

AN ALTERNATIVE TO VINBERG'S ALGORITHM

DANIEL ALLCOCK

ABSTRACT. Vinberg's algorithm is the main method for finding fundamental domains for reflection groups acting on hyperbolic space. Experience shows that it can be slow. We explain why this should be expected, and prove this slowness in some cases. And we provide an alternative algorithm that should be much faster. It depends on an algorithm for finding vectors with small positive norm in indefinite binary quadratic forms, of independent interest.

1. INTRODUCTION

The Weyl group $W(L)$ of a Lorentzian lattice L is generated by reflections and acts properly discontinuously on hyperbolic space. Any one of its chambers serves as a fundamental domain, and computing $W(L)$ amounts to explicitly describing a chamber. Since the early 1970's, Vinberg's algorithm has been the tool for doing this. In this paper we present a new method, which has several advantages. In very general terms, Vinberg's algorithm is like solving Pell's equation $x^2 - ny^2 = 1$ by brute force, first trying $y = 1$, then $y = 2$, etc. Ours is like the exponentially faster solution using continued fractions. We hope to apply it to the problem of classifying all Lorentzian lattices which are reflective (ie have chambers of finite hyperbolic volume).

Vinberg's algorithm requires a choice of a point k of hyperbolic space, a choice of chamber C containing k , and a way to enumerate all reflections in $W(L)$ whose mirrors lie at any chosen hyperbolic distance from k . The output is the list of all simple roots of C . If there are only finitely many then the algorithm recognizes this and terminates. If not, then the algorithm might not terminate, but will still find all the simple roots if left to run forever.

Our algorithm requires k to be a vertex, and only finds the simple roots corresponding to k 's component of the topological boundary ∂C . Neither of these limitations is serious. First, k has always been chosen

Date: August 31, 2021.

2010 Mathematics Subject Classification. Primary: 20F55, 11Y16.

Supported by Simons Foundation Collaboration Grant 429818.

to be a vertex of C , except in a few special situations—the main example being Conway’s group $W(\text{II}_{25,1})$, whose chamber has infinite volume and infinitely many facets [6]. Second, ∂C is connected unless C has infinite volume. The usual method of recognizing that C has infinite volume is to find enough vertices of C that one can construct infinitely many diagram automorphisms of C . This applies without change in our situation (see remark 9.6). On the other hand, our method only works for lattices over \mathbb{Z} . Also, we do use Vinberg’s algorithm in a limited way, for working out the simple roots at a corner of the chamber, given the corner.

When comparing algorithms, one usually analyzes the asymptotics of the time or space required, as a function of the size of the input. This does not make sense in this case, because there is no upper bound on the size of the answer, in terms of the size of the input. Furthermore, Nikulin showed that there are only finitely many reflective lattices [11][7]. So there are constant upper bounds on the run-times of both algorithms (ours and Vinberg’s), when they are applied to reflective lattices. And applied to non-reflective lattices, both algorithms run forever. So there is no asymptotic way to compare them. (In a few nonreflective examples like $\text{II}_{25,1}$, Vinberg’s algorithm “terminates” after finding infinitely many simple roots in finitely many batches [6][4]. There are only finitely many examples like this [12].)

So our main goal is to prove the correctness of our algorithm. See remark 8.3 for a heuristic argument that it is much faster than Vinberg’s.

Experience shows that for some lattices, Vinberg’s algorithm can take a very long time. It also slows down dramatically as the dimension of L increases. In section 4 we prove that Vinberg’s algorithm can require astronomical time for fairly simple lattices. Our algorithm was motivated by the observed slowness, this proven slowness, and our pessimism about the performance of existing implementations of Vinberg’s algorithm in large dimensions ($\dim(L) > 6$ or so).

Our basic idea is to “walk along the edges of C ”. That is, given a vertex, we find the vertices at the other ends of the edges emanating from it. This amounts to searching for small-but-positive-norm vectors in 2-dimensional Lorentzian lattices, which we do using a variation of the standard method for solving Pell’s equation $x^2 - ny^2 = 1$. This appears in section 8; we hope it is of independent interest. After finding these neighboring vertices, we find their neighbors, and then their neighbors’ neighbors, etc. We stop when the endpoints of all edges are known (or all the walking tires us out). The stopping criterion, ie for

recognizing that C has finite volume, amounts to running out of edges to walk along (Theorem 9.5).

In section 2 we review some mostly-standard notation and language. In sections 3–4 we review Vinberg’s algorithm and establish the claimed slowness. Section 5 introduces “normed Dynkin diagrams”, an almost trivial variation on ordinary Dynkin diagrams. They are useful when finding the “batch 0” simple roots in sections 6 and 7, at ordinary vertices and cusps respectively, and in edgewalking in section 9, which also relies on our algorithms for rank 2 Lorentzian lattices in section 8.

2. PRELIMINARIES

We recall some standard language and notation. The only non-standard content is the notation $O^\uparrow(L)$ and the concepts of constraint lattices and almost-roots. The notation $\langle \dots \rangle$ means the sublattice/subgroup/subspace generated by \dots , depending on context.

A lattice L means a free abelian group equipped with a \mathbb{Q} -valued symmetric bilinear form, called the inner product or dot product. For $v, w \in L$, it is written $v \cdot w$, and v^2 is an abbreviation for $v \cdot v$, called the norm of v . We call v isotropic if its norm is 0, and L isotropic if it contains nonzero isotropic vectors. L is called Euclidean if it is positive definite. If $X \subseteq L$ then X^\perp means the sublattice consisting of all $v \in L$ having zero inner product with all members of X .

L is called integral if all inner products are integers. It is called nondegenerate if $L^\perp = 0$. In this case the dual lattice $L^* = \text{Hom}(L, \mathbb{Z})$ can be identified with the set of vectors in $L \otimes \mathbb{Q}$ having integral inner products with all elements of L . When L is integral and nondegenerate, then $L \subseteq L^*$, and L^*/L is called the discriminant group of L .

L is called Lorentzian if its signature is $(n, 1)$ for some $n \geq 1$, meaning that $L \otimes \mathbb{R}$ is the orthogonal direct sum of an n -dimensional positive definite subspace and a 1-dimensional negative definite subspace. In this case we make the following definitions. A vector of positive resp. negative norm is called spacelike, and a nonzero isotropic vector is called lightlike. The lightlike vectors fall into two connected components, called the two light cones. When one of them is distinguished, we sometimes refer to its convex hull as the future cone and call its elements future-directed. These terms are borrowed from relativity.

An isometry from one lattice to another is a group isomorphism that respects inner products. The orthogonal group $O(L)$ means the group of isometries from L to itself. Every isometry has determinant ± 1 , and $SO(L)$ means the index 2 subgroup with determinant 1. If L is

Lorentzian, then $O^\uparrow(L)$ and $SO^\uparrow(L)$ mean the index 2 subgroups of $O(L)$ and $SO(L)$ whose elements preserve each of the two light cones.

If $\alpha^2 > 0$ and $\alpha \cdot L \subseteq \frac{\alpha^2}{2}\mathbb{Z}$, then we call α an almost-root of L . The importance of almost-roots is that the reflection $x \mapsto x - 2(x \cdot \alpha)\alpha/\alpha^2$ in α is an element of $O^\uparrow(L)$. If α is also primitive (ie $L/\langle\alpha\rangle$ is torsion-free) then we call α a root. Every almost-root is either a root or twice a root. The ones which are not roots are not very interesting, but it is convenient to keep them around in the intermediate steps of our algorithms and then filter them out later. The convenience comes from the following lemma, which follows immediately from the definitions.

Definition/Lemma 2.1 (Constraint lattices). *Suppose L is a lattice and $N > 0$. Then the norm N vectors in the constraint lattice $L_N = L \cap \frac{N}{2}L^*$ are exactly the norm N almost-roots of L . \square*

Lemma 2.2 (Root norms). *Suppose L is a nondegenerate integral lattice and e is the exponent of the discriminant group L^*/L . Then the norm of any root of L divides $2e$. \square*

Proof. Suppose α is a root, and define $M = \alpha^\perp$. By $L \cdot \alpha \subseteq \frac{\alpha^2}{2}\mathbb{Z}$, $M \oplus \langle\alpha\rangle$ has index 1 or 2 in L . If the index is 1, then the discriminant group $\mathbb{Z}/\alpha^2\mathbb{Z}$ of $\langle\alpha\rangle$ is a direct summand of that of L , so $\alpha^2|e$.

Now suppose the index is 2, so L is generated by $M \oplus \langle\alpha\rangle$ and a vector $(u, \alpha/2)$ with $u \in M^*$. Integrality forces α^2 to be even. From $[L : M \oplus \langle\alpha\rangle] = 2$ follows $2u \in M$. But $u \notin M$, because otherwise $\alpha/2$ would lie in L , contradicting the hypothesis that α is a root. So the image of u in M^*/M has order 2.

Because M is integral, the \mathbb{Q} -valued inner product on M^* descends to a \mathbb{Q}/\mathbb{Z} -valued bilinear pairing on M^*/M . It is standard that this is nondegenerate. Since $u \notin M$, there exists $m \in M^*$ with $m \cdot u \not\equiv 0 \pmod{1}$. Since u has order 2 (mod M), the only possibility is $m \cdot u \equiv \frac{1}{2} \pmod{1}$. Therefore $\lambda = (m, \alpha/\alpha^2)$ has integral inner product with $(u, \alpha/2)$. Since λ also has integral inner products with α and all elements of M , we have $\lambda \in L^*$. Considering its second component shows that $n\lambda$ can only lie in L if $\frac{\alpha^2}{2}|n$. Therefore the exponent of L^*/L is divisible by $\alpha^2/2$. That is, $\alpha^2|2e$. \square

We defined reflections in roots (and almost-roots) above. One can define them more generally, for example allowing reflections in timelike vectors. But for us, reflections mean reflections in roots. A reflection group means a group generated by reflections. As a special case, the group $W(L)$ generated by all reflections of L is called the Weyl group of L , and lies in $O^\uparrow(L)$.

We write $\mathbb{R}^{n,1}$ for $(n + 1)$ -dimensional Minkowski space; $L \otimes \mathbb{R}$ is isometric to $\mathbb{R}^{n,1}$ for any Lorentzian lattice L . Hyperbolic space H^n means the image in $P\mathbb{R}^{n,1}$ of the timelike vectors, and its boundary ∂H^n means the image of the lightlike vectors. When we have one light cone distinguished, we sometimes identify H^n with the norm -1 hyperboloid in the future cone. In that case, we use future-pointing vectors to represent points of $\overline{H^n} := H^n \cup \partial H^n$.

