

Orbits in the Leech Lattice

Daniel Allcock

CONTENTS

- 1. Introduction
- 2. Notation and Background
- 3. Orbits Under Co_0
- 4. Orbits under Co_0 : Performance
- 5. Orbits Under $2^{12}:M_{24}$
- 6. Orbits Under M_{24}
- 7. Remarks
- Acknowledgments
- References

We provide an algorithm for determining whether two vectors in the Leech lattice are equivalent under its isometry group, the Conway group Co_0 of order $\sim 8 \times 10^{18}$. Our algorithm reduces the test of equivalence to at most four tests under the subgroup $2^{12}:M_{24}$ and a test under this subgroup to at most 12 tests under M_{24} . We also give algorithms for testing equivalence under these two subgroups. We describe our intended applications to the symmetry groups of Lorentzian lattices and the enumeration of lattices of dimension ~ 24 with good properties such as having small determinant. Our methods rely on and develop the work of R. T. Curtis.

1. INTRODUCTION

The Leech lattice Λ is a lattice in 24-dimensional Euclidean space with many remarkable properties, for us the most important of which is that its isometry group (modulo $\{\pm I\}$) is one of the sporadic finite simple groups. The isometry group is called the Conway group Co_0 , and our purpose is to present an algorithm for a computer to determine whether two given vectors of Λ are equivalent under Co_0 . The group is fairly large, of order $> 8 \times 10^{18}$, so the obvious try-every-isometry algorithm is useless. Instead, we use the geometry of Λ to build a faster algorithm; along the way we build algorithms for the problem of equivalence of vectors in \mathbb{R}^{24} under the subgroups $2^{12}:M_{24}$ and M_{24} , where M_{24} is the Mathieu group permuting the 24 coordinates. One can implement and use the algorithms without any deep familiarity with Co_0 and M_{24} , although more background is needed for the proofs of correctness and the bounds on performance.

The motivation for the algorithm is the problem of extending results from [Vinberg 72, Vinberg and Kaplinskaja 78, Borchers 87] on the isometry groups of the unimodular Lorentzian lattices $I_{n,1}$ to the case $n = 24$. The idea (following [Conway and Sloane 82, Borchers 87]) is that $\text{Aut } I_{n,1}$ is best understood when embedded in the isometry group of the even unimodular Lorentzian lattice $II_{25,1}$. This larger group has a particularly simple

2000 AMS Subject Classification: Primary 20D08, 11H06

Keywords: Leech lattices, S-lattices, Golay code, orbit

structure, discovered in [Conway 83]: its reflection subgroup has a fundamental domain Δ with one facet for each element of Λ , and the group Co_∞ of affine isometries of Λ is the subgroup of $\text{Aut } II_{25,1}$ preserving Δ . This means that the equivalence problem of two points in hyperbolic 25-space can be determined by applying reflections to bring them into Δ and then testing their equivalence under Co_∞ . Since $Co_\infty = \Lambda:Co_0$, this essentially reduces to equivalence testing under Co_0 .

Borcherds [Borcherds 88, Borcherds 84] used these techniques involving $II_{25,1}$ to enumerate the unimodular lattices in dimensions 24 and 25 and the bimodular lattices of dimension 25. Similar calculations with Co_∞ were used by [Conway et. al. 82] to enumerate the deep holes of Λ (see [Borcherds 85] for an easier approach) and by [Borcherds et. al. 88] to enumerate the shallow holes. All of these enumerations required extremely lengthy hand computation and detailed familiarity with Λ and Co_0 . There is no reason to doubt these enumerations, but our algorithm could be used to verify them. Also, it might be useful for working with the fake monster lie algebra (see [Borcherds 90]), whose root lattice is $II_{25,1}$.

Background on Λ , M_{24} , and certain of its sublattices, called \mathcal{S} -lattices, appears in Section 2. The algorithm for Co_0 appears in Section 3 and relies essentially on a theorem of Curtis [Curtis 73]: every isometry between two \mathcal{S} -lattices extends to an isometry of Λ . Our algorithm proceeds by constructing a tree of data, so the natural worry is that branching in the tree could require an exponential amount of computation. In Section 4, we prove that this is not a problem; regardless of the height of the tree its width is bounded by 16. The algorithm reduces a test of equivalence of two vectors either to an application of Curtis' theorem or to a few tests (at most four) of equivalence under $2^{12}:M_{24}$, which is much easier to deal with because it consists of signed permutations. In Section 5, we show how to reduce a test of equivalence under $2^{12}:M_{24}$ to at most 12 tests under M_{24} , and, in Section 6, we provide an algorithm for testing equivalence under M_{24} . Any other algorithms for $2^{12}:M_{24}$ and M_{24} could be used as "subroutines" in place of ours when testing equivalence under Co_0 . Section 7 contains a few remarks.

This paper is a development of part of my dissertation [Allcock 96], and I am grateful to my "unofficial" thesis advisor R. Borcherds for suggesting the problem. The paper has been rewritten from scratch, and the results of Section 4 are completely new. The algorithm for M_{24} is also new; the original one involved fewer special

cases but was much more intricate. The original preprint received limited circulation under the title "Recognizing Equivalence of Vectors in the Leech Lattice."

2. NOTATION AND BACKGROUND

We use the ATLAS notation for groups [Conway et. al. 85], so we say that a group has structure $G:H$ if it has a normal subgroup G , the quotient of which is H . If we write $G:H$ then the extension splits, and if we write $G \cdot H$ then it does not. We write p^n for the direct product of n copies of the cyclic group of order p .

To describe M_{24} , we follow [Conway 88]. Let Ω be a fixed set of size 24, and consider 2^Ω as a vector space over the Galois field \mathbb{F}_2 , addition being given by symmetric difference. We often refer to a size n set as an n -ad, or a monad, duad, triad, tetrad, etc. The Golay code \mathcal{C} is a subspace of 2^Ω of dimension 12 and weight (= set size) distribution $0^1 8^{759} 12^{2576} 16^{759} 24^1$. Elements of \mathcal{C} are called codewords or \mathcal{C} -sets. The octads and 16-ads of \mathcal{C} are called *special* and the dodecads are called *umbral*; every octad, dodecad, and 16-ad appearing in this paper is special (respectively, umbral) unless otherwise noted, so we will drop the qualifier except for emphasis. Every pentad of Ω lies in a unique octad. A partition of Ω into three special octads is called a trio (*three octads*). Hand calculations involving \mathcal{C} are best done using Curtis' Miracle Octad Generator (MOG), which arranges the 24 points of Ω into a 4×6 array, which falls into three 4×2 bricks. Each brick is an octad and the set of all three is called the standard trio. We refer to the leftmost MOG column as the standard tetrad. See [Conway 88] for a wealth of information about the MOG array. (Conway's MOG differs by a reflection from Curtis' original array [Curtis 76].)

The Mathieu group M_{24} is the group of permutations of Ω that preserve \mathcal{C} ; it is simple of order 244,823,040 and acts 5-transitively on Ω . Therefore, the pointwise stabilizers M_{23} , M_{22} , M_{21} , M_{20} , and M_{19} of a monad, duad, triad, tetrad, and pentad are well defined up to conjugacy. M_{24} is also transitive on octads and dodecads; the stabilizer of a dodecad is M_{12} , another of Mathieu's sporadic simple groups. M_{24} is also transitive on trios.

The Golay cocode \mathcal{C}^* is $2^\Omega/\mathcal{C}$, also of dimension 12. Any element of \mathcal{C}^* has either a unique smallest representative in 2^Ω (in which case this subset of Ω has size ≤ 3 and is usually identified with the element of \mathcal{C}^*) or else has six smallest representatives (in which case they form a sextet, so named because the *six* representatives are mutually disjoint *tetrads*). Every tetrad lies in a unique

sextet. If the image in \mathcal{C}^* of a set $X \subseteq \Omega$ is a sextet, then we say that X represents that sextet. We call a subset of Ω small if it has size ≤ 4 and large otherwise. If $X \subseteq \Omega$ does not represent a sextet, then the unique small set to which X is congruent modulo \mathcal{C}^* is called the small representative of X . By transitivity on tetrads, M_{24} is transitive on sextets. We define the standard sextet to be the sextet containing the standard tetrad; this consists of the six MOG columns. The subgroup of M_{24} preserving the standard sextet is called the sextet group and has structure $2^6:3 \cdot S_6$ and order 138,240.

A basic tool in our algorithm for M_{24} is reducing a set $X \subseteq \Omega$ modulo \mathcal{C} . By this, we mean finding a small set to which X is congruent modulo \mathcal{C} . When this set has size ≤ 3 , then it is uniquely determined and we denote it by \bar{X} , and when it has size 4, then X represents a sextet. Computing a small representative is just a rephrasing of the problem of finding a codeword nearest to X , for which there are several algorithms in the literature, e.g., [Conway and Sloane 86, Vardy and Be'ery 91].

We equip $\mathbb{R}^{24} = \mathbb{R}^\Omega$ with the inner product $x \cdot y = \frac{1}{8} \sum_{i \in \Omega} x_i y_i$, and by the norm of $x \in \mathbb{R}^{24}$ we mean $x^2 = x \cdot x$. The Leech lattice Λ is the set of vectors x with integral coordinates x_i satisfying

- (i) the coordinates are all congruent modulo 2; write m for their common congruence class;
- (ii) the set of i for which x_i takes any given value modulo 4 is a \mathcal{C} -set;
- (iii) the sum of the x_i is congruent to $4m$ modulo 8.

One can check that Λ is an even unimodular lattice that contains vectors of all even norms except 2. The type of a lattice vector is half its norm. (This is not important, but it is part of the vocabulary of Λ .) M_{24} acts on Λ by permuting coordinates, and \mathcal{C} acts by negation of coordinates on \mathcal{C} -sets. Together they generate a group $2^{12}:M_{24}$. The full group of isometries of Λ is called Co_0 in honor of Conway, who found additional symmetries (see below); it has order 8,315,553,613,086,720,000 and, modulo its center $\{\pm I\}$, it is a simple group. It acts transitively on lattice vectors of each of the types 2, 3, and 4 (and 5 and 7 too, though we will not need this). A sublattice L of Λ is called primitive if $L = (L \otimes \mathbb{Q}) \cap \Lambda$.

A beautiful property of Λ is the simple structure of its residue classes mod 2. Each class has a representative of type 0, 2, 3, or 4; such vectors are called short. The only congruences among short vectors are that each of type 2 or 3 is congruent to its negative, and that vectors of type 4 fall into sets of 48, called frames. Each

frame consists of 24 mutually orthogonal pairs of antipodal vectors, for example, the standard frame consists of the vectors $(0, \dots, 0, \pm 8, 0, \dots, 0)$, where the ± 8 may occur in any position. Co_0 acts transitively on frames, and the stabilizer of the standard frame is the group $2^{12}:M_{24}$ discussed above. We also refer to the element of $\Lambda/2\Lambda$ represented by the members of a frame as a frame.

If $x \in \Lambda$, then we write \bar{x} for the image of x in $\Lambda/2\Lambda$, and when \bar{x} is not a frame, we write \hat{x} for a short representative for \bar{x} , which is well defined up to sign. If L is a sublattice of Λ , then we write \bar{L} for its image in $\Lambda/2\Lambda$.

A basic tool in our Co_0 algorithm is reducing a vector $x \in \Lambda$ modulo 2, by which we mean finding a short vector s to which x is congruent modulo 2Λ . This can be accomplished using an algorithm for decoding Λ , which means finding a lattice point closest to any given element of \mathbb{R}^{24} . Such algorithms appear in [Conway and Sloane 86, Vardy and Be'ery 93]. To find s , decode $x/2$ to obtain $\lambda \in \Lambda$. By [Conway et. al. 82] (see also [Borchers 85]), every point of \mathbb{R}^{24} lies within $\sqrt{2}$ of Λ , so $(\lambda - x/2)^2 \leq 2$ and $s = x - 2\lambda$ is short. If $s^2 < 8$, then we know that s is actually well defined up to sign, and $\pm s$ are the possibilities for \hat{x} . We wonder whether there is a faster algorithm for reducing an element of $\Lambda \bmod 2$; the problem seems so much more specialized than the general decoding problem that perhaps a more specialized algorithm could be faster. The speed of decoding Λ is the main factor in the speed of the Co_0 algorithm.

Curtis introduced the concept of an \mathcal{S} -lattice in [Curtis 73]; the idea is to consider sublattices L of Λ none of whose elements represent frames and for which there is no obvious enlargement. Precisely, L is an \mathcal{S} -lattice if no element of L represents a frame, and if for each $x \in L$, L contains \hat{x} and $(x - \hat{x})/2$. The type of an \mathcal{S} -lattice is the formal symbol $2^a 3^b$, where a (respectively, b) is the number of antipodal pairs of type 2 (respectively, 3) vectors in L . One has $a + b + 1 = 2^{\dim L}$, and Curtis completely classified the \mathcal{S} -lattices, finding 11 orbits under Co_0 . (He used a slightly different definition of \mathcal{S} -lattice, but his arguments work just as well for this definition, which is from [Conway et. al. 85].) While we do not need the classification, we remark that the largest \mathcal{S} -lattice, of type $2^{27} 3^{36}$, plays a major role in the performance analysis in Section 4.

Also, two \mathcal{S} -lattices of the same type are Co_0 -equivalent, so referring to an \mathcal{S} -lattice by its type is unambiguous (up to Co_0). Finally, the following theorem is implicit in [Curtis 73] and stated explicitly in [Conway et. al. 85]. It is essential for the validity of our Co_0 algorithm.

