We now have an analogue of the upper halfplane $\mathfrak{H}$, namely our globally symmetric space $D = G/A_G K$. To build locally symmetric spaces, the analogues of the open modular curves, we need to get discrete subgroups into the picture. This brings us to arithmetic groups.

Let $\mathbf{G}$ be a linear algebraic group. Then a subgroup $\Gamma \subset \mathbf{G}$ is an *arithmetic group* if it is commensurable with $\mathbf{G}(\mathbb{Z})$. This means the intersection $\Gamma \cap \mathbf{G}(\mathbb{Z})$ has finite index in both Γ and $\mathbf{G}(\mathbb{Z})$.[4]

For instance, suppose $\mathbf{G}$ is GL_n. Then $\mathbf{G}(\mathbb{Z})$ is $\mathrm{GL}_n(\mathbb{Z})$, the group of invertible integral matrices. If we put $\mathbf{G} = R_{F/\mathbb{Q}}\mathrm{GL}_n$, then we find $\mathbf{G}(\mathbb{Z}) = \mathrm{GL}_n(\mathscr{O})$. We can make further examples by taking quotients. For instance, if we choose an ideal $I \subset \mathscr{O}$, we can consider the map $\mathrm{GL}_n(\mathscr{O}) \to \mathrm{GL}_n(\mathscr{O}/I)$. The kernel is a subgroup of finite index in $\mathrm{GL}_n(\mathscr{O})$ called a *congruence subgroup*.

Given an arithmetic group Γ we can form the quotient

$$Y_\Gamma = \Gamma\backslash D = \Gamma\backslash G/A_G K.$$

The space Y_Γ is a *locally symmetric space*. This is our replacement for the open modular curve. We propose to study

[4] This definition of arithmetic group suffices for our purposes because we have defined our algebraic groups as subgroups of GL_n. If one works more abstractly, then the correct condition is that $\Gamma \subset \mathbf{G}(\mathbb{Q})$ is arithmetic if for any $\mathbb{Q}$-embedding $\iota\colon \mathbf{G} \to \mathrm{GL}_n$, the group $\iota(\Gamma)$ is commensurable with $\iota(\mathbf{G}) \cap \mathrm{GL}_n(\mathbb{Z})$.

$$H^*(Y_\Gamma; \mathbb{C}); \tag{14}$$

classes in these spaces will be our analogue of holomorphic modular forms of weight two.

What about higher weight modular forms? We can find analogues of these as well, if we're willing to work with fancier cohomology. Let (ρ, M) be a finite-dimensional (complex) rational representation of G. The reader should think of the case of G a classical matrix group and ρ a classical polynomial representation. We get a representation of Γ in M. If Γ is torsionfree, then the fundamental group of Y_Γ is Γ. The representation $\rho\colon \Gamma \to \mathrm{GL}(M)$ thus induces a local coefficient system $\mathscr{M}$ on Y_Γ. We can then form the cohomology spaces

$$H^*(Y_\Gamma; \mathscr{M}). \tag{15}$$

For more details about local coefficients see [13, 60, 63]. This construction works even if Γ has torsion, although the quotient is an orbifold, not a manifold. Nevertheless cohomology with coefficients in a local system still makes sense for such objects. If $G = \mathrm{SL}_2(\mathbb{R})$ and $\Gamma \subset \mathrm{SL}_2(\mathbb{Z})$, this construction is exactly what we did to express higher weight forms in terms of cohomology of the modular curve, cf. (4)–(5). In that case the degree k homogeneous polynomials are an incarnation of the standard representation of SL_2 of dimension $k+1$.

We claim that the cohomology spaces (15), which include those in (14) as a special case, provide a means to compute certain automorphic forms explicitly. Certainly it is not clear that cohomology has anything to do with automorphic forms, although (4)–(5) give some evidence. Justifying this relationship in detail would take us well beyond the scope of these lectures; for an excellent discussion of the connection, we refer to [43, 55, 64]. What we can say is the following:

1. According to a deep theorem of Franke [28], which proved a conjecture of Borel, the cohomology groups $H^*(Y_\Gamma; \mathscr{M})$ can be directly computed in terms of certain automorphic forms (those that are "cohomological," also known as those with "nonvanishing $(\mathfrak{g}, K)$ cohomology" [65]).
2. There is a direct sum decomposition

$$H^*(Y_\Gamma; \mathscr{M}) = H^*_{\mathrm{cusp}}(Y_\Gamma; \mathscr{M}) \oplus \bigoplus_{\{\mathbf{P}\}} H^*_{\{\mathbf{P}\}}(Y_\Gamma; \mathscr{M}), \tag{16}$$

where the sum is taken over the set of classes of *associate proper* $\mathbb{Q}$-*parabolic subgroups of* $\mathbf{G}$. (cf. [42, Chapter 2])

The summand $H^*_{\mathrm{cusp}}(Y_\Gamma; \mathscr{M})$ of (16) is called the *cuspidal cohomology*; this is the subspace of classes represented by cuspidal automorphic forms. The remaining summands constitute the *Eisenstein cohomology* of Γ [38]. In particular the summand indexed by $\{\mathbf{P}\}$ is constructed using Eisenstein series attached to certain cuspidal automorphic forms on lower rank groups; one should compare (16) with (4). Hence $H^*_{\mathrm{cusp}}(Y_\Gamma; \mathscr{M})$ is in some sense the most important part of the cohomology: all the rest can be built systematically from cuspidal cohomology on lower rank groups.[5]

We emphasize that the cohomological automorphic forms are a very special subset of all the automorphic forms, and that in some sense the typical automorphic form will not contribute to the cohomology of an arithmetic group. For $\mathrm{SL}_2/\mathbb{Q}$, for example, it is only the holomorphic modular forms of weights ≥ 2 that appear. Neither the (real-analytic) Maass forms nor the weight 1 holomorphic forms are cohomological.

The underlying reason comes from the infinite-dimensional representation theory of $\mathrm{SL}_2(\mathbb{R})$. We can only sketch the connection here; for more details, including undefined terms, see [14, 30]. Any automorphic form f on $\mathrm{SL}_2/\mathbb{Q}$ gives rise to an *automorphic representation* π. This is a certain subquotient of $L^2(\mathrm{SL}_2(\mathbb{Q})\backslash\mathrm{SL}_2(\mathbb{A}_\mathbb{Q}))$, where $\mathbb{A}_\mathbb{Q}$ is the adele ring of $\mathbb{Q}$. The representation π factors as a restricted tensor product of local representations

$$\pi_\infty \otimes \bigotimes_{p \text{ prime}} \pi_p.$$

The factor π_∞ is a unitary representation of $\mathrm{SL}_2(\mathbb{R})$. Apart from the trivial representation, the irreducible unitary representations of $\mathrm{SL}_2(\mathbb{R})$ come in four families:

1. The principal series.
2. The discrete series.
3. The limits of discrete series.
4. The complementary series.

Which of these occur as π_∞ depends on what f is. If f is a Maass form, then π_∞ is principal series. If f is holomorphic, then π_∞ is either discrete series ($k \geq 2$) or a limit of discrete series ($k = 1$). The complementary series do not appear as π_∞. Only the discrete series are cohomological, which is why we only see holomorphic modular forms of weights ≥ 2 is the cohomology of the modular curves.

Since many—indeed most—automorphic forms are not cohomological, why do we study cohomological forms? Here is one answer. Our ultimate goal is not to study automorphic forms for their own sake, but instead to pursue links between automorphic forms and arithmetic. The standard example occurs in every course on modular forms: the mysterious connection between counting points on elliptic curves over prime fields and computing Hecke eigenvalues of weight two holomorphic cusp forms. In general one expects that certain automorphic forms on a reductive group $\mathbf{G}$ should have connections to arithmetic geometry (Galois representations). These connections are revealed through

[5] In fact the cuspidal cohomology can itself come from groups of lower rank, through functorial liftings. The paper [6] contains evidence of cohomological lifts of paramodular forms on $\mathrm{Sp}_4/\mathbb{Q}$ to $\mathrm{SL}_4/\mathbb{Q}$.

the Hecke eigenvalues. The cohomology of arithmetic groups gives us a way to get our hands on some automorphic forms, and these forms are among those predicted to be related to arithmetic. Some forms we'd like to see will be missing (e.g. weight 1 holomorphic forms), but in any event the cohomological automorphic forms are a natural and tractable class to investigate.

6 Reduction theory I: the rational numbers

At this point we have found our analogues of the modular curves, the locally symmetric spaces. We need to compute their cohomology. As everyone learns in algebraic topology courses, to compute cohomology one needs a cochain complex, and a first step to finding this is forming a cell decomposition of the underlying space. For the spaces $Y(N), X(N)$, this can be done using the Farey tessellation of the upper halfplane. In fact, since the Farey tessellation is $\mathrm{SL}_2(\mathbb{Z})$-invariant, it leads to a nice triangulation of the quotient by *any* finite-index torsionfree subgroup of $\mathrm{SL}_2(\mathbb{Z})$. Even if a subgroup has torsion, as $\Gamma_0(N)$ typically does, one can still use the Farey tessellation to compute cohomology. One can take the barycentric subdivision, or can work with more sophisticated techniques that incorporate the torsion (cf. §8). Thus, for modular curves, one has a powerful tool to compute cohomology.

Unfortunately, for general locally symmetric spaces $\Gamma\backslash D$ the situation is not as nice. In fact we don't know a good way to construct Γ-invariant subdivisions of D for an arbitrary symmetric space! But all is not lost: we have one general tool, a tool that works in the important case of GL_n over number fields. The construction has its origin in Voronoi's work on reduction theory for positive-definite quadratic forms [68], which we discuss in this section as a warm-up. We treat the case of general number fields in §7.

Let us reconsider the setting of §§2–3 and show a different way to build the Farey tessellation. Let $V = \mathrm{Sym}_2(\mathbb{R})$ be the three-dimensional vector space of 2×2 real symmetric matrices. Inside V we have the subset C of positive-definite matrices. The set C is a convex cone: if $x \in C$ then so is ρx for any $\rho \in \mathbb{R}_{>0}$, and if $x, y \in C$ so is $x + y$. The vector space V comes with an inner product

$$\langle x, y\rangle = \mathrm{Trace}(xy), \tag{17}$$

and C is self-adjoint with respect to this product, namely

$$C = C^* = \{y \in V \mid \langle x, y\rangle > 0 \text{ for all } x \in V\}.$$

The group $G = \mathrm{SL}_2(\mathbb{R})$ acts on V by $(g, x) \mapsto gxg^t$, and this action preserves C. The stabilizer of any point is a conjugate of $K = \mathrm{SO}(2)$. The G-action is not transitive, so we can't identify C with $G/K = \mathfrak{H}$, but the action is transitive after we mod out C by homotheties. This leads to an identification

$$C/\mathbb{R}_{>0} \xrightarrow{\sim} \mathfrak{H}, \tag{18}$$

which is compatible with the action of $\mathrm{SL}_2(\mathbb{Z})$ on both sides. The cone C is an example of a real *self-adjoint homogeneous cone* [26]. Figure 2 shows C, where we have used the coordinates $\left(\begin{smallmatrix} x & z \\ z & y \end{smallmatrix}\right)$. The cone is determined by the inequalities $xy - z^2 > 0, x > 0$. We have also indicated a few forms in C and its closure.

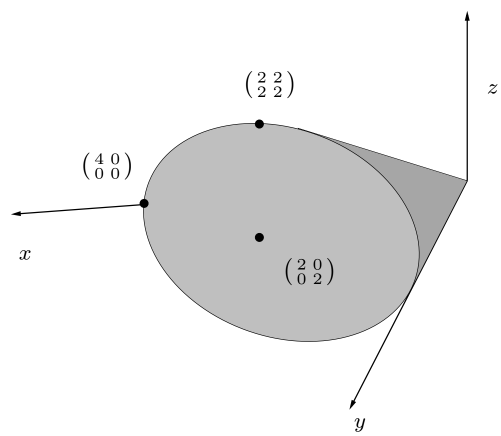

Fig. 2 The cone of positive definite binary quadratic forms. The ellipse is the slice of constant trace 4.

Now consider the closure $\bar{C}$ of C. This consists of certain rank 1 symmetric matrices, as well as the unique rank 0 symmetric matrix. Consider the lattice $\mathbb{Z}^2$, which we write as column vectors. We have a map $q\colon \mathbb{Z}^2 \to \bar{C}$ given by $q(x) = xx^t$. Restricting q to $\mathbb{Z}^2 \smallsetminus \{0\}$, we obtain a collection of nonzero points Ξ in the boundary $\partial C = \bar{C} \smallsetminus C$. The image is discrete since it lies in the lattice of integral symmetric matrices. Furthermore the action of $\mathrm{SL}_2(\mathbb{Z})$ on V induces an action on these points. As we shall see, the points Ξ are almost exactly the vertices of the Farey tessellation.

What makes (18) so useful is that C gives a *linear* model of $\mathfrak{H}$, apart from the mild complication of the homotheties. In particular, given the linear structure and convexity of C and the collection of points Ξ, the geometer's next step is irresistible: take the convex hull Π of Ξ. The result is a huge polyhedron equipped with an action of $\mathrm{SL}_2(\mathbb{Z})$. Of course, there is no reason a priori that Π is a nice object, or is even computable in any reasonable sense.

Fortunately for us, this is not the case. The polyhedron Π is very nice, with a beautiful combinatorial structure. It has the deficiency of not being locally finite (each vertex meets infinitely many edges), but its facets (top-dimensional proper faces) are finite polytopes, in fact triangles. Moreover, after modding out by homotheties and applying (18), the proper faces of Π become the vertices, arcs, and triangles of the Farey tessellation. The connection can be understood as in §3. Suppose $(a,b)^t$ and $(c,d)^t$ are primitive vectors giving cusps at the endpoints of an arc in $\mathscr{T}$. Then this arc is exactly the image of the edge of Π between $q(a,b)$ and $q(c,d)$. Note how useful the homotheties are in (18). The map q "lifts" the nonzero integral points $(a,b)^t$, primitive or not, up along ∂C (cf. Figure 3). When forming Π by taking the convex hull, we group the lifts nontrivially into the faces of Π. The projection back down to $\bar{C}/\mathbb{R}_{>0} \xrightarrow{\sim} \mathfrak{H} \cup \mathbb{P}^1(\mathbb{R})$ then recovers the Farey triangles. From this many properties of Π become clear:

- Modulo the action of $\mathrm{SL}_2(\mathbb{Z})$, there are only finitely many vertices, edges, and triangles in Π.
- Every edge meets finitely many triangles (namely two), but every vertex meets infinitely many edges. Thus the polyhedron fails to be locally finite, but only "at infinity."

