Integer Matrices

Ben Marlin

Introduction

The goal of this document is to compile scattered results throughout the literature on the similarity of integer matrices. We will mostly follow David Husert's dissertation "Similarity of Integer Matrices," but will appeal to other sources as necessary. Although there will be some standard module or number theoretic results for which we do not include a proof, references will be provided that contain full proofs. We will not focus on fully reproducing the (mostly quite lucid) proofs of Husert, and rather we will focus on filling in the trickier details of such proofs and adding supplementary notes. Other papers we will discuss leave many more details out, so we will give full proofs of these results with the details filled in.

Let R be a commutative ring. Matrices $A, B \in \mathrm{Mat}_n(R)$ are said to be similar (over R) if there exists a matrix $C \in \mathrm{GL}_n(R)$ such that $A = CBC^{-1}$. We will focus on the case of $R = \mathbb{Z}$, where invertible matrices are precisely those with determinant equal to ± 1. Our goal will be to understand when two given integral matrices are similar over the integers, and understand how to algorithmically test for similarity. We will also want to understand the number of similarity classes of integral matrices of a given minimal and characteristic polynomial (so that we can know when we have found a set of representatives of the similarity classes). Although similarity is completely understood over a field, the problem is much harder over $\mathbb{Z}$ and involves nontrivial algebraic number theory.

Contents

1 Preliminaries

We cover the algebraic preliminaries needed for the later sections on integer matrices. We focus specifically on the results that we need and do not attempt to give a comprehensive review of algebra. Although the connection of some results to matrices may not be immediately clear, their usefulness will be apparent in later sections.

1.1 Orders and Lattices

Let R be a Noetherian integral domain with field of fractions K. For any K-space V, an R-lattice in V is a finitely generated R-submodule M of V such that $KM = V$. We may also consider R-lattices without specifying an ambient space. Generally, an R-lattice M is a finitely generated torsion-free R-module. Indeed, M may be viewed as a lattice in KM (its localization at $R \setminus \{0\}$). For such a lattice, we define its rank to be $\dim_K(KM)$.

If A is a finite dimensional algebra over K, then we define an R-order in A to be a subring Λ of A that is also an R-lattice in A. We will most often be considering $\mathbb{Z}$-orders, but occasionally this more general perspective will be useful. Many results about $\mathbb{Z}$-orders also apply to orders over Dedekind domains.

We will frequently be concerned with full modules over orders. To be explicit, let A be a finite dimensional K-algebra and Λ an R-order in A. Let L^* be an A-module (which we may also refer to as a representation of A). Then, a full Λ-module in L^* is a finitely generated Λ-module $\mathcal{U}$ such that $K\mathcal{U} = L^*$. We might also call a full Λ-module a Λ-lattice. In particular, following Husert's notation, we will care about the situation where our finite dimensional algebra is $A = \mathcal{K} = \bigoplus_{i=1}^{s} \mathcal{K}_i$, $L^* = \mathcal{K}^{\mathbf{n}}$, and $\Lambda = \mathcal{O}$ is an equation order in $\mathcal{K}$ corresponding to the characteristic polynomial of a given matrix.

There are several important observations about this case. First, since the classes of free and torsion-free modules coincide over a principal ideal domain R, any R-order $\mathcal{O}$ is R-free. Indeed, such an order is a finitely generated R-module, and as a submodule of a (necessarily torsion-free) vector space, it is torsion-free. Similarly, any full $\mathcal{O}$-module $\mathcal{U}$ in $\mathcal{K}^{\mathbf{n}}$ is $\mathbb{Z}$-free of rank $\dim_{\mathbb{Q}} \mathcal{K}^{\mathbf{n}}$. By extending maps of full $\mathcal{O}$-modules to $\mathcal{K}$-linear maps $\mathcal{K}^{\mathbf{n}} \to \mathcal{K}^{\mathbf{n}}$, we find $f \in \mathrm{Hom}_{\mathcal{O}}(\mathcal{U}, \mathcal{B})$ satisfies that for all $u \in \mathcal{U}$, $f(u) = \Gamma(u)$, for some $\Gamma \in \mathcal{M}(\mathbf{n}, \mathcal{K})$ such that $\Gamma\mathcal{U} = \mathcal{B}$. Essentially, all maps of full modules may be realized as block diagonal matrices (and vice versa). See Proposition (1.1) in Husert for details.

Second, when $\mathcal{K}$ is a number field (so not a proper direct sum of number fields), the full $\mathcal{O}$-modules in $\mathcal{K}$ are precisely the nonzero fractional ideals of $\mathcal{O}$ in $\mathcal{K}$ (for any order in a number field, its field of fractions is the number field). This will become important later because in certain cases we will be able to use the well-developed theory of class groups of rings of integers to understand full modules.

Third, the number of isomorphism classes of full $\mathcal{O}$-module classes in $\mathcal{K}^{\mathbf{n}}$ is finite, owing to the Jordan-Zassenhaus theorem, which is a vast generalization of the classical result on the finiteness of the class number of a number field. We will discuss this theorem in depth in a later section.

1.2 Algebraic Integers and Dedekind Domains

In this section, we will address the main properties of algebraic integers and Dedekind domains which we will use throughout the paper. Since there are so many properties to cover, we save space by not writing in the usual theorem-proof format.