The mirror of a reflection means its fixed-point set in H^n . Removing all mirrors from H^n leaves a disconnected set. The closure of any one of them in $\overline{H^n}$ is called a Weyl chamber (or just a chamber). In particular, chambers contains their ideal vertices. When we have a chamber in mind, we usually write C for it. The Weyl group $W(L)$ acts simply transitively on the chambers. The dihedral angles of C are integral submultiples of π ; in particular C has no obtuse dihedral angles. Sometimes we pass without comment between C and its preimage in the future cone.

Because $W(L)$ is discrete in $O^\uparrow(L \otimes \mathbb{R})$, its mirrors form a locally finite subset of H^n , so there is subtlety when speaking of faces of C of dimension > 1 . There is a little subtlety when speaking of vertices, if C has infinitely many facets. This is only relevant in section 9, where we give a precise definition of “corners”.

A root α is called a simple root of a chamber C if it is orthogonal to a facet of C and has negative inner product with vectors in the interior of C . Another way to say this is that C is the projectivization of the set of future-directed vectors that have negative inner product with all simple roots. (The convention for Lorentzian lattices is the opposite of the convention in Lie theory, where simple roots are inward-pointing not outward-pointing.) A standard consequence of the absence of obtuse dihedral angles is that the simple roots of a (chamber of a) positive-definite lattice are linearly independent.

3. VINBERG’S ALGORITHM

In this section we describe Vinberg’s algorithm [16, §3.2], at first geometrically, and then algebraically in the special case of reflection groups of lattices. Suppose W is a discrete group of isometries of H^n , generated by reflections. In this general setting, the two unit vectors in $\mathbb{R}^{n,1}$ orthogonal to the mirror of a reflection are called the roots of that reflection. Suppose $k \in \mathbb{R}^{n,1}$ is timelike; it is called the “control vector”, and we use the same symbol for its image in hyperbolic space H^n . Suppose known a set of simple roots B_0 (“batch 0”) for its W -stabilizer W_k . Let C_k be the corresponding chamber of W_k , ie the cone in $\mathbb{R}^{n,1}$

of vectors having nonpositive inner products with the batch 0 roots. There is a unique chamber C of W that lies in C_k and contains k . Vinberg's algorithm extends B_0 to the set B of simple roots for C .

The construction is inductive, so for each $D \geq 0$ we write B_D ("the distance D batch") for the set of simple roots for C whose mirrors lie at hyperbolic distance D from k . Obviously $B = \cup_{D \geq 0} B_D$, so it will suffice to find each B_D . We assumed B_0 known, which allows the induction to begin.

Suppose we have already found $B_{\leq D} = \cup_{d \leq D} B_d$. Consider all mirrors of W lying at distance $> D$ from k , and let D' be the distance to the closest of these. (D' exists because W is discrete.) For each mirror at that distance, consider its outward-pointing root α (ie, $\alpha \cdot k < 0$). We call α "accepted" if it has nonpositive inner products with all roots in $B_{\leq D}$. Vinberg's algorithm amounts to the statement that $B_{D'}$ is this set of accepted roots. (Of course, all batches between D and D' are empty.)

The algorithm might not terminate. This is in the nature of things, because some Weyl chambers have infinitely many facets. But even in this case, $B = \cup_{D \geq 0} B_D$. Also, in [16, Prop. 5] Vinberg established a criterion on $B_{\leq D}$: if the polytope they define has finite hyperbolic volume, then all later batches are empty. In this case we regard the algorithm as terminating, with $B = B_{\leq D}$. When one suspects that there are infinitely many simple roots, some ingenuity is needed to establish this rigorously. This is beyond the scope of this paper; the main method is to show that C has infinitely many isometries; see [3, Theorem 5.5], [13, II§3], [2, p. 20] and [15, Prop. 3.2] for examples of this.

Another difficulty is that one can only call it an algorithm if one has an algorithm for finding D' and all roots α of W with α^\perp at distance D' from k . Both problems are straightforward in the case of Lorentzian lattices. Namely, suppose L is a Lorentzian lattice and $W \subseteq \text{Aut}(L)$. One scales the control vector to be a primitive lattice vector. Also, one changes the definition of the roots of a mirror: now we use the two primitive lattice vectors orthogonal to it. This is convenient and clearly has no effect on the algorithm.

The hyperbolic distance between k and α^\perp is $\sinh^{-1}(-k \cdot \alpha / \sqrt{-k^2 \alpha^2})$. Because \sinh^{-1} is increasing and k^2 is constant, considering roots α in increasing order of D is the same as considering them in increasing order of $-k \cdot \alpha / \sqrt{\alpha^2}$, which we call the priority of α . The algorithm examines roots with smaller priority first. There are only finitely many possibilities for α^2 (Lemma 2.2). For fixed α^2 , the possibilities for

the priority p lie in a scaled copy of \mathbb{Z} . For fixed α^2 and p , one can enumerate all roots of L with those parameters. This boils down to iterating over all the elements of the lattice k^\perp , or one of finitely many translates of it inside $k^\perp \otimes \mathbb{Q}$, whose norm is specified in terms of α^2 and p . Together these facts allow one to find D' and $B_{D'}$, making concrete the inductive step of Vinberg's algorithm.

All of Vinberg's geometric arguments apply perfectly well with hyperbolic space H^n replaced by the sphere S^n . The changes needed to the above are the following: replace $\mathbb{R}^{n,1}$ by Euclidean space \mathbb{R}^{n+1} and \sinh^{-1} by \sin^{-1} , and assume $k \neq 0$ rather than that k is timelike.

Conway [6] found another variation of the classical form of the algorithm. He observed that the same algorithm works perfectly well with k allowed to be timelike, provided that we remain in the Lorentzian lattice setting. Two observations are needed. First, one uses a notion of "distance" to the ideal point of H^n , whose level sets are the horospheres centered there. Happily, this requires no change to the algorithm (when it is formulated in terms of the priority). Second, we can no longer use the discreteness of W to conclude that D' exists. However, the possible priorities are constrained just as before, so one can examine them all in increasing order. When non-empty, the set of roots of L with given norm and priority is now infinite, but falls into finitely many orbits under the unipotent subgroup of the $O(L)$ -stabilizer of k . Finding these orbits boils down to examining finitely many cosets of $k^\perp/\langle k \rangle$ (or a sublattice of it) in $(k^\perp/\langle k \rangle) \otimes \mathbb{Q}$. So again the search for D' and $B_{D'}$ can be made explicit.

At several points it will be convenient to use a minor variation on Vinberg's algorithm:

Lemma 3.1. *Suppose L is a Euclidean or Lorentzian lattice, W is a subgroup of $O(L)$ that is generated by reflections, and $k \in L \otimes \mathbb{R} - \{0\}$. Call k the control vector, and if L is Lorentzian then make the extra assumption that k is timelike or lightlike. Suppose C is a chamber of W that contains k , and B_0 is the set of simple roots of C that are orthogonal to k .*

Suppose β_1, β_2, \dots is a finite or infinite sequence of almost-roots of W , having negative inner products with k , such that

- (1) *they have nonpositive inner products with all members of B_0 ;*
- (2) *they include all the simple roots of C other than those in B_0 ;*
- (3) *they appear in nondecreasing order of priority;*
- (4) *those with equal priority appear in nondecreasing order of norm.*

Inductively define β_i as “approved” if it has nonpositive inner product with all of its predecessors $\beta_{j < i}$ which were approved.

Then the approved β_i are the simple roots of C that are not in B_0 .

Proof. Fix $i \geq 1$ and suppose that each β_j with $j < i$ is approved if and only if it is a simple root of C . We will prove the same for β_i .

First, supposing that β_i is a simple root of C , we will show that it is approved. All the simple roots of C with (nonzero) priority less than than β_i occur among the $\beta_{j < i}$, by hypotheses (2) and (3). By induction they were approved, and no other predecessors of β_i were. Since β_i is a simple root, it has nonpositive inner products with all other simple roots. Therefore it has nonpositive inner products with all approved predecessors. So β_i is approved.

Next, supposing β_i is a root and is approved, we will prove that it is a simple root of C . All simple roots of C with (nonzero) smaller priority occur among the $\beta_{j < i}$ by (2) and (3), and were approved, by induction. Since β_i is approved, it has nonpositive inner products with all of them. So β_i is accepted by Vinberg’s algorithm, and is therefore a simple root.

Finally, supposing that β_i is not a root, we must prove that it is not approved. Because β_i is an almost-root but not a root, $\beta_i/2$ is a root. Note that $\beta_i/2$ has the same priority as β_i , but strictly smaller norm.

If $\beta_i/2$ is a simple root, then it equals some β_j , by (2). In fact it is a predecessor of β_i , by (4). By induction it was approved. From $\beta_i \cdot \beta_i/2 > 0$ follows that β_i is not approved.

On the other hand, if $\beta_i/2$ is not a simple root, then it is not accepted by Vinberg’s algorithm. So $\beta_i/2$ has positive inner product with some simple root with strictly smaller priority than $\beta_i/2$. By (2) this simple root is some β_j , and considering priority shows that $j < i$. By induction β_j was approved. But $\beta_j \cdot \beta_i$ has the same sign as $\beta_j \cdot \beta_i/2 > 0$. So β_i is not approved. \square

4. SLOWNESS

In our generalization [2] of Nikulin’s classification of rank 3 reflective Lorentzian lattices [13, I], we ran Vinberg’s algorithm on 204 520 candidate lattices. Most candidates took a fraction of a second. But some took many hours, examining long sequences of batches that all turned out empty. There is a phenomenon in the rank 2 case that predicts this behavior, and persists in higher dimensions.

(The proof in [2] examines 857 candidates. Later we simplified the overall logic at the cost of testing more candidates, and improved our

implementation of Vinberg’s algorithm. The phenomenon of long sequences of empty batches was already apparent for the original 857.)

The slowness of Vinberg’s algorithm for some lattices is closely related to the well-known fact that the fundamental solution (x, y) to Pell’s equation $x^2 - ny^2 = 1$ may be large, even when n is small. For example, if $n = 106$ then $(x, y) = (32\,080\,051, 3\,115\,890)$. The naive way to find (x, y) is to start with the known solution $(1, 0)$, then look for solutions $(?, 1)$, then for solutions $(?, 2)$, and so on. This amounts to checking that $106 \cdot 1^2 + 1 = 107$ is not a square, then that $106 \cdot 2^2 + 1 = 425$ is not a square, etc. This process takes 3 115 890 iterations to find a solution. On the other hand, the continued fractions method of solving Pell’s equation takes 17 steps and can be done by hand.

The main point of this section is that Vinberg’s algorithm amounts to the naive approach. To make this precise, take L to have inner product matrix $\begin{pmatrix} 1 & 0 \\ 0 & -106 \end{pmatrix}$, take the control vector k to be $(0, 1)$, take batch 0 to consist of the root $(-1, 0)$, and take W to be the subgroup of $\text{Aut}(L)$ generated by the reflections in norm 1 roots. Then Vinberg’s algorithm exactly matches the previous paragraph, finding 3 115 889 consecutive empty batches. We made the artificial restriction to norm 1 roots in order to make this simple statement.