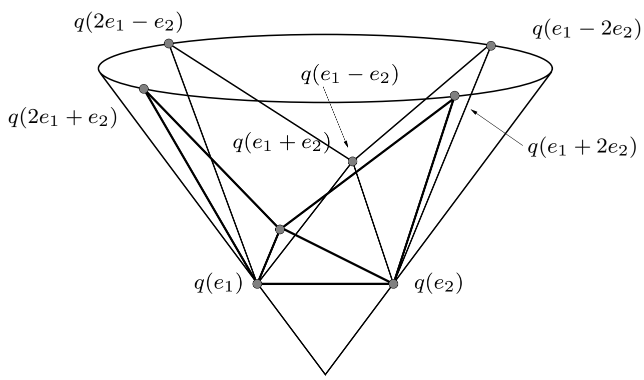

Fig. 3 A few facets of the Voronoi polyhedron Π for $\mathrm{SL}_2(\mathbb{Z})$.

Now we turn to higher rank. Let $n \geq 2$, let V be $\mathrm{Sym}_n(\mathbb{R})$, the real vector space of $n \times n$ symmetric matrices. Let $C \subset V$ be the convex cone of positive definite matrices. Everything we did before goes through without trouble. Again the group $G = \mathrm{SL}_n(\mathbb{R})$ acts on C by $(g,x) \mapsto gxg^t$. The quotient of C by homotheties is isomorphic to the symmetric space $D = \mathrm{SL}_n(\mathbb{R})/\mathrm{SO}(n)$, where G acts on D by left translations. We have a map $q\colon \mathbb{Z}^n \smallsetminus \{0\} \to \bar{C}$ determining a point set $\Xi \subset \partial C$. The convex hull of Ξ is a polyhedron Π, called the *Voronoi polyhedron*. By construction $\mathrm{SL}_n(\mathbb{Z})$ acts on Π. The cones on the faces of Π descend to form cells in $\bar{D}$, where the latter is a certain natural compactification of D.

Voronoi defined and studied Π because he was interested in the reduction theory of positive-definite quadratic forms, which essentially boils down to finding a nice fundamental domain of $\mathrm{SL}_n(\mathbb{Z})$ acting on C. To explain his results we need the concept of a perfect form. A recent treatment of Voronoi's work can be found in [54, Chapter 3].

Let $A \in C$, and let Q_A be the corresponding positive-definite quadratic form. Given a point $x \in \mathbb{R}^n$, regarded as a column vector, we can evaluate Q_A on it in a variety of ways:

$$Q_A(x) = \sum_{i,j} A_{ij} x_i x_j = x^t A x = \langle x x^t, A \rangle.$$

Thus if $x \in \mathbb{Z}^n$, then the inner product $\langle q(x), A \rangle$ gives the value of Q_A on the lattice point x. The *minimum* $m(A)$ of Q_A is by definition the minimum of $\langle q(x), A \rangle$ as x ranges over all points in $\mathbb{Z}^n \smallsetminus \{0\}$. The set $M(A)$ of *minimal vectors* is the subset of $\mathbb{Z}^n$ on which the minimum is attained. A quadratic form is called *perfect* if it can be reconstructed from the knowledge of its minimum and its minimal vectors. For instance, the binary quadratic form $x^2 + xy + y^2$ is perfect, whereas $x^2 + y^2$ is not. Whether or not a form is perfect is unchanged under homothety.

Voronoi proved that modulo $\mathrm{SL}_n(\mathbb{Z})$, the polyhedron Π has finitely many faces. He also proved that the facets of Π are in bijection with the homothety classes of perfect quadratic forms. Under this bijection, if F is a facet of Π with vertices $\xi_1, \dots, \xi_k$, then the inverse images of the ξ_i in $\mathbb{Z}^n$ are the minimal vectors of a form in the corresponding class.

He even gave an algorithm that, starting with an initial perfect form, produces a list of perfect forms modulo $\mathrm{SL}_n(\mathbb{Z})$, and used it to compute perfect forms for $n \leq 5$. Today we have a good understanding of the combinatorics of Π up to $n = 7$. For $n = 8$ we know the $\mathrm{SL}_8(\mathbb{Z})$-orbits of the facets of Π, but a notorious bugaboo living in dimension 8 challenges further progress: the E_8 root lattice [22]. This lattice gives rise to a perfect form whose corresponding facet of Π contains more that 2.5×10^{14} maximal faces!

Voronoi's algorithm produces, as a by-product, an explicit reduction theory for C. Consider the collection of cones Σ in $\bar{C}$ obtained by taking the cones on the faces of Π. Modulo $\mathrm{SL}_n(\mathbb{Z})$ there are only finitely many cones in Σ, and one can prove that if a cone meets C then its stabilizer in $\mathrm{SL}_n(\mathbb{Z})$ is finite. Thus the top dimensional cones in Σ are very close to fundamental domains of $\mathrm{SL}_n(\mathbb{Z})$; we saw this already for $n = 2$, where each Farey triangle was a union of three fundamental domains of $\mathrm{SL}_2(\mathbb{Z})$. In particular, if $\Gamma \subset \mathrm{SL}_n(\mathbb{Z})$ is torsionfree, then one can make a fundamental domain for Γ by taking a union of finitely many closed top-dimensional cones in Σ. In any case, whether Γ has torsion or not, one can show that any point $x \in C$ lies in a unique cone $\sigma(x) \in \Sigma$. It turns out that Voronoi's algorithm to enumerate perfect forms leads to an algorithm that finds $\sigma(x)$.

7 Reduction theory II: general number fields

Voronoi's theory is fantastic for the rational numbers, but what about other number fields? The good news is that we have tools for explicit reduction theory there as well, thanks to work of Ash [2] and Koecher [41]. The former was developed in the context of compactifying hermitian locally symmetric spaces [1], and is in some ways more similar to Voronoi's original theory. The latter sacrifices some of the structure of the former, but has the advantage that it works over any number field: it can be Galois or not, CM or not, totally real or not, and so on.

Let $F/\mathbb{Q}$ be a number field of degree $d = r+2s$, where $F \otimes \mathbb{R} \simeq \mathbb{R}^r \times \mathbb{C}^s$. Let $\mathscr{O} = \mathscr{O}_F$ be the ring of integers of F. Our goal is to compute the cohomology of $\mathrm{GL}_n(\mathscr{O})$ and its congruence subgroups by building cell decompositions of locally symmetric spaces as in §6. It turns out that this can be done in a straightforward way. There are just a few differences from the rational case.

First we need a vector space and a cone. The field F has r real embeddings and s complex conjugate pairs of complex embeddings. For each pair of complex conjugate embeddings, choose and fix one. We can then identify the infinite places of F with its real embeddings and our choice of complex embeddings. For each infinite place v of F, let V_v be the real vector space of $n \times n$ real symmetric (respectively, of complex Hermitian) matrices $\mathrm{Sym}_n(\mathbb{R})$ (resp., $\mathrm{Herm}_n(\mathbb{C})$) if v is real (resp., complex). Let C_v be the corresponding cone of positive definite (resp., positive Hermitian) forms. Put $V = \prod_v V_v$ and $C = \prod_v C_v$, where the products are taken over the infinite places of F. We equip V with the inner product

$$\langle x, y \rangle = \sum_v c_v \,\mathrm{Trace}(x_v y_v), \tag{19}$$

where the sum is taken over the infinite places of F, and c_v equals 1 for v real and 2 for v complex. Once again, the cone C is self-adjoint with respect to this inner product.

Koecher calls C a *positivity domain*; one can regard it as the cone of real-valued positive quadratic forms over F in n-variables. Specifically, if $A \in C$ is a tuple (A_v), then A determines a quadratic form Q_A on F^n by

$$Q_A(x) = \sum c_v x_v^* A_v x_v,$$

where c_v is defined in (19) and $*$ denotes transpose if v is real, and conjugate transpose if v is complex. Such forms are sometimes called *Humbert forms* in the literature. Note that we do not require that (A_v) arises from a matrix with entries in F via the embedding $F \to F \otimes \mathbb{R}$. Instead each A_v is an independent matrix in its C_v.[6]

[6] Although this construction sounds strange, we shall see that it is a reasonable notion of forms over F. Not every quadratic form of interest comes from a matrix A that is the

The group $G = \mathrm{GL}_n(\mathbb{R})^r \times \mathrm{GL}_n(\mathbb{C})^s$ acts on V by

$$(g \cdot y)_v = \begin{cases} g_v y_v g_v^t & v \text{ real,} \\ g_v y_v \bar{g}_v^t & v \text{ complex.} \end{cases}$$

This action preserves C, and exhibits G as the full automorphism group of C. In fact, we can identify the quotient $C/\mathbb{R}_{\geq 0}$ of C by homotheties with the globally symmetric space $D = G/KA_G$, where $K \simeq \mathrm{O}(n)^r \times \mathrm{U}(n)^s$ is a maximal compact subgroup of G and A_G is the split component, cf. §4.

Now we construct a subset $\Xi \subset \partial\bar{C}$. We can do almost exactly what we did before; we just have to take the different embeddings of F into account. In particular, the nonzero (column) vectors $\mathscr{O}^n \smallsetminus \{0\}$ determine points in V via

$$q\colon x \longmapsto (x_v x_v^*). \tag{20}$$

We have $q(x) \in \partial C$ for all $x \in \mathscr{O}^n$, and we define

$$\Xi = \bigl\{ q(x) \bigm| x \in \mathscr{O}^n \smallsetminus \{0\} \bigr\}.$$

These points play the same role for the forms in C that the lattice points did for positive-definite quadratic forms in §6. In particular, they lead to a notion of perfection. Given $A \in C$ we define its *minimum* to be

$$m(A) = \inf_{\xi \in \Xi} \langle \xi, A \rangle$$

and its *minimal vectors* by

$$M(A) = \{ \xi \in \Xi \mid \langle \xi, A \rangle = m(A) \}.$$

A form is called *perfect* if it can be recovered from the knowledge of $m(A)$ and $M(A)$.

With this set up, Koecher proves that perfect forms exist, and that every perfect form has finitely many minimal vectors. Given a perfect form A, let $\sigma(A) \subset \bar{C}$ be the cone

$$\sigma(A) = \bigl\{ \sum \rho_\xi \xi \bigm| \xi \in M(A), \rho_\xi \geq 0 \bigr\}.$$

Koecher calls $\sigma(A)$ a *perfect pyramid*, and proves that they behave almost identically to Voronoi's perfect cones:

1. Any compact subset of C meets only finitely many perfect pyramids.
2. Two different perfect pyramids have no interior point in common.

---

image of a matrix from F under the embeddings; in particular the perfect forms defined in this section usually do not come from a matrix over F.

3. Given any perfect pyramid σ, there are only finitely many perfect pyramids σ' such that $\sigma \cap \sigma'$ contains a point of C (which, by item (2), must lie on the boundaries of σ, σ').
4. The intersection of any two perfect pyramids is a common face of each.
5. Let σ be a perfect pyramid a τ and codimension one face of σ. If τ meets C, then there is another perfect pyramid σ' such that $\sigma \cap \sigma' = \tau$.
6. We have $\bigcup_{\sigma \in \Sigma} \sigma \cap C = C$.

Now we bring our discrete group into the picture. The group $\mathrm{GL}_n(\mathscr{O})$ acts on C and takes Ξ into itself. It clearly acts on the set of perfect pyramids and thus on the cones in Σ. Koecher proves that there are finitely many $\mathrm{GL}_n(\mathscr{O})$-orbits in Σ, and that each $\sigma \in \Sigma$ that meets C has at worst a finite stabilizer. Quotienting the entire picture out by homotheties, we wind up with a picture exactly analogous to the Farey tessellation of the upper halfplane, and we can use the resulting decomposition of the symmetric space D to compute cohomology of finite-index subgroups of $\mathrm{GL}_n(\mathscr{O})$.

We conclude this discussion with two points. First, it is essential that we use GL_n instead of SL_n. To pass from C to the SL-symmetric space D_{SL}, we would need to divide *each factor* C_v by $\mathbb{R}_{>0}$, not just the product. In fact, one passes from D_{GL} to D_{SL} by dividing out by the group of real points of the unit group. This kills the flat factor (and explains why it has the same dimension as the rank of $\mathscr{O}^{\times}$). But we only know how to do the explicit reduction theory for the full cone C, and so we have to keep the flat factor. Second, when $F = \mathbb{Q}$, we can construct the cones in Σ by taking cones on the faces of the Voronoi polyhedron Π. For general F, we can define the *Koecher polyhedron* to be the convex hull of Ξ. It's not hard to see that any perfect form gives rise to a facet of Π, but the converse is not clear. One needs to know that the perfect pyramids have no *dead ends* (cf. [23, §3]). More discussion can be found in [36, §2].

8 The cohomological dimension and spines

In §7 we explained how to find the analogues of the Farey tessellations for GL_n over number fields. In this section we want to explain how to use them to compute cohomology.

Let Γ be an arithmetic group in a reductive $\mathbb{Q}$-group $\mathbf{G}$. Assume for the moment that Γ is torsionfree. Let $D = G/KA_G$ be the global symmetric space, and let Y_Γ be the locally symmetric space $\Gamma\backslash D$. Let M be a $\mathbb{Z}[\Gamma]$-module and let $\mathscr{M}$ be the associated local system on Y_Γ. Then the *cohomological dimension* of Γ is defined to be the smallest integer ν such that $H^i(Y_\Gamma; \mathscr{M}) = 0$ for all $\mathscr{M}$ [56, §1.2]. We extend this to Γ with torsion by defining the *virtual cohomological dimension* $\mathrm{vcd}(\Gamma)$ to be the cohomological dimension of any finite-index torsionfree subgroup of Γ [56, §1.8]. One can show that this is well-defined.

It turns out that one can compute the virtual cohomological dimension for any arithmetic group. By a result of Borel–Serre [10, Theorem 11.4.4], we have

$$\operatorname{vcd}(\Gamma) = \dim(D) - r_{\mathbb{Q}}(\mathbf{G}/R(\mathbf{G})), \tag{21}$$

where $R(\mathbf{G})$ is the radical. In general this is less than the dimension of D, which is the same as the dimension of Y_Γ. For instance, if $\mathbf{G} = \mathrm{SL}_2/\mathbb{Q}$, then $D = \mathfrak{H}$, which has (real) dimension 2, and $r_{\mathbb{Q}}(\mathbf{G}) = 1$ (the radical is trivial since $\mathbf{G}$ is semisimple). Thus the cohomology of the open modular curve $\Gamma\backslash\mathfrak{H}$ vanishes in degrees > 1. For $\mathbf{G} = R_{F/\mathbb{Q}}\mathrm{GL}_n$, where $F \otimes \mathbb{R} \simeq \mathbb{R}^r \times \mathbb{C}^s$, we have

$$\dim(D) = r \cdot \frac{n(n+1)}{2} + s \cdot n^2 - 1, \quad r_{\mathbb{Q}}(\mathbf{G}/R(\mathbf{G})) = n - 1.$$

For some examples of these numbers, see Table 1 on page 40.