Theorem 2.1. *Any isometry between two \mathcal{S} -lattices extends to an isometry of Λ .*

Our final prerequisite is a small tool used in the Co_0 algorithm; we will need to be able to find an element of Co_0 carrying any given frame to the standard one. It is enough to carry any given type 4 vector into the standard frame. We need the extra automorphism η of Λ discovered in [Conway 69]. We write e_i for the vector $(0, \dots, 0, 1, 0, \dots, 0)$ with the 1 in the i th place; for each i we write $T(i)$ for the MOG column containing i and $e_{T(i)}$ for the vector with coordinates equal to 1 on $T(i)$ and 0 elsewhere. The isometry η is defined by first negating the leftmost MOG column (the standard tetrad) and then sending each e_i to $e_i - \frac{1}{2}e_{T(i)}$.

Algorithm 2.2. Given $v \in \Lambda$ of type 4, this algorithm carries v to a member of the standard frame.

Step 1.	2222	$2\bar{2}\bar{2}\bar{2}$	$\bar{4}000$
Step 2.	4222	$4\bar{2}\bar{2}\bar{2}$	$5\bar{1}\bar{1}\bar{1}$
Step 3.	3333	$3\bar{3}\bar{3}\bar{3}$	$\bar{6}000$
Step 4.	5331	$5\bar{3}\bar{3}\bar{1}$	$\bar{6}220$
Step 5.	6222	$6\bar{2}\bar{2}\bar{2}$	$\bar{4}044$
Step 6.	4444	$4\bar{4}\bar{4}\bar{4}$	$\bar{8}000$

By a step “AAAA BBBB CCCC” we mean the following operation. Search for a tetrad on which the absolute values of the coordinates of w are the numbers AAAA; if no such tetrad is found, then proceed to the next step. If one is found, then find $\sigma \in 2^{12}:M_{24}$ such that the coordinates of $\sigma(w)$ on the standard tetrad are exactly the numbers BBBB, replace w by $\eta \circ \sigma(w)$, whose coordinates on the standard tetrad are then the numbers CCCC, and then proceed to the next step. We have written \bar{m} for $-m$.

The proof is a straightforward examination of the list of type 4 vectors in Λ [Conway and Sloane 88, page 133]. In order to use the algorithm, one must be able to find an element of M_{24} carrying any given tetrad to the standard one and be able to find an element of \mathcal{C} that changes the signs on the standard tetrad in any desired way. Both tasks are accomplished with a few precomputed elements of M_{24} and \mathcal{C} .

3. ORBITS UNDER Co_0

The idea of our algorithm for detecting equivalence of $v, w \in \Lambda$ under Co_0 is simple and best motivated by considering an example of how the algorithm might work. Suppose for simplicity that neither v nor w lies in 2Λ .

If one represents a frame and the other does not, then they are clearly inequivalent. If both represent frames (say ϕ, ψ), then we find $g, h \in Co_0$ carrying ϕ and ψ to the standard frame. It is easy to see that v and w are Co_0 -equivalent if and only if $g(v)$ and $h(w)$ are equivalent under the stabilizer $2^{12}:M_{24}$ of the standard frame. This is a reduction to a much smaller problem, since the subgroup acts by signed permutations. If neither represents a frame, then \hat{v} and \hat{w} are well defined up to sign; we suppose for this example that $\hat{v} \cdot v \neq 0$ and $\hat{w} \cdot w \neq 0$, which is usually the case. We choose \hat{v} and \hat{w} so that $\hat{v} \cdot v$ and $\hat{w} \cdot w$ are positive. Then, v and w are Co_0 -equivalent if and only if the lattice spanned by v and \hat{v} is Co_0 -equivalent to that spanned by w and \hat{w} , by an isometry carrying v to w and \hat{v} to \hat{w} . We can enlarge these two lattices by adjoining the lattice vectors $(v + \hat{v})/2$ and $(w + \hat{w})/2$ and then reducing these new vectors modulo 2. If we get frames, then we can reduce the problem to $2^{12}:M_{24}$ as before; otherwise, we can enlarge our lattices by adjoining the short representatives of $(v + \hat{v})/2$ and $(w + \hat{w})/2$. We can repeat the process until either we find frames or the lattices stop growing. In the former case we reduce to $2^{12}:M_{24}$, and in the latter case we will use the theorem of Curtis on \mathcal{S} -lattices (Theorem 2.1) to determine equivalence or inequivalence.

The main point of the example is that one should keep track of sublattices of Λ , not just vectors. It also exhibits various exceptional cases we must deal with, such as what to do when reducing a vector mod 2 yields the 0 class or a short vector orthogonal to v .

The basic object in our construction is what we call a marked lattice, which is a nonempty ordered list of linearly independent elements of Λ . The language reflects the idea that a marked lattice is a sublattice of Λ equipped with a distinguished basis. Our algorithm determines whether two given marked lattices are equivalent under Co_0 ; this solves the equivalence problem for nonzero $v, v' \in \Lambda$, because we can apply the algorithm to the one-dimensional marked lattices they span. Given a marked lattice, we will iteratively construct new marked lattices from it until this process halts and then study the last ones constructed in order to obtain information about L . This information is enough to determine the Co_0 -equivalence or -inequivalence of marked lattices.

We begin with the iterative step. Given a marked lattice L with basis e_1, \dots, e_n , we define $\ell_y = \sum y_i e_i$ for all $y \in \{0, 1\}^n$. These are a set of representatives for $L/2L$, and we regard the y s as ordered under the lexicographic order of $\{0, 1\}^n$. Recall that for $v \in \Lambda$, we write \bar{v} for the image of v in $\Lambda/2\Lambda$ and when \bar{v} is not a

frame, we write \hat{v} for a short representative of v , which is well defined up to sign.

Algorithm 3.1. Given a marked lattice L , this algorithm either (i) produces a frame ϕ , (ii) produces a marked lattice M , (iii) produces a set $\{M_{\pm}\}$ of two marked lattices, or (iv) terminates without producing anything.

- Step 1. (Applies if there exists y with $\bar{\ell}_y$ a frame.) Write y for the first such, set ϕ equal to the frame $\bar{\ell}_y$, and quit, yielding case (i).
- Step 2. (Applies if there exists $y \neq (0, \dots, 0)$ with $\bar{\ell}_y = 0$.) Write y for the first such, define i by the condition that the first nonzero entry of $y \in \{0, 1\}^n$ is the i th, define M as the marked lattice with basis $e_1, \dots, e_{i-1}, \ell_y/2, e_{i+1}, \dots, e_n$, and quit, resulting in case (ii).
- Step 3. (Applies if there exists y with $\hat{\ell}_y$ in neither $L \otimes \mathbb{Q}$ nor L^\perp .) Write y for the first such, define s to be whichever of $\pm \hat{\ell}_y$ has positive inner product with the first of e_1, \dots, e_n with which it has nonzero inner product, define M to be the marked lattice with basis e_1, \dots, e_n, s , and quit, resulting in case (ii).
- Step 4. (Applies if there exists $y \neq (0, \dots, 0)$ with $\hat{\ell}_y \perp L$.) Write y for the first such, define M_{\pm} as the marked lattices with bases $e_1, \dots, e_n, \pm \hat{\ell}_y$, and quit, resulting in case (iii).
- Step 5. (Applies in all other cases.) Quit without producing anything, resulting in case (iv).

There is nothing to prove except that our constructions make sense. In Step 2, $\ell_y/2$ lies in Λ because of the hypothesis $\bar{\ell}_y = 0$, and the linear independence of $e_1, \dots, e_{i-1}, \ell_y/2, e_{i+1}, \dots, e_n$ follows from the fact that their integral span contains e_i . In Steps 3 and 4, $\hat{\ell}_y$ is determined up to sign—if $\bar{\ell}_y$ were zero or a frame then the algorithm would have stopped at Step 1 or 2. In Step 3, one of $\pm \hat{\ell}_y$ is distinguished, but in Step 4 neither is—the two marked lattices M_{\pm} must be treated equally, and only the set of both of them together is natural. (They are two different markings of the same underlying lattice.) Finally, Step 5 does not really quit without producing anything, because it produces a proof that $(L \otimes \mathbb{Q}) \cap \Lambda$ is an \mathcal{S} -lattice; see Lemma 3.4 below.

The output of the algorithm is natural in the sense that if L and L' are marked lattices and $g \in Co_0$ carries L

to L' , then L produces a frame ϕ if and only if L' produces a frame ϕ' , in which case $g(\phi) = \phi'$, and similarly for the other cases. To see this, observe that every criterion and construction in the algorithm uses only the geometry of Λ and the given ordering of the basis e_1, \dots, e_n .

A marked lattice that occurs in the output of the algorithm applied to L is called a child of L . A descendant of L is defined to be L itself, a child of L , or a child of a child of L , and so on. We are abusing the usual meaning of the word by regarding L as a descendant of itself. We write Ω_L for the set of childless descendants of L . Since the construction of children is natural, we see by induction that Ω_L is also natural, in the sense that any $g \in Co_0$ carries Ω_L to $\Omega_{g(L)}$.

Lemma 3.2 assures us that Ω_L is finite and nonempty and that it can be computed by repeatedly applying Algorithm 3.1 to find the “family tree” of L . It may appear that this computation of Ω_L involves an exponential explosion, on account of the opportunity for Step 4 to introduce branching into the tree. In fact, this does not happen: in Section 4, we prove that $|\Omega_L| \leq 4$ except when L lies in an \mathcal{S} -lattice, when $|\Omega_L|$ is still bounded by 16.

Lemma 3.2. *A child of a marked lattice L strictly contains L .*

Proof: A child of L constructed by one of Steps 3 or 4 of Algorithm 3.1 is obtained by adjoining a new linearly independent vector to the basis for L , so the assertion is obvious. A child M of L constructed by Step 2 is obtained by replacing e_i by $\ell_y/2$ in the basis (e_1, \dots, e_n) for L . Since $\ell_y = e_i + \sum_{j \neq i} y_j e_j$, we have $e_i = 2 \cdot (\ell_y/2) - \sum_{j \neq i} y_j e_j$, so that M contains L . Also, $\ell_y/2$ lies in M but not in L —if it lay in L , then ℓ_y would represent the zero element of $L/2L$, contrary to the assumption $y \neq (0, \dots, 0)$. \square

Next, we consider a childless descendant M of L . If Algorithm 3.1 applied to M stops at Step 1, with output the frame ϕ , then we say that M yields, or produces, or determines ϕ . Lemma 3.4 asserts that if M determines no frame then M lies in an \mathcal{S} -lattice.

Lemma 3.3. *If the short representatives of the elements of a d -dimensional frame-free subspace of $\Lambda/2\Lambda$ lie in a d -dimensional subspace of \mathbb{R}^{24} , then their integral span is an \mathcal{S} -lattice and is primitive in Λ .*

Proof: Write M_0 for the integral span of the short representatives and note that $(M_0 \otimes \mathbb{Q}) \cap \Lambda = \bar{M}_0$, since

\bar{M}_0 has cardinality $\geq 2^d$ and lies in $\overline{(M_0 \otimes \mathbb{Q}) \cap \Lambda}$, which has cardinality $\leq 2^d$. Obviously, $0 \in M_0$. If $v \in (M_0 \otimes \mathbb{Q}) \cap \Lambda$ is nonzero, then \hat{v} lies in M_0 . Also, $(v - \hat{v})/2 \in (M_0 \otimes \mathbb{Q}) \cap \Lambda$ has smaller norm than v if we choose \hat{v} such that $v \cdot \hat{v} \geq 0$. By induction on the norm, $(v - \hat{v})/2$ lies in M_0 , hence, $v = \hat{v} + 2 \cdot (v - \hat{v})/2$ does too. This proves that M_0 is primitive, and it follows that it is an \mathcal{S} -lattice. \square

Lemma 3.4. *Suppose M is a childless marked lattice. If it determines a frame, then M does not lie in an \mathcal{S} -lattice. Otherwise, the short representatives for $\bar{M} \subseteq \Lambda/2\Lambda$ span an \mathcal{S} -lattice that contains M , has the same rational span as M , and is the smallest \mathcal{S} -lattice containing M .*

Proof: We write m_y for the standard representatives of $M/2M$, just as we wrote ℓ_y for those of $L/2L$. If the algorithm applied to M yields a frame, then some m_y represents a frame. No member of any \mathcal{S} -lattice can represent a frame, so M cannot lie in an \mathcal{S} -lattice.

Now, suppose M is a childless marked lattice that determines no frame. Since the algorithm does not halt at Step 1, the \hat{m}_y are well defined up to sign. Since it does not halt at Step 2, the map $M/2M \rightarrow \Lambda/2\Lambda$ is injective, so that M and \bar{M} have the same dimension. Since the algorithm does not halt at Step 3 or 4, all the \hat{m}_y lie in the rational span of M ; we write M_0 for the lattice they generate. By Lemma 3.3, M_0 is an \mathcal{S} -lattice and is primitive, so it contains M . Finally, any \mathcal{S} -lattice containing M must contain the short representatives for all elements of M and hence must contain M_0 . \square

Next, we claim that either all the marked lattices $M \in \Omega_L$ yield frames (in which case we say that L ultimately determines frames) or none of them do. By considering the vectors adjoined to L to obtain its children, one sees that any \mathcal{S} -lattice containing L also contains the children. By induction, any \mathcal{S} -lattice containing L contains every descendant of L . By Lemma 3.4, if any $M \in \Omega_L$ determines a frame, then M does not lie in any \mathcal{S} -lattice. Therefore, L does not lie in an \mathcal{S} -lattice, so no member of Ω_L lies in an \mathcal{S} -lattice, so every member of Ω_L determines a frame.