For the more natural case $W = W(L)$, theorem 4.1 shows that Vinberg’s algorithm will stop when it finds the root $r = (41234, 4005)$ of norm 106. This lies in the 4005th batch of norm 106 roots (taking into account the fact that 106 must divide $r \cdot k$). So the algorithm examines at least this many batches before stopping.

It is easy to give examples with much worse performance. In the rest of this section we suppose

- n is a square-free integer bigger than 1;
- \mathcal{O} is the ring of algebraic integers in $\mathbb{Q}(\sqrt{n})$;
- L is the lattice underlying \mathcal{O} (at the scale where 1 has norm 1);
- $W = W(L)$;
- k (the control vector) is a fixed square root \sqrt{n} of n ;
- $\alpha_0 = -1$ (this constitutes batch 0);
- C is the chamber of W containing k and having α_0 as a simple root;
- α is the other simple root of C .

We will show how to find α and give a lower bound for the number of batches Vinberg’s algorithm examines before finding α . The example above was the $n = 106$ case, and some other examples appear in table 1.

Theorem 4.1 (Slowness in rank 2). *In the setting above, take u to be a fundamental unit for \mathcal{O} , and $v = u$ or u^2 according to whether u has*

n	α^2	coefficient of 1	coefficient of \sqrt{n}	batch number
2	2	2	1	1
3	6	3	1	1
5	5	5/2	1/2	1
6	3	3	1	1
7	2	3	1	1
19	38	57	13	13
67	134	1809	221	221
73	73	9125	1068	2 136
97	97	55193	5604	11 208
193	193	24508105	1764132	3 528 264
241	241	1102388225	71011068	142 022 136
337	337	18648111017	1015827336	2 031 654 672
409	409	2263478264165	111921796968	223 843 593 936
601	601	3419107492676845	139468303679532	278 936 607 359 064
769	769	453881183125633513	16367374077549540	32 734 748 155 099 080

TABLE 1. The second simple root α of the lattice L underlying the ring of algebraic integers in $\mathbb{Q}(\sqrt{n})$. Running Vinberg's algorithm, with \sqrt{n} as control vector and $\{-1\}$ as batch 0, finds α in the indicated batch of roots of that norm.

norm 1 or -1 . Suppose u is chosen so that v has positive trace, and negative inner product with k . Then α is the smallest positive multiple of $1 + v$ that lies in \mathcal{O} .

Proof. The reflection R in the root 1, which acts by the composition of Galois conjugation and multiplication by -1 , lies in $O^\dagger(L)$. By the positive trace hypothesis, v lies in the same branch of the norm 1 hyperbola as 1. So its multiplication operator V also lies in $O^\dagger(L)$. Furthermore, by its definition in terms of u , v generates the multiplicative group of all lattice points in that branch of the hyperbola. It follows that $\langle V, R \rangle$ acts transitively on the norm 1 elements of L . Since the $O^\dagger(L)$ -stabilizer of 1 is trivial, it follows that $\langle V, R \rangle = O^\dagger(L)$. This group is infinite dihedral, because R inverts V . Because $V \circ R$ is also a reflection, we have $W(L) = O^\dagger(L)$.

Because Galois conjugation exchanges v and v^{-1} , it follows that $V \circ R$ sends 1 to $-v$ and v to -1 . Therefore $V \circ R$ negates $v + 1$, and it follows that the smallest positive multiple α' of $v + 1$ that lies in L is a root of L . Because α' has negative inner product with α_0 , $\{\alpha_0, \alpha'\}$ is a set of simple roots for the reflection group they generate. We have seen that this group is all of $W(L)$, so they are simple roots for L . Finally, by construction their chamber contains k . So $\alpha' = \alpha$. \square

We used the PARI/GP package [14] to work out many examples, a few of which appear in table 1. The last column says which batch of roots of that norm contains α . If \mathcal{O} has no elements with half-integral \sqrt{n} -component, then this is just the \sqrt{n} -component of α . Otherwise, the batch number is twice this. Vinberg's algorithm inspects at least this many batches to find α . Usually it will also inspect some (empty) batches of roots of other norms. So the later entries of the table show that Vinberg's algorithm can take a very long time.

Each of these troublesome lattices can be embedded in higher-dimensional Lorentzian lattices, dashing any hopes one might have that this is 2-dimensional phenomenon. For example, $(41234, 4005, 0, 0)$ is a simple root of $\begin{pmatrix} 1 & 0 \\ 0 & -106 \end{pmatrix} \oplus \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$. Therefore running Vinberg's algorithm with $(-1, 0, 0, 0)$ as control vector takes at least 4005 batches to find it. (Working out each batch in rank 4 also requires more work than in rank 2.)

5. NORMED DYNKIN DIAGRAMS

There are several variations on the definition of a Dynkin diagram. The one we use is the following: a graph in which each pair of nodes are joined by at most one bond, which may have any one of four types: single, double, triple and heavy. Double and triple bonds are oriented, and heavy bonds may be oriented or not. (We interpret triple bonds as in Lie theory, indicating an angle $\pi/6$, rather than the $\pi/5$ of [16]. Geometrically, heavy bonds indicate an angle " π/∞ ", meaning parallelism. For example, the bond in the affine diagram \tilde{A}_1 is a heavy bond.)

Suppose L is a lattice, $\alpha_1, \dots, \alpha_n$ are roots of it, and any two of them have nonpositive inner product and positive-definite or positive-semidefinite span. Then their Dynkin diagram is defined in the usual way: its nodes are $\alpha_1, \dots, \alpha_n$ with α_i, α_j joined (or not) as follows:

- (1) if $\alpha_i \perp \alpha_j$ then they are not joined;
- (2) if the inner product matrix of α_i and α_j is a rational multiple of

$$\begin{pmatrix} 2 & -1 \\ -1 & 2 \end{pmatrix}, \begin{pmatrix} 2 & -1 \\ -1 & 1 \end{pmatrix}, \begin{pmatrix} 6 & -3 \\ -3 & 2 \end{pmatrix}, \begin{pmatrix} 4 & -2 \\ -2 & 1 \end{pmatrix} \text{ resp. } \begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix}$$

then they are joined by a single bond, double bond, triple bond, oriented heavy bond, resp. unoriented heavy bond. In the oriented cases, the bond points from the longer root to the shorter.

Under the positive definiteness/semidefiniteness hypothesis, these alternatives cover all cases.

A norm on a Dynkin diagram means the assignment of a positive number to each node (its “norm”), subject to constraints corresponding to the previous paragraph:

- (1) two nodes joined by a single bond, or by an unoriented heavy bond, have equal norms;
- (2) if two nodes are joined by a double, triple, resp. oriented heavy bond, then the norm of the node at the tail of the arrow is 2, 3, resp. 4 times the norm of the node at the tip of the arrow.

A normed Dynkin diagram means a Dynkin diagram together with a norm on it. Given L and $\alpha_1, \dots, \alpha_n$ as above, the basic example of a norm on their Dynkin diagram is got by defining the norm of each node to be the norm of the corresponding root. Conversely, if Δ is a normed Dynkin diagram, then it arises in this way from a unique inner product on the free abelian group generated by the nodes of Δ : the norm on Δ gives the norms of the generators, and then the edges of Δ determine their pairwise inner products.

A Dynkin diagram Δ is called spherical if it admits a norm for which the inner product on A is positive definite. (An equivalent definition is to require this for every norm.) This is equivalent to each component of Δ having one of the classical types $A_n, B_n, C_n, D_n, E_6, E_7, E_8, F_4, G_2$. We call Δ cuspidal if it has a norm which is positive semidefinite but not definite. If Δ is connected, then this is equivalent to Δ appearing in the well-known list of “affine” Dynkin diagrams \tilde{A}_n, \dots (The nomenclature is less standard than in the spherical case; see [1],[8],[10].) If Δ is disconnected, then it is cuspidal if and only if every component is spherical or affine, with at least one affine component being present.

Now suppose Δ is a normed spherical Dynkin diagram and N is a positive number. By a norm N extension of Δ we mean an inclusion of Δ into a normed Dynkin diagram Δ' with one more node than Δ , the extending node having norm N . We call the extension spherical resp. cuspidal if Δ' is. These are the only extensions that will be important for us.

In our root-finding algorithms for Lorentzian lattices, we will need to find all such extensions of Δ . The main application is to find all possibilities for the root system at a vertex, given the root system at an incident edge. The classification of spherical Dynkin diagrams shows that every norm N spherical extension occurs in one of the following ways, where α denotes the extending node:

- (1) α is not joined with any node of Δ .
- (2) α is singly joined with 1, 2 or 3 nodes of Δ of norm N .

- (3) α is doubly joined to/from a node of Δ of norm $N/2$ or $2N$, and singly joined with 0 or 1 nodes of Δ of norm N .
- (4) α is triply joined to/from a node of Δ of norm $N/3$ or $3N$.

To find the norm N spherical extensions, one just constructs all these extensions and discards the ones that are not spherical. The size of this enumeration is fairly small, for example if Δ is an A_1^{20} diagram with all norms equal, then it has $1 + 20 + \binom{20}{2} + \binom{20}{3} = 1351$ spherical extensions with that norm. In our intended application, the classification of reflective Lorentzian lattices, this is a worst-case scenario, because Esselmann proved that no such lattices exist of rank > 22 . Typically there are just a few spherical extensions.

Enumerating the cuspidal norm N extensions of Δ is similar; every one occurs in one of the following ways:

- (1) α is singly joined with 1, 2, 3 or 4 nodes of norm N .
- (2) α is doubly joined to/from a node with norm $N/2$ or $2N$, and doubly joined to/from another node with norm $N/2$ or $2N$.
- (3) α is doubly joined to/from one node which has norm $N/2$ or $2N$, and singly joined with 0, 1 or 2 nodes of norm N .
- (4) α is triply joined to/from one node which has norm $N/3$ or $3N$, and singly joined with 0 or 1 nodes with norm N .
- (5) α is joined by an unoriented heavy bond with a node of Δ having norm N .
- (6) α is joined by an oriented heavy bond to/from a node of Δ having norm $N/4$ or $4N$.

One constructs all these extensions and discards the non-cuspidal ones.

6. BATCH 0 (SPHERICAL CASE)

In the setting which motivates this section, one is given a Lorentzian lattice L and a timelike lattice vector k . To begin Vinberg's algorithm, with k as control vector, one must find "batch 0", meaning a system of simple roots for the root system Π_k consisting of all roots of L that are orthogonal to k . Our goal in this section is to explain how to do this systematically. We assume given enough members of Π_k to span k^\perp rationally. Geometrically this corresponds to k being a vertex of a Weyl chamber. The Lorentzian hypothesis is not relevant in this section; all we need is a positive definite sublattice.