For our main groups of interest, there is way to understand geometrically why $\operatorname{vcd}(\Gamma)$ should be given by (21). We first consider the simplest case: $\mathbf{G} = \mathrm{SL}_2/\mathbb{Q}$. Consider the Farey tessellation of $\mathfrak{H}$. Inside the tessellation one can find a regular 3-tree W that's dual to the tessellation (Figure 4). The vertices of W lie at the $\mathrm{SL}_2(\mathbb{Z})$-translates of $\omega = e^{2\pi i/3}$, and the edges meet the edges of the Farey tessellation at the $\mathrm{SL}_2(\mathbb{Z})$-translates of i. The tree, like the tessellation, is an $\mathrm{SL}_2(\mathbb{Z})$-equivariant collection of cells, but it has an advantage over the tessellation: modulo any finite index subgroup $\Gamma \subset \mathrm{SL}_2(\mathbb{Z})$, the tree is *compact*. For instance, if one takes $\Gamma = \Gamma(N)$ for $N = 3, 4, 5$ and computes $\Gamma\backslash W$, the result is once again very familiar (cf. [52]).

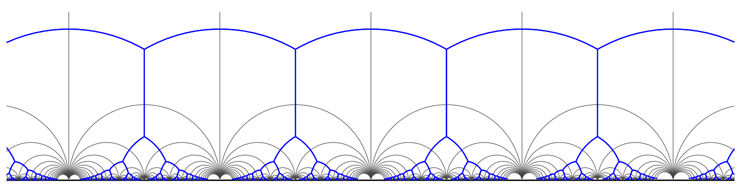

Fig. 4 The retract inside $\mathfrak{H}$.

But there is still more: not only is W modulo Γ compact, it is actually an $\mathrm{SL}_2(\mathbb{Z})$-equivariant deformation retract of $\mathfrak{H}$. In other words, there is a continuous map $r\colon \mathfrak{H} \to W$ that is $\mathrm{SL}_2(\mathbb{Z})$-equivariant and is the identity when restricted to W. This means

$$H^*(Y_\Gamma; \mathscr{M}) \xrightarrow{\sim} H^*(\Gamma\backslash W; \mathscr{M})$$

for any finite-index subgroup Γ. Since $\Gamma\backslash W$ is compact and of lower dimension than Y_Γ, it is easier to work with.

This motivates the following definition. Let D be a symmetric space acted on by an arithmetic group Γ. We say that a subspace $W \subset D$ is a *spine* for Γ if the following hold:

1. There is a Γ-equivariant deformation retraction of D onto W.
2. W is a locally finite regular cell complex of dimension $\mathrm{vcd}(\Gamma)$.
3. Γ acts on W with finite stabilizers, and modulo Γ there are only finitely many cells in W.

Spines are only known for a few symmetric spaces, and almost all of those are some form of GL_n. For $\mathrm{GL}_2/\mathbb{Q}$ the existence of a spine is classical. For $\mathrm{GL}_n/\mathbb{Q}$, $n \geq 3$, spines were constructed by Ash, Soulé, and Lannes–Soulé [2, 3, 59]. For GL_2/F when F is imaginary quadratic, spines were built by Mendoza [49], Flöge [27], and Vogtmann [66]. The most general construction along these lines is due to Ash, and is known as the *well-rounded retract* [3]. It treats $\mathbf{G}$ such that $\mathbf{G}(\mathbb{Q}) = \mathrm{GL}_n(A)$, where A is a division algebra over $\mathbb{Q}$. This includes $R_{F/\mathbb{Q}}\mathrm{GL}_n$. Outside of these cases, spines are only known sporadically. Yasaki proved a spine exists when $\mathbf{G}$ has $\mathbb{Q}$-rank one and gave complete details for $\mathrm{SU}(2,1)$ over $\mathbb{Q}(\sqrt{-1})$ [69, 70]. McConnell–MacPherson [46, 47] constructed a spine for $\mathrm{Sp}_4/\mathbb{Q}$.

Ash's construction [3] provides a spine for $R_{F/\mathbb{Q}}\mathrm{GL}_n$; in fact he builds an $(h-1)$-parameter family of spines, where h is the class number of $\mathscr{O}$. For our purposes we prefer to follow an idea in Ash's earlier paper [2], which gives a spine beginning from a Koecher-like decomposition. This has the advantage that the resulting spine is clearly dual to the cones in Koecher fan, just like the tree is dual to the Farey tessellation. We only sketch the construction here, and leave the details to the reader.

Consider the 1-dimensional cones in the Koecher fan Σ. Each contains a distinguished point, namely the first point in Ξ that lies on it. We call this point a *spanning point* of the 1-cone, and thus for any cone in Σ we can speak of its spanning points. Using the spanning points, we can form the barycentric subdivision of the cones in Σ to make a new fan $\tilde{\Sigma}$. Note that $\mathrm{GL}_n(\mathscr{O})$ acts on $\tilde{\Sigma}$.

The fan Σ has the property that any cone of dimension $n-1$ cannot meet C, i.e. such cones lie in ∂C. We define a subcone $W' \subset C$ by taking the union of all cones in $\tilde{\Sigma}$ that are contained entirely in C. We claim that $W = W'/\mathbb{R}_{>0}$ gives the spine in D.

Figure 5 illustrates this for $\mathrm{SL}_2/\mathbb{Q}$, after we've modded out by homotheties; thus this represents a "cross-section" of what's happening in the cone C. The triangle on the left is taken from the Farey tessellation and has vertices at infinity. The triangle in the middle has been barycentrically subdivided. The heavy lines in the triangle on the right are pieces of the retract W. Note that the edges in the spine are actually unions of cells from the barycentric subdivision. This is what happens in general: the cells in W will be glued together from cells arising from $\tilde{\Sigma}$ mod homotheties.

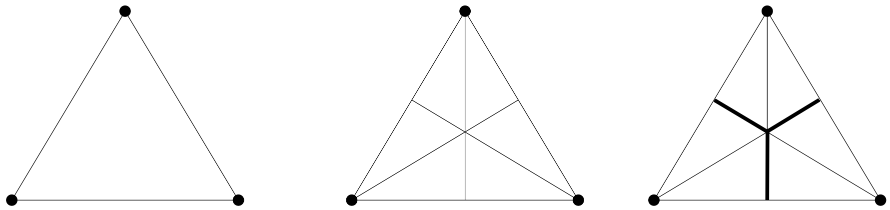

Fig. 5 Subdividing the Farey tessellation to make a spine.

From this construction it is not hard to see that W meets all the criteria to be a spine. For instance, the retraction is done piecewise-linearly, simply by appropriately projecting within each cone in $\tilde{\Sigma}$ (cf. [2, §3]). The stabilizer of a cell in W is the same as the stabilizer of the dual cone in Σ. These stabilizers are finite, since the dual cones meet C.

We end this section with a few words about how one can use W to compute the cohomology of Y_Γ. If Γ is torsionfree, there is not much to say. One simply takes the regular cell complex $\Gamma\backslash W$ and proceeds as usual in algebraic topology courses. But if Γ has torsion, as most of the Γ do that we care about, the situation is more complicated. We can't simply divide out by the action of Γ, since there can be nontrivial stabilizer subgroups.

One solution to this dilemma would be to pass to the barycentric subdivision $\tilde{W}$ of W. It's not hard to see that all Γ-stabilizers in $\tilde{W}$ are trivial for any Γ, so $\Gamma\backslash\tilde{W}$ is a regular cell complex. But this is not always such a useful path to take. The number of Γ-orbits in $\tilde{W}$, for instance, will be much greater than the number in W. It's also less clear how to use $\tilde{W}$ to compute the action of the Hecke operators on the cohomology, cf. §11.

Another solution is to compute the *equivariant cohomology* $H^*_\Gamma(W, \mathscr{M})$. This is a ramped-up version of cohomology that takes into account the stabilizers. Usually when a group acts on a space, the equivariant cohomology computes different information from the cohomology of the quotient, but we're lucky in this case: we have an isomorphism $H^*_\Gamma(W; \mathscr{M}) \simeq H^*(\Gamma\backslash W; \mathscr{M})$. In particular, we have $H^*_\Gamma(W; \mathbb{C}) \simeq H^*(\Gamma\backslash W; \mathbb{C})$. To prove this one uses a spectral sequence relating the two cohomology theories; for more information see [13, Chapter VII] (in the language of homology) or [5, §3]. What makes everything work is that (1) the stabilizers on W are finite and (2) the local systems we consider come from *complex* representations of our discrete group. In particular, the orders of the stabilizers are invertible in the ring over which the coefficient modules are defined.

The paper [5] explains in great detail how to compute the boundary maps one needs to compute $H^*_\Gamma(W; \mathbb{C})$. Another presentation can be found in [24]. The amazing fact is that, after all the dust settles, the boundary map is essentially what one would make from the Koecher fan! In other words, there is a natural chain complex one could build from the cones in Σ. One simply takes the free abelian groups generated on the oriented cones that meet C, and then the boundary map is induced from passing from a cone to the codimension one cones on its boundary. This gives a chain complex that mod Γ computes the cohomology of Y_Γ, at least if Γ is torsionfree. If Γ has torsion, then the boundary maps must include information about the stabilizers in its definition and things get more complicated (cf. [5]); otherwise the construction is the same. The indexing in this scheme is very pleasant: the cones of *codimension* k in Σ induce the chain group that captures H^k.

9 Hecke operators and modular symbols

It is now time to talk about Hecke operators. They are a collection of linear maps on the cohomology, and their eigenvalues reveal the arithmetic lurking in the cohomology.

We fix a reductive group $\mathbf{G}$ and an arithmetic group Γ. The symmetric space (resp., locally symmetric space) is denoted D (resp., Y_Γ) as usual. Let $g \in \mathbf{G}$ have the property that Γ and $g^{-1}\Gamma g$ have finite index in $\Gamma' = \Gamma \cap g^{-1}\Gamma g$. We get a diagram

$$\begin{tikzcd} & \Gamma'\backslash D \arrow[dl, "s"'] \arrow[dr, "t"] & \\ \Gamma\backslash D & & \Gamma\backslash D \end{tikzcd}, \tag{22}$$

where the map t is the composition of $\Gamma\backslash D \to g^{-1}\Gamma g\backslash D$ with the diffeomorphism $g^{-1}\Gamma g\backslash D \to \Gamma\backslash D$ given by multiplication on the left by g:

$$g^{-1}\Gamma g x \longmapsto \Gamma g x.$$

The diagram (22) is called a *Hecke correspondence*. The condition on g ensures that s, t are finite-to-one maps.

The diagram (22) induces a map on cohomology, namely

$$t_* s^* \colon H^*(Y_\Gamma; \mathscr{M}) \to H^*(Y_\Gamma; \mathscr{M}). \tag{23}$$

The map s^* is just the usual map s induces on cohomology, but t_* only makes sense because t has finite fibers. This kind of map is sometimes called a "wrong-way" map. It can be built using integration over the fibers, if one uses de Rham cohomology [11], or via the transfer map in group cohomology [13, III.9]. The map (23) is called an *Hecke operator* and is denoted T_g. For instance if one takes $\mathbf{G} = \mathrm{SL}_2/\mathbb{Q}$ and $g = \left(\begin{smallmatrix} 1 & 0 \\ 0 & p \end{smallmatrix}\right)$, one gets the classical operator T_p.

The Hecke operators satisfy many properties. For instance, the correspondence and the operator depend only on the double coset $\Gamma g \Gamma$. The operators form an algebra by composition. Given any g as above, we can write

$$\Gamma g \Gamma = \bigsqcup_{h \in \Omega} \Gamma h, \tag{24}$$

where the Ω is a finite set (i.e., the double coset $\Gamma g \Gamma$ is a finite union of single cosets). Thus the diagram (22) can be thought of as a "multi-valued function" on Y_Γ:[7] we have

$$\Gamma x \longmapsto \{\Gamma h x\}_{h \in \Omega}. \tag{25}$$

Many details about the Hecke operators on $\mathrm{GL}_n/\mathbb{Q}$, especially the algebra structure, can be found in [57, Chapter 3].

We are keenly interested in using topological tools to determine eigenvalues and eigenclasses of the Hecke operators in cohomology. But unfortunately we cannot *directly* use the cell decompositions we constructed in §§6–7 to achieve this. The problem is that the Hecke correspondences do not act cellularly on our decompositions, and so we cannot easily write the Hecke operator as a linear map on our chain complexes. This is already visible for $\mathrm{SL}_2(\mathbb{Z})$: the Farey tessellation is not taken into itself under the "map" (25). To see this, take $g = \left(\begin{smallmatrix} 1 & 0 \\ 0 & 2 \end{smallmatrix}\right)$ and consider the edge from 0 to ∞, which is colored green in Figure 6. We have $\Omega = \{\left(\begin{smallmatrix} 1 & 0 \\ 0 & 2 \end{smallmatrix}\right), \left(\begin{smallmatrix} 2 & 0 \\ 0 & 1 \end{smallmatrix}\right), \left(\begin{smallmatrix} 1 & 1 \\ 0 & 2 \end{smallmatrix}\right)\}$. After applying (25), this edge is taken to itself with multiplicity two and the red edge in Figure 6, which runs from $1/2$ to ∞. The red edge isn't an edge in the tessellation.

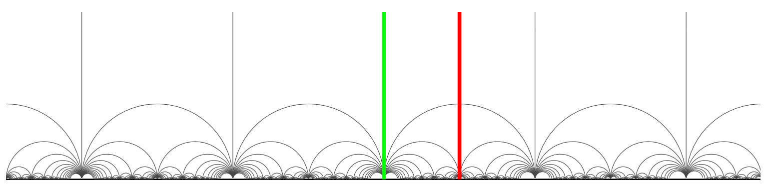

Fig. 6 Hecke correspondences don't preserve the tessellation.

How to deal with this problem? One solution is to simply refine the cell decomposition to include these new edges. In fact for T_2 on the upper halfplane this is not an unreasonable approach. It's not hard to see that essentially all one needs to do is to add all the $\mathrm{SL}_2(\mathbb{Z})$-translates of the red edge and to take the common refinement of them with the original Farey triangles. A similar strategy can be applied to T_p. The problem is that this will only allow one to compute a single Hecke operator at a time, with the geometric complexity

[7] According to R. MacPherson, the great geometers of old were perfectly comfortable with multi-valued functions and would have embraced such a perspective. It is only modern mathematicians who have the paucity of imagination to insist that functions be single-valued.

rapidly increasing as p increases. Usually one wants to be able to compute *many* Hecke operators, for as many p as possible.