If one marked lattice ultimately determines frames and another does not, then they are not Co_0 -equivalent. We next give necessary and sufficient conditions for the equivalence of two marked lattices that both ultimately determine frames. (This is the generic case.) Then, we will treat the case where neither does. If $L = (e_1, \dots, e_n)$ is a marked lattice that ultimately determines frames,

let $\{\phi_1, \dots, \phi_k\}$ be the set of the frames determined by the childless descendants of L . For each $i = 1, \dots, k$, let g_i be an element of Co_0 carrying ϕ_i to the standard frame; these can be found with Algorithm 2.2. Then, let $e_{i,j} = g_i(e_j)$ for $j = 1, \dots, n$. We use a similar notation for a second marked lattice L' of the same dimension n .

Theorem 3.5. *With the notation above, L and L' are Co_0 -equivalent if and only if there exist $g \in 2^{12}:M_{24}$ and $i \in \{1, \dots, k\}$ such that $g(e'_{1,j}) = e_{i,j}$ for all $j = 1, \dots, n$.*

Proof: The “if” direction is trivial. For the other direction, suppose $h \in Co_0$ carries L' to L ; then, it carries the frames ultimately determined by L' to those ultimately determined by L . In particular, $h(\phi'_1) = \phi_i$ for some $i = 1, \dots, k$. Then, $g = g_i \circ h \circ (g'_1)^{-1}$ carries $g'_1(L') = (e'_{1,1}, \dots, e'_{1,n})$ to $g_i(L) = (e_{i,1}, \dots, e_{i,n})$. It also carries the standard frame to itself and hence lies in $2^{12}:M_{24}$. \square

Theorem 3.5 reduces a test of equivalence under Co_0 to $k \leq |\Omega_L|$ tests of equivalence under $2^{12}:M_{24}$, so the main worry from a performance perspective is that Ω_L might be large. This turns out not to be a problem; in Section 4, we show that $|\Omega_L| \leq 4$ when L ultimately determines frames.

Now, suppose $L = (e_1, \dots, e_n)$ is a marked lattice that does not ultimately determine frames; in this case L lies in an \mathcal{S} -lattice. Suppose $M = (f_1, \dots, f_d)$ is a childless descendant with mod 2 classes $\bar{m}_y, y \in \{0, 1\}^d$. For each $y \neq (0, \dots, 0)$, let \hat{m}_y be the unique short representative of \bar{m}_y that has positive inner product with the first f_i with which it has nonzero inner product. (Existence of such an f_i follows from the condition $\hat{m}_y \in M \otimes \mathbb{Q}$; if this fails then Step 3 or 4 of Algorithm 3.1 would have produced at least one child of M .) Define

$$\begin{aligned} A(M, y, z) &= \hat{m}_y \cdot \hat{m}_z, & \text{for all } y, z \in \{0, 1\}^d, \\ B(M, y, i) &= \hat{m}_y \cdot e_i, & \text{for all } y \in \{0, 1\}^d \\ & & \text{and } i \in \{1, \dots, n\}. \end{aligned}$$

As M varies over Ω_L , we obtain a family of arrays of integers.

Theorem 3.6 shows that a straightforward comparison of these arrays determines whether L is Co_0 -equivalent to another marked lattice L' that lies in some \mathcal{S} -lattice. From a performance perspective, the main worry is that the list of A s might be immense. In fact, it can be immodestly large but not huge; because M lies in an \mathcal{S} -lattice, we have $d \leq 6$.

Theorem 3.6. *Suppose L and L' are marked lattices of dimension n that lie in \mathcal{S} -lattices, M' is a fixed childless descendant of L' , of dimension say d' , and that notation is otherwise as above. Then, L and L' are Co_0 -equivalent if and only if there is a childless descendant M of L , of dimension $d = d'$, such that*

$$A(M, y, z) = A(M', y, z), \quad \text{for all } y, z \in \{0, 1\}^d$$

and

$$B(M, y, i) = B(M', y, i), \quad \text{for all } i \in \{1, \dots, n\}$$

and $y \in \{0, 1\}^d$.

Proof: If $g \in Co_0$ carries L' to L , then it carries M' to some childless descendant M of L , and the naturality of our constructions imply the stated equalities. On the other hand, suppose these equalities hold for some childless descendant M of L . The first batch of equalities implies that the map $\hat{m}'_y \rightarrow \hat{m}_y$ defines an isometry from the integer span M'_0 of the \hat{m}'_y to the integer span M_0 of the \hat{m}_y . Since M'_0 and M_0 are \mathcal{S} -lattices, Theorem 2.1 implies that this isometry between M'_0 and M_0 extends to an isometry T of Λ . The second batch of equalities implies that T carries $e'_1, \dots, e'_n \in M'_0$ to $e_1, \dots, e_n \in M_0$. □

4. ORBITS UNDER Co_0 : PERFORMANCE

Most of this section is devoted to proving the following upper bound on the amount of branching in the family tree of a marked lattice. This is very important, because it rules out an exponential explosion in the computation. However, this section is not needed for the correctness of any of the algorithms in this paper.

Theorem 4.1. *A marked lattice that ultimately determines frames has at most four childless descendants. A marked lattice that lies in an \mathcal{S} -lattice has at most 16 childless descendants.*

The fact that we have a stronger bound in the first case is good, because this is the generic case. More precisely, it is easy to estimate the number of lattice vectors of norm $\leq N$ that lie in \mathcal{S} -lattices. The dominant term comes from the \mathcal{S} -lattices of type $2^{27}3^{36}$. Each is a scaled copy of the dual of the E_6 root lattice and has determinant 3^5 . Therefore, it contains $\sim V(6, \sqrt{N})/3^{5/2}$ vectors of norm $\leq N$, where $V(d, r)$ is the volume of a radius- r ball in \mathbb{R}^d . Similarly, Λ contains $\sim V(24, \sqrt{N})/1$ vectors of norm $\leq N$. The stabilizer G of one of these \mathcal{S} -lattices is $2.(3^5.2).U_4(2)$ by [Curtis 76, page 573], so they are

$|Co_0|/|G|$ in number. Putting these data together shows that a random lattice vector of norm $\leq N$ lies in an \mathcal{S} -lattice with probability

$$\sim \frac{|Co_0|}{|G|} \cdot \frac{V(6, \sqrt{N})}{3^{5/2}} \bigg/ V(24, \sqrt{N}) = \frac{2^{23}3^{25}5^7 7^3 11^2 13 \cdot 23}{\pi^9 \sqrt{3}} N^{-9} \sim (5.67 \times 10^{13}) N^{-9}.$$

For small N , this is a poor estimate because of the \mathcal{S} -lattices' intersections; indeed, it does not drop below 1 until $N = 34$. However, it is easy to treat vectors of small norm directly, with the following results; we denote orbits of type ≤ 11 (i.e., $N \leq 22$) using the notation of [Conway 69]. A vector of type 2 or 3 spans an \mathcal{S} -lattice, and one of type 4, 6_{32} , 8_{42} , or 10_{52} represents a frame. One of type 8_{22} is twice a type 2 vector. A vector of type 5, 6_{22} , 7, 8_{32} , 9_{32} , or 10_{33} has one childless descendant, which appears after two generations and is an \mathcal{S} -lattice of type $2^2 3^1$, $2^3 3^0$, $2^1 3^2$, $2^2 3^1$, $2^0 3^3$, or $2^1 3^2$, respectively. A vector of type 9_{42} , 10_{42} , or 11_{43} has one childless descendant; it appears after two generations and determines a frame. The most complicated case is a vector of type 11_{52} , which has two childless descendants, both three-dimensional \mathcal{S} -lattices of type $2^5 3^2$. Considering the sizes of all these orbits shows that a random vector of norm ≤ 22 lies in an \mathcal{S} -lattice with probability $\sim .092$. This probability drops off rapidly as N increases.

We will also obtain the following bound on the number of generations in the family tree:

Theorem 4.2. *Suppose $v \in \Lambda - \{0\}$ and $L_0 \subseteq L_1 \subseteq \dots \subseteq L_n$ are marked lattices, L_0 being the span of v and L_i a child of L_{i-1} for $i = 1, \dots, n$. Then,*

$$n \leq \log_4(6^6 v^2) = \log_4 v^2 + 7.75\dots$$

To prove Theorems 4.1 and 4.2 we will need the following definitions and notation. $\Lambda/2\Lambda$ is equipped with the quadratic form Q obtained by reducing vectors' types modulo 2. We will consider the restriction of Q to subspaces of $\Lambda/2\Lambda$, so we need notation for quadratic forms over \mathbb{F}_2 . We write S (split) and A (anisotropic) for the 2-variable quadratic forms xy and $x^2 + xy + y^2$ and \square (square) and 0 (zero) for the 1-variable forms x^2 and 0. Any \mathbb{F}_2 quadratic form is a direct sum of copies of these forms, and there are isomorphisms $S \oplus S \cong A \oplus A$, $S \oplus \square \cong A \oplus \square$, and $\square \oplus \square \cong \square \oplus 0$. Any \mathbb{F}_2 quadratic form q has an associated bilinear form $b(x, y) = q(x+y) - q(x) - q(y)$. We write B for the bilinear form on $\Lambda/2\Lambda$, which is given

by reducing lattice vectors' inner products modulo 2. We call a subspace isotropic if Q vanishes identically on it; note that this is a stronger condition than the vanishing of B . Vectors of Λ of types 2 and 3 are called 2-vectors and 3-vectors, and we call their images in $\Lambda/2\Lambda$ "2-classes" and "3-classes." If L is a sublattice of Λ , then we call it reduced if $L/2L \rightarrow \Lambda/2\Lambda$ is injective; its reduction is defined as $(L \otimes \mathbb{Z}[\frac{1}{2}]) \cap \Lambda$ and is reduced.

Now, suppose L is reduced and \bar{L} frame free. We write L_0 for the integer span of $\{\hat{v}|\bar{v} \in \bar{L} \text{ and } \hat{v} \in L \otimes \mathbb{Q}\}$. Almost all of our arguments refer to L_0 rather than L itself. If $\bar{v} \in \bar{L} - \{0\}$ has $\hat{v} \perp L$, then we call \bar{v} an ambiguous class and \hat{v} an ambiguous vector. The point of these definitions is that Algorithm 3.1 produces a child of L if some $\bar{v} \in \bar{L}$ has no representative in L_0 , and it produces two children when every such class is ambiguous. The ambiguity is whether to extend L to a larger marked lattice by adjoining \hat{v} or $-\hat{v}$.

Lemma 4.3. *Every seven-dimensional subspace of $\Lambda/2\Lambda$ contains a frame, as does every three-dimensional isotropic subspace.*

Proof: The first claim follows from the second, since any \mathbb{F}_2 -quadratic form of dimension ≥ 7 contains a three-dimensional isotropic subspace. To prove the second claim, choose short vectors $e_1, e_2,$ and e_3 representing the members of a basis of the subspace and write L for their integral span. If any e_i has type 4, then we are done. Otherwise, they all have type 2 (norm 4) and their pairwise inner products are in $\{0, \pm 2\}$. It is well known that any family of norm 2 vectors having inner products in $\{0, \pm 1\}$ spans a direct sum of $A_n, D_n,$ and E_n root lattices, so L is a scaled version of $A_1^3, A_2 \oplus A_1,$ or A_3 . Each of these root systems contains a pair of orthogonal roots, so L contains a norm 8 (type 4) vector and \bar{L} contains a frame. \square

Proof of Theorem 4.2: If a child of a marked lattice has the same dimension, its determinant is that of its parent divided by 4. If it has larger dimension, its determinant is at most that of its parent multiplied by 6. By Lemma 4.3, the latter can occur at most six times. The theorem follows because L_0 has determinant v^2 and the determinants of all the L_i are integers. \square

The next three lemmas deal with special arrangements of vectors and are difficult to motivate without first looking at the proof of Lemma 4.8. We recommend jumping straight to Lemma 4.7.