Problem 6.1 (Spherical Batch 0). *Suppose given a lattice L and a positive-definite subspace V of $L \otimes \mathbb{R}$. Suppose given some roots of L that span V . Find a set of simple roots for the (spherical) root system consisting of all roots of L in V .*

Problem 6.2 (Spherical Batch 0, iterative step). *Suppose L is an integral lattice and Π is its root system. Suppose $\alpha_1, \dots, \alpha_m \in \Pi$ are linearly independent with positive definite span. Write V_{m-1} for the real span of $\alpha_1, \dots, \alpha_{m-1}$, and similarly for V_m . Assuming that $\alpha_1, \dots, \alpha_{m-1}$ is a system of simple roots for $\Pi \cap V_{m-1}$, find the unique $\alpha'_m \in \Pi$ satisfying*

- (1) $\alpha_1, \dots, \alpha_{m-1}, \alpha'_m$ is a system of simple roots for $\Pi \cap V_m$, and
- (2) $\alpha'_m \in V_{m-1} + \mathbb{R}_{>0}\alpha_m$.

Solution. Write Δ for the normed Dynkin diagram Δ formed by $\alpha_1, \dots, \alpha_{m-1}$. Write k for the primitive lattice vector in $V_m \cap V_{m-1}^\perp$ that lies in the halfspace $V_{m-1} + \mathbb{R}_{<0}\alpha_m$.

Step 1: find a list \mathcal{N} of positive numbers, that contains the norms of all roots of L . (For example, appeal to Lemma 2.2.)

Step 2: for each $N \in \mathcal{N}$, find the set \mathcal{D}_N of all norm N spherical extensions $\Delta \rightarrow \Delta'$ of normed Dynkin diagrams.

Step 3: for each $N \in \mathcal{N}$, and each extension $\Delta \rightarrow \Delta'$ in \mathcal{D}_N , define $\alpha_{N,\Delta'}$ to be the (unique) element of $V_{m+1} + \mathbb{R}_{<0}k$ which has norm N and whose inner products with $\alpha_1, \dots, \alpha_{m-1}$ are as specified by Δ' . Call $\alpha_{N,\Delta'}$ a candidate if it lies in the constraint lattice $L_N = L \cap \frac{N}{2}L^$.*

Step 4: compute the priority $-k \cdot \alpha_{N,\Delta'} / \sqrt{N}$ of each candidate. Then α'_m is a candidate of minimal priority, indeed the (unique) one of them with smallest norm.

Proof. We will write Π_m for $\Pi \cap V_m$, and similarly for Π_{m-1} . These root systems have rank m and $m-1$ by the linear independence hypothesis. First we prove the existence and uniqueness of α'_m . Because $\alpha_1, \dots, \alpha_{m-1}$ is a set of simple roots for $\Pi_m \cap k^\perp = \Pi_{m-1}$, it has exactly two extensions to a system of simple roots for Π_m . Of the corresponding two chambers for Π_m , one contains k and the other $-k$. By its definition, k has negative inner product with α_m . Therefore condition (2) implies $k \cdot \alpha'_m < 0$. So the Weyl chamber for $\alpha_1, \dots, \alpha_{m-1}, \alpha'_m$ is the one which contains k . This proves existence and uniqueness.

We will prove that α'_m is a candidate in the sense of step 3. Writing N for $\alpha'_m{}^2$, we have $N \in \mathcal{N}$ by the construction of \mathcal{N} . Similarly, $\Delta \rightarrow \Delta'$ lies in \mathcal{D}_N , where Δ' is the normed Dynkin diagram of $\alpha_1, \dots, \alpha_{m-1}, \alpha'_m$. We claim $\alpha'_m = \alpha_{N,\Delta'}$. These vectors both have norm N , and they have the same inner products with $\alpha_1, \dots, \alpha_{m-1}$. To prove they are equal it is enough to prove that the signs of their inner products with k are equal. By its definition, $\alpha_{N,\Delta'}$ has negative inner product with k , and we saw above that α'_m does too. So $\alpha'_m = \alpha_{N,\Delta'}$. In particular, $\alpha_{N,\Delta'}$ is a norm N root of L , and therefore lies in the constraint lattice L_N . So α'_m is a candidate.

On the other hand, every candidate β is an almost-root of L , because it is a norm β^2 vector in the constraint lattice L_{β^2} . Because it lies in $V_{m-1} + \mathbb{R}_{<0}k$, β has negative inner product with k . Furthermore, its inner products with $\alpha_1, \dots, \alpha_{m-1}$ are given by a normed Dynkin diagram, and are therefore nonpositive. We have verified that the sequence of candidates, ordered first by priority and second by norm, satisfies the hypotheses of Lemma 3.1, with k as control vector.

The conclusion of that lemma is that step 4 reveals α'_m . In detail: We know that only one candidate is a simple root, and we already named it α'_m . Therefore α'_m is the only candidate approved by Lemma 3.1. The first candidate is automatically approved, so it is α'_m . This holds for any ordering of the candidates satisfying the ordering hypotheses (3)–(4) of Lemma 3.1. So only α'_m has smallest norm, among all candidates of smallest priority. \square

Solution to Problem 6.1. In the situation of Problem 6.1, suppose $\alpha_1, \dots, \alpha_n$ are the given roots that span V . By discarding those which are linear combinations of their predecessors, we may suppose they are linearly independent. For each m , we write V_m for the rational span of $\alpha_1, \dots, \alpha_m$. Obviously α_1 forms a system of simple roots for $\Pi \cap V_1$. For consistency with the inductive step, define $\alpha'_1 = \alpha_1$.

Now suppose $m > 1$ and that $\alpha'_1, \dots, \alpha'_{m-1}$ is a system of simple roots for $\Pi \cap V_{m-1}$. In particular, $\alpha'_1, \dots, \alpha'_{m-1}$ have the same rational span as $\alpha_1, \dots, \alpha_{m-1}$, so $\alpha'_1, \dots, \alpha'_{m-1}, \alpha_m$ are linearly independent with positive definite span. Using a solution to Problem 6.2, find the unique $\alpha'_m \in \Pi \cap (V_{m-1} + \mathbb{R}_{>0}\alpha_m)$, which together with $\alpha'_1, \dots, \alpha'_{m-1}$ forms a system of simple roots for $\Pi \cap V_m$. Repeating this argument through the $m = n$ case shows that $\alpha'_1, \dots, \alpha'_n$ are a system of simple roots for $\Pi \cap V$. \square

Remark 6.3. $\Pi \cap V$ has more than one system of simple roots, and we seldom care which one is used. But the proof distinguishes one, in terms of the initial roots $\alpha_1, \dots, \alpha_n$. One can use this to detect whether $\alpha_1, \dots, \alpha_n$ form a system of simple roots for $\Pi \cap V$. Namely, they do if and only if $\alpha'_m = \alpha_m$ for all m .

7. BATCH 0 (CUSPIDAL CASE)

This section has the same motivation as section 6, except with the control vector k being an ideal point of hyperbolic space. The method is conceptually similar to the iterative step of the spherical case, but some details are more complicated.

Problem 7.1 (Cuspidal Batch 0). *Suppose L is a Lorentzian lattice, Π is its root system, k is a primitive lightlike lattice vector, and V is its orthogonal complement in $L \otimes \mathbb{R}$. Suppose U is a complement to $\mathbb{R}k$ in V , and u_1, \dots, u_n are a system of simple roots for $\Pi \cap U$ and span U . Let $C \subseteq H^{n+1} \cup \partial H^{n+1}$ be the unique Weyl chamber for $W(\Pi \cap V)$ that contains k and whose simple roots include u_1, \dots, u_n . Find the remaining simple roots of C .*

Remark 7.2. Suppose that instead of u_1, \dots, u_n we are given enough members of $\Pi \cap V$ to span V modulo $\mathbb{R}k$, with no other hypotheses. Then by discarding some of them, and applying the spherical batch 0 case, we could construct u_1, \dots, u_n with the properties stated in the problem.

Solution to Problem 7.1. *Write Δ for the normed Dynkin diagram of u_1, \dots, u_n . Take the future cone to be the convex hull of the light cone that contains k . Choose k' to be any future-directed vector in the plane U^\perp , but not in $\mathbb{R}k$.*

Step 1: Find a list \mathcal{N} of positive numbers, that contains the norms of all the roots of L . (For example, appeal to Lemma 2.2.)

Step 2: For each $N \in \mathcal{N}$, find the set \mathcal{D}_N of all norm N cuspidal extensions $\Delta \rightarrow \Delta'$ of normed Dynkin diagrams. (See Section 5.)

Step 3: For each $N \in \mathcal{N}$, and each extension $\Delta \rightarrow \Delta'$ in \mathcal{D}_N , consider the element $u_{N, \Delta'}$ of U whose inner products with u_1, \dots, u_n are as specified by Δ' . If the constraint lattice $L_N = L \cap \frac{N}{2}L^$ contains a vector of the form $u_{N, \Delta'} + ck$ with $c \in \mathbb{Q}$, then define $v_{N, \Delta'}$ as the vector of this form with c as small as possible subject to being positive. Call these vectors the candidates.*

Step 4: Compute the priority $-k' \cdot v_{N, \Delta'} / \sqrt{N}$ of each candidate, using k' as control vector. Order the candidates first by priority and second by norm.

Step 5: Inductively define a candidate to be approved if it has non-positive inner products with all of its predecessors that were approved. Then the simple roots of C are u_1, \dots, u_n and the approved candidates.

Proof. Because u_1, \dots, u_n are among the simple roots of C , and k' is orthogonal to them, k' lies in an edge of C . So all remaining simple roots of C have negative inner products with k' .

We begin by claiming that every simple root α of C , other than u_1, \dots, u_n , is a candidate in the sense of step 3. First, $\alpha \notin U$, because u_1, \dots, u_n are already a set of simple roots for the n -dimensional spherical root system $\Pi \cap U$. Therefore u_1, \dots, u_n, v are a basis for V . Second, write N for α^2 ; by definition \mathcal{N} contains N . Third, write Δ'

for the normed Dynkin diagram got by adjoining α to Δ . Because Δ' forms a basis for V , which is positive semidefinite but not positive definite, $\Delta \rightarrow \Delta'$ is a norm N cuspidal extension. Therefore $u_{N,\Delta'}$ is defined. Its inner products with u_1, \dots, u_n are determined by the diagram extension, so they are the same as the inner products of α with u_1, \dots, u_n . It follows that $u_{N,\Delta'}$ is the projection of α to U .