Another solution is to build a bigger complex with Γ-action that also computes $H^*(Y_\Gamma)$ but has the additional property of admitting a Hecke action. Such a complex can't possibly be finite modulo Γ, so a priori is not computationally useful.

The example of $\mathrm{SL}_2/\mathbb{Q}$ is instructive. We recall from §3 how we used the Farey triangulation to compute cohomology. There we took a congruence subgroup $\Gamma \subset \mathrm{SL}_2(\mathbb{Z})$ and let X_Γ be the compactified modular curve. The chain complex built from the Farey tessellation allowed us to compute the relative homology $H_1(X_\Gamma, \partial X_\Gamma; \mathbb{C})$. We now need to add edges to this triangulation to account for the Hecke images. A little experimentation with (25) quickly convinces one that one needs to add the images (under the projection $\mathfrak{H}^* \to X_\Gamma$) of *all* geodesics in $\mathfrak{H}^*$ from cusp to cusp, if one wants to include all possible Hecke images of Farey edges. This is now a huge collection of geodesics. It's clearly not Γ-finite, since the absolute value of the determinant of a pair of cusps is preserved by Γ (cf. (7)). Nevertheless, we'll see that this is a good idea.

Thus one is naturally led to a new model for $H_1(X_\Gamma, \partial X_\Gamma; \mathbb{C})$. One considers the $\mathbb{C}$-vector space U generated by symbols $\mathbf{u} = [x_1, x_2]$, where each x_i is a cusp. We think of the symbol $\mathbf{u}$ as corresponding to the class in $H_1(\mathfrak{H}^*, \partial\mathfrak{H}^*; \mathbb{C})$ of the oriented geodesic from x_1 to x_2. As such, these symbols have to satisfy some obvious relations. For instance, $[x_2, x_1] = -[x_1, x_2]$, since the geodesics have opposite orientations. Certainly $[x_1, x_1] = 0$, since the geodesic is a point. The most complicated one is the *three-term* relation, which says

$$[x_1, x_2] + [x_2, x_3] + [x_3, x_1] = 0. \tag{26}$$

This simply says that if three geodesics form the oriented boundary of a triangle, the sum of their classes should vanish. The space U is the space of *modular symbols*. It has a Γ-action since Γ acts on the cusps, and we let U_Γ be the quotient U by relations of the form $\mathbf{u} - \gamma \cdot \mathbf{u}$ (i.e., the space of coinvariants.) Then U_Γ is isomorphic to $H_1(X_\Gamma, \partial X_\Gamma; \mathbb{C})$ and admits a Hecke action: if $\Gamma g \Gamma$ has the decomposition (24), then

$$T_g(\mathbf{u}) = \sum_{h \in \Omega} h \cdot \mathbf{u}. \tag{27}$$

By (6) this corresponds to a Hecke action on $H^1(Y_\Gamma; \mathbb{C})$, and thus on the weight two modular forms.

Now we connect this back to our original model for the relative H_1, which came from the Farey tessellation. Clearly the Farey tessellation determines a subspace $U' \subset U$, the subspace generated by *unimodular symbols*, which by definition are the symbols with determinant ± 1 (cf. (7)). This is an easier space to work with since the space of coinvariants U'_Γ is finite. So our computational problem becomes the following:

1. Start with a cycle $\eta \in U'_\Gamma$, representing a class in $H_1(X_\Gamma, \partial X_\Gamma; \mathbb{C})$.
2. Lift η to sum of symbols $\tilde\eta = \sum a_{\mathbf{u}} \mathbf{u}$ and thus to an element of U', which determines an element of U.
3. Now apply the Hecke operator (27) to compute $T_g(\tilde\eta)$. Thus will lie in U and *not* U' in general.
4. Somehow push $T_g(\tilde\eta)$ to an equivalent element in U'. Here "equivalent" means that we are allowed to rewrite $T_g(\tilde\eta)$ using the defining relations for U, such as (26). The goal is to obtain a sum $\tilde\theta = \sum b_{\mathbf{u}} \mathbf{u}$, where each $\mathbf{u}$ is unimodular.
5. Then one projects $\tilde\theta$ back down to U'_Γ.

Of course step (4) is the subtle part of this process; this is where all the action takes place. For modular symbols on $\mathrm{SL}_2/\mathbb{Q}$, step (4) is done through a version of the continued fraction algorithm and is known as *Manin's trick*. See [61, Proposition 3.11] for an exposition. Since our goal is to treat more complicated groups, here we content ourselves with showing what happens to the red modular symbol from Figure 6. We can push it back to the unimodular subspace by applying one three-term relation, which writes it as the sum of the two orange modular symbols in Figure 7.

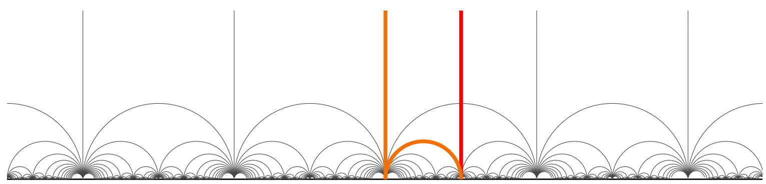

Fig. 7 Making the red modular symbol unimodular.

10 The sharbly complex

We return to $\mathbf{G} = R_{F/\mathbb{Q}}\mathrm{GL}_n$. Our goal is to describe a complex $\mathscr{S}_*$ with Hecke action that computes $H^*(Y_\Gamma; \mathbb{C})$. This will take the place of the modular symbols U from §9. In fact it will turn out that U is the quotient $\mathscr{S}_0/\partial\mathscr{S}_1$, where ∂ is the boundary map in the complex. Like the modular symbols, this complex will be built from tuples of "cusps," but now the cusps will be equivalence classes of the points we used in the construction of the Koecher fan.

Recall that for any $x \in \mathscr{O}^n \smallsetminus \{0\}$, we have constructed a point $q(x) \in \bar C$ (see (20)). Write $x \sim y$ if there exists $\lambda \in \mathbb{R}_{>0}$ such that $q(x) = \lambda q(y)$. Thus x is equivalent to y if they determine the same ray in $\bar C$. Let A_k be the set of formal $\mathbb{C}$-linear sums of symbols $\mathbf{u} = [x_1, \ldots, x_{k+n}]$, where each x_i is in $\mathscr{O}^n \smallsetminus \{0\}$. Let C_k be the submodule generated by the elements

1. $[x_{\sigma(1)}, \ldots, x_{\sigma(k+n)}] - \operatorname{sgn}(\sigma)[x_1, \ldots, x_{k+n}]$ for any permutation σ on $(k+n)$ letters,
2. $[x, x_2, \cdots, x_{k+n}] - [y, y_2, \ldots, y_{k+n}]$ if $x \sim y$, and
3. $\mathbf{u}$ if $x_1, \cdots, x_{k+n}$ are contained in a hyperplane (we say $\mathbf{u}$ is *degenerate*).

The quotient $\mathscr{S}_k = A_k/C_k$ is called the space of *k-sharblies*. We define a boundary map $\partial\colon \mathscr{S}_{k+1} \to \mathscr{S}_k$ by linearly extending

$$\partial[x_1, \cdots, x_{k+n}] = \sum_{i=1}^{k+n} (-1)^i [x_1, \cdots, \hat{x}_i, \ldots, x_{k+n}], \tag{28}$$

where $\hat{x}_i$ means omit x_i. The resulting complex $\mathscr{S}_*$ is called the *sharbly complex*. Note that $\mathscr{S}_*$ is a homological complex—the boundary maps decrease degrees.

For example, let $\mathbf{G} = \mathrm{SL}_2/\mathbb{Q}$. Two points $x, y \in \mathbb{Z}^2 \smallsetminus \{0\}$ satisfy $x \sim y$ if and only if they determine the same cusp of the upper halfplane. The defining relations for C_0, together with the additional relations obtained from the image of $\partial(\mathscr{S}_1)$, are exactly the relations used to build U.

The complex $\mathscr{S}_*$ has a left Γ-action for any subgroup $\Gamma \subset \mathrm{GL}_n(\mathscr{O})$: if $g \in \Gamma$ and $\mathbf{u} = [x_1, \ldots, x_{k+n}]$, then $g \cdot \mathbf{u} = [gx_1, \ldots, gx_{k+n}]$. The Γ-action commutes with the boundary, so we can form the complex $(\mathscr{S}_*)_\Gamma$ of coinvariants. We claim

$$H^{\mathrm{vcd}(\Gamma)-k}(Y_\Gamma; \mathbb{C}) \xrightarrow{\sim} H_k((\mathscr{S}_*)_\Gamma).$$

This follows from *Borel–Serre duality*, as we now explain. We first need to recall the Steinberg module.

Let $V = F^n$ be an n-dimensional vector space over F. We build a simplicial complex $\mathscr{T}$, called the *Tits building*, from this vector space as follows. The vertices are the proper nonzero subspaces of V. Subspaces $V_1, \ldots, V_{k+1}$ determines a k-simplex if they can be arranged into a flag

$$\{0\} \subsetneq V_1 \subsetneq V_2 \subsetneq \cdots \subsetneq V_{k+1} \subsetneq V.$$

By the Solomon–Tits theorem, $\mathscr{T}$ has the homotopy type of a bouquet of $(n-2)$-spheres.[8] In particular the reduced homology groups $\tilde{H}_*(\mathscr{T})$ are nonzero only in degree $(n-2)$.

One can construct classes in $\tilde{H}_{n-2}(\mathscr{T})$ by taking the fundamental classes of apartments: one chooses a basis $E = \{v_1, \ldots, v_n\}$ of V and considers all the possible flags that can be constructed from E by taking spans of permutations of subsets. By appropriately choosing signs one obtains a class $\langle v_1, \ldots, v_n \rangle \in \tilde{H}_{n-2}(\mathscr{T})$. It is known that such classes span the homology.

[8] A *bouquet of spheres* is wedge sum of a set of spheres.

We have an action of $\mathbf{G}(\mathbb{Q})$, and by definition, the *Steinberg module* St_n is the $\mathbf{G}(\mathbb{Q})$-module $\tilde{H}_{n-2}(\mathscr{T})$.

For instance, suppose $n = 3$ and $F = \mathbb{Q}$. Then the vertices of $\mathscr{T}$ come in two types, namely those indexed by lines and those indexed by planes. Two vertices are joined by an edge if one corresponds to a line and one to a plane, and the line is contained in the plane. There are no higher-dimensional simplices. Thus $\mathscr{T}$ is a graph, which certainly has the homotopy type of a bouquet of circles. If we fix a triple of lines in $\mathbb{Q}^3$ only meeting at the origin, we can determine 3 planes by taking their pairwise spans. Thus we obtain 6 different flags whose edges can be grouped together into a hexagon (Figure 8; the white (resp., black) dots correspond to the lines (resp., planes)). The classes of all such hexagons span $\tilde{H}_1(\mathscr{T})$.

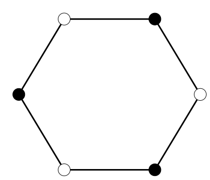

Fig. 8 A cycle in the Tits building.

Now we are ready to connect the Steinberg module to the sharbly complex. The Borel–Serre duality theorem [10] states that for any arithmetic group $\Gamma \subset \mathbf{G}(\mathbb{Q})$, we have

$$H^{\mathrm{vcd}(\Gamma)-k}(Y_\Gamma;\mathbb{C}) \xrightarrow{\sim} H_k(\Gamma;\mathrm{St}_n \otimes \mathbb{C}).$$

A more general result holds for the local coefficient systems $\mathscr{M}$; one simply replaces $\mathrm{St}_n \otimes \mathbb{C}$ with $\mathrm{St}_n \otimes M$. Thus to compute the cohomology of Γ, we need to take a resolution of the Steinberg module. This is what the sharbly complex gives us. We have a map

$$\varepsilon\colon \mathscr{S}_0 \longrightarrow \mathrm{St}_n$$

gotten by taking the 0-sharbly $\mathbf{u} = [x_1, \dots, x_n]$ to the class $\langle x_1, \dots, x_n\rangle$. It is not hard to see that the composition $\partial \circ \varepsilon\colon \mathscr{S}_1 \to \mathrm{St}_n$ vanishes, so we get a map of complexes $\mathscr{S}_* \to \mathrm{St}_n$ (we regard St_n as being a complex concentrated in degree 0.) The sharbly complex itself is easily seen to be acyclic, and thus it furnishes us with a resolution of St_n.

Hence we can use the sharbly complex to compute the cohomology $H^*(Y_\Gamma;\mathbb{C})$. By extending the coefficients of the sharbly complex, we can even use it to compute the cohomology with coefficients $H^*(Y_\Gamma;\mathscr{M})$. The sharbly complex admits an action of the Hecke operators. Suppose $\eta \in (\mathscr{S}_k)_\Gamma$ is a cycle. Then we can lift η to a k-sharbly chain $\tilde{\eta} = \sum a_{\mathbf{u}}\mathbf{u} \in \mathscr{S}_k$, where

$a_{\mathbf{u}} \in \mathbb{C}$ and almost all coefficients are zero. If a Hecke operator T has coset representatives Ω, as in (24), then we put

$$\tilde{\mu} = T(\tilde{\eta}) = \sum a_{\mathbf{u}}(h \cdot \mathbf{u}).$$

In general $\tilde{\mu}$ will depend on our choice of Ω, but the image μ of $\tilde{\mu}$ in $(\mathscr{S}_k)_\Gamma$ will not. Thus we can define $T(\eta)$ to be μ.

At this point we find ourselves in a very similar position to that in §9. We have a complex $\mathscr{S}_*$ that allows to compute the cohomology of Y_Γ, a complex that is the analogue of the modular symbols U. The only glitch is that each $\mathscr{S}_k$ isn't finite modulo Γ. This is easy to see. If $\mathbf{u} = [x_1, \ldots, x_n]$ is a 0-sharbly, we can compute how "big" it is using determinants. For any $x \in \mathscr{O}^n \smallsetminus \{0\}$, let $x' \in \mathscr{O}^n \smallsetminus \{0\}$ denote the unique point such that $x' \sim x$ and $q(x')$ is closest to the origin in the ray through $q(x)$. We define $\operatorname{Size}(\mathbf{u}) = |\operatorname{N}_{F/\mathbb{Q}} \det(x'_1, \ldots, x'_n)| \in \mathbb{Z}_{>0}$. If is clear that Size is constant on $\mathrm{GL}_n(\mathscr{O})$-orbits in $\mathscr{S}_0$, and that Size is unbounded on $\mathscr{S}_0$.