Lemma 4.4. *Consider a two-dimensional subspace V of $\Lambda/2\Lambda$ with one class τ of type 2 and two classes θ_1 and θ_2 of type 3. Then, either every pair among $\hat{\tau}, \hat{\theta}_1,$ and $\hat{\theta}_2$ is orthogonal or no pair is.*

Proof: By computing $B|_V$ from $Q|_V$, we see that all inner products among the three vectors are even. If $\hat{\tau}$ makes a nonzero inner product with one of the $\hat{\theta}$, say $\hat{\theta}_1$, then, without loss of generality, we may take $\hat{\tau} \cdot \hat{\theta}_1 = 2$ and $\hat{\theta}_2 = \hat{\theta}_1 - \hat{\tau}$. Direct computation shows that no two are orthogonal. We have $\hat{\theta}_1 \cdot \hat{\theta}_2 \neq \pm 2$, because otherwise $\tau = \theta_1 - \theta_2$ would be a frame, not a 2-class. If $\hat{\theta}_1 \cdot \hat{\theta}_2 \neq 0$, then, without loss of generality, we have $\hat{\theta}_1 \cdot \hat{\theta}_2 = 4$ and $\hat{\tau} = \hat{\theta}_1 - \hat{\theta}_2$, and again direct computation shows that no two are orthogonal. \square

Lemma 4.5. *Suppose W is a three-dimensional frame-free subspace of $\Lambda/2\Lambda$ with $Q|_W \cong \square \oplus 0 \oplus 0$. Then, the 2-vectors representing the 2-classes of W span an \mathcal{S} -lattice S of type $2^3 3^0$. Writing $\theta_1, \dots, \theta_4$ for the 3-classes of V , if two of the $\hat{\theta}_i$ are orthogonal to S , then all of them are. In this case, $\hat{\theta}_i \perp \hat{\theta}_j$ for all $i \neq j$, and there is an \mathcal{S} -lattice of type $2^{27} 3^{36}$ containing S and all the $\hat{\theta}_i$.*

Proof: The span S of the 2-vectors is clearly an \mathcal{S} -lattice of the specified type, since they have nonzero even inner products with each other. It will be convenient to name the 2-classes $\tau_1, \tau_\omega,$ and $\tau_{\bar{\omega}}$, the subscripts denoting the nonzero elements of the finite field \mathbb{F}_4 ; we may take $\hat{\tau}_1 + \hat{\tau}_\omega + \hat{\tau}_{\bar{\omega}} = 0$. By considering $Q|_W$, one sees that all inner products of representatives for elements of W are even. The inner product of a 2-vector and a 3-vector in Λ lies in $\{0, \pm 1, \pm 2, \pm 3\}$. Therefore, if some $\hat{\theta}_i$, say $\hat{\theta}_1$, is not in S^\perp , then, without loss of generality, we may suppose $\hat{\theta}_1 \cdot \hat{\tau}_1 = 0, \hat{\theta}_1 \cdot \hat{\tau}_\omega = 2,$ and $\hat{\theta}_1 \cdot \hat{\tau}_{\bar{\omega}} = -2$. Then, $\hat{\theta}_1 + \hat{\tau}_{\bar{\omega}}$ and $\hat{\theta}_1 - \hat{\tau}_\omega$, also 3-vectors, are also not in S^\perp and represent two of the remaining 3-classes of W . We have proven that if any of the $\hat{\theta}_i$ is nonorthogonal to S , then at least three of them are, which is a restatement of our second claim.

Now, suppose all the $\hat{\theta}_i$ lie in S^\perp ; we use new subscripts, writing $\theta_0, \theta_1, \theta_\omega,$ and $\theta_{\bar{\omega}}$ for the 3-classes of W . We may do this in such a way that the addition table for W is $\tau_a + \tau_b = \tau_{a+b}, \tau_a + \theta_b = \theta_{a+b},$ and $\theta_a + \theta_b = \tau_{a+b}$. (For purposes of this notation, we take $\tau_0 = 0 \in W$.) The claim $\hat{\theta}_i \cdot \hat{\theta}_j = 0$ for $i \neq j$ follows from Lemma 4.4, because $\hat{\theta}_i$ is orthogonal to $\widehat{\theta_i + \theta_j} = \hat{\tau}_{i+j}$. All that remains is to construct the desired \mathcal{S} -lattice. For this we note that besides the $\hat{\tau}_a$ and $\hat{\theta}_b, \Lambda$ also contains the 16 3-vectors $(\pm \hat{\theta}_0 \pm \hat{\theta}_1 \pm \hat{\theta}_\omega \pm \hat{\theta}_{\bar{\omega}})/2,$ since

$\theta_0 + \theta_1 + \theta_\omega + \theta_{\bar{\omega}} = 0 \in \Lambda/2\Lambda$. Similarly, it contains the 48 2-vectors $(\pm\hat{\tau}_a \pm \hat{\theta}_b \pm \hat{\theta}_{a+b})/2$ with $a, b \in \mathbb{F}_4$ and $a \neq 0$, the 24 3-vectors $(\pm(\hat{\tau}_a + \hat{\tau}_b) \pm \hat{\theta}_a \pm \hat{\theta}_b)/2$ with $a, b \in \mathbb{F}_4 - \{0\}$, and the 24 3-vectors $(\pm(\hat{\tau}_a + \hat{\tau}_b) \pm \hat{\theta}_0 \pm \hat{\theta}_c)/2$ where a, b , and c are the nonzero elements of \mathbb{F}_4 in any order. We have exhibited a six-dimensional lattice with 54 2-vectors and 72 3-vectors, which is an \mathcal{S} -lattice by Lemma 3.3. \square

Lemma 4.6. *Co_0 acts transitively on ordered triples (u, v, w) of vectors in Λ satisfying $u^2 = v^2 = 4, w^2 = 6, u \cdot v = u \cdot w = v \cdot w = 2$. The lattice K spanned by three such vectors is primitive in Λ and has exactly one ambiguous class \bar{x} . The pointwise stabilizer of K exchanges the two short representatives $\pm\hat{x}$, the lattice obtained adjoining them to K is not reduced, and its reduction represents a frame.*

Proof: By [Curtis 73, page 562], we may take $u = (4, -4, 0^{22})$ and $w = (5, 1^{23})$. The simultaneous stabilizer of these vectors is the Higman-Sims group HS of order 44,352,000, and one can enumerate the 2-vectors having inner product 2 with each of u and w . There are 1,100 of them, one of which is $v = (4, 0, -4, 0^{21})$. We will prove transitivity on the 1,100 by showing that the stabilizer G of v in HS has order $|HS|/1100$. (It obviously has at least this order.)

We write K for the lattice spanned by u, v , and w . G preserves each of $w = (5, 1, 1, 1^{21})$, $w - u = (1, 5, 1, 1^{21})$, and $w - v = (1, 1, 5, 1^{21})$, so it preserves their sum, say $x = (7, 7, 7, 3^{21}) = (3, 3, 3, (-1)^{21}) + 2(2, 2, 2, 2^{21})$. The 3-vector $\hat{x} = (3^3, -1^{21})$ is orthogonal to K , so \bar{x} is an ambiguous class. No other class of \bar{K} is ambiguous because we have seen that K contains $u, v, u - v, w, w - u$, and $w - v$.

The stabilizer $G_{\hat{x}}$ of \hat{x} fixes (2^{24}) , which represents the standard frame. Therefore, $G_{\hat{x}} \subseteq 2^{12}:M_{24}$, and it is easy to see that $G_{\hat{x}} = M_{21}$, of order 20,160. Since $|M_{21}| = |HS|/2200$, and every element of G fixes or negates \hat{x} , we have $|G| \leq |HS|/1100$. This gives $|G| = |HS|/1100$, proving the claimed transitivity. We also see $[G : G_{\hat{x}}] = 2$, so some element of G negates \hat{x} . The four-dimensional lattice spanned by K and \hat{x} is not reduced since we have already seen that its reduction contains (2^{24}) , which represents the standard frame.

To prove the primitivity, compute the determinant of K , which is 56. Since the only square dividing 56 is 4, any three-dimensional enlargement of K contains K of index 2. No such enlargement exists in Λ because K is reduced. \square

Lemma 4.7. *Suppose L is a marked lattice with two children. Then,*

- (a) \bar{L} is frame-free;
- (b) L is reduced;
- (c) for all $\ell \in L$, $\hat{\ell}$ either lies in L_0 or is ambiguous;
- (d) \bar{L} contains an ambiguous class;
- (e) \bar{L} contains no isotropic 3-space;
- (f) the only elements of \bar{L} that can be ambiguous are those in $\ker(B|_{\bar{L}})$;
- (g) if L_0 contains a 2-vector, then it contains the short representatives of all the 2-classes of \bar{L} ;
- (h) if L_0 contains no 2-vectors, then any two independent 3-vectors in L_0 have inner product 0 or ± 3 , and L_0 is a scaled copy of a direct sum of A_n, D_n , and E_n root lattices.

Proof: (a), (b), and (c) hold because Algorithm 3.1 applied to L does not stop at any of its first three steps. We have (d), because Step 4 of that algorithm does apply to L . Because of Lemma 4.3, (a) implies (e). If $\bar{\ell} \in \bar{L}$ does not lie in $\ker(B|_{\bar{L}})$, then $\hat{\ell}$ has odd inner product with some element of L , so (c) implies $\hat{\ell} \in L_0$. This proves (f). If \hat{v} is a 2-vector of L_0 and $\bar{w} \in \bar{L}$ is an ambiguous 2-class, then $\hat{w} \perp \hat{v}$, so that $\hat{w} + \hat{v}$ has type 4 and \bar{L} contains a frame, contradicting (a). This proves (g).

Now, we prove the first part of (h). If $v, w \in L_0$ are independent 3-vectors, then their inner product is 0, ± 1 , ± 2 , ± 3 , or ± 4 , so it suffices to rule out the cases $v \cdot w = 1, 2$, or 4. In the last case, $v - w$ is a 2-vector of L_0 , contrary to hypothesis. In the case $v \cdot w = 2$, $v - w$ has type 4, contrary to (a). Now, we show that $v \cdot w = 1$ leads to a contradiction; since Co_0 acts transitively on pairs of 3-vectors with inner product 1, the following claims can be verified by checking a single example. The sum $v + w$ has type 5, and the two possibilities for $\widehat{v + w}$ are $\pm u$, where u has type 3 and makes inner product 1 with each of v and w . Therefore, $u \in L_0$ by (f). Now, $u + v + w \in 2\Lambda$, and we write x for $(u + v + w)/2$, which lies in L_0 by (b). Then, $x - u, x - v$, and $x - w$ are 2-vectors of L_0 , contrary to hypothesis. The second part of (h) follows from the first by mimicking the argument for Lemma 4.3. \square

Lemma 4.8. *Suppose L is a marked lattice of dimension d that has two children. Then one of the following holds:*

- (i) $d \leq 2$;
- (ii) L lies in an \mathcal{S} -lattice;
- (iii) L_0 is a copy of the three-dimensional lattice K of Lemma 4.6 and contains L ; or
- (iv) $d = 3$ and $\dim L_0 \leq 1$.

Proof: It suffices to assume $d \geq 3$ and show that one of (ii), (iii), and (iv) applies. We will refer to the assertions of Lemma 4.7 simply by their letter. By (b), $\dim \bar{L} = d$. By (e), \bar{L} contains no isotropic 3-space. By (d) and (f), $\ker(B|_{\bar{L}}) \neq 0$. One can enumerate the \mathbb{F}_2 quadratic forms satisfying these conditions, with the result that one of following holds:

- $d = 5$ and $Q|_{\bar{L}} \cong S \oplus A \oplus 0$ or $S \oplus S \oplus \square$;
- $d = 4$ and $Q|_{\bar{L}} \cong S \oplus \square \oplus 0$ or $A \oplus 0 \oplus 0$;
- $d = 3$ and $Q|_{\bar{L}} \cong S \oplus \square$, $S \oplus 0$, $A \oplus 0$, or $\square \oplus 0 \oplus 0$.

For cases $S \oplus A \oplus 0$ and $S \oplus 0$, L_0 contains a 2-vector by (f); since the nonzero element of $\ker(B|_{\bar{L}})$ is a 2-class, its short representatives lie in L_0 by (g). But then there are no ambiguous classes, contrary to (d). We now treat the remaining cases one by one, splitting the last case into two parts. To prove that L lies in an \mathcal{S} -lattice, it suffices to prove that L_0 does and that $\dim L_0 = \dim L$, because then L lies in the rational span of an \mathcal{S} -lattice. Since \mathcal{S} -lattices are primitive (their images in $\Lambda/2\Lambda$ satisfy the hypotheses of Lemma 3.3), L lies in the \mathcal{S} -lattice itself. In all cases except $\square \oplus 0 \oplus 0$, the equality $\dim L_0 = \dim L$ is automatic because the elements of $\bar{L} - \ker(B|_{\bar{L}})$ span \bar{L} .

Case $S \oplus S \oplus \square$. We will show that L_0 lies in an \mathcal{S} -lattice of type $2^{27}3^{36}$. Only the 3-class in $\ker(B|_{\bar{L}})$ is ambiguous, so L_0 contains 2-vectors representing the 15 2-classes of \bar{L} . These classes may be identified with the 15 duads from $\{1, \dots, 6\}$ in such a way that classes have even inner product if and only if the corresponding duads are disjoint. In this notation, if i, j, k, l, m , and n are $1, \dots, 6$ in any order, then ij, kl , and mn lie in a two-dimensional isotropic subspace. (The fact that such a labeling is possible stems from the fact that S_6 is the isometry group of $S \oplus S \oplus \square$.) The isotropic subspace containing $12, 34$, and 56 lifts to an \mathcal{S} -lattice of type $2^3 3^0$; we choose lifts $\widehat{12}, \widehat{34}$, and $\widehat{56}$ summing to 0. Every other duad is disjoint from one of these three and meets the other two, and we choose the short representative \widehat{ij} of ij which makes inner product -2 with whichever one of $\widehat{12}, \widehat{34}$, and $\widehat{56}$ it has even inner product with.

We have now chosen representatives for all the 2-classes of \bar{L} , and we claim that any two with even inner

product have inner product -2 . If $\widehat{13} \cdot \widehat{24} = 2$, then $\widehat{13} - \widehat{24}$ and $\widehat{56}$ would be orthogonal 2-vectors and \bar{L} would contain a frame. If $\widehat{13} \cdot \widehat{45} = 2$, then $\widehat{13} - \widehat{45}$ and $\widehat{56}$ would be orthogonal 2-vectors and \bar{L} would contain a frame. Up to symmetry of our notation, these are the only cases to check, so the claim is proven.

Next, we claim that any pair of our vectors with odd inner product have inner product 1. If i, j, k, ℓ, m , and n are $1, \dots, 6$ in any order then $\widehat{ij} + \widehat{k\ell} + \widehat{mn} = 0$, $\widehat{ik} \cdot \widehat{mn} = -2$, and $\widehat{ik} \cdot \widehat{ij}$ and $\widehat{ik} \cdot \widehat{k\ell}$ lie in $\{\pm 1\}$. Therefore, $\widehat{ik} \cdot \widehat{ij} = 1$ for any distinct i, j , and k , proving the claim.