Therefore $\alpha = u_{N,\Delta'} + ck$ for some $c \in \mathbb{Q}$. From $\alpha \cdot k' < 0$, $k \cdot k' < 0$ and $\alpha \cdot k' = ck \cdot k'$, we see $c > 0$. We have $\alpha \in L_N$ because α is a norm N root. So the candidate $v_{N,\Delta'}$ is defined. We claim that α coincides with it. By definition, $v_{N,\Delta'} \in L_N$ has the form $u_{N,\Delta'} + c_0k$ with $0 < c_0 \leq c$. So it is enough to prove $c_0 = c$. We suppose $c_0 < c$ and derive a contradiction. Since $v_{N,\Delta'}$ is a norm N vector in L_N , it is an almost-root, so it is either a root or twice a root. In either case, the mirror $v_{N,\Delta'}^\perp$ strictly separates k' from α^\perp . Therefore no chamber having α as a simple root can contain k' , which is a contradiction.

Next we observe that every candidate β is an almost-root, since it is a norm β^2 element of L_{β^2} . By the $c > 0$ condition in the definition of β , we have $k' \cdot \beta < 0$. The inner products of β with u_1, \dots, u_n are given by the normed Dynkin diagram Δ' and are therefore nonpositive. Now Lemma 3.1, applied to the root system $\Pi \cap V$, finishes the proof. \square

8. SHORT VECTORS IN 2-DIMENSIONAL LORENTZIAN LATTICES

Throughout this section, take L to be a 2-dimensional Lorentzian lattice and $k \in L \otimes \mathbb{Q}$ to be timelike or lightlike. Write P for the plane $L \otimes \mathbb{R}$, and suppose $\frac{1}{2}P$ is an open halfplane in P bounded by $\mathbb{R}k$, that contains timelike vectors arbitrarily close to k . We think of k as defining a point (ordinary or ideal) of hyperbolic 1-space H^1 , and $\frac{1}{2}P$ as representing a ray emanating from it. The condition on the existence of spacelike vectors near k corresponds to the fact that there is only one “ray” emanating from an ideal point of H^1 , rather than two from an ordinary point. The goal of this section is to search for short spacelike lattice vectors $r \in \frac{1}{2}P$ whose orthogonal complements are near k .

We write S for the sector of vectors in $\frac{1}{2}P$ with nonnegative norm. We illustrate the situation in figure 8.1, according to whether k is timelike or lightlike. We order the rays emanating from 0 into S according to their slope in the figure (arrows indicate the direction of increase). This induces a partial order on S and a total order on the set of primitive lattice vectors that lie in S . We write Ω for the set of lightlike vectors along the top of S . Take r to be the primitive lattice vector in S that is orthogonal to k . In section 9, L will be the projection of a higher-dimensional lattice to a plane P , and r will be got by projecting




FIGURE 8.1. The plane $P = L \otimes \mathbb{R}$, with timelike or lightlike vertex k . The diagonal lines are lightlike, the shading indicates $\frac{1}{2}P$, and the circular arcs indicate the ordering on the sector S .

a root of it to P and then scaling to make it primitive. Section 9 is also the source of the label “ray/line to walk along”.

Problem 8.1. *Given a primitive lattice vector $r \in S - \Omega$, and $M > 0$, find the primitive lattice vectors of norm $\leq M$ in S that come after r , in increasing order.*

Algorithm 8.2 below contains the essential idea of the solution. The name “Promised” refers to the hypothesis that S contains a primitive lattice vector of norm $\leq M$, after r . There are two natural settings in which this is automatic. First, if L is isotropic, then the primitive lattice vector in Ω comes after r (by the hypothesis $r \notin \Omega$) and has norm 0. Second, if L is anisotropic and $r^2 \leq M$, then S contains the images of r under $\text{SO}^\uparrow(L) \cong \mathbb{Z}$, half of which come after r .

We distinguish an orientation on P by taking a basis whose first member lies in the interior of $\frac{1}{2}P$ and whose second member is k . An “oriented basis” will mean a basis for L representing this orientation.

Suppose $M > 0$ and $r \in S$ is a primitive lattice vector. Then any $l \in L$ is called an M -supplement of r if two conditions hold. First, (r, l) must be an oriented basis. Second, l must not lie in the interior of the convex hull of the norm M hyperbola in S . We think of r as standing for “right” because it points into S , and l for “left”. If the first condition holds, then by subtracting a multiple of r from l we can arrange for the second condition to also hold.

Algorithm 8.2 (Promised(M, r, l)). *Suppose $r \in S - \Omega$ is a primitive lattice vector, $M > 0$, and l is an M -supplement of r . Suppose also that S contains a primitive lattice vector of norm $\leq M$ that comes after r .*

Then this recursive algorithm returns a pair (r', l') , where r' is the first such vector, and l' is an M -supplement of r' .

- (1) [Find middle] Set $m \leftarrow r + l$.
- (2) [Go right?] If $m^2 \leq M$ or $m \cdot l < 0$ then return $\text{Isotropic}(M, r, m)$.
- (3) [Done?] If $l^2 \geq 0$ and $r \cdot l > 0$ then return $(l, -r)$.
- (4) [Go left] Return $\text{Isotropic}(M, m, r)$.

We first remark that r' exists. By hypothesis, S contains a primitive lattice vector of norm $\leq M$ after r . The set of such vectors is discrete in P , and it follows easily that there is a first one.

We also remark that in the case of L isotropic, this algorithm suffices to solve Problem 9.2. Given r , one finds an M -supplement for it, and then repeats the algorithm finitely many times, using the output of each iteration as input for the next. The process stops once the primitive lattice vector in Ω is found.

Proof. First we note that r' lies in the open half-plane $\mathbb{R}r + \mathbb{R}_{>0}l$, because it comes after r in the ordering on S . We claim that r' even lies in the half-open sector $\mathbb{R}_{\geq 0}r + \mathbb{R}_{>0}l$. If $l^2 < 0$ or $l \cdot r < 0$, then the intersection of the half-plane with S lies in this sector, so certainly r' does too. On the other hand, if $l^2 \geq 0$ and $l \cdot r \geq 0$, then l lies in S , and it has norm $\leq M$ by the M -supplement hypothesis. By definition, $r' \leq l$ with respect to the ordering on S , so again r' lies in $\mathbb{R}_{\geq 0}r + \mathbb{R}_{>0}l$.

The proof that the algorithm terminates, and gives the claimed results, uses induction on the sum of the (just-proven-to-be nonnegative) coefficients of r' with respect to r, l . The base of the induction is step 3. For orientation we remark that the primitive lattice vector m lies in the “middle” of this sector, and subdivides it into the “right” subsector $\mathbb{R}_{\geq 0}r + \mathbb{R}_{>0}m$ and “left” subsector $\mathbb{R}_{\geq 0}m + \mathbb{R}_{>0}l$. The “go right” and “go left” steps search for r' in these subsectors.

First suppose step 2 applies. Then r' lies in the right subsector $\mathbb{R}_{\geq 0}r + \mathbb{R}_{>0}m$, by the same argument we used to prove $r' \in \mathbb{R}_{\geq 0}r + \mathbb{R}_{>0}l$. Furthermore, the conditions for step 2 say that m is an M -supplement of r , so M, r, m satisfy the hypotheses of $\text{Isotropic}(\cdot, \cdot, \cdot)$. The m -coefficient of r' with respect to (r, m) is the same as the l -coefficient of r' with respect to (l, m) . And the r -coefficient of r' with respect to (r, m) is strictly smaller than with respect to (r, l) . By induction, $\text{Isotropic}(M, r, m)$ returns (r', l') where l' is an M -supplement to r' . This completes the proof if step 2 applies.

Now suppose step 3 applies. It will be enough to prove $r' = l$, because $(l, -r)$ is an oriented basis, and $-r$ is an M -supplement to l by $-r \notin S$. By the conditions for step 3, l lies in S , so the M -supplement hypothesis forces $l^2 \leq M$. Because $r' \in \mathbb{R}_{\geq 0}r + \mathbb{R}_{>0}l$, to show $r' = l$ it suffices to

show that every $v \in \mathbb{Z}_{>0}r + \mathbb{Z}_{>0}l$ has norm $> M$. Note that v can be expressed as the sum of m and some nonnegative multiples of r and l . Therefore $v^2 > M$ follows from the facts $r^2 \geq 0$ (since $r \in S$), $l^2 \geq 0$ and $l \cdot r > 0$ (since this step does apply), $m^2 > M$ (since step 2 didn't apply), and $m \cdot l \geq 0$ and $m \cdot r \geq 0$ (since $m = r + l$). This completes the proof if step 3 applies.

Finally, suppose step 4 applies. The calculation in the previous paragraph shows that every $v \in \mathbb{Z}_{\geq 0}r + \mathbb{Z}_{>0}m$ has norm $> M$. Therefore r' cannot lie in the right subsector, and must lie in the left subsector $\mathbb{R}_{\geq 0}m + \mathbb{R}_{>0}l$. Since step 2 didn't apply, m lies in S , but not in Ω . Obviously (m, l) is an oriented basis for L . Since l is an M -supplement to r , it is also one for m . Therefore (M, m, l) satisfy the hypotheses of $\text{Isotropic}(\cdot, \cdot, \cdot)$. Arguing as for step 2 shows that the coefficients of r' with respect to (m, l) have smaller sum than do its coefficients with respect to (r, l) . By induction, $\text{Isotropic}(M, m, l)$ returns (r', l') where l' is an M -supplement for r' . This completes the proof. \square

Remark 8.3 (Performance). Write (r_n, l_n) for the second and third arguments to the n th call to **Promised**, starting with (r_0, l_0) . Write $\begin{pmatrix} a_n \\ b_n \end{pmatrix}$ for the coefficients of the final answer r' with respect to (r_n, l_n) . Because $(r_{n+1}, l_{n+1}) = (r_n, r_n + l_n)$ or $(r_n + l_n, l_n)$, $\begin{pmatrix} a_{n+1} \\ b_{n+1} \end{pmatrix}$ is got from $\begin{pmatrix} a_n \\ b_n \end{pmatrix}$ by left-multiplying by $\begin{pmatrix} 1 & -1 \\ 0 & 1 \end{pmatrix}$ or $\begin{pmatrix} 1 & 0 \\ -1 & 1 \end{pmatrix}$. All coefficients arising this way are positive. Therefore the algorithm amounts to running the “subtract the smaller from the larger” version of the Euclidean algorithm on the coefficients $\begin{pmatrix} a_0 \\ b_0 \end{pmatrix}$. Of course, we don't know ahead of time what these are. Running the Euclidean algorithm backwards amounts to computing the continued fraction expansion of a_0/b_0 . This is the sense in which our alternative to Vinberg's algorithm resembles the continued fraction expansion approach to Pell's equation.

A simple modification of our algorithm makes it analogous to the usual “replace the larger by its remainder upon division by the smaller” version of the Euclidean algorithm. We group together consecutive applications of step 6, by calling $\text{Promised}(M, r, \text{Canonical}(M, r, l))$ instead of $\text{Promised}(M, r, r + l)$. Here $\text{Canonical}(M, r, l) = l + ?r$ is the canonical M -supplement discussed below. We group together consecutive applications of step 4 in a similar way.