Hence we have an analogue of the space of modular symbols. What's missing is the analogue of the unimodular subspace U'. This subspace U' has two characterizations: it can be defined as the space spanned by (i) determinant one modular symbols, or (ii) modular symbols whose support is an edge in the Farey tessellation. By coincidence these are the same condition for $\mathrm{SL}_2(\mathbb{Z})$, but this is not necessarily true in other settings. Our perspective is that the second criterion is more robust, and works better for other groups.

Thus the decompositions we constructed in §§6–7 return to the stage. Recall that Σ is the fan of Koecher cones in the closed cone $\bar{C}$, and that the quotient $C/\mathbb{R}_{>0}$ is a model for the symmetric space associated to $\mathbf{G}$. Recall also that we can compute the cohomology $H^*(Y_\Gamma; \mathbb{C})$ by computing the equivariant cohomology $H^*_\Gamma(W; \mathbb{C})$ of the retract W dual to Σ, and that the complex used to do this is essentially the chain complex on the Koecher cones. Since the Koecher cones are encoded by tuples of nonzero points in $\mathscr{O}^n$, just as the sharblies are, it is irresistible to try to build a "Koecher" subcomplex of $\mathscr{S}_*$ corresponding to the Koecher cones. In fact this is basically what the unimodular subspace is: it comes from the subcomplex of $\mathscr{S}_0$ built from the Farey edges, i.e. the Voronoi 2-cones.

This is the idea that will eventually allow us to compute the action of the Hecke operators on the cohomology, but implementing it is not as straightforward as one might hope. An immediate problem is that the Koecher cones need not be simplicial. Hence for some Koecher cones there is not an obvious way to build a corresponding sharbly chain.

One way to deal with this is as follows. A basis k-sharbly $\mathbf{u} = [x_1, \ldots, x_{k+n}]$ induces a collection of rays $\mathbb{R}_{>0}q(x_1), \ldots, \mathbb{R}_{>0}q(x_{k+n})$ in $\bar{C}$. We say $\mathbf{u}$ is *reduced* if there is a top-dimensional Koecher cone that contains these rays. We say a k-sharbly chain is reduced if all its basis sharblies are reduced. Note that a sharbly being reduced is not the same as saying that these rays span a cone in the fan Σ, since the cones in Σ aren't necessarily simplicial.

Reduced sharblies are our analogue of unimodular symbols. It is clear that the reduced sharblies form a subcomplex of $\mathscr{S}_*$, since the boundary of any reduced sharbly chain is reduced. Let $\mathscr{R}_*$ be this complex. It is finite modulo Γ, since there are only finitely many Koecher cones modulo Γ.

Thus $(\mathscr{R}_*)_\Gamma$ looks like a good candidate to compute $H^*(Y_\Gamma; \mathbb{C})$, but unfortunately it doesn't work. The problem is that because the Koecher fan isn't simplicial, the complex $(\mathscr{R}_*)_\Gamma$ could be missing some identifications necessary to capture the cohomology we want. Rather than pursuing a complete presentation, which is not necessary for our purposes, we give an example to illustrate what's going on.

Suppose that $n = 2$ and that Σ contains a 3-cone σ that is a cone on a square; this is the simplest way Σ can fail to be non-simplicial. There are four reduced 1-sharblies "associated" to σ, namely those corresponding the four simplicial cones obtained by drawing the two different diagonals across the square (Figure 9 shows a cross section). Let's call these 1-sharblies $\mathbf{u}_1, \dots, \mathbf{u}_4$. Certainly we want to have the relation $\mathbf{u}_1 + \mathbf{u}_2 = \mathbf{u}_3 + \mathbf{u}_4$. If the stabilizer of σ in Γ includes rotation by 90°, then we can pick it up when we pass to the coinvariants. But if we don't have this rotation we may miss a relation we clearly want, and Y_Γ might appear to have extra cohomology.

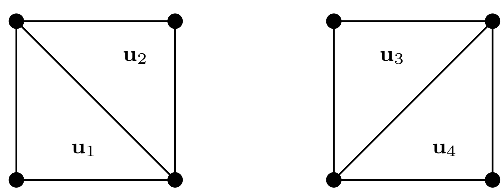

Fig. 9 Reduced 1-sharblies from a cone on a square.

There are three ways out of this problem:

1. *We can ask for less* by restricting ourselves to torsionfree Γ. Then we can take each cone $\sigma \in \Sigma$ and can simplicially subdivide it *without adding new rays*. Since Γ is torsionfree we can perform this subdivision Γ-equivariantly. We can define a subcomplex $\mathscr{R}'_* \subset \mathscr{R}_*$ corresponding to these cones. Then $(\mathscr{R}'_*)_\Gamma$ is finite and computes the cohomology of Y_Γ. But this possibility is rather unappealing, since our main groups of interest (analogues of $\Gamma_0(N) \subset \mathrm{SL}_2(\mathbb{Z})$) have torsion in general.
2. *We can steel ourselves* and can work with $\mathscr{R}_*$. We just add all the extra relations we need. In particular, if two unions of cones corresponding to two sets of reduced k-sharblies form the same cone in $\bar{C}$, then we pick up a relation, exactly as Figure 9 tells us that $\mathbf{u}_1 + \mathbf{u}_2 = \mathbf{u}_3 + \mathbf{u}_4$. This can be somewhat painful to work out, but there are only finitely many possibilities to worry about, and often the stabilizers of the Koecher cones can help.

3. *We can gamble* and only compute cohomology in the degrees where the Koecher fan is simplicial. In particular suppose we are interested in the cohomology group $H^{\mathrm{vcd}(\Gamma)-i}$. Then we want to work with k-sharblies for $k = i-1, i, i+1$. If the Koecher fan is simplicial in these dimensions, then we can use $\mathscr{R}_*$ without having to muck around with subdivisions.

The last idea sounds somewhat crazy, but it turns out to work very well in practice: experience teaches us that in many examples the Koecher fan is simplicial up to cones of relatively large dimension. Even better, for various reasons we may only care about the small-dimensional cones in Σ, where Σ is often simplicial. For instance this is what happens in the computations described in §12. In the next section, we will try to explain this.

11 Hecke operators and sharbly reduction

In this section we finally describe how to use the sharbly complex to compute the action of the Hecke operators. We will explain what happens in the mysterious step (4). However, we must come clean at the very beginning, and confess that we don't actually have a *proof* that our techniques to compute Hecke operators work. Nevertheless, the techniques are robust enough that they have worked in every attempt, without fail. We frame the discussion in a sequence of heuristics.

Before we begin, we must discuss the cuspidal range. Recall the decomposition (16) from §5:

$$H^*(Y_\Gamma;\mathbb{C}) = H^*_{\mathrm{cusp}}(Y_\Gamma;\mathbb{C}) \oplus \bigoplus_{\{\mathbf{P}\}} H^*_{\{\mathbf{P}\}}(Y_\Gamma;\mathbb{C}). \tag{29}$$

Here we have taken cohomology with trivial coefficients. The *cuspidal cohomology* H^*_{cusp} is the most interesting summand from our point of view, since it corresponds to certain cuspidal automorphic forms; these are our analogues of the weight two holomorphic modular forms.

This decomposition only tells us that there is a summand for the cusp forms. It doesn't say anything about which cohomological degrees actually contain cuspidal cohomology, so where does it live? It turns out that cuspidal cohomology is a rather picky beast: outside of a certain interval, called the *cuspidal range*, the cuspidal cohomology vanishes. One can get an estimate for this range using the structure theory of $G = \mathbf{G}(\mathbb{R})$ [43, §2.C]. Assume that $\mathbf{G} = R_{F/\mathbb{Q}}\mathrm{SL}_n$, so that $G \simeq \mathrm{SL}_n(\mathbb{R})^r \times \mathrm{SL}_n(\mathbb{C})^s$, with maximal compact subgroup $K \simeq \mathrm{SO}(n)^r \times \mathrm{SU}(n)^s$ and associated symmetric space $D = G/K$. Let $r(G) = r_{\mathbb{C}}(G)$ be the absolute rank of G, and let $l_0(G) = r(G) - r(K)$. Then for any coefficient system $\mathscr{M}$ as before and any arithmetic group Γ, we have

$$H^i_{\mathrm{cusp}}(Y_\Gamma;\mathscr{M}) = 0$$

unless

$$b(\Gamma) := \frac{1}{2}(\dim D - l_0(G)) \leq i \leq \frac{1}{2}(\dim D + l_0(G)) =: t(\Gamma). \tag{30}$$

For instance, if $F = \mathbb{Q}$ then $r(\mathrm{SL}_2(\mathbb{R}))$ is the rank of $\mathrm{SL}_2(\mathbb{C})$, which is 1, and $r(\mathrm{SO}(2))$ is the rank of $\mathrm{SO}(2, \mathbb{C})$, which is 1. Thus $l_0 = 0$, and the cuspidal cohomology occurs in degree $1 = (\dim \mathfrak{H})/2$. This is why we see the holomorphic modular forms in H^1. For SL_3 we have $r(\mathrm{SL}_3(\mathbb{R})) = 2$, but $r(\mathrm{SO}(3)) = 1$ (since this is the rank of $\mathrm{SO}(3, \mathbb{C})$, which has Cartan-Killing type $B_1 \simeq A_1$). Thus $l_0 = 1$. The space D has dimension 5, and so cuspidal cohomology can only show up in degrees $2 \leq i \leq 3$.

The case of F imaginary quadratic is also interesting. We have $G = \mathrm{SL}_2(\mathbb{C})$, thought of as a *real* Lie group, not a complex Lie group. Thus $r(\mathrm{SL}_2(\mathbb{C})) = 2$, since $G(\mathbb{C})$ is *two* copies of $\mathrm{SL}_2(\mathbb{C})$. On the other hand $r(\mathrm{SU}(2)) = 1$, since $\mathrm{SU}(2)$ is the compact form of $\mathrm{SL}_2(\mathbb{C})$. Hence $l_0 = 1$, and since $\dim \mathfrak{H}_3 = 3$, cuspidal cohomology can only occur in degrees $1 \leq i \leq 2$.

What about the difference between $\mathbf{G} = R_{F/\mathbb{Q}}\mathrm{GL}_n$ and $\mathbf{G}' = R_{F/\mathbb{Q}}\mathrm{SL}_n$? The only difference in the locally symmetric spaces attached to these two groups is that the GL-space Y_Γ is a torus bundle over the SL-space Y'_Γ. One can find a $\Gamma \subset \mathrm{GL}_n(\mathscr{O})$ such that this bundle is trivial, which means (by the Künneth theorem) that $H^*(Y_\Gamma) = H^*(T) \otimes H^*(Y'_\Gamma)$, where $T \simeq (S^1)^{r+s-1}$. From this it is clear that the formula (30) becomes

$$b(\Gamma) := \frac{1}{2}(\dim D - l(G)) \leq i \leq \frac{1}{2}(\dim D + l(G)) =: t(\Gamma), \tag{31}$$

where $l(G) = l_0(\mathbf{G}'(\mathbb{R})) + r + s - 1$. Notice that the lower bound $b(\Gamma)$ doesn't change. The upper bound $t(\Gamma)$ grows, but the difference $\mathrm{vcd}(\Gamma) - t(\Gamma)$ doesn't change. Table 1 gives some examples of these numbers.

Thus the cuspidal cohomology only occurs in a restricted range. Furthermore, if a given cusp form contributes to any cohomology group in this interval, then it does so to all of them, and in an easily understood way. Therefore if one wants to compute cuspidal cohomology, one might as well pick a single group in the cuspidal range to study.

Now we work with k-sharbly cycles. There are two reasons we prefer to make k as small as possible. First, in many examples the Koecher fan is simplicial in low dimensions; thus it is easy to map the chain complex coming from the reduction theory into the sharbly complex. Second, k-sharbly cycles become more difficult to handle as k increases, since our main tool is to modify k-sharbly cycles $\sum a[x_1, \ldots, x_{k+n}]$ by fiddling with subtuples of the x_i of order n. This leads to our first heuristic:

(A)*It is better to work with sharbly cycles in low degree, and thus with cohomology in high degree.*

In fact, to date we have focussed on k-sharbly cycles when $k = 0$ or 1, for the number fields F where the cuspidal cohomology contributes either to

$\mathrm{vcd}(\Gamma)$ or $\mathrm{vcd}(\Gamma)-1$. Some of these cases have been previously studied using other techniques. Indeed this prior work was extremely important to us; our work would not have been possible without it.

11.1 0*-sharblies*

We begin with 0-sharblies and $F = \mathbb{Q}$. Let $\mathbf{u} = [x_1, \dots, x_n]$ be a 0-sharbly. We may assume each x_i is primitive, which is equivalent to $x_i = x_i'$. Call $\mathbf{u}$ *unimodular* if $\mathrm{Size}(\mathbf{u}) = 1$. Unimodular 0-sharblies are reduced in the sense of §10, since they are all equivalent to $\mathbf{v} = [e_1, \dots, e_n]$ modulo $\mathrm{GL}_n(\mathbb{Z})$, where $\{e_i\}$ is the standard basis of $\mathbb{Z}^n$; the vectors e_i are some of the minimal vectors of the A_n perfect form. We have the following fundamental result of Ash–Rudolph [9], which generalizes Manin's trick to higher dimensions:

Theorem 1. [9] *If* $\mathrm{Size}([x_1, \dots, x_n]) > 1$*, then there exists* $x \in \mathbb{Z}^n \smallsetminus \{0\}$ *such that*

$$0 \leq \mathrm{Size}([x, x_1, \dots, \hat{x}_i, \dots, x_n]) < \mathrm{Size}([x_1, \dots, x_n]), \quad (i = 1, \dots, n).$$

Thus $H^{\mathrm{vcd}(\Gamma)}(Y_\Gamma; \mathbb{C})$ *is spanned by unimodular* 0*-sharblies.*

We call any integral point that can play the role of x in the theorem a *reducing point* for $\mathbf{u} = [x_1, \dots, x_n]$. The set of all reducing points for $\mathbf{u}$ is denoted $\mathrm{Red}\,\mathbf{u}$. The proof of Theorem 1 is constructive, and gives an algorithm for finding a reducing point for $\mathbf{u}$. The algorithm is a higher-dimensional version of continued fractions.

Theorem 1 has been applied to compute the Hecke action on $H^3(Y_\Gamma)$ when $\Gamma \subset \mathrm{SL}_3(\mathbb{Z})$ [4,62]. Variants have been used for $R_{F/\mathbb{Q}}\mathrm{SL}_2$ where F is an imaginary quadratic field with ring of integers $\mathscr{O}$ norm-Euclidean [17,25,32] or a PID [21], and even in some cases of nontrivial class group [16,44]. These algorithms compute the Hecke action on $H^2(Y_\Gamma)$ for $\Gamma \subset \mathrm{SL}_2(\mathscr{O})$.