We have proven that the configuration of the 2-vectors is uniquely determined. That is, we may choose coordinates on $L_0 \otimes \mathbb{R}$ such that $\widehat{12}$ is the vector $\frac{1}{\sqrt{3}}(2, 2, -1, -1, -1, -1)$ and similarly for the other \widehat{ij} , with the 2s appearing in the i th and j th positions. (We refer to the standard metric on \mathbb{R}^6 and note that all our vectors have coordinate sum zero.) Therefore, L_0 contains the vectors $\alpha_i = \frac{1}{\sqrt{3}}(1, \dots, 1, -5, 1, \dots, 1)$, with the -5 in the i th position. Since these have even inner product with all the \widehat{jk} , all represent the ambiguous class of \bar{L} , say \bar{v} . This class has type 3 and \bar{v} is orthogonal to L_0 , so we may suppose $\hat{v} = (1, \dots, 1)$ in our coordinates. Since $\alpha_i \equiv \hat{v}$ modulo 2Λ , Λ contains the six vectors $a_i = (\alpha_i + \hat{v})/2$ and also the six vectors $b_i = (\alpha_i - \hat{v})/2$. The a_i and b_i have type 2. Adjoining them to the \widehat{ij} yields the configuration of 27 2-vectors spanning Curtis' \mathcal{S} -lattice $2^{27}3^{36}$. We have proven that L_0 lies in this \mathcal{S} -lattice.

Case $S \oplus \square \oplus 0$. We will show that L_0 lies in an \mathcal{S} -lattice $2^{27}3^{36}$. By (f), L_0 contains short representatives for the classes not in $\ker(B|_{\bar{L}})$; since one of these is a 2-class, (g) implies that L_0 also contains a short representative a for the unique 2-class in $\ker(B|_{\bar{L}})$. There are six other 2-classes in \bar{L} , which fall into three pairs, each pair summing to \bar{a} . Since each pairs trivially with \bar{a} under B , we may choose short representatives having inner product 2 with a . Consideration of $B|_{\bar{L}}$ shows that vectors of different pairs have inner product ± 1 . A contradiction arises if any of these inner products is -1 , so all are $+1$. It follows that our seven vectors may be taken to be $a = (2, 0, 0, 0)$, $(1, \pm\sqrt{3}, 0, 0)$, $(1, 0, \pm\sqrt{3}, 0)$, and $(1, 0, 0, \pm\sqrt{3})$. A simple argument shows that every point in their real span lies at distance < 2 of their integral span S . (The key is that S contains an A_3 lattice orthogonal to a , scaled to have minimal norm 6; this scaled A_3 has covering radius $\sqrt{3}$ by [Conway and Sloane 88, page 112].) This implies that S is primitive in Λ , since Λ has minimal norm 4. Therefore, $L_0 = S$. No 3-vector in S represents θ_1 or

θ_2 , so $\hat{\theta}_1$ and $\hat{\theta}_2$ lie in L^\perp . Now, we can use Lemma 4.5; consider the three-dimensional subspace of \bar{L} containing a , θ_1 , θ_2 , and any other 2-class. By Lemma 4.5, the short representatives for this subspace span a six-dimensional lattice that lies in (hence, rationally spans) the rational span of L_0 , $\hat{\theta}_1$, and $\hat{\theta}_2$. Also, this 6-space meets Λ in an \mathcal{S} -lattice $2^{27}3^{36}$, so L_0 lies in the \mathcal{S} -lattice.

Case $A \oplus 0 \oplus 0$. We will show that L_0 lies in an \mathcal{S} -lattice $2^{27}3^{36}$. By (f), L_0 contains the short representatives for the 12 elements of $\bar{L} - \ker(B|_{\bar{L}})$, all of type 3. If L_0 contained the short representative for any element of $\ker(B|_{\bar{L}})$, then it would contain all of them by (g), so no class would be ambiguous, contrary to (d). Therefore, L_0 contains no 2-vectors. By (h), the 3-vectors of L_0 form a (scaled) root system of dimension ≤ 4 . Since there are 24 roots, it can only be of type D_4 . This D_4 is of course orthogonal to the \mathcal{S} -lattice $2^{33}3^0$ spanned by the short representatives of $\ker(B|_{\bar{L}})$, and Lemma 4.5 implies that some \mathcal{S} -lattice $2^{27}3^{36}$ contains both this \mathcal{S} -lattice and the D_4 . In particular, it contains L_0 .

Case $S \oplus \square$. We will show that L_0 lies in an \mathcal{S} -lattice 2^93^6 . There is only one ambiguous class, which has type 3. Considering $Q|_{\bar{L}}$, we see that L_0 contains three linearly independent 2-vectors, any two having odd inner product. After negating one of them, we may suppose that either all their inner products are 1 or that two are 1 and the third is -1 . In the latter case, the three vectors form the configuration of figure I.8 of [Curtis 73], and Curtis shows that the three vectors span an \mathcal{S} -lattice 2^33^4 . This would mean that no class is ambiguous, contrary to (d). Therefore, all the inner products are 1, and the vectors form the configuration of Curtis' figure I.7. In this case, Curtis shows that L_0 lies in a four-dimensional \mathcal{S} -lattice 2^93^6 .

Case $A \oplus 0$. We will show that L_0 lies in an \mathcal{S} -lattice of type $2^{27}3^{36}$. By (f), only the 2-class could be ambiguous, so L_0 contains 12 short representatives of the six 3-classes of \bar{L} . Also, L_0 does not contain any 2-vectors, for otherwise no class of \bar{L} would be ambiguous. By (h), the 3-vectors form a (scaled) root system of dimension ≤ 3 ; since there are 12 roots, it must have type A_3 . We choose coordinates using the standard inner product on \mathbb{R}^4 , with $a_{12} = (\sqrt{3}, -\sqrt{3}, 0, 0)$ and similarly for a_{ij} ($i, j = 1, \dots, 4$ with $i \neq j$), the $\sqrt{3}$ appearing in the i th place and the $-\sqrt{3}$ appearing in the j th. We write b for a 2-vector representing the 2-class of \bar{L} , which is of course orthogonal to the 3-vectors. In our coordinates we may take $b = (1, 1, 1, 1)$. If i, j, k , and ℓ are 1, 2, 3, and 4, in any order, then $a_{ij} + a_{k\ell}$ has even type. Since it is too short

to lie in 2Λ , it is congruent modulo 2Λ to b , so Λ contains the vectors $(a_{ij} + a_{k\ell} + b)/2$. Consider the vectors a_{13} , a_{24} , b , and $b' = (a_{12} + a_{34} + b)/2$. Since $a_{13} + a_{24} \equiv b$, their images in $\Lambda/2\Lambda$ span a three-dimensional space. Since a_{13} and a_{24} are orthogonal to the \mathcal{S} -lattice $2^{33}3^0$ spanned by b and b' , Lemma 4.5 implies that all four vectors lie in an \mathcal{S} -lattice $2^{27}3^{36}$. Therefore, L_0 lies in this \mathcal{S} -lattice.

Case $\square \oplus 0 \oplus 0$, with L_0 containing a 2-vector. We will show that one of Cases (ii) and (iii) applies. By (g), L_0 contains the short representatives for all three 2-classes of \bar{L} ; together these span an \mathcal{S} -lattice $2^{33}3^0$, say S . By Lemma 4.5, either all the short representatives of the 3-classes are orthogonal to S or only one is. In the second case, L_0 contains a triple of vectors as in Lemma 4.6, hence it contains a copy of the lattice K treated there. Since K is primitive, it equals L_0 , and since $\dim L_0 = \dim L$, K contains L . We have proven that (iii) applies. Now, we treat the case in which all short representatives of the 3-classes are orthogonal to S . By Lemma 4.5, representatives $\hat{a}_1, \dots, \hat{a}_4$ for the four 3-classes are mutually orthogonal; we write V for their real span and note that $V \cap \Lambda$ also contains the vectors $(\pm\hat{a}_1 \pm \hat{a}_2 \pm \hat{a}_3 \pm \hat{a}_4)/2$. Now, L contains a vector z with $\bar{z} = a_1$, and we claim that z has nonzero projection into V . If $z \perp V$, then the vectors $(\hat{a}_1 + \hat{a}_2 + \hat{a}_3 + \hat{a}_4)/2$ and $(\hat{a}_1 + z)/2$ of Λ have nonintegral inner product, which is impossible. Since the projection of z is nonzero, some \hat{a}_i is not in L^\perp and, therefore, lies in L_0 . Then, L_0 has dimension 3 and, hence, the same rational span as L . Then, the usual argument applies: since L_0 lies in an \mathcal{S} -lattice, so does L . We note for use in the proof of Theorem 4.1 that L_0 is the orthogonal direct sum of S with the span of a 3-vector.

Case $\square \oplus 0 \oplus 0$, with L_0 containing no 2-vectors. Clearly, either (iv) holds or else L_0 contains two linearly independent short vectors; we will show that under the latter condition, L lies in an \mathcal{S} -lattice $2^{27}3^{36}$. We write b_0 and b_1 for two 3-classes with short representatives in L_0 . We write a_1 , a_ω , $a_{\bar{\omega}}$, b_ω , and $b_{\bar{\omega}}$ for the other elements of \bar{L} , using the notation in the proof of Lemma 4.5. We choose short representatives \hat{a}_1 , \hat{a}_ω , and $\hat{a}_{\bar{\omega}}$ summing to zero. Since \hat{b}_0 and \hat{b}_1 are orthogonal to the \hat{a} , Lemma 4.5 implies that \hat{b}_ω and $\hat{b}_{\bar{\omega}}$ are also orthogonal to the \hat{a} , that the \hat{b} are all mutually orthogonal, and that the rational span of all the \hat{a} and \hat{b} meets Λ in an \mathcal{S} -lattice $2^{27}3^{36}$. We write V for the real span of \hat{b}_ω , $\hat{b}_{\bar{\omega}}$, and the \hat{a} . Now, L contains some z with $\bar{z} = b_\omega$; we claim that z has nonzero projection to V —otherwise, the vectors $(z + \hat{b}_\omega)/2$ and $(\hat{a}_1 + \hat{b}_\omega + \hat{b}_{\bar{\omega}})/2$ of Λ would have nonintegral inner prod-

uct. Therefore, one of $\hat{a}_1, \hat{a}_\omega, \hat{a}_{\bar{\omega}}, \hat{b}_\omega,$ and $\hat{b}_{\bar{\omega}}$ is not in L^\perp , hence lies in L_0 . Since L_0 contains no 2-vectors, either \hat{b}_ω or $\hat{b}_{\bar{\omega}}$ lies in L_0 . Then, L_0 has the same dimension as L ; since L_0 lies in an \mathcal{S} -lattice, L does too. For use in the proof of Theorem 4.1, we note that L_0 is spanned by three mutually orthogonal 3-vectors, since no other element of \bar{L}_0 has its short representative in $L_0 \otimes \mathbb{Q}$. \square

Proof of Theorem 4.1:

Frames case. Suppose a marked lattice T does not lie in any \mathcal{S} -lattice and has more than four childless descendants. Then, it has a descendant U with two children, one of which, say V , has a descendant W with two children, and one of which, say X , has a descendant Y with two children. Since T does not lie in any \mathcal{S} -lattice, neither does any descendant. Lemma 4.8 implies $\dim Y \leq 3$, and since we have

$$1 \leq \dim T \leq \dim U < \dim V \leq \dim W < \dim X \leq \dim Y \leq 3,$$

it follows that

$$\begin{aligned} \dim T &= \dim U = 1, \\ \dim V &= \dim W = 2, \\ \dim X &= \dim Y = 3. \end{aligned}$$

Since V is obtained from U (and X from W) by adjoining a short vector, Y contains two linearly independent short vectors. In particular, $\dim Y_0 > 1$, so by Lemma 4.8, Y_0 is a copy of K and contains Y . Now, U has a single basis vector, say a , and V has basis (a, b) with $a \equiv b$ modulo 2Λ and $a \perp b$. Then, W is the reduction of V , and Y is obtained by adjoining a short vector c that is orthogonal to both a and b . We have deduced that there is a vector $a \in K$ that is orthogonal to two mutually orthogonal short vectors of K and congruent to one of them modulo 2Λ (hence modulo $2K$). Now, such a exists and is unique up to isometry of K and multiplication by an odd number. To see this, one considers the short vectors of K , all of which are enumerated in the proof of Lemma 4.6, and finds all orthogonal pairs of such vectors. Then, one checks that $\text{Aut } K$ acts transitively on such pairs, so, without loss of generality, we may suppose our pair is $u - v$ and w in the notation of that lemma. Then, a must be an odd multiple of the generator of their orthogonal complement in K , which is $(-14, 14, 14, 2^{21})$. We have shown that if any element of Λ ultimately determines frames and has more than 4 childless descendants then this one does. An explicit calculation yields only 4,

which proves (the frames case of) the theorem and also the optimality of the upper bound 4.