It is standard, but usually phrased differently, that the size of the arguments to the (usual) Euclidean algorithm grows exponentially in the number of steps the algorithm takes. The same reasoning shows that the time required to write down the result of our modified Algorithm 8.2 (in decimal notation) grows at least linearly in the number of steps taken. This does not prove that the run time is linear in the

size of the answer, because as the numbers grow larger, the arithmetic operations in each step take longer. But it does suggest that Algorithm 8.2 will perform very quickly. Since the rest of the algorithms in this section, and our alternative to Vinberg's algorithm (section 9), are just wrappers around it, they should be similarly fast.

See [5, §5.7.1–5.7.2] and references there for more precise estimates of the running time of the continued fraction algorithm, taking into account the growth of the coefficients. See also the survey [9] for other approaches to Pell's equation.

In the rest of this section we suppose L is anisotropic. The main complication is that $\text{SO}^\uparrow(L) \cong \mathbb{Z}$, and our algorithms must search for a generator at the same time they search for vectors of small norm. To do this we introduce what we call the canonical M -supplement to a primitive lattice vector $r \in S$, for a given $M > 0$. This is the unique M -supplement l for r , for which $l + r$ is not an M -supplement for r . The point of this construction is that $\text{SO}^\uparrow(L)$ respects canonical M -supplementation: if $\phi \in \text{SO}^\uparrow(L)$, then the canonical M -supplement of $\phi(r)$ is the ϕ -image of the canonical M -supplement of r .

To find the canonical M -supplement, start with any M -supplement and repeatedly add r to it. Alternately, one can give a formula for it:

Lemma 8.4 (Canonical(M, r, l)). *Suppose L is integral and anisotropic, M is a positive integer and (r, l) is an oriented basis for P with $r \in S$. Then the canonical M -supplement of r is $l + Kr$ where*

$$K = \left\lfloor \frac{-r \cdot l + \lfloor \sqrt{(r \cdot l)^2 - r^2(l^2 - M)} \rfloor}{r^2} \right\rfloor$$

□

We explained above that if $r^2 \leq M$, then **Promised** is enough to find the next primitive lattice vector $r' \in S$ of norm $\leq M$. But if $r^2 > M$, then such a vector might not even exist. To approach the problem we start with a more modest goal: finding a shorter spacelike vector, or determining that none exist.

Algorithm 8.5 (Shorter(r, l)). *Suppose L is anisotropic, $r \in S$ is a primitive lattice vector, and l is its canonical r^2 -supplement.*

If L has no spacelike vectors of norm $< r^2$, then this algorithm says so. Otherwise it returns (r', l') , where r' is the first primitive such vector of S after r , and l' is the canonical $(r')^2$ -supplement of r' .

- (1) [Constants] Set $M \leftarrow r^2$ and then $A \leftarrow \begin{pmatrix} M & r \cdot l \\ r \cdot l & l^2 \end{pmatrix}$.
- (2) [Enter main loop] Set $(r, l) \leftarrow \text{Promised}(r^2, r, l)$
- (3) [Canonicalize] Set $l \leftarrow \text{Canonical}(r^2, r, l)$.

- (4) [Found?] If $r^2 < M$ then return (r, l) .
- (5) [Nonexistence?] If the inner product matrix of (r, l) is equal to A , then stop: L contains no spacelike vectors of norm $< M$.
- (6) [Repeat] Go back to step 2.

Proof. We write (r_n, l_n) for the value of (r, l) at the n th entry to step 2. We will prove the following inductively. First, r_n lies in S and has norm M , and l_n is its canonical M -supplement. Second, 0 is the only lattice vector of norm $< M$ in the sector $\mathbb{R}_{\geq 0}r_1 + \mathbb{R}_{\geq 0}r_n$. These properties are trivial for (r_1, l_1) . Now suppose $n \geq 1$; we will work through the steps and see what the algorithm does.

Step 2 makes sense because S contains vectors of norm r_n^2 that come after r_n , for example some of the $\text{SO}^\uparrow(L)$ -images of r_n . After this step and the next, r is the first vector of norm $\leq M$ in S after r_n , and l is the canonical r^2 -supplement to r .

If step 4 applies, then by induction, r is the first vector of norm $< M$ in S after r_1 , and we already know that l is its canonical r^2 -supplement. So the algorithm halts with the claimed result. So suppose step 4 does not apply. That is, $r^2 = M$. By induction, 0 is the only lattice vector of norm $< M$ in the sector $\mathbb{R}_{\geq 0}r_1 + \mathbb{R}_{\geq 0}r$.

If step 5 applies, then (r, l) has the same inner product matrix as (r_1, l_1) , so the linear transformation sending (r_1, l_1) to (r, l) is a member of $\text{SO}^\uparrow(L)$. Therefore S is the union of the $\text{SO}^\uparrow(L)$ -images of the sector $\mathbb{R}_{\geq 0}r_1 + \mathbb{R}_{\geq 0}r$, which we just saw contains no lattice vectors of norm $< M$ except for 0. So the same holds for S , and the algorithm halts with the claimed result.

If step 5 does not apply, then we return to step 2, so the pair (r_{n+1}, l_{n+1}) is defined. We have already verified our inductive hypotheses on it. It remains only to show that the algorithm terminates. Otherwise we would have an infinite sequence $(r_1, l_1), (r_2, l_2), \dots$, where the r_n are consecutive norm M lattice vectors in S , and each l_n is the canonical M -supplement to r_n . Because there are only finitely many $\text{SO}^\uparrow(L)$ -orbits of norm M lattice vectors, some $g \in \text{SO}^\uparrow(L)$ sends r_1 to some $r_{n>1}$. By the nature of the canonical M -supplement, g sends l_1 to l_n . So (r_n, l_n) has inner product matrix A , and the algorithm would have halted via step 5. \square

Algorithm 8.6 ($\text{NotPromised}(M, r, l)$). *Suppose L is anisotropic, $r \in S$ is a primitive lattice vector, l is the canonical r^2 -supplement to r , and $M > 0$.*

If L has no spacelike vectors of norm $\leq M$, then this algorithm says so. Otherwise, it returns (r', l') where r' is the first primitive such vector in S after r , and l' is the canonical M -supplement of r' .

- (1) **[Easy case]** If $r^2 \leq M$, then set $(r, l) \leftarrow \text{Promised}(M, r, l)$ and then $l \leftarrow \text{Canonical}(M, r, l)$, and return (r, l) .
- (2) **[Nonexistence?]** Set $(r, l) \leftarrow \text{Shorter}(r, l)$, unless **Shorter** (r, l) reports nonexistence. In the latter case, stop: S contains no vectors of norm $\leq M$.
- (3) **[Done?]** If $r^2 \leq M$, then set $l \leftarrow \text{Canonical}(M, r, l)$ and return (r, l) .
- (4) **[Repeat]** Go back to step 2.

Proof. If case 1 applies, then S contains primitive lattice vectors of norm $\leq M$ after r , for example some of the $\text{SO}^\uparrow(L)$ -images of r . The only hypothesis of **Promised** (\cdot, \cdot, \cdot) that requires checking is that l is a not-necessarily-canonical M -supplement of r . This follows from the fact that it is an r^2 -supplement and $r^2 \leq M$. We refer to **Promised** to justify the result in this case. (**Promised** does not canonicalize the M -supplement in its output, which is why we do it here.) So suppose case 1 does not apply. That is, $r^2 > M$.

We write (r_n, l_n) for the value of (r, l) at the n th entry to step 2. We will prove by induction that the r_i form an increasing sequence of primitive lattice vectors in S , the norms of the r_i exceed M and form a strictly decreasing sequence, and that there are no lattice vectors of norm $< r_n^2$ except 0 in the sector $\mathbb{R}_{\geq 0}r_1 + \mathbb{R}_{\geq 0}r_n$. When $n = 1$ these claims are trivial. Now suppose $n \geq 1$; we will work through the steps and see what the algorithm does.

In step 2, if **Shorter** (r_n, l_n) reports nonexistence, then S contains no lattice vectors of norm $< r_n^2$, hence none of norm $\leq M$, as claimed. Otherwise, r is the first primitive lattice vector after r_n in S with norm $< r_n^2$, and l is the canonical r^2 -supplement of r .

In particular, if $r^2 \leq M$, then r is the first primitive lattice vector in S after r_n with norm $\leq M$. By inductive hypothesis the same statement holds with r_1 in place of r_n . This justifies step 3 if it applies. (Before the canonicalization in this step, l is the canonical r^2 -supplement of r , which might not be the canonical M -supplement.)

So suppose $r^2 > M$. Then we return to step 2, so (r_{n+1}, l_{n+1}) is defined. Along the way we have proven that it satisfies the inductive hypotheses. All that remains is to show that the algorithm terminates. This holds because r_1^2, r_2^2, \dots is strictly decreasing, and L has only finitely many possible norms $\leq r_1^2$ of spacelike vectors. \square

Finally we can solve the anisotropic case of Problem 9.2.

Algorithm 8.7 (*Anisotropic*(M, r, l)). Suppose L is anisotropic, $r_0 \in S$ is a primitive lattice vector, $M > 0$, and l_0 is the canonical r_0^2 -supplement of r_0 . If L has no spacelike vectors of norm $\leq M$, then this algorithm says so. Otherwise, it returns a generator g of $\text{SO}^\uparrow(L)$, and a nonempty sequence $\mathcal{R} = (r_1, r_2, \dots, r_n)$ of vectors, such that:

$$r_1, \dots, r_n, g(r_1), \dots, g(r_n), g^2(r_1), \dots, g^2(r_n), \dots$$

is the sequence, in ascending order, of all primitive lattice vectors in S that have norm $\leq M$ and come after r_0 .

- (1) [Nonexistence?] Set $(r_1, l_1) \leftarrow \text{NotPromised}(M, r_0, l_0)$, unless $\text{NotPromised}(M, r_0, l_0)$ reports nonexistence. In the latter case, stop: L has no spacelike vectors with norm $\leq M$.
- (2) [Initialize]
 - (a) set \mathcal{R} to be the one-term sequence (r_1) ,
 - (b) set $(r, l) \leftarrow (r_1, l_1)$, and
 - (c) set $A \leftarrow \begin{pmatrix} r^2 & r \cdot l \\ r \cdot l & l^2 \end{pmatrix}$.
- (3) [Enter main loop] Set $(r, l) \leftarrow \text{Promised}(M, r, l)$ and then $l \leftarrow \text{Canonical}(M, r, l)$.
- (4) [Done?] If the inner product matrix of (r, l) equals A , then set g to be the linear transformation sending (r_1, l_1) to (r, l) , and return g and \mathcal{R} .
- (5) [Append & repeat] Append r to \mathcal{R} , and go to step 3.