Theorem 1 provides a beautiful way to reduce 0-sharblies to unimodular, but its beauty is also its fatal flaw when one goes from $\mathbb{Q}$ to other number fields: it relies on the Euclidean algorithm in an essential way. Thus the method breaks down on fields that aren't norm-Euclidean, in particular for fields with nontrivial class number. Other techniques must then be tried. This leads us to our next heuristic:

(B) *Choose candidates for* 0*-sharblies using the geometry of Koecher fan, and not using continued fractions.*

This is based on ideas that go back to [33,34]. This idea and variations of it have been used in [5,35,36,71]. Here's how it works. To make the discussion more accessible we restrict to $R_{F/\mathbb{Q}}\mathrm{GL}_2$, which has $\mathbb{Q}$-rank 1. A 0-sharbly $\mathbf{u}$ is then a pair $[x_1, x_2]$. In the compactified symmetric space $\bar{D} = \bar{C}/\mathbb{R}_{>0}$, we can realize $\mathbf{u}$ as a geodesic φ running from the image of $q(x_1)$ to that of $q(x_2)$. It will meet the images of the Koecher cones in various ways, and $\mathbf{u}$ will be reduced if and only if φ is completely contained in the image of a Koecher cone.

We want to find a reducing point $x \in \mathscr{O}^n$ such that in the three-term relation

$$[x_1, x_2] = [x_1, x] + [x, x_2] =: \mathbf{u}_1 + \mathbf{u}_2, \tag{32}$$

the 0-sharblies $\mathbf{u}_1, \mathbf{u}_2$ will be closer to being reduced (for instance, we might ask that their sizes be smaller than that of $\mathbf{u}$). The idea is that *the reducing point x should be selected from the vertices of the Koecher cone containing the center of φ*. Here the center means the ray through the points $q(x_1 + x_2)$; note that this is only well-defined since we've chosen specific points in the rays $\mathbb{R}_{>0}q(x_1), \mathbb{R}_{>0}q(x_2)$.

Why should an x chosen in this way be a reducing point? We must confess that we don't know. Geometric motivation comes from looking at the Voronoi cones for $\mathrm{SL}_2(\mathbb{Z})$. Consider Figure 10, which shows a cross-section of the cone C. The triangles in the middle look rather large compared to the triangles on the outside, so the center of a "balanced" 0-sharbly will tend to land there. Any of the three points w, w', w'' appear to be good choices for a reducing point, since the two 0-sharblies on the right of (32) look like they will cut across fewer top-dimensional cones and will be closer to being reduced.

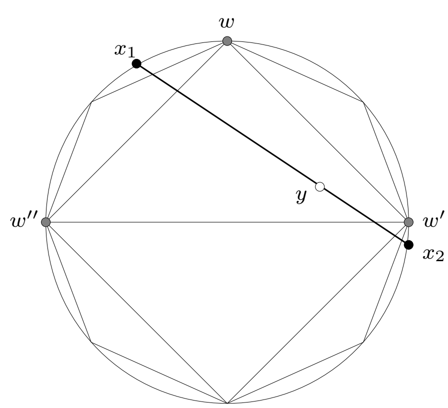

Fig. 10 Using the tessellation to find reducing points. The point y is the center, and w, w', w'' are the potential reducing points.

Of course, motivation is not a proof, and sadly we don't have a proof that this will work. Thus we don't know if selecting reducing points in this way eventually allows one to write a 0-sharbly as a sum of reduced ones. For $\mathrm{SL}_2(\mathbb{Z})$, it's not hard to engineer a geometric proof that realizes the motivation above, but the general case is unknown. Nevertheless, the heuristic works extremely well in practice. In fact, the 0-sharblies on the right of (32) tend to be *much* closer to being reduced than the original. We also expect this idea to work in cases where Γ is cocompact, such as [31, 51, 67].

11.2 1-*sharblies*

Now we go to 1-sharblies. We continue to take $\mathbf{G} = R_{F/\mathbb{Q}}\mathrm{GL}_2$. We will also assume that Γ is torsionfree to avoid some complications. We assure the reader that these restrictions are only for convenience: we have also applied the ideas in this subsection to groups of higher $\mathbb{Q}$-rank, in particular $\mathrm{GL}_n/\mathbb{Q}$, $n = 3, 4$, and $R_{F/\mathbb{Q}}\mathrm{GL}_3$ where F is imaginary quadratic. We have also treated Γ with torsion. The reader curious about dealing with torsion and higher $\mathbb{Q}$-rank can consult [5, 34]. The papers [36, 37] also contain more details of this method. Our goal here is to explain its geometry and combinatorics.

Let $\mathbf{u} = [x_1, x_2, x_3]$ be a basis 1-sharbly. We call the three 0-sharblies that appear in $\partial\mathbf{u}$ the *submodular symbols* of $\mathbf{u}$. We denote the set of all submodular symbols appearing in $\mathbf{u}$ by $Z(\mathbf{u})$, and extend this notation to sharbly chains in the obvious way.

Let $\eta = \sum a_{\mathbf{u}}\mathbf{u}$ be a 1-sharbly chain that is cycle modulo Γ. Note that this is now a nontrivial condition, unlike for 0-sharblies. In particular, any 0-sharbly chain in $\mathscr{S}_0$ automatically a cycle, even without passing to $(\mathscr{S}_0)_\Gamma$. For a 1-sharbly chain to have zero boundary modulo Γ, there must be nontrivial identifications among the submodular symbols $Z(\eta)$ appearing in its boundary.

Suppose η is not reduced. How can we reduce it? One criterion for whether or not η is reduced involves its submodular symbols. Since we have heuristics for reducing 0-sharblies, it makes sense to try to reduce η by somehow reducing $Z(\eta)$. This is in fact the approach we take, although there are two subtleties:

1. We can try to reduce the modular symbols in $Z(\eta)$ by choosing a bunch of reducing points. But what's the best way to do this?
2. And if we select reducing points, what are we supposed to do with them? In other words, how do we combine the candidates and the points in η into a new 1-sharbly cycle that is somehow better?
3. Certainly if η is reduced, then is submodular symbols are as well. Unfortunately the converse is not true in general: it is possible for $Z(\eta)$ to consist entirely of reduced 0-sharblies and for η to not be reduced. What do we do about this?

Let's treat these one at a time. First, we should use heuristic (A) to pick reducing points, just as we did for 0-sharblies on their own. However we have to take more care in this case, since the submodular symbols in $Z(\eta)$ are the boundary $\partial\eta \in \mathscr{S}_0$, and for η to be a cycle modulo Γ we need the image of $\partial\eta$ to vanish in $(\mathscr{S}_0)_\Gamma$. This means we need to choose reducing points Γ-equivariantly. In other words, suppose $\mathbf{v}, \mathbf{v}'$ appear in $\partial\mathbf{u}$ and satisfy $\mathbf{v} = \gamma\cdot\mathbf{v}'$ for some $\gamma \in \Gamma$. Then if we choose $w \in \operatorname{Red}\mathbf{v}$ for $\mathbf{v}$, we must take γw for $\mathbf{v}'$. We also insist that we choose the *best possible* reducing points, in the sense that the sizes of the resulting 0-sharblies are as small as possible (for instance, we might want to minimize the sum of sizes).

The second issue is more interesting. Assume that all submodular symbols of $\mathbf{u} = [x_1, x_2, x_3]$ need reducing (the general case is an easy variation of this). Choose three reducing points w_1, w_2, w_3 using (A); we label them such that w_i goes with $[x_j, x_k]$ if any only if i, j, k are distinct. We need to assemble them into a 2-sharbly chain whose boundary contains $\mathbf{u}$ and some 1-sharblies that have a chance of being closer to reduced. This gives us our next heuristic:

(C) *Three points and a triangle make an octahedron.*

Figure 11 illustrates what's happening. We use the six points $x_1, \ldots, w_3$ to make an octahedron O. We can subdivide O any way we like into tetrahedra to make a 2-sharbly chain; its boundary in $\mathscr{S}_2$ will be eight 1-sharblies, which in $\mathscr{S}_1$ induces the relation

$$\mathbf{u} = \mathbf{u}_1 + \mathbf{u}_2 + \mathbf{u}_3 - \mathbf{u}_{12} - \mathbf{u}_{13} - \mathbf{u}_{23} + \mathbf{u}_{123}. \tag{33}$$

In (33) the notation $\mathbf{u}_S$ for $S \subset \{1, 2, 3\}$ means the 1-sharbly in the boundary that contains $\{w_i \mid i \in S\}$. When we mod out by Γ, a miracle happens: (33) becomes

$$\mathbf{u} = -\mathbf{u}_{12} - \mathbf{u}_{13} - \mathbf{u}_{23} + \mathbf{u}_{123}. \tag{34}$$

That is, the three 1-sharblies represented by the shaded triangles in Figure 11 disappear. The reason is simple. Consider the 1-sharbly $\mathbf{u}_1 = [w_1, x_2, x_3]$. Since $\mathbf{u}$ was part of a cycle mod Γ, there must be some 0-sharbly $\mathbf{v}'$ that cancels $\mathbf{v} = [x_2, x_3]$ when $\partial\eta$ is taken. This 0-sharbly cannot be reduced since $\mathbf{v}$ isn't, so we must have chosen a reducing point for it. Since we made these choices Γ-equvariantly, there must be some *other* 1-sharbly $\mathbf{u}'$ with $\mathbf{v}'$ in its boundary. When we build an octahedron O' over $\mathbf{u}'$, a triangle in $\partial O'$ will cancel $\mathbf{u}_1$ in ∂O.

Thus over η we replace each 1-sharbly with four new 1-sharblies and obtain a new cycle η' mod Γ. Why is η' better than η? Consider (34). Certainly $\mathbf{u}_{12}, \mathbf{u}_{13}, \mathbf{u}_{23}$ look better than $\mathbf{u}$, since the reducing points have improved some edges. But why is $\mathbf{u}_{123}$ better? Notice that it is built from the reducing points. When choosing them we only look at the 0-sharblies; we don't consider them collectively until we package them into $\mathbf{u}_{123}$. In fact, upon reflection it's clear that we *can't* choose the w_i with the intent of making $\mathbf{u}_{123}$ good, since we have to choose them Γ-equivariantly over the whole cycle η, and we have no control over what this cycle looks like. All we can do is look at the submodular symbols and pick reducing points locally, in other words without considering what 1-sharblies contain them.

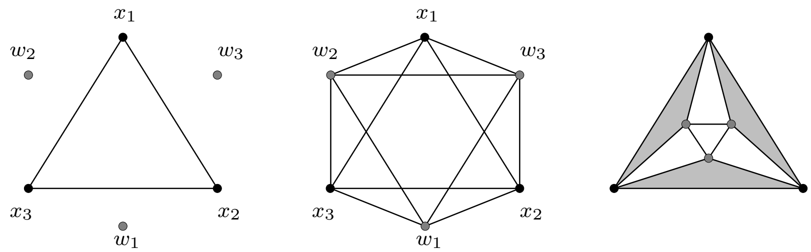

Fig. 11 The 1-sharbly $\mathbf{u}$ and its reducing points (left) assembled into an octahedron (middle). The shaded triangles (right) don't appear in $(\mathscr{S}_1)_\Gamma$, so $\mathbf{u}$ is transformed to four new 1-sharblies.

We claim that $\mathbf{u}_{123}$ will be much closer to being reduced than $\mathbf{u}$. In fact, in practice $\mathbf{u}_{123}$ will be far better than the other 1-sharblies in (34). We have no proof of this, but again we can provide some geometric motivation. Up to a flat factor, the symmetric space D is of noncompact type and is thus nonpositively curved. This means the centers of the facets of a simplex tend to be close to the center of the simplex itself; think of what triangles look like in the hyperbolic plane. The hyperbolic plane is not an entirely accurate picture of what happens, since D may have high-dimensional flat subspaces if the $\mathbb{R}$-rank of $\mathbf{G}$ is large, but nevertheless this picture is compelling.

Thus the three reducing points for the submodular symbols of $\mathbf{u}$ tend to be taken from the same Koecher cone, or at least from cones that are very close together. Therefore the potentially bad 1-sharbly $\mathbf{u}_{123} = [w_1, w_2, w_3]$ tends to be very close to reduced. Since $\mathbf{u}_{123}$ has good edges, so do the other three 1-sharblies on the right of (34). Figure 12 illustrates this principle in the cone model of $\mathfrak{H}$. The reducing points w_i are the vertices of two adjacent triangles, which means $\mathbf{u}_{123}$ can't be too bad.

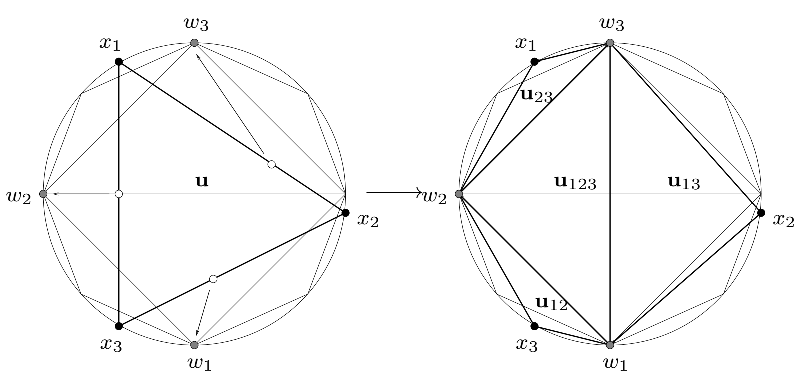

Fig. 12 The transformation of Figure 11 viewed in the cone. The new 1-sharblies are closer to being reduced.

Finally we come to the last issue: what do we do if a 1-sharbly has reduced edges but is not itself reduced? This phenomenon does occur, quite often in fact, although not for $\mathbf{G} = \mathrm{SL}_2/\mathbb{Q}$. This phenomenon is an artifact of the flat factor and reflects the presence of the unit group $\mathscr{O}^\times$. What happens is that we wind up with a 1-sharbly $[x_1, x_2, x_3]$ when we really want $[\varepsilon_1 x_1, \varepsilon_2 x_2, \varepsilon_3 x_3]$ for some $\varepsilon_i \in \mathscr{O}^\times$ of infinite order. The fix in this case involves "subdividing edges at infinity." The process is similar to what we describe above, but is in fact easier since it is an abelian analogue. For more details, we refer to [36, §7] for an example.