\mathcal{S} -lattice case. If a marked lattice L lies in an \mathcal{S} -lattice and has more than 16 childless descendants, then the family tree must branch five times. Since \mathcal{S} -lattices have dimension at most 6, we deduce the existence of marked lattices L_i of dimension i for each $i = 1, \dots, 5$, each with two children, L_1 a descendant of L , and each other L_i a descendant of L_{i-1} . We consider L_3 ; by the argument for Lemma 4.8, one of the following is true:

- (i) $Q|_{\bar{L}} \cong S \oplus \square$ and L_3 lies in a four-dimensional \mathcal{S} -lattice;
- (ii) $Q|_{\bar{L}} \cong A \oplus 0$ and $L_{3,0}$ is a copy of the A_3 lattice, scaled to have minimal norm 6;
- (iii) $Q|_{\bar{L}} \cong \square \oplus 0 \oplus 0$ and $L_{3,0}$ is spanned by an \mathcal{S} -lattice M of type $2^3 3^0$ and a 3-vector orthogonal to M ;
- (iv) $Q|_{\bar{L}} \cong \square \oplus 0 \oplus 0$ and $L_{3,0}$ is spanned by three mutually orthogonal 3-vectors.

Case (i) clearly cannot arise. In each Case (ii)–(iv), the argument for the frames case shows that $L_{3,0}$ contains orthogonal short vectors b and c such that L_1 is spanned by a vector v in $\langle b, c \rangle^\perp \subseteq L_{3,0}$, with $v \equiv b$ or c modulo 2Λ . In each case, $\{b, c\}$ is unique up to isometry of $L_{3,0}$, and L_1 is spanned by an odd multiple of the generator a of their orthogonal complement in $L_{3,0}$, which is easy to find. In Cases (ii) and (iv), a is congruent to neither b nor c . In Case (iii), a is a vector of type 6_{22} , which has only one childless descendant. \square

Theorem 4.1 is the best possible because the following considerations construct a marked lattice with 16 childless descendants. Let E be an \mathcal{S} -lattice $2^{27} 3^{36}$ and L_3 the orthogonal complement in E of three mutually orthogonal 3-vectors $e_1, e_2,$ and e_3 . ($\text{Aut}(E)$ acts transitively on ordered 4-tuples of mutually orthogonal 3-vectors in E , so there are no choices to make.) Then, all but three classes of \bar{L}_3 have short representatives in L_3 , and these three classes have representatives $e_1, e_2,$ and e_3 . Adjoining one of the e_i and reducing yields a lattice L_4 , with all but two classes of \bar{L}_4 having short representatives in L_4 . These two classes are represented by the two remaining e_i . Adjoining and reducing again yields L_5 , and every class of \bar{L}_5 has a short representative in L_5 except one, represented by the last of the e_i . Adjoining and reducing one last time yields E . It is easy to find a vector in L_3 with two children and L_3 among its descendants. Such a vector has 16 childless descendants.

5. ORBITS UNDER $2^{12}:M_{24}$

To test the equivalence of vectors under $2^{12}:M_{24}$, we reduce the problem to several tests under M_{24} . Recall that the normal subgroup 2^{12} of $2^{12}:M_{24}$ is a copy of \mathcal{C} acting on \mathbb{R}^{24} by negating signs on \mathcal{C} -sets.

Theorem 5.1. *Suppose $v, w \in \mathbb{R}^{24}$, V (respectively, W) is the set of vectors in $\mathcal{C} \cdot v$ (respectively, $\mathcal{C} \cdot w$) with fewest possible negative coordinates, and $w' \in W$. Then, v and w are $2^{12}:M_{24}$ -equivalent if and only if w' is M_{24} -equivalent to some element of V .*

The proof is trivial, and the factor determining performance is clearly the size of V ; in fact, $|V| \leq 12$ in all cases. To verify this, we used the enumeration [Conway and Sloane 88, page 284] of M_{24} -orbits of subsets X of Ω . For each X , we considered the subgroup of $\{\pm 1\}^X$ obtained by restricting \mathcal{C} -sets to X and found all its coset representatives in $\{\pm 1\}^X$ with the fewest possible -1 s. We also verified $|V| \leq 12$ by using an implementation of our original M_{24} algorithm to enumerate M_{24} -orbits of vectors having all coordinates in $\{0, \pm 1\}$. For each such vector, we computed V by the method given below. The extreme case $|V| = 12$ occurs only when V has support an umbral dodecad and oddly many negative coordinates, or when V has support a 13-ad containing an umbral dodecad and has oddly many negative coordinates on the dodecad. The next-worst case gives $|V| = 8$.

To use Theorem 5.1, we need to be able to find V and W . We may do this by applying one of the Golay decoding algorithms in [Conway and Sloane 86] or [Vardy and Be'ery 91]. These are ‘‘soft’’ decoding algorithms, which means that they view \mathcal{C} as a subset of \mathbb{R}^{24} , with a \mathcal{C} -set γ corresponding to the vector with coordinates -1 on γ and $+1$ elsewhere. Given a vector of \mathbb{R}^{24} , these algorithms return a \mathcal{C} -set with maximal inner product with that vector. Given v , we define the vector v_0 whose i th coordinate is the sign ($0, +1$, or -1) of the i th coordinate of v . After applying the Golay decoder to v_0 to obtain a codeword γ , set $v' = \gamma \cdot v$. A little thought verifies that v' has the fewest negative coordinates of any element of $\mathcal{C} \cdot v$. The Golay decoders of [Conway and Sloane 86, Vardy and Be'ery 91] proceed by computing the inner products of codewords with a given vector and returning a codeword with the largest inner product. (The cleverness of these algorithms lies in doing this with as few computations as possible.) We may run a modified version of the decoder, which returns *every* codeword with maximal inner product. Denoting this set of codewords by C , we have $V = C \cdot v$.

If the support $S \subseteq \Omega$ of v is disjoint from some code-words, then finding V this way is inefficient, because C will be larger (possibly much larger) than V . This is only a problem when S is not very big, and such cases may be handled by the following shortcut, which works when $|S| < 12$. Let γ be a codeword nearest S . If γ is not an octad or does not lie in S , then V consists of just one vector, with all coordinates positive. If γ is an octad lying in S and v has an even number of negative coordinates on γ , then again V contains a single vector, with all coordinates positive. If γ is an octad lying in S and v has an odd number of negative coordinates on γ , then V contains eight vectors, each with all coordinates positive except for one of the eight points of γ , which is negative.

6. ORBITS UNDER M_{24}

The problem of equivalence of vectors $v, w \in \mathbb{R}^{24}$ under M_{24} is not really a question about vectors in \mathbb{R}^{24} . Namely, we write $C_v = (c_1, \dots, c_k)$ where $c_1 < \dots < c_k$ are the distinct values taken by the coordinates of v , P_i for the subset of Ω where the coordinates of v equal c_i , and \mathcal{P}_v for the corresponding (ordered) partition (P_1, \dots, P_k) of Ω . Then, v and w are M_{24} -equivalent if and only if $C_v = C_w$ and \mathcal{P}_v is M_{24} -equivalent to \mathcal{P}_w . Therefore, we restrict ourselves to the problem of equivalence of ordered partitions \mathcal{P} and \mathcal{P}' of Ω . All partitions in our analysis are ordered, so we will suppress the qualifier except for emphasis. The idea is to search for sextets by reducing members of \mathcal{P} and \mathcal{P}' modulo \mathcal{C} and by several other means. Our procedure will almost always find a sextet, reducing the equivalence problem to the sextet group $2^6:3 \cdot S_6$ of order only 138,240. Of course, we also treat the special case where no sextets are found. We do not provide an algorithm for the sextet group; various combinations of brute force and cleverness are possible.

Suppose \mathcal{P} is an ordered partition of Ω . If some member P of \mathcal{P} represents a sextet, then testing the M_{24} -equivalence of \mathcal{P} with some other ordered partition \mathcal{P}' can be reduced to the sextet group (see Theorem 6.6 below). Otherwise, the unique small representative \bar{P} of P modulo \mathcal{C} will typically cut the members of \mathcal{P} non-trivially, so that \mathcal{P} determines a finer partition, whose members may themselves be reduced modulo \mathcal{C} , and so on. We will try to refine \mathcal{P} as much as possible in this manner, interrupting the refinement process if a sextet appears at any point. To make this precise, suppose $\mathcal{P} = (P_1, \dots, P_n)$ is an ordered partition and X a subset of Ω . We say that X refines \mathcal{P} if X is not a union

• 3.3.3.3.3.3.3.3	<u>5.1.2.6.10</u>	<u>6.1.1.8.8</u>	<u>8.13.1.1.1</u>
• <u>9.3.3.3.3.3</u>	◦ <u>5.1.1.1.6.10</u>	6.2.10.3.3	<u>8.14.2</u>
• <u>10.2.3.3.3.3</u>	• <u>5.3.8.3.3.2</u>	◦ <u>6.2.13.3</u>	<u>8.14.1.1</u>
• <u>12.3.3.3.3</u>	<u>5.3.8.8</u>	<u>6.2.15.1</u>	<u>8.15.1</u>
• <u>13.3.3.3.2</u>	<u>5.2.1.8.8</u>	<u>6.1.1.15.1</u>	<u>8.15.1</u>
<u>16.3.3.2</u>	<u>5.1.1.1.8.8</u>	<u>6.2.16</u>	• <u>9.3.9.3</u>
21.3	• <u>5.3.13.3</u>	<u>6.1.1.16</u>	◦ <u>9.3.10.2</u>
21.2.1	<u>5.3.14.2</u>	<u>7.1.7.1.7.1</u>	<u>9.3.12</u>
21.1.1.1	<u>5.3.16</u>	<u>7.1.7.1.8</u>	<u>9.2.1.12</u>
22.2	<u>5.2.1.16</u>	<u>7.1.8.8</u>	<u>9.1.1.1.12</u>
22.1.1	<u>5.1.1.1.16</u>	<u>7.1.14.2</u>	<u>10.2.11.1</u>
23.1	<u>6.2.6.10</u>	<u>7.1.14.1.1</u>	<u>10.1.1.11.1</u>
24	<u>6.1.1.6.10</u>	<u>7.1.15.1</u>	<u>10.2.12</u>
• <u>5.3.5.3.3.3.2</u>	<u>6.2.6.9.1</u>	<u>7.1.16</u>	<u>10.1.1.12</u>
• <u>5.3.5.3.5.3</u>	◦ <u>6.1.1.6.9.1</u>	<u>8.8.8</u>	<u>11.1.11.1</u>
• <u>5.3.5.3.6.2</u>	<u>6.2.7.1.8</u>	• <u>8.8.3.3.2</u>	<u>11.1.12</u>
• <u>5.3.5.3.8</u>	<u>6.1.1.7.1.8</u>	<u>8.13.3</u>	<u>12.12</u>
• <u>5.3.6.2.8</u>	<u>6.2.8.8</u>	<u>8.13.38.13.2.1</u>	

TABLE 1. Possible shapes of refined partitions; see Lemma 6.2. Each underline means that the union of the underlined sets is a codeword. A bullet • indicates that a partition of that shape determines a sextet in Algorithm 6.3, and a circle ◦ indicates that partitions of that shape are treated by one of the special cases of Theorem 6.7. Partitions of other listed shapes are treated by the general case of Theorem 6.7.

of members of \mathcal{P} . In this case, suppose P_{i_1}, \dots, P_{i_m} ($i_1 < \dots < i_m$) are the members of \mathcal{P} that meet X but do not lie in it; we define the refinement of \mathcal{P} by X to be

$$(Q_1, \dots, Q_n, P_{i_1} \cap X, \dots, P_{i_m} \cap X),$$

where $Q_j = P_j - X$ for $j \in \{i_1, \dots, i_m\}$ and $Q_j = P_j$ otherwise. We say that \mathcal{P} is refined if no member of \mathcal{P} represents a sextet and no member's small representative refines \mathcal{P} .

Algorithm 6.1. (Refinement.) This algorithm accepts an ordered partition \mathcal{P} of Ω and returns either a sextet S or a refined partition \mathcal{R} of Ω ; in the latter case, we call \mathcal{R} the refinement of \mathcal{P} .

Step 1. Set $\mathcal{R} = \mathcal{P}$.

Step 2. If any member of \mathcal{R} represents a sextet, set S to be the sextet represented by the first such member, and quit.

Step 3. If any member of \mathcal{R} of size 5 or more has small representative that refines \mathcal{R} , replace \mathcal{R} by its refinement by \bar{R} , where R is the first such member, and go back to Step 2.

Step 4. Quit, returning \mathcal{R} .

Proof: The algorithm terminates because each repetition of Step 3 increases the number of members of \mathcal{R} , of which there may be at most 24. All we must show is that \mathcal{R} is refined when defined. If \mathcal{R} is defined, the algorithm reaches Step 4, so neither Step 2 nor 3 applies to \mathcal{R} . Since Step 2 does not apply, no member of \mathcal{R} represents a sextet, so, in particular, no member of \mathcal{R} has size 4. Since Step 3 does not apply, no member of \mathcal{R} of size > 4 has a small representative that refines \mathcal{R} . Since any member of \mathcal{R} of size < 4 is its own small representative, \mathcal{R} is refined. \square

Because our criteria and constructions refer only to \mathcal{C} and the ordering of \mathcal{P} , the sextet is natural when defined, in the sense that if the algorithm applied to \mathcal{P} produces a sextet S and $g \in M_{24}$, then the algorithm applied to $g(\mathcal{P})$ produces the sextet $g(S)$. Similarly, if \mathcal{P} has refinement \mathcal{R} and $g \in M_{24}$, then $g(\mathcal{P})$ has refinement $g(\mathcal{R})$. Refined partitions are rather special:

Lemma 6.2. *Suppose \mathcal{R} is a refined ordered partition of Ω . Then, either some tetrad is a union of members of \mathcal{R} or else \mathcal{R} has one of the 71 shapes listed in Table 1. (A line beneath some of the listed sets indicates that their union is a codeword.)*

Proof: The proof is an uninspiring slog through many cases. We assume throughout that no tetrad is a union

of members of \mathcal{R} . What gets the enumeration off the ground is that a member R of \mathcal{R} of size ≥ 5 is disjoint from its small representative—otherwise, \bar{R} would refine \mathcal{R} . This means that the only sets that can appear in \mathcal{R} are monads, duads, triads, and sets of the form

(a Golay codeword) – (0, 1, 2, or 3 of its points).