Proof. If NotPromised reports nonexistence in step 1, then L has no spacelike vectors of norm $\leq M$, as claimed. Otherwise, step 2 sets r to be the first primitive lattice vector of S with norm $\leq M$ after r_0 , and sets l to be its canonical M -supplement.

Now consider the n th entry to step 3, at which point the initial subsequence r_1, \dots, r_n of \mathcal{R} has been defined, r is another name for r_n , l and l_1 are the canonical M -supplements of r and r_1 , and the following inductive hypotheses hold. First, each $r_{i \geq 1}$ is the first primitive lattice vector in S of norm $\leq M$ that comes after r_{i-1} . Second, no element of $\text{SO}^\uparrow(L)$ sends r_1 to any of r_2, \dots, r_n . $\text{Promised}(M, r, l)$ makes sense in step 3 because $r_n^2 \leq M$.

Upon entering step 4, r is the first primitive lattice vector in S of norm $\leq M$ that comes after r_n , and l is its canonical M -supplement. If step 4 applies, then g is an isometry. It lies in $\text{SO}(L)$ because (r_1, l_1) and (r, l) are oriented bases. Together with $r_1, r \in S$ this shows that $g \in \text{SO}^\uparrow(L)$. It is a generator for $\text{SO}^\uparrow(L)$ because otherwise a generator would send r_1 to some primitive lattice vector in S that lies strictly between r_1 and r . This would contradict the fact that r_1, \dots, r_n are all the norm $\leq M$ primitive lattice vectors that come strictly after r_0

and strictly before r . Since $g(r_1) = r$, this also justifies the rest of our claims.

If step 4 does not apply, then step 5 defines $r_{n+1} = r$, and l is already its canonical M -supplement. No element of $\mathrm{SO}^\uparrow(L)$ sends r_1 to r_{n+1} , because it would also send l_1 to l and the algorithm would have stopped at the previous step. So the induction continues after returning to step 3. The algorithm terminates, because otherwise the infinite sequence r_1, r_2, \dots would contain $g(r_1)$, where g is a generator of $\mathrm{SO}^\uparrow(L)$. But then the algorithm would have stopped via step 4. \square

9. EDGEWALKING

This section contains the heart of the paper. Informally stated, we solve the following problem: given a corner of a Weyl chamber, and an edge emanating from it, find the corner at the other end of that edge. If the chamber has finite volume, then repeating this process will find all the corners and simple roots.

Before making a formal statement, we say what we mean by a corner. Suppose L is a Lorentzian lattice of dimension $n + 1$, Π is its root system, C is a chamber for $W(L)$, and $\{\alpha_i\}$ are simple roots for C . An ordinary corner of C means a point $k \in C \cap H^n$ for which $k^\perp \cap \Pi$ has rank n . An ideal corner means a point $k \in C \cap \partial H^n$, such that every horosphere centered there meets C in a compact set. An edge means the obvious thing; we only mention that we take edges to contain their ideal endpoints. Topologically each edge is a closed segment, and if it is incident at a corner, then the corner is an endpoint of the segment.

Lemma 9.1 (Edges end at corners). *In the notation above, suppose C is a chamber for $W(L)$, k is a corner of C , and e is an edge incident to k . Then the other end of e is also a corner of C .*

Proof. Write u_1, \dots, u_{n-1} for the simple roots of C orthogonal to e , and k' for the other end of e . If k' is an ordinary point, then the fact that e ends there forces the existence of a simple root β which is orthogonal to k' but not e . So $k'^\perp \cap \Pi$ contains $u_1, \dots, u_{n-1}, \beta$, whose span has rank n .

Now suppose k' is ideal. We use the upper half-space model, with k' at vertical infinity. Each horosphere H centered at k' appears as a horizontal plane, and the mirrors u_i^\perp appear as vertical halfplanes whose intersection is a vertical half-line. The $O(L)$ -stabilizer of k' contains a normal subgroup \mathbb{Z}^n consisting of horizontal translations, that acts cocompactly on H . The \mathbb{Z}^n -images of the u_i^\perp are roots of L . Their orthogonal complements cut H into bounded pieces. $C \cap H$ lies in the closure of one of these pieces, and is closed. So it is compact. \square

Problem 9.2 (Edgewalking). *Suppose L is a Lorentzian lattice of rank $n + 1$, Π is its root system, $W = W(L)$, C is a Weyl chamber of W , k is a corner of C , and e is an edge of C incident to k . Find the corner k' at the other end of e , and the simple roots of C that are orthogonal to k' .*

Our solution uses the following notation:

- u_1, \dots, u_{n-1} are the simple roots of C that are orthogonal to e ;
- Δ is the normed Dynkin diagram of u_1, \dots, u_{n-1} ;
- U is the real span of u_1, \dots, u_{n-1} ;
- P is the plane in $L \otimes \mathbb{R}$ orthogonal to U ;
- π_U, π_P are the orthogonal projection maps to U and P ;
- $\frac{1}{2}P$ is the open half-plane in P corresponding to e ;
- S is the sector of spacelike and lightlike vectors in $\frac{1}{2}P$;
- r_0 is the primitive lattice vector in S that is orthogonal to k ;
- L_N , for any given $N > 0$, is the constraint lattice $L \cap \frac{N}{2}L^*$.

In section 8 we defined a partial order on S . We refine it here to a total order, by declaring that if two members of S are proportional, then the one closer to 0 precedes the one further from 0.

Solution to Problem 9.2.

Step 1: Find a list \mathcal{N} of positive numbers, that contains the norms of all the roots of L . (For example, appeal to Lemma 2.2.)

Step 2: For each $N \in \mathcal{N}$, find the set \mathcal{D}_N of all norm N spherical extensions $\Delta \rightarrow \Delta'$ of normed Dynkin diagrams. (See Section 5.)

Step 3: For each $N \in \mathcal{N}$, and each extension $\Delta \rightarrow \Delta'$ in \mathcal{D}_N , define the following (if the stated conditions are met).

- (1) $u_{N,\Delta'}$ is the element of U having inner products with u_1, \dots, u_{n-1} as specified by the extension $\Delta \rightarrow \Delta'$.
- (2) The “residual norm” $R_{N,\Delta'}$ is $N - u_{N,\Delta'}^2$.
- (3) If some $v \in L_N$ satisfies $\pi_U(v) = u_{N,\Delta'}$, then the “residual coset” $C_{N,\Delta'}$ is $\pi_P(v) + (L_N \cap P)$. (This is independent of v .)
- (4) If $C_{N,\Delta'}$ is defined, and some vector of $S \cap C_{N,\Delta'}$ has norm $R_{N,\Delta'}$ and comes after $\mathbb{R}_{>0}r_0 \subseteq S$, then the “root residue” $r_{N,\Delta'}$ is the first such vector.
- (5) If $r_{N,\Delta}$ is defined, then the “candidate” $\alpha_{N,\Delta'}$ is $u_{N,\Delta'} + r_{N,\Delta'}$.

Step 4: If no candidates exist, then stop: the other endpoint k' of e is the ideal point at the end of the ray emanating from k along e . Our solution to the cuspidal batch 0 problem (problem 7.1) extends $\{u_1, \dots, u_{n-1}\}$ to a set of simple roots for $\Pi \cap k'^{\perp}$.

Step 5: Define α as the (unique) candidate $\alpha_{N,\Delta'}$ satisfying:

- (1) among all candidates, it minimizes $-k \cdot r_{N,\Delta'}/\sqrt{R_{N,\Delta'}}$,
- (2) among all candidates satisfying (1), it minimizes the priority $-k \cdot \alpha_{N,\Delta'}/\sqrt{N} = -k \cdot r_{N,\Delta'}/\sqrt{N}$.
- (3) among all candidates satisfying (1) and (2), it has smallest norm.

Then $\{u_1, \dots, u_{n-1}, \alpha\}$ is the set of simple roots for $\Pi \cap k'^{\perp}$ corresponding to the chamber C . In particular, the other endpoint k' corresponds to $\alpha^{\perp} \cap P \subseteq \mathbb{R}^{n,1}$.

The “residual” objects in step 3 refer to properties of the projection to P of a (hypothetical) root α with norm N and diagram extension Δ' . The point is that it is the part of α not already described by $u_{N,\Delta'}$.

Remark 9.3 (Simpler implementation, better performance). As written, we compute each candidate separately. This makes it easier to explain the proof, but gives up a great deal because the algorithms in section 8 are perfectly suited to searching for all candidates with fixed N simultaneously. Furthermore, most candidates have no chance of satisfying (1) in step 2, and their computation can be avoided.

Suppose $N \in \mathcal{N}$ is fixed, and write M for the largest residual norm among all $\Delta' \in \mathcal{D}_N$. Apply the algorithms in section 8 to the lattice $\pi_P(L_N)$. That is, if $P \cap L_N$ is isotropic then use **Promised** repeatedly; otherwise use **Anisotropic** once. This gives the list \mathcal{V} of all primitive vectors $r \in S \cap \pi_P(L_N)$ that have norm $\leq M$ and come after r_0 , in increasing order. For each such r , consider all of its positive integer multiples s that have norm $\leq M$. For each s , check whether there exists $\Delta' \in \mathcal{D}_N$ with $s^2 = R_{N,\Delta'}$ and $s \in C_{N,\Delta'}$. If so, then we have found $r_{N,\Delta'} = s$, hence the candidate $\alpha_{N,\Delta'}$. If any candidates arise from r in this way, then we halt the search for norm N candidates after finding them. The point is that if r' comes after r in \mathcal{V} , then r'^{\perp} is further from k than r^{\perp} is. Therefore any candidates arising from r' would be discarded in step 5(1).

Furthermore, once a candidate has been found, it can be used to abandon the searches for candidates of other norms. These searches should be abandoned when any additional candidates they might find would automatically be discarded in (1) or (2) of step 5.

Remark 9.4. Using **Anisotropic** in the previous remark returns a list of vectors $r_1, \dots, r_m \in S \cap L_N$ and a generator g for $\text{SO}^{\uparrow}(\pi_P(L_N))$, such that

$$\mathcal{V} = (r_1, \dots, r_m, g(r_1), \dots, g(r_m), g^2(r_1), \dots, g^2(r_m), \dots)$$

One cannot examine just r_1, \dots, r_m when checking whether they could be root residues $r_{N,\Delta'}$, because for example $g(r_1)$ may lie in a different

coset of $P \cap L_N$ in $\pi_P(L_N)$ than r_1 does. Instead, one must examine all terms of \mathcal{V} through $g^{m-1}(r_m)$, where g^m is a power of g that is the restriction of an isometry of L_N . Sufficient conditions for this are that g^m preserves $P \cap L_N$ and acts trivially on $\pi_P(L_N)/(P \cap L_N)$. Therefore m can be found from g . If one of $r_1, \dots, r_m, \dots, g^{m-1}(r_1), \dots, g^{m-1}(r_N)$ has a positive integer multiple s that lies in $C_{N,\Delta'}$ and has norm $R_{N,\Delta'}$, then $r_{N,\Delta'}$ is the first such multiple (under our total ordering on S). Otherwise $r_{N,\Delta'}$ does not exist.