	$\dim D$	$\mathrm{vcd}(\Gamma)$	$t(\Gamma)$	$b(\Gamma)$
$\mathrm{SL}_2/\mathbb{Q}$	2	1	1	1
$\mathrm{SL}_3/\mathbb{Q}$	5	3	3	2
$\mathrm{SL}_4/\mathbb{Q}$	9	6	5	4
$\mathrm{SL}_5/\mathbb{Q}$	14	10	8	6
$R_{F_{0,1}/\mathbb{Q}}(\mathrm{SL}_2)$	3	2	2	1
$R_{F_{2,0}/\mathbb{Q}}(\mathrm{SL}_2)$	4	3	2	2
$R_{F_{2,0}/\mathbb{Q}}(\mathrm{GL}_2)$	5	4	3	2
$R_{F_{1,1}/\mathbb{Q}}(\mathrm{SL}_2)$	5	4	3	2
$R_{F_{1,1}/\mathbb{Q}}(\mathrm{GL}_2)$	6	5	4	2
$R_{F_{0,2}/\mathbb{Q}}(\mathrm{SL}_2)$	6	5	4	2
$R_{F_{0,2}/\mathbb{Q}}(\mathrm{GL}_2)$	7	6	5	2

Table 1 Examples of dimensions of symmetric spaces (§4), virtual cohomological dimensions (§8), and cuspidal ranges (§11). $F_{r,s}$ means that F is a number field with $F \otimes \mathbb{R} \simeq \mathbb{R}^r \times \mathbb{C}^s$.

12 Computational examples

As proof of concept, we conclude by presenting some examples of computations done with these techniques [35, 36]. Other examples can be found in [5–8, 37]. Here we consider GL_2 over two number fields:

- F_1 is the quartic field $\mathbb{Q}(\zeta)$, where ζ is a primitive fifth root of unity. Thus $F_1 \simeq \mathbb{Q}[x]/(x^4 + x^3 + x^2 + x + 1)$. F_1 has discriminant 5^3, is Galois with Galois group isomorphic to $\mathbb{Z}/4\mathbb{Z}$, and is a CM extension of its totally real subfield $F_1^+ = \mathbb{Q}(\eta)$, $\eta = \zeta + 1/\zeta$. The ring of integers $\mathscr{O}_1$ has class number one and as a $\mathbb{Z}$-module equals $\mathbb{Z}[\zeta]$.
- F_2 is the nonreal cubic field of discriminant -23. Thus $F_2 = \mathbb{Q}(t)$ with t a root of $x^3 - x^2 + 1$. The field F_2 is not Galois over $\mathbb{Q}$ (obviously), and has no subfields other than $\mathbb{Q}$. The ring of integers $\mathscr{O}_2$ has class number one and admits a power basis: $\mathscr{O}_2 = \mathbb{Z}[t]$.

In some sense these two fields are as different as possible. F_1 is highly symmetric, in fact much more symmetric than a complex quartic field "deserves" to be, since it is a cyclotomic extension. F_2, on the other hand, has no symmetry at all. But this does not mean F_2 is devoid of charm. For example, it is the first cubic field in databases of such fields ordered by absolute value of discriminant. Its Galois closure is the Hilbert class field of $\mathbb{Q}(\sqrt{-23})$, which means F_2 often appears in algebraic number theory courses as an appealing example. The field F_2 is the invariant trace field of the Weeks–Mateo's–Fomenko manifold, which is the contender for the orientable hyperbolic 3-manifold of minimal volume [45]. This field even appears in architecture and design: the unique real root of $x^3 - x - 1$, which lives in F_2, is known in some circles as the *plastic number*, and apparently functions as a generalization of the golden ratio [50]

The focus of the papers [35, 36] is testing the relationship between automorphic forms on GL_2 and elliptic curves. In particular, we wanted to test whether elliptic curves over these number fields were modular, in the weakest possible sense that still has content: matching of partial L-functions on both sides, at least as far as we could compute. We recall what this means.

Let F be a number field with ring of integers $\mathscr{O}$. Let E be an elliptic curve defined over F with conductor $\mathfrak{n} \subset \mathscr{O}$. Given any prime $\mathfrak{p}$ not dividing $\mathfrak{n}$, one defines an integer $a_\mathfrak{p}(E)$ by

$$a_\mathfrak{p}(E) = \mathrm{N}(\mathfrak{p}) + 1 - N_\mathfrak{p}. \tag{35}$$

Here $\mathrm{N}(\mathfrak{p})$ is the cardinality of the residue field $\mathbb{F}_\mathfrak{p} = \mathscr{O}/\mathfrak{p}$, and $N_\mathfrak{p}$ is the number of points E has in $\mathbb{F}_\mathfrak{p}$. These numbers can be assembled into a (partial) L-function $L(s, E)$ by making an Euler product

$$L(s, E) = \prod_{\mathfrak{p} \nmid \mathfrak{n}} \frac{1}{1 - a_\mathfrak{p}(E)\,\mathrm{N}(\mathfrak{p})^{-s} + \mathrm{N}(\mathfrak{p})^{1-2s}}. \tag{36}$$

The product (36) can be completed with certain local factors at the "bad primes"—those $\mathfrak{p}$ that divide $\mathfrak{n}$—and with Gamma factors for the archimedian places of F, so that the resulting L-function has a functional equation of the form $s \mapsto 2 - s$ and has analytic continuation to the complex plane. For more details, see [58].

On the automorphic side, let f be an cuspidal automorphic form on GL_2/F, also of conductor $\mathfrak{n}$. Assume f is an eigenform for the Hecke operators. Then f produces a collection of eigenvalues $a_\mathfrak{p}(f)$, one for each for each prime $\mathfrak{p}$ not dividing $\mathfrak{n}$, and we can make an L-function $L(s, f)$ using (36), with $a_\mathfrak{p}(E)$ replaced with $a_\mathfrak{p}(f)$.

Now a modularity result predicts, at the lowest level, that there should be a tight relationship between the L-functions constructed from elliptic curves and modular forms. In particular, given E, one hopes to find f such that

$$L(s,E) = L(s,f). \tag{37}$$

In particular one should have $a_{\mathfrak{p}}(f) = a_{\mathfrak{p}}(E)$ for all $\mathfrak{p} \nmid \mathfrak{n}$. We say f is *attached* to E if (37) is true. Conversely, given a cuspidal automorphic form with rational Hecke eigenvalues, one expects to find an elliptic curve attached to it in this sense, with matching conductor. This is motivated by what happens when $F = \mathbb{Q}$. In this case, thanks to work of many people, we know that this is true. Hence one might hope the same phenomenon happens over general number fields, or at least might check the extent to which it does. As it turns out, there are some subtleties; we indicate some below.

The papers [35, 36] test this relationship in the following way. First we generate a small database of elliptic curves over F in the most ingenuous way: we search over a family of Weierstrass equations

$$y^2 + a_1xy + a_3 = x^3 + a_2x^2 + a_4x + a_6$$

by taking a_1, a_2, a_3, a_4, a_6 in some bounded subset of $\mathscr{O}$. On the automorphic side, we look for appropriate automorphic forms by computing cohomology of the subgroups $\Gamma_0(\mathfrak{n}) \subset \mathrm{GL}_2(\mathscr{O})$ using techniques described above. In particular we can use the reduction algorithm in §11 to decompose the cohomology into eigenspaces for the Hecke operators $T_{\mathfrak{p}}$ for a range of primes $\mathfrak{p}$. We found excellent agreement between the arithmetic and automorphic sides:

- For each elliptic curve E with norm conductor within the range of our cohomology computations, we found a cuspidal cohomology class with rational Hecke eigenvalues that matched the point counts for E as in (35), for every Hecke operator that we checked.[9]
- For each cuspidal cohomology class with rational Hecke eigenvalues, we found a corresponding elliptic curve whose point counts matched every eigenvalue we computed, with just one exception: over the CM field F_1, we found an eigenclass with rational eigenvalues that corresponds to an abelian surface over F_1^+.[10]

We now present more information for our specific fields.[11]

---

[9] For general number fields, it is not expected that every elliptic curve should correspond to a cusp form in this way. For instance, suppose F is complex quadratic. Then if E is defined over F and has complex multiplication by an order in $\mathscr{O}_F$, then E should correspond to an Eisenstein series, cf. [25].

[10] Similar phenomena happen over complex quadratic fields [18], and can be expected to happen over any CM field.

[11] The cohomology computations in the following are simplified by the fact that F_1 and F_2 each have class number 1. One can still perform these computations for fields with higher class numbers, although it is best to work adelically. In practice this means that one has to work with several copies of the locally symmetric spaces instead of one, each equipped with its own Koecher decomposition. However, such complications are not always necessary. For F imaginary quadratic with odd class number, for instance, Lingham [44] developed a technique to work with a single connected component.

12.1 The field $F_1 = \mathbb{Q}(\zeta)$

The positivity domain C_1 is $\mathrm{Herm}_2(\mathbb{C})^2$, which has real dimension 8. The Koecher polyhedron and the perfect forms for F_1 were determined by D. Yasaki [72]. (For GL_2 over F_1 and F_2, the facets of the Koecher polyhedron are in bijection with the perfect forms.) Modulo the action of $\mathrm{GL}_2(\mathscr{O}_1)$ there is just one perfect form with 240 minimal vectors. Any two minimal vectors that differ by multiplication by a torsion element of $\mathscr{O}_1^\times$ determine the same vertex of the Koecher polyhedron. Thus the Koecher fan Σ_1 contains one top-dimensional cone; the corresponding facet of the Koecher polyhedron has $24 = 240/10$ vertices. One can easily compute the rest of the cones in Σ_1 modulo $\mathrm{GL}_2(\mathscr{O}_1)$. One finds 5 $\mathrm{GL}_2(\mathscr{O}_1)$-classes of 7-cones, 10 classes of 6-cones, 11 classes of 5-cones, 9 classes of 4-cones, 4 classes of 3-cones, 2 classes of 2-cones, and 1 class of 1-cones.

The symmetric space X_1 attached to $\mathbf{G}_1 = R_{F_1/\mathbb{Q}}\mathrm{GL}_2$ has dimension 7. Since the derived subgroup $\mathbf{G}_1'$ has $\mathbb{Q}$-rank one, the virtual cohomological dimension is 6. By Table 1 the cuspidal cohomology occurs in degrees $3, 4, 5$. Using 1-sharblies we compute with $H^5(\Gamma;\mathbb{C})$.[12] The Koecher fan is simplicial in dimensions $2, 3, 4$, so we can identify $\mathscr{R}_*$ in these degrees with a subcomplex of $\mathscr{S}_*$.

We were able to compute H^5 for all levels $\mathfrak{n}$ with $\mathrm{N}(\mathfrak{n}) \leq 4941$. For $\mathfrak{n}$ prime we were able to carry the computations further to $\mathrm{N}(\mathfrak{n}) \leq 7921$. We also computed Hecke operators on H^5. For all levels we were able to compute at least up to $T_\mathfrak{q}$ with $\mathfrak{q} \subset \mathscr{O}_1$ prime satisfying $\mathrm{N}(\mathfrak{q}) \leq 41$; at some smaller levels we computed much further.

We found a variety of phenomena:

1. We found classes that seemed to correspond to elliptic curves defined over F_1 that were not base changes from any subfield of F_1. The smallest conductor we found was $\mathfrak{p}$ with $\mathrm{N}(\mathfrak{p}) = 701$. The elliptic curve had equation $(a_1, a_2, a_3, a_4, a_6) = (-\zeta-1, \zeta^2-1, 1, -\zeta^2, 0)$. Altogether we found 13 such examples.[13]
2. We found classes that seemed to correspond to base changes of curves from $\mathbb{Q}$ and curves/abelian surfaces from F_1^+.
3. We found "old" cohomology classes, namely eigenclasses whose eigenvalues matched those of eigenclasses appearing at smaller level norms.
4. One eigenclass ξ at level norm 3025 deserves some extra discussion. Let $\mathfrak{m} \subset \mathscr{O}_{F_1^+}$ be the ideal $\mathfrak{p}_5^2\mathfrak{p}_{11}$. The space of parallel weight 2 Hilbert modular newforms of level $\mathfrak{m}$ contains an eigenform g with Hecke eigenvalues $a_\mathfrak{q}$ in the field F_1^+. For any prime $\mathfrak{q} \subset \mathscr{O}_{F^+}$, let $q \in \mathbb{Z}$ be the prime under $\mathfrak{q}$. Then we have $a_\mathfrak{q}(g) = 0$ if $q = 5$, and

[12] Actually, to avoid precision problems we work with the large finite field $\mathbb{F}_{12379}$ instead of $\mathbb{C}$.

[13] One curve was found by Mark Watkins using the method of Cremona–Lingham [20]; see the appendix to [35].

$$a_{\mathfrak{q}}(g) \in \begin{cases} \mathbb{Z} & \text{if } q = 1 \bmod 5, \\ \mathbb{Z}\cdot\sqrt{5} & \text{if } q = 2,3,4 \bmod 5. \end{cases} \tag{38}$$

These conditions (38), together with the Hecke eigenvalues of g, imply that the L-series $L(s,g)L(s,g\otimes\varepsilon)$ agrees with the L-series attached to our eigenclass ξ, where ε is the unique quadratic character of $\mathrm{Gal}(F_1/F_1^+)$. Thus ξ has rational Hecke eigenvalues, but does not correspond to an elliptic curve over F_1. Instead, it can be attributed to an abelian surface over F_1^+ with extra symmetry, and thus gives an example over F_1 of phenomena first seen in [18] for complex quadratic fields.

Other than item 4, we found perfect matching between elliptic curves and cuspidal cohomology classes with rational eigenvalues.

12.2 The field $F_2 = \mathbb{Q}(t)$

For this field the positivity domain C_2 is $\mathrm{Sym}_2(\mathbb{R})\times\mathrm{Herm}_2(\mathbb{C})$, which has real dimension 7. Modulo the action of $\mathrm{GL}_2(\mathscr{O}_2)$ there are nine perfect forms. Of these, seven give simplicial cones in the Koecher fan Σ_2; for the other two the facets of the Koecher polyhedron have eight and nine vertices respectively. [14] For the rest of Σ_2, one finds 35 $\mathrm{GL}_2(\mathscr{O}_2)$-classes of 6-cones, 47 classes of 5-cones, 31 classes of 4-cones, 10 classes of 3-cones, 2 classes of 2-cones, and 1 class of 1-cones.

The symmetric space X_2 attached to $\mathbf{G}_2 = R_{F_2/\mathbb{Q}}\mathrm{GL}_2$ has dimension 6. As before the derived subgroup $\mathbf{G}_2'$ has $\mathbb{Q}$-rank one; the virtual cohomological dimension is 5. The cuspidal cohomology occurs in degrees $3, 4$, and 1-sharblies enable us to compute $H^4(\Gamma;\mathbb{C})$.[15] As before the Koecher fan is simplical in the dimensions we care about, so in these degrees we can identify $\mathscr{R}_*$ with a subcomplex of $\mathscr{S}_*$.