Also, if R is a member of \mathcal{R} , then \bar{R} refines \mathcal{R} if and only if the codeword $R + \bar{R}$ refines \mathcal{R} . Table 1 was obtained by first enumerating those partitions with no large sets, then those with exactly one large set, then those containing a pentad and at least one other large set, then those containing a hexad and at least one other large set, but no pentad, and so on.

As an example of the argument, we enumerate the partitions containing exactly one pentad p and also some other large set. If there is a hexad h , then the octads \mathcal{O}_p and \mathcal{O}_h containing p and h meet in 0, 2, or 4 points, because they are \mathcal{C} -sets. The last case is impossible, because then p and h would overlap. If they are disjoint, then \mathcal{R} has shape $\underline{5.3.6.2.8}$, where (n) indicates some partition of n points. If the (3) is not a triad or the (2) is not a duad, then some tetrad would be a union of members of \mathcal{R} ; therefore, \mathcal{R} has shape $\underline{5.3.6.2.8}$. By hypothesis, the (8) contains no pentad, and it cannot contain a monad, duad, or tetrad. Therefore, the (8) is a single octad, so \mathcal{R} has shape $\underline{5.3.6.2.8}$, which appears in the table.

We summarize and continue this argument in a compact notation.

$5.6.(13)$ and $\mathcal{O}_p \cap \mathcal{O}_h = \emptyset \Rightarrow \underline{5.3.6.2.8} \Rightarrow \underline{5.3.6.2.8} \Rightarrow \underline{5.3.6.2.8}$.

$5.6.(13)$ and $|\mathcal{O}_p \cap \mathcal{O}_h| = 2 \Rightarrow \underline{5.1.2.6.10} \Rightarrow \underline{5.1.2.6.10}$, so both $\underline{5.1.2.6.10}$ and $\underline{5.1.1.1.6.10}$ appear in the table.

$5.7.(12)$ and no 6s $\Rightarrow \underline{5.3.7.1.8} \Rightarrow$ tetrad.

$5.8.(11)$ and no 6s or 7s $\Rightarrow \underline{5.3.8.8} \Rightarrow \underline{5.3.8.8}$ or $\underline{5.3.8.3.3.2}$; in the latter case, we must have $\underline{5.3.8.3.3.2}$.

$5.9.(10)$ and no 6s, 7s, or 8s $\Rightarrow \mathcal{O}_p$ meets the 9-ad's dodecad in 2, 4, or 6 points; only the first case is possible. Therefore, we have $\underline{5.1.2.1.9.6} \Rightarrow$ tetrad.

$5.10.(9)$ and no 6s through 9s $\Rightarrow \underline{5.1.2.10.6} \Rightarrow$ tetrad.

$5.N.(19 - N)$ for $N = 11$ or $12 \Rightarrow$ a contradiction, since \mathcal{O}_p would meet the N -ad's dodecad in at least two points, hence refine \mathcal{R} .

$5.N.(19 - N)$ for $N = 13, \dots, 16 \Rightarrow$ the N -ad's 16-ad meets \mathcal{O}_p in 0, 4, 6, or 8 points; only the first is possible without refining \mathcal{R} . Therefore, we have one of $\underline{5.3.13.3}$, $\underline{5.3.14.2}$, $\underline{5.3.15.1}$, and $\underline{5.3.16}$. All cases except those with tetrads appear in the table.

If \mathcal{R} contains a pentad, then it cannot contain a set larger than a 16-ad, for else \mathcal{O}_p would meet this set and hence refine \mathcal{R} . □

Algorithm 6.3. (Sextet search.) Suppose \mathcal{P} is an ordered partition of Ω . This algorithm returns either a sextet S or else the refinement \mathcal{R} of \mathcal{P} exists and the algorithm returns it.

Step 1. Apply Algorithm 6.1 (refinement), obtaining either a sextet (in which case we set S to be this sextet and quit) or the refinement \mathcal{R} of \mathcal{P} .

Step 2. (Applies if some tetrad is the union of members of \mathcal{R} .) Set S to be the sextet containing the first such tetrad and quit. See below for the meaning of “first.”

Step 3. (Applies if \mathcal{R} has three triads.) Let t_1, t_2 , and t_3 be the first three triads, in order. If $t_1 \cup t_2$ represents a sextet, then set S to be this sextet and quit. Otherwise, if $t_2 \cup t_3$ represents a sextet, then set S to be this sextet and quit. Otherwise, the octads containing $t_1 \cup t_2$ and $t_2 \cup t_3$ meet in a tetrad; set S to be the corresponding sextet and quit.

Step 4. (Applies if \mathcal{R} has shape $2.3^2.5^2.6$, $3^2.5^2.8$, $3^2.5.13$, or $3^2.9^2$.) Write t_1 and t_2 for the two triads, in order. If $t_1 \cup t_2$ represents a sextet, then set S to be this sextet and quit. Otherwise, the octad containing $t_1 \cup t_2$ meets the first pentad of \mathcal{R} (or the first nonad in case of shape $3^2.9^2$) in one point, which with t_1 forms a tetrad; set S to be the corresponding sextet and quit.

Step 5. (Applies if \mathcal{R} has shape $2.3.5.6.8$.) The octad containing the union of the duad and triad meets the pentad in a single point, which with the triad forms a tetrad; set S to be the corresponding sextet and quit.

Step 6. (Applies if \mathcal{R} has shape $2.3^2.8^2$.) Let t be the first triad, \mathcal{O} the union of the small sets, and T the trio consisting of \mathcal{O} and the two octads

of \mathcal{R} ; set S to be the sextet determined by the containments $t \subseteq \mathcal{O} \in T$ (see Lemma 6.4 below) and quit.

Step 7. (Applies in all other cases.) Quit, returning \mathcal{R} .

Beyond the fact that the constructions make sense, there is nothing to prove. For Step 2, we need a notion of the first tetrad T that is a union of members of \mathcal{R} . Here is one possibility. If \mathcal{R} contains a tetrad, then T is the first such (with respect to the ordering of \mathcal{R}). If \mathcal{R} contains a triad and a monad, then T is the union of the first triad and first monad. If \mathcal{R} contains two duads, T is the union of the first two. If \mathcal{R} contains a duad and two monads, T is the union of the duad and the first two monads. Otherwise, \mathcal{R} contains four monads and T is the union of the first four.

For Step 3, we must show that the octads containing $t_1 \cup t_2$ and $t_2 \cup t_3$ meet in four points (when such octads exist); this is because they meet in at least three (both contain t_2) but cannot coincide (else $t_1 \cup t_2 \cup t_3$ would lie in an octad). For Step 4, under the condition that $t_1 \cup t_2$ does not represent a sextet, we claim that the octad \mathcal{O} containing it meets each pentad (or nonad in the case of shape $3^2.9^2$) once. For the case $2.3^2.5^2.6$, the sets assemble into \mathcal{C} -sets as [5.3.5.3.6.2](#), according to Table 1. Then, it is clear that \mathcal{O} meets each octad [5.3](#) in at least three points, hence exactly four, and, therefore, meets each pentad once.

The same argument applies essentially verbatim for the case of shapes $3^2.5^2.8$ and $3^2.5.13$. For the case $3^2.9^2$, \mathcal{R} has shape [9.3.9.3](#) by Table 1, and \mathcal{O} meets each dodecad [9.3](#) in at least three points, hence exactly four. The argument for Step 5, with shape [5.3.6.2.8](#), is the same. For Step 6, we note that \mathcal{R} has shape [8.8.3.3.2](#) by Table 1, so \mathcal{O} and the two octads are special, and we apply the following lemma:

Lemma 6.4. *If t is a triad, \mathcal{O} a special octad, and T a trio, with $t \subseteq \mathcal{O} \in T$, then there are octads $\mathcal{O}' \in \mathcal{C}$ meeting \mathcal{O} in exactly four points, including those of t , and disjoint from one of the other octads of T . All such octads \mathcal{O}' meet \mathcal{O} in the same tetrad.*

Proof: M_{24} acts transitively on trios, the stabilizer of a trio acts transitively on its octads, and the stabilizer of a trio and one of its octads acts 3-transitively on the octad. Therefore, it suffices to verify the claim for one particular inclusion $t \subseteq \mathcal{O} \in T$. If we take T to be the standard trio, then the ‘‘Turyn’’ description of \mathcal{C} in [Conway and

Sloane 88, Chapter 11, Section 12] makes this obvious. (The tetrad is the unique tetrad containing t that lies in the ‘‘line code.’’) \square

Remark 6.5. Finding τ algorithmically is easy; simply transform T to the standard trio in such a way that t is carried into one of the MOG columns. Then, τ is just that column and the sextet is the standard one. Performing this operation requires finding a suitable permutation in M_{24} , which is easy using a few precomputed permutations. Also, such a permutation will be needed anyway in Theorem 6.6 below.

We say that \mathcal{P} determines a sextet S if S is the output of Algorithm 6.3 applied to \mathcal{P} . All the constructions of the algorithm are natural in the sense that they use only the ordering on \mathcal{R} and the structure of \mathcal{C} . Therefore, if $g \in M_{24}$ and \mathcal{P} determines the sextet S , then $g(\mathcal{P})$ determines the sextet $g(S)$. On the other hand, if \mathcal{P} does not determine a sextet, then neither does $g(\mathcal{P})$, and g carries the refinement of \mathcal{P} to that of $g(\mathcal{P})$. Given ordered partitions \mathcal{P} and \mathcal{P}' of Ω , if one determines a sextet and the other does not, then they are not M_{24} -equivalent. If both determine sextets, then whether they are M_{24} -equivalent is determined by Theorem 6.6, and if neither does, then it is determined by Theorem 6.7.

Theorem 6.6. *Suppose \mathcal{P} and \mathcal{P}' are ordered partitions of Ω that determine sextets S and S' , and let $g, g' \in M_{24}$ carry S and S' to the standard sextet. Then, \mathcal{P} and \mathcal{P}' are M_{24} -equivalent if and only if $g(\mathcal{P})$ and $g'(\mathcal{P}')$ are equivalent under the sextet group $2^6:3.S_6$ of order 138, 240.*

Proof: The ‘‘if’’ part is trivial. For the converse, suppose $h \in M_{24}$ carries \mathcal{P} to \mathcal{P}' , so that it carries S to S' . Then, $g' \circ h \circ g^{-1}$ carries $g(\mathcal{P})$ to $g'(\mathcal{P}')$; it also carries the standard sextet to itself, so it lies in the sextet group. (This is essentially the argument we used for Theorem 3.5.) \square

Finding the permutations g and g' needed to apply Theorem 6.6 is easy: we can just refer to tables of coset representatives for M_{23} in M_{24} , M_{22} in M_{23} , M_{21} in M_{22} , and M_{20} in M_{21} .

Theorem 6.7. *Suppose \mathcal{P} and \mathcal{P}' are ordered partitions of Ω that do not determine sextets, and let \mathcal{R} and \mathcal{R}' be their refinements. If \mathcal{R} and \mathcal{R}' are not S_{24} -equivalent, then \mathcal{P} and \mathcal{P}' are not M_{24} -equivalent. If \mathcal{R} and \mathcal{R}' are S_{24} -equivalent, then we have the following cases:*

1. (Applies if both \mathcal{R} and \mathcal{R}' have shape $1^3.5.6.10$ or both have shape $1^3.6^2.9$.) The octad containing the first hexad of \mathcal{R} contains two of the monads and the remaining monad is the first, second, or third monad of \mathcal{R} ; define $n = 1, 2, \text{ or } 3$ in these cases, and define n' similarly using \mathcal{R}' in place of \mathcal{R} . Then, \mathcal{P} and \mathcal{P}' are M_{24} -equivalent if and only if $n = n'$.
2. (Applies if both \mathcal{R} and \mathcal{R}' have shape $2.3.6.13$.) Let h be the hexad of \mathcal{R} and \mathcal{O} the octad containing the duad and triad, and similarly for h' and \mathcal{O}' . Then, \mathcal{P} and \mathcal{P}' are M_{24} -equivalent if and only if $|h \cap \mathcal{O}| = |h' \cap \mathcal{O}'|$.
3. (Applies if both \mathcal{R} and \mathcal{R}' have shape $2.3.9.10$.) Let n be the monad of \mathcal{R} and \mathcal{O} the octad containing the duad and triad, and similarly for n' and \mathcal{O}' . Then, \mathcal{P} and \mathcal{P}' are M_{24} -equivalent if and only if $|n \cap \mathcal{O}| = |n' \cap \mathcal{O}'|$.
4. (Applies in all other cases.) \mathcal{P} and \mathcal{P}' are M_{24} -equivalent.

Proof: Since \mathcal{P} and \mathcal{P}' are M_{24} -equivalent if and only if \mathcal{R} and \mathcal{R}' are, we may, without loss of generality, replace \mathcal{P} and \mathcal{P}' by \mathcal{R} and \mathcal{R}' throughout; then, the first claim is obvious. In Cases (1)–(3), the criteria used to conclude inequivalence involve only the orderings on \mathcal{R} and \mathcal{R}' and the structure of \mathcal{C} , so any conclusion of inequivalence is justified. We now treat Cases (1)–(4) individually, proving that any conclusion of equivalence is justified.