Proof. First we claim that if a candidate $\alpha_{N,\Delta'}$ is defined, then it is an almost-root of L . This follows from Lemma 2.1, because $\alpha_{N,\Delta'}$ has norm $R_{N,\Delta'} + u_{N,\Delta'}^2 = N$ and lies in $C_{N,\Delta'} + u_{N,\Delta'} \subseteq L_N$.

Furthermore, $\alpha_{N,\Delta'}^\perp$ meets the hyperbolic line containing e , because $\alpha_{N,\Delta'}$ extends u_1, \dots, u_{n-1} to the spherical Dynkin diagram Δ' . Together with the fact that $r_{N,\Delta'}$ comes after $\mathbb{R}_{>0}r_0$ under the ordering on S , this shows that $\alpha_{N,\Delta'}^\perp$ meets the ray E emanating from k along e , and has negative inner product with k . We can now finish the proof in the case that k' is ideal. In this case, no mirror of $W(L)$ can meet E in hyperbolic space. Therefore no $\alpha_{N,\Delta'}$ can exist, and the algorithm recognizes in step 4 that k' is the ideal endpoint of E .

So suppose k' is timelike, and write β for the root which extends u_1, \dots, u_{n-1} to the system of simple roots for $\Pi \cap k'^\perp$ whose Weyl chamber contains C . We claim β is a candidate. Define $r = \pi_P(\beta)$ and $N = \beta^2$, and let Δ' be the normed Dynkin diagram of $u_1, \dots, u_{n-1}, \beta$. That r comes after $\mathbb{R}_{>0}r_0 \subseteq S$ follows from the fact that β^\perp meets E but not its initial endpoint k . Obviously $k \cdot r = k \cdot \beta$, which is negative by the definition of β . Because $\beta \in L_N$ and $\pi_U(\beta) = u_{N,\Delta'}$, the residual coset $C_{N,\Delta'}$ is defined and contains r . Because β can serve a ‘‘some vector’’ in step 3(4), $r_{N,\Delta'}$ and $\alpha_{N,\Delta'}$ are defined. Now, $r_{N,\Delta'}$ cannot precede r in the ordering on S , because then the mirror $\alpha_{N,\Delta'}^\perp$ would cut the hyperbolic segment $\overline{kk'} = e$. On the other hand, the definition of $r_{N,\Delta'}$ shows that it cannot come after r . Because r and $r_{N,\Delta'}$ have the same norm, and neither precedes the other, they are equal. Therefore β is the candidate $\alpha_{N,\Delta'}$.

Now suppose γ is any candidate, and $w = \pi_P(\gamma)$. The hyperbolic distance from k to $E \cap \gamma^\perp$ is at least as large as the hyperbolic distance from k to k' , or else k' would not be the other endpoint of e . Therefore

$$-k \cdot w / \sqrt{w^2} \geq -k \cdot r / \sqrt{r^2}$$

with equality just if $\gamma \perp k'$. So the candidates satisfying (1) in step 5 are exactly the candidates whose mirrors pass through k' . Obviously β is among them.

We saw that these candidates are almost-roots. So we may apply Lemma 3.1 to them, taking $W = W(\Pi \cap k^\perp)$ and using k as the control vector. Arguing as in the last step in the proof of our solution to Problem 6.2 shows that β is the unique candidate satisfying (1)–(3) of step 5. \square

We organize repeated use of the algorithm as follows. We maintain lists of corners found, simple roots found, and rays remaining to explore. Initially, these consist of the given corner k , the given simple roots for $\Pi \cap k^\perp$, and the rays emanating from k along edges. (These edges can be read from the Dynkin diagram of $\Pi \cap k^\perp$.) Now we repeat the following process: walk along the first unexplored ray E , using our solution to Problem 9.2, and then remove E from the list of rays to explore. If the corner k' at the other end of E is not already known, then do the following. First, add k' to the list of known corners. Second, append to the list of known roots all the simple roots orthogonal to k' that are not already known. Third, to the list of unexplored rays append all the rays emanating from k' , along edges.

Theorem 9.5. *The chamber C has finite volume if and only if this process terminates (meaning that after some finite number of steps the list of unexplored rays becomes empty). In this case, it finds all the corners and simple roots of C .*

Proof. If walking along an edge leads to a corner already known, then the list of unexplored rays shrinks, because one is removed and none are added. Therefore, if C has finite volume, hence finitely many corners, then the algorithm terminates.

Now suppose the process terminates. To prove that all simple roots and corners are found, we will use the projective model of hyperbolic space. Suppose $\alpha_1, \dots, \alpha_k$ are the simple roots found, and let $B = \{v \in L \otimes \mathbb{R} \mid v \cdot \alpha_i \forall i\}$ be the cone they define. We regard PB as a convex polytope in the affine subspace of $P(L \otimes \mathbb{R})$, got by discarding a suitable hyperplane of $P(L \otimes \mathbb{R})$. (For example, discard the orthogonal complement of a vector in the interior of B .) Note that PB contains C . We claim that all vertices of PB are corners of C . Certainly one vertex is, namely the initial corner (call it v_1).

Now suppose v_2, \dots, v_l are given, each v_i being a vertex of B adjacent to v_{i-1} . We claim that v_l is a corner of C , not just a vertex of B . By induction, v_{l-1} is a corner of C . (The base case v_1 is already known.) Whenever we find a new vertex, we also find all simple roots of C orthogonal to it. It follows that every edge of PB emanating from v_{l-1} contains an edge of C emanating from v_{l-1} . Because the algorithm has

terminated, we have already walked along the ray from v_{l-1} toward v_l , found the second corner c of C on it, and found all simple roots of C orthogonal to c . In particular, c is also a vertex of B , which forces $c = v_l$, finishing the proof of the inductive step.

Because PB is a finite-sided polytope, every vertex of PB lies at the end of such a sequence v_1, \dots, v_l . Therefore all vertices of PB lie in $\overline{H^n}$. So their convex hull PB does too, and it follows that B contains no spacelike vectors. No unfound simple root of C can exist, because it would be a spacelike element of B . So $C = PB$. Since C is the convex hull of finitely many points of $\overline{H^n}$, it has finite hyperbolic volume. \square

Remark 9.6 (Termination in the infinite-volume case). If C has infinite volume, then our algorithm can fail to find all vertices and simple roots of C , even if left to run forever. This happens whenever ∂C is disconnected; for examples of this see [13, III]. In this case, the algorithm finds exactly the (infinitely many) vertices in the component K of ∂C containing k , and the simple roots orthogonal to each of them.

When using Vinberg's algorithm, the usual strategy for detecting that C has infinite volume, effective in principle and often in practice, is to keep searching for simple roots until enough are found to detect some elements of $O^\uparrow(L)$ that preserve C . If enough are found to generate an infinite group, then C must have infinite volume. For example, whether two vertices are $O^\uparrow(L)$ -equivalent can be determined by trying to identify the simple roots at one with simple roots at the other.

This strategy works just as well, and probably better, in our setting, because finding vertices is a basic part of our approach. If C has infinite volume, then the component of ∂C containing k has infinitely many vertices, so the $O^\uparrow(L)$ -stabilizer of that component of ∂C is infinite. As one finds vertices, one can test them for equivalence with known vertices. Once enough vertices are found, infinitely many automorphisms of C will be visible, so C must have infinite volume.

REFERENCES

- [1] Allcock, Daniel, Presentation of affine Kac-Moody groups over rings. *Algebra Number Theory* **10** (2016), no. 3, 533–556.
- [2] Allcock, Daniel, The reflective Lorentzian lattices of rank 3. *Mem. Amer. Math. Soc.* **220** (2012), no. 1033, x+108 pp. ISBN: 978-0-8218-6911-6.
- [3] Borcherds, Richard, Automorphism groups of Lorentzian lattices. *J. Algebra* **111** (1987), no. 1, 133–153.
- [4] Borcherds, Richard E., Lattices like the Leech lattice. *J. Algebra* **130** (1990), no. 1, 219–234.
- [5] Cohen, Henri, *A course in computational algebraic number theory*. Graduate Texts in Mathematics, 138. Springer-Verlag, Berlin, 1993. xii+534 pp. ISBN: 3-540-55640-0.

- [6] Conway, J. H., The automorphism group of the 26-dimensional even unimodular Lorentzian lattice. *J. Algebra* **80** (1983), no. 1, 159–163.
- [7] Esselmann, Frank, Über die maximale Dimension von Lorentz-Gittern mit coendlicher Spiegelungsgruppe. *J. Number Theory* **61** (1996), no. 1, 103–144.
- [8] Kac, Victor G., *Infinite-dimensional Lie algebras (3rd edition)*. Cambridge University Press, Cambridge, 1990. xxii+400 pp. ISBN: 0-521-37215-1; 0-521-46693-8.
- [9] Lenstra, H. W., Jr., Solving the Pell equation. *Notices Amer. Math. Soc.* **49** (2002), no. 2, 182–192.
- [10] Moody, Robert V.; Pianzola, Arturo, *Lie algebras with triangular decompositions*. Canadian Mathematical Society Series of Monographs and Advanced Texts. A Wiley-Interscience Publication. John Wiley & Sons, New York, 1995. xxii+685 pp. ISBN: 0-471-63304-6.
- [11] Nikulin, V. V., On the arithmetic groups generated by reflections in Lobachevsky spaces. *Math. USSR-Izv.* **16** (1981), no. 3, 573–601.
- [12] Nikulin, V. V., Reflection groups in Lobachevsky spaces and an identity for the denominator of Lorentzian Kac-Moody algebras. *Izv. Math.* **60** (1996), no. 2, 305–334
- [13] Nikulin, V. V., On the classification of hyperbolic root systems of rank three. *Proc. Steklov Inst. Math.* **2000**, no. 3(230), 1–241
- [14] The PARI Group, PARI/GP version 2.13.1, Univ. Bordeaux, 2019, <http://pari.math.u-bordeaux.fr/>.
- [15] Scharlau, Rudolf; Walhorn, Claudia, Integral lattices and hyperbolic reflection groups. *Journées Arithmétiques*, 1991 (Geneva). *Astérisque* No. 209 (1992), 15–16, 279–291.
- [16] È. B. Vinberg, On groups of units of certain quadratic forms, *Math. USSR-Sbornik* **87** (1972) 17–35.

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF TEXAS, AUSTIN

Email address: allcock@math.utexas.edu

URL: <http://www.math.utexas.edu/~allcock>