We computed the cohomology at 308 different levels, including all ideals with level norm ≤ 835, and the Hecke operators:

1. As for F_1, we found examples of elliptic curves over F_2 that are not base changes from $\mathbb{Q}$ and that are apparently attached to eigenclasses of the appropriate levels. The first curve appears at level norm 89, and has equation $(a_1,a_2,a_3,a_4,a_6) = (t-1, -t^2-1, t^2-t, t^2, 0)$. Altogether 43 such curves were found in the range of our computations.
2. We found "old" cohomology classes.

[14] One cannot help but notice the contrast between the highly symmetric perfect cone for F_1 and the minimally symmetric perfect cones for F_2.

[15] See footnote (12).

3. We found one base change from $\mathbb{Q}$ to F: at level norm $529 = 23^2$, there is an eigenclass with eigenvalues in $\mathbb{Q}(\sqrt{5})$. This class is accounted for by the level 23 abelian surface over $\mathbb{Q}$ with real multiplication by $\mathbb{Q}(\sqrt{5})$.

References

1. A. Ash, D. Mumford, M. Rapaport, and Y. Tai., *Smooth compactifications of locally symmetric varieties*, Math. Sci. Press, Brookline, Mass., 1975.
2. A. Ash, *Deformation retracts with lowest possible dimension of arithmetic quotients of self-adjoint homogeneous cones*, Math. Ann. **225** (1977), no. 1, 69–76.
3. A. Ash, *Small-dimensional classifying spaces for arithmetic subgroups of general linear groups*, Duke Math. J. **51** (1984), no. 2, 459–468.
4. A. Ash, D. Grayson, and P. Green, *Computations of cuspidal cohomology of congruence subgroups of* $\mathrm{SL}(3, \mathbf{Z})$, J. Number Theory **19** (1984), no. 3, 412–436.
5. A. Ash, P. E. Gunnells, and M. McConnell, *Cohomology of congruence subgroups of* $\mathrm{SL}_4(\mathbb{Z})$, J. Number Theory **94** (2002), no. 1, 181–212.
6. A. Ash, P. E. Gunnells, and M. McConnell, *Cohomology of congruence subgroups of* $\mathrm{SL}(4, \mathbb{Z})$*. II*, J. Number Theory **128** (2008), no. 8, 2263–2274.
7. A. Ash, P. E. Gunnells, and M. McConnell, *Cohomology of congruence subgroups of* $\mathrm{SL}_4(\mathbb{Z})$*. III*, Math. Comp. **79** (2010), no. 271, 1811–1831.
8. A. Ash, P. E. Gunnells, and M. McConnell, *Torsion in the cohomology of congruence subgroups of* $\mathrm{SL}(4, \mathbb{Z})$ *and Galois representations*, J. Algebra **325** (2011), 404–415.
9. A. Ash and L. Rudolph, *The modular symbol and continued fractions in higher dimensions*, Invent. Math. **55** (1979), no. 3, 241–250.
10. A. Borel and J.-P. Serre, *Corners and arithmetic groups*, Comment. Math. Helv. **48** (1973), 436–491, Avec un appendice: Arrondissement des variétés à coins, par A. Douady et L. Hérault.
11. R. Bott and L. W. Tu, *Differential forms in algebraic topology*, Graduate Texts in Mathematics, vol. 82, Springer-Verlag, New York, 1982.
12. R. Brooks, *Platonic surfaces*, Comment. Math. Helv. **74** (1999), no. 1, 156–170.
13. K. S. Brown, *Cohomology of groups*, Graduate Texts in Mathematics, vol. 87, Springer-Verlag, New York, 1994, Corrected reprint of the 1982 original.
14. D. Bump, *Automorphic forms and representations*, Cambridge Studies in Advanced Mathematics, vol. 55, Cambridge University Press, Cambridge, 1997.
15. D. Bump, *Lie groups*, Graduate Texts in Mathematics, vol. 225, Springer-Verlag, New York, 2004.
16. J. Bygott, *Modular forms and modular symbols over imaginary quadratic fields*, Ph.D. thesis, Exeter University, 1999.
17. J. E. Cremona, *Hyperbolic tessellations, modular symbols, and elliptic curves over complex quadratic fields*, Compositio Math. **51** (1984), no. 3, 275–324.
18. J. E. Cremona, *Abelian varieties with extra twist, cusp forms, and elliptic curves over imaginary quadratic fields*, J. London Math. Soc. (2) **45** (1992), no. 3, 404–416.
19. J. E. Cremona, *Algorithms for modular elliptic curves*, second ed., Cambridge University Press, Cambridge, 1997.
20. J. E. Cremona and M. P. Lingham, *Finding all elliptic curves with good reduction outside a given set of primes*, Experiment. Math. **16** (2007), no. 3, 303–312.
21. J. E. Cremona and E. Whitley, *Periods of cusp forms and elliptic curves over imaginary quadratic fields*, Math. Comp. **62** (1994), no. 205, 407–429.
22. M. Dutour, F. Vallentin, and A. Schürmann, *Classification of perfect forms in dimension* 8, talk at Oberwolfach meeting *Sphere packings: Exceptional structures and relations to other fields*, November 2005.

23. M. Dutour Sikirić, A. Schürmann, and F. Vallentin, *A generalization of Voronoi's reduction theory and its application*, Duke Math. J. **142** (2008), no. 1, 127–164.
24. P. Elbaz-Vincent, H. Gangl, and C. Soulé, *Perfect forms and the cohomology of modular groups*, eprint arXiv:1001.0789 (2010), arXiv:1001.0789.
25. J. Elstrodt, F. Grunewald, and J. Mennicke, *On the group* $\mathrm{PSL}_2(\mathbf{Z}[i])$, Number theory days, 1980 (Exeter, 1980), London Math. Soc. Lecture Note Ser., vol. 56, Cambridge Univ. Press, Cambridge, 1982, pp. 255–283.
26. J. Faraut and A. Korányi, *Analysis on symmetric cones*, Oxford Mathematical Monographs, The Clarendon Press Oxford University Press, New York, 1994, Oxford Science Publications.
27. D. Flöge, *Zur Struktur der* PSL_2 *über einigen imaginär-quadratischen Zahlringen*, Math. Z. **183** (1983), no. 2, 255–279, URL: `http://dx.doi.org/10.1007/BF01214824`, doi:10.1007/BF01214824.
28. J. Franke, *Harmonic analysis in weighted* L_2*-spaces*, Ann. Sci. École Norm. Sup. (4) **31** (1998), no. 2, 181–279.
29. E. Freitag, *Hilbert modular forms*, Springer-Verlag, Berlin, 1990.
30. S. S. Gelbart, *Automorphic forms on adèle groups*, Princeton University Press, Princeton, N.J., 1975, Annals of Mathematics Studies, No. 83.
31. M. Greenberg and J. Voight, *Computing systems of Hecke eigenvalues associated to Hilbert modular forms*, Math. Comp. **80** (2011), no. 274, 1071–1092.
32. F. Grunewald, H. Helling, and J. Mennicke, SL_2 *over complex quadratic number fields. I*, Algebra i Logika **17** (1978), no. 5, 512–580, 622.
33. P. E. Gunnells, *Modular symbols for* $\mathbb{Q}$*-rank one groups and Voronoĭ reduction*, J. Number Theory **75** (1999), no. 2, 198–219.
34. P. E. Gunnells, *Computing Hecke eigenvalues below the cohomological dimension*, Experiment. Math. **9** (2000), no. 3, 351–367.
35. P. E. Gunnells, F. Hajir, and D. Yasaki, *Modular forms and elliptic curves over the field of fifth roots of unity*, Experimental Math. (to appear).
36. P. E. Gunnells and D. Yasaki, *Modular forms and elliptic curves over the cubic field of discriminant* -23, submitted, 2012.
37. P. E. Gunnells and D. Yasaki, *Hecke operators and Hilbert modular forms*, Algorithmic number theory, Lecture Notes in Comput. Sci., vol. 5011, Springer, Berlin, 2008, pp. 387–401.
38. G. Harder, *Eisenstein cohomology of arithmetic groups and its applications to number theory*, Proceedings of the International Congress of Mathematicians, Vol. I, II (Kyoto, 1990) (Tokyo), Math. Soc. Japan, 1991, pp. 779–790.
39. Harish-Chandra, *Eisenstein series over finite fields*, Functional analysis and related fields (Proc. Conf. M. Stone, Univ. Chicago, Chicago, Ill., 1968), Springer, New York, 1970, pp. 76–88.
40. S. Helgason, *Differential geometry, Lie groups, and symmetric spaces*, Graduate Studies in Mathematics, vol. 34, American Mathematical Society, Providence, RI, 2001, Corrected reprint of the 1978 original.
41. M. Koecher, *Beiträge zu einer Reduktionstheorie in Positivitätsbereichen. I*, Math. Ann. **141** (1960), 384–432.
42. R. P. Langlands, *On the functional equations satisfied by Eisenstein series*, Springer-Verlag, Berlin, 1976, Lecture Notes in Mathematics, Vol. 544.
43. J.-S. Li and J. Schwermer, *Automorphic representations and cohomology of arithmetic groups*, Challenges for the 21st century (Singapore, 2000), World Sci. Publ., River Edge, NJ, 2001, pp. 102–137.
44. M. Lingham, *Modular forms and elliptic curves over imaginary quadratic fields*, Ph.D. thesis, Nottingham, 2005.
45. C. Maclachlan and A. W. Reid, *The arithmetic of hyperbolic 3-manifolds*, Graduate Texts in Mathematics, vol. 219, Springer-Verlag, New York, 2003.

46. R. MacPherson and M. McConnell, *Classical projective geometry and modular varieties*, Algebraic analysis, geometry, and number theory (Baltimore, MD, 1988), Johns Hopkins Univ. Press, Baltimore, MD, 1989, pp. 237–290.
47. R. MacPherson and M. McConnell, *Explicit reduction theory for Siegel modular threefolds*, Invent. Math. **111** (1993), no. 3, 575–625.
48. J. I. Manin, *Parabolic points and zeta functions of modular curves*, Izv. Akad. Nauk SSSR Ser. Mat. **36** (1972), 19–66.
49. E. R. Mendoza, *Cohomology of* PGL_2 *over imaginary quadratic integers*, Bonner Mathematische Schriften [Bonn Mathematical Publications], 128, Universität Bonn Mathematisches Institut, Bonn, 1979, Dissertation, Rheinische Friedrich-Wilhelms-Universität, Bonn, 1979.
50. R. Padovan, *Proportion: Science, philosophy, architecture*, Taylor& Francis, 1999.
51. A. Page, *Computing arithmetic Kleinian groups*, ArXiv e-prints (2012), arXiv:1206.0087.
52. Plato, *The Collected Dialogues of Plato*, Princeton University Press, 2005.
53. V. Platonov and A. Rapinchuk, *Algebraic groups and number theory*, Pure and Applied Mathematics, vol. 139, Academic Press Inc., Boston, MA, 1994, Translated from the 1991 Russian original by Rachel Rowen.
54. A. Schürmann, *Computational geometry of positive definite quadratic forms*, University Lecture Series, vol. 48, American Mathematical Society, Providence, RI, 2009, Polyhedral reduction theories, algorithms, and applications.
55. J. Schwermer, *Geometric cycles, arithmetic groups and their cohomology*, Bull. Amer. Math. Soc. (N.S.) **47** (2010), no. 2, 187–279.
56. J.-P. Serre, *Cohomologie des groupes discrets*, Prospects in mathematics (Proc. Sympos., Princeton Univ., Princeton, N.J., 1970), Princeton Univ. Press, Princeton, N.J., 1971, pp. 77–169. Ann. of Math. Studies, No. 70.
57. G. Shimura, *Introduction to the arithmetic theory of automorphic functions*, Publications of the Mathematical Society of Japan, No. 11. Iwanami Shoten, Publishers, Tokyo, 1971, Kanô Memorial Lectures, No. 1.
58. J. H. Silverman, *The arithmetic of elliptic curves*, second ed., Graduate Texts in Mathematics, vol. 106, Springer, Dordrecht, 2009.
59. C. Soulé, *Cohomologie de* $SL_3(Z)$, C. R. Acad. Sci. Paris Sér. A-B **280** (1975), no. 5, Ai, A251–A254.
60. N. Steenrod, *Homology with local coefficients*, Ann. of Math. (2) **44** (1943), 610–627.
61. W. Stein, *Modular forms, a computational approach*, Graduate Studies in Mathematics, vol. 79, American Mathematical Society, Providence, RI, 2007, With an appendix by Paul E. Gunnells.
62. B. van Geemen, W. van der Kallen, J. Top, and A. Verberkmoes, *Hecke eigenforms in the cohomology of congruence subgroups of* $\mathrm{SL}(3, \mathbf{Z})$, Experiment. Math. **6** (1997), no. 2, 163–174.
63. J. W. Vick, *Homology theory*, second ed., Graduate Texts in Mathematics, vol. 145, Springer-Verlag, New York, 1994, An introduction to algebraic topology.
64. D. A. Vogan, Jr., *Cohomology and group representations*, Representation theory and automorphic forms (Edinburgh, 1996), Proc. Sympos. Pure Math., vol. 61, Amer. Math. Soc., Providence, RI, 1997, pp. 219–243.
65. D. A. Vogan, Jr. and G. J. Zuckerman, *Unitary representations with nonzero cohomology*, Compositio Math. **53** (1984), no. 1, 51–90.
66. K. Vogtmann, *Rational homology of Bianchi groups.*, Math. Ann. **272** (1985), no. 3, 399–419.
67. J. Voight, *Computing automorphic forms on Shimura curves over fields with arbitrary class number*, Algorithmic number theory, Lecture Notes in Comput. Sci., vol. 6197, Springer, Berlin, 2010, pp. 357–371.
68. G. Voronoĭ, *Nouvelles applications des paramétres continus à la théorie des formes quadratiques, I. Sur quelques propriétés des formes quadratiques positives parfaites*, J. Reine Angew. Math. **133** (1908), 97–178.

69. D. Yasaki, *On the existence of spines for $\mathbb{Q}$-rank 1 groups*, Selecta Math. (N.S.) **12** (2006), no. 3-4, 541–564.
70. D. Yasaki, *An explicit spine for the Picard modular group over the Gaussian integers*, J. Number Theory **128** (2008), no. 1, 207–234.
71. D. Yasaki, *Bianchi modular forms*, package for the Magma Computer Algebra System, 2009.
72. D. Yasaki, *Binary Hermitian forms over a cyclotomic field*, J. Algebra **322** (2009), 4132–4142.