For the case $1^3.5.6.10$, both \mathcal{R} and \mathcal{R}' have shape $5.1.1.1.6.10$ and $n = n'$. Write D (respectively, D') for the dodecad $5.1.6$ of \mathcal{R} (respectively, \mathcal{R}'). By hypothesis, D (respectively, D') contains the n th monad of \mathcal{R} (respectively, \mathcal{R}'). Since M_{24} is transitive on dodecads, we may suppose $D' = D$. Since the stabilizer M_{12} of D is 5-transitive on its complement, we may suppose that for $i \neq n$, the i th monad of \mathcal{R} coincides with that of \mathcal{R}' ; write δ for the duad consisting of these two monads. Now, only two octads of \mathcal{C} meet $\Omega - D$ exactly in δ , one of which meets D in the hexad of \mathcal{R} and one of which (possibly the same one) meets D in the hexad of \mathcal{R}' . The pointwise stabilizer of δ in M_{12} is M_{10} , which contains a permutation exchanging these two octads, so we may suppose the hexads of \mathcal{R} and \mathcal{R}' coincide. Finally, the setwise stabilizer of this hexad in M_{10} is A_6 , acting transitively on the remaining six points of D . Therefore, we may suppose that the n th monad of \mathcal{R} coincides with that of \mathcal{R}' , i.e., $\mathcal{R}' = \mathcal{R}$.

For the case $1^3.6^2.9$, both \mathcal{R} and \mathcal{R}' have shape $6.1.1.6.1.9$ and the argument is the same except that the last step is replaced by A_6 , acting transitively on the decad.

For the case $2.3.6.13$, both \mathcal{R} and \mathcal{R}' have shape $6.2.13.3$. The essential fact is the following: the group $2^4:S_6$ that simultaneously stabilizes a special octad \mathcal{O} and a duad δ contained in it acts with two orbits on triads t in the complementary 16-ad. These orbits are distinguished by whether the octad containing $t \cup \delta$ meets \mathcal{O} in two or four points.

For the case $2.3.9.10$, both \mathcal{R} and \mathcal{R}' have shape $9.3.10.2$, and the essential fact is the following: the group $M_{10}.2$ that simultaneously stabilizes a dodecad D and a duad δ in it acts with two orbits on triads in the complementary dodecad. These orbits are distinguished by whether the octad containing $t \cup \delta$ meets D in two or four points.

Now we treat all other cases. The common shape of \mathcal{R} and \mathcal{R}' is not one of those marked with a bullet \bullet in Table 1, because otherwise they would have determined sextets in Algorithm 6.3. The shapes marked with circles \circ are those we have just treated. For all remaining shapes in the table, M_{24} acts transitively on ordered partitions having that shape and forming \mathcal{C} -sets in the indicated way. The analysis is easier than the cases above and uses the following facts. M_{24} acts 5-transitively on Ω and transitively on octads, dodecads and trios. The octad stabilizer acts on the octad as A_8 and acts $(3+2)$ - and $(1+3)$ -transitively on the octad and its complement. The trio group permutes the three octads of the trio as S_3 , and the subgroup preserving each of them acts $(3+1+0)$ - and $(2+1+1)$ -transitively on the octads in any order. The stabilizer M_{12} of a dodecad D acts $(5+0)$ -, $(2+1)$ -, $(1+2)$ -, and $(0+5)$ -transitively on D and its complement. The setwise stabilizer of a duad δ not meeting D is a subgroup $M_{10}.2$; this group contains a permutation preserving each of the two octads that meet $\Omega - D$ exactly in δ , say \mathcal{O} and \mathcal{O}' , and swapping the points of δ . It also contains a permutation swapping \mathcal{O} and \mathcal{O}' and preserving each point of δ . The subgroup preserving each point of δ and each of the two octads is A_6 , acting in the natural way on each of the hexads $\mathcal{O} - \delta$ and $\mathcal{O}' - \delta$ and transitively on the 10-ad $\Omega - (D \cup \delta)$. Most of these assertions appear explicitly in [Conway and Sloane 88, Chapter 10], and the rest implicitly. \square

7. REMARKS

We close with a few remarks, first on equivalence testing under the infinite group $Co_\infty = \Lambda:Co_0$ and then on

possible analogues of our algorithm for the complex and quaternionic versions of Λ .

Since we developed these algorithms in order to sort negative-norm elements of $II_{25,1}$ into orbits, we offer here a sketch of how to do this. Following the line of reasoning in the introduction reduces this to the problem of determining equivalence of vectors in $\Lambda \otimes \mathbb{Q}$ under Co_∞ . Borchers ([Borchers 95] has implemented an algorithm to quickly find all lattice vectors lying within a given distance of a given $v \in \Lambda \otimes \mathbb{Q}$. Coupling this with a decoder for Λ , we may find the set N_v of all nearest neighbors to v . Then, given $v, w \in \Lambda \otimes \mathbb{Q}$ and $w' \in N_w$, v and w are Co_∞ -equivalent if and only if $w - w'$ and $v - v'$ are Co_0 -equivalent for some $v' \in N_v$. Of course, $w - w'$ and $v - v'$ are in $\Lambda \otimes \mathbb{Q}$ not Λ , but after scaling we may apply our Co_0 -algorithm. This reduces the test under Co_∞ to at most $|N_v|$ tests under Co_0 .

Happily, we never need to worry about N_v having size larger than 25, because of the following considerations. The diagram Δ_v of v is defined as the graph whose vertices are the elements of N_v ; two are unjoined (respectively, joined, doubly joined) if their difference has type 2 (respectively, 3, 4). See [Conway et. al. 82] for more information about these diagrams. Δ_v is always the disjoint union of spherical and affine Dynkin diagrams of types a_n, d_n, e_n, A_n, D_n , and E_n and hence is a very simple combinatorial object. It is also true that $|N_v| \leq 48$, and if $|N_v| > 25$, then v is a deep hole of Λ and its orbit under Co_∞ is completely determined by the combinatorial type of Δ_v . This means that a test for equivalence under Co_∞ may be reduced to 25 tests for equivalence under Co_0 . We will usually be able to do even better than this by using some feature of Δ . For example, suppose v and w are shallow holes, each with diagram of type $a_{17}d_7a_1$. Then, any transformation carrying w to v must carry the branch point w' of Δ_w to the branch point v' of Δ_v . Therefore, v and w are Co_∞ -equivalent if and only if $v - v'$ and $w - w'$ are Co_0 -equivalent, so only a single equivalence test is required. There are two orbits of shallow holes with this diagram (see [Borchers et. al. 88]), so we cannot determine equivalence by simply inspecting Δ_v and Δ_w .

An extension of our ideas in a different direction concerns the complex and quaternionic versions of Λ , whose automorphism groups are $6.Suz$ and $2G_2(4)$, where Suz is Suzuki's sporadic finite simple group. The complex Leech lattice (see [Wilson 83]) is a lattice over $\mathbb{Z}[\omega = \sqrt[3]{1}]$, whose underlying real lattice is Λ , scaled to have minimal norm 6. It has 3^{12} congruence classes modulo $\theta = \omega - \bar{\omega} = \sqrt{-3}$, and these have minimal represen-

tatives of norm ≤ 9 . The only congruences modulo θ among these vectors are that each norm 6 vector is congruent to its multiples by powers of ω and that the norm 9 vectors fall into classes (frames) of 36 vectors, any two of which are orthogonal or proportional by $\omega^{\pm 1}$. The stabilizer of a frame is $3^5:M_{11}$ (M_{11} being the Mathieu group of order 7,920). In a manner similar to $2^{12}:M_{24}$, M_{11} acts by permuting 12 complex coordinates, and 3^5 acts by multiplying the coordinates by various scalars. If one proved an analogue of Curtis' theorem on 8-lattices, then one could devise an algorithm for equivalence of vectors in the complex Leech lattice under $6.Suz$. If such an analogue is true, then its proof should be easier than that of Curtis' original theorem.

The quaternionic Leech lattice (see [Wilson 82]) has an even simpler structure: modulo $1 + i$, every vector is congruent to either 0 or a minimal vector. The minimal vectors fall into 4,095 classes (again called frames), each consisting of 48 vectors, any two of which are either proportional by an element of $\{\pm 1, \pm i, \pm j, \pm k\}$ or orthogonal. The stabilizer of a frame is the group $2^6.(2^2 \times A_5).2$ of order 30,720, and so an analogue of the Co_0 algorithm is obvious: supposing that $1 + i$ divides neither v nor w , reduce v and w modulo $1 + i$, carry each of the resulting frames to the standard one, and then check for equivalence under the group stabilizing the frame. The reader should note that each of Wilson's frames [Wilson 82] contains three of ours, with $3 \cdot 48 = 144$ vectors, and has slightly larger stabilizer.

ACKNOWLEDGMENTS

The author was partly supported By NSF grants DMS-0070930 and DMS-0231585.

REFERENCES

- [Allcock 96] D. J. Allcock. "The Leech Lattice and Hyperbolic Geometry." Ph.D. diss., U.C. Berkeley, 1996.
- [Borchers 84] R. E. Borchers. "The Leech Lattice and Other Lattices." PhD diss., Cambridge University, 1984.
- [Borchers 85] R. E. Borchers. "The Leech Lattice." *Proceedings of the Royal Society of London* A398 (1985), 365–376.
- [Borchers 87] R. E. Borchers. "Automorphism Groups of Lorentzian Lattices." *Journal of Algebra* 111 (1987), 133–153.
- [Borchers 88] R. E. Borchers. "The 24-Dimensional Odd Unimodular Lattices." In *Sphere Packings, Lattices, and Groups*, edited by J. H. Conway and N. J. A Sloane, pp. 421-430. Berlin: Springer-Verlag, 1988.

- [Borcherds 90] R. E. Borcherds, “The Monster Lie Algebra.” *Advances in Mathematics* 53 (1990), 30–47.
- [Borcherds 95] R. E. Borcherds. Personal Communication, 1995.
- [Borcherds et. al. 88] R. E. Borcherds, J. H. Conway, and L. Queen. “The Cellular Structure of the Leech Lattice.” In *Sphere Packings, Lattices, and Groups*, edited by J. H. Conway and N. J. A Sloane, pp. 513–521. Berlin: Springer-Verlag, 1988
- [Conway 69] J. H. Conway. “A Group of Order 8,315,553,613,086,720,000.” *Bulletin of the London Mathematical Society* 1 (1969), 79–88.
- [Conway 83] J. H. Conway. “The Automorphism Group of the 26-Dimensional Even Unimodular Lorentzian Lattice.” *Journal of Algebra* 80 (1983), 159–163.
- [Conway 88] J. H. Conway. “The Golay Codes and the Mathieu Groups.” In *Sphere Packings, Lattices, and Groups*, edited by J. H. Conway and N. J. A Sloane, pp. 299–330. Berlin: Springer-Verlag, 1988.
- [Conway and Sloane 82] J. H. Conway and N. J. A. Sloane. “Leech Roots and Vinberg Groups.” *Proceedings of the Royal Society of London* A384 (1982), 233–258.
- [Conway and Sloane 86] J. H. Conway and N. J. A. Sloane. “Soft Decoding Techniques for Codes and Lattices, Including the Golay Code and the Leech Lattice.” *IEEE Transactions on Information Theory* 32 (1986), 41–50.
- [Conway and Sloane 88] J. H. Conway and N. J. A. Sloane, editors. *Sphere Packings, Lattices and Groups*. New York: Springer-Verlag 1988.
- [Conway et. al. 82] J. H. Conway, R. A. Parker, and N. J. A. Sloane. “The Covering Radius of the Leech Lattice.” *Proceedings of the Royal Society of London* A380 (1982), 261–290.
- [Conway et. al. 85] J. H. Conway, R. T. Curtis, S. P. Norton, R. A. Parker, and R. A. Wilson. *ATLAS of Finite Groups*. Oxford, UK: Oxford University Press, 1985.
- [Curtis 73] R. T. Curtis. “On Subgroups of $\cdot 0$, I: Lattice Stabilizers.” *Journal of Algebra* 27 (1973), 549–573.
- [Curtis 76] R. T. Curtis. “A New Combinatorial Approach to M_{24} .” *Proceedings of the Cambridge Philosophical Society* 79 (1976), 25–42.
- [Vardy and Be’ery 91] A. Vardy and Y. Be’ery. “More Efficient Soft Decoding of the Golay Codes.” *IEEE Transactions on Information Theory* 37:3 (1991), 667–672.
- [Vardy and Be’ery 93] A. Vardy and Y. Be’ery. “Maximum Likelihood Decoding of the Leech Lattice.” *IEEE Transactions on Information Theory* 39:4 (1993), 1435–1444.
- [Vinberg 72] E. B. Vinberg. “On the Groups of Units of Certain Quadratic Forms.” *Math. Sb.* 87 (1972), 18–36.
- [Vinberg and Kaplinskaja 78] E. B. Vinberg and I. M. Kaplinskaja. “On the Groups $O_{18,1}(Z)$ and $O_{19,1}(Z)$.” *Dokl. Akad. Nauk.* 238 (1978), 1273–1275.
- [Wilson 82] R. A. Wilson. “The Quaternionic Lattice for $2G_2(4)$ and Its Maximal Subgroups.” *Journal of Algebra* 77 (1982), 449–466.
- [Wilson 83] R. A. Wilson. “The Complex Leech Lattice and Maximal Subgroups of the Suzuki Group.” *Journal of Algebra* 84 (1983), 151–188.

Daniel Allcock, Department of Mathematics, University of Texas at Austin, Austin, TX 78712 (allcock@math.utexas.edu)

Received December 22, 2004; accepted June 7, 2005.