1.2.1 Algebraic Integers

Suppose that $R \subseteq S$ is an extension of commutative rings. We call such an extension finite provided that S is a finitely generated R-module. An element $s \in S$ is called integral over R (or an algebraic integer) if it satisfies a monic polynomial over R. If every element of S is integral, then S is an integral extension. This terminology mirrors the familiar terminology associated with field extensions.

A fundamental result in algebraic number theory is that all finite extensions of commutative rings are integral. This is a consequence of the Cayley-Hamilton theorem. A general form of the Cayley-Hamilton theorem asserts that for each commutative ring R, any endomorphism of a finitely generated R-module satisfies a monic polynomial over R.

Then, for any $s \in S$, we consider the map φ given by left multiplication by s. It follows that $p(s) = p(s) \cdot 1 = p(\varphi) \cdot 1 = 0$ for some monic polynomial $p(x) \in R[x]$.

It is easily verified that $s \in S$ is integral if and only if the extension $R[s]$ is finitely generated. Since finite ring extensions are integral, this easily implies that the set of all elements of S integral in R forms a ring called the integral closure of R in S and denoted $\bar{R}$. We say that a ring is integrally closed if its integral closure in its field of fractions is itself.

1.2.2 Dedekind Domains

Rings of integral elements will form the motivating example for the general notion of a Dedekind domain, which generalizes principal ideal domains to situations where unique factorization may fail.

Let $\mathcal{K}$ be a number field and consider the ring of algebraic integers $\mathcal{O}_{\mathcal{K}}$ in $\mathcal{K}$. We will show that $\mathcal{O}_{\mathcal{K}}$ is a $\mathbb{Z}$-order in K which is Noetherian, integrally closed domain, and every nonzero prime ideal of $\mathcal{O}_{\mathcal{K}}$ is maximal. The fact that $\mathcal{O}_{\mathcal{K}}$

is a $\mathbb{Z}$-order follows immediately from the fact that $\mathcal{O}_\mathcal{K}$ is a finitely generated $\mathbb{Z}$-module and contains a basis of $\mathcal{K}$.

Rings of integers are Dedekind: Once we know $\mathcal{O}_\mathcal{K}$ is a finitely generated $\mathbb{Z}$-module, we immediately know it is Noetherian since finitely generated modules over Noetherian rings are Noetherian. Since $\mathbb{Z}$ is easily verified to be integrally closed by Euclid's lemma (in fact, the standard argument holds for an arbitrary GCD domain), and by induction we may prove integrality is transitive, it follows that $\mathcal{O}_\mathcal{K}$ is integrally closed.

It remains to show that nonzero prime ideals are maximal. Suppose that P is a nonzero prime ideal of $\mathcal{O}_\mathcal{K}$. Let $\alpha \in P$ be a nonzero element with minimal polynomial $m_\alpha(x) \in \mathbb{Z}[x]$ (recall that an algebraic number $z \in \bar{\mathbb{Q}}$ has an integer minimal polynomial if and only if z is an algebraic integer by Gauss' lemma). By definition, $m_\alpha(\alpha) = 0$, so $a_0 = -(\alpha^n + \cdots + a_1\alpha) \in P$. Since $m_\alpha(x)$ is irreducible, $a_0 \neq 0$, and a_0 annihilates $\mathcal{O}_\mathcal{K}/P$ as a $\mathbb{Z}$-module. Thus, $\mathcal{O}_\mathcal{K}/P$ is an integral domain which is finitely generated and torsion as a $\mathbb{Z}$-module. It follows that $\mathcal{O}_\mathcal{K}/P$ is a finite integral domain (i.e., a field), so P is forced to be maximal as desired.

Equivalent definitions: The three properties mentioned above (Noetherian integrally closed domain with nonzero primes being maximal) is the definition of a Dedekind domain. There are many important equivalent characterizations of a Dedekind domain.

It requires a fair amount of effort to establish the equivalence of these definitions. Although we will not prove the equivalence, let us describe a road-map for the proof. One can first prove that for any integral domain a nonzero fractional ideal is invertible if and only if it is projective. One then proves that a domain is Dedekind if and only if every fractional ideal is invertible. To build up to the desired result that every nonzero proper ideal factors uniquely as the product of primes, one can prove that a domain is Dedekind if and only if every nonzero proper ideal is the product of maximal ideals. With this preliminary result in hand, one then proves the powerful result that a domain is Dedekind if and only if each nonzero proper ideal factors uniquely as the product of prime ideals.

Because each nonzero fractional ideal is invertible, the collection of nonzero fractional ideals of a Dedekind domain R form an abelian group with ideal multiplication. The principal fractional ideals form a normal subgroup, and quotienting by this subgroup we obtain the so-called ideal class group of R. We will show later that if $\mathcal{K}$ is a number field, then the ideal class group of $\mathcal{O}_\mathcal{K}$ is finite. We call the cardinality of the ideal class group of $\mathcal{O}_\mathcal{K}$ the class number of $\mathcal{K}$.

Localization of Dedekind domains are DVRs: A quick corollary of this final result is that the localization of a Dedekind domain at a maximal ideal is a discrete valuation ring, i.e., a non-field local PID. Since the localization of a Dedekind

domain at any multiplicative subset is Dedekind, it suffices to show that a local Dedekind domain is a non-field PID. Let R be a local Dedekind domain with M its unique maximal ideal. For each $x \in M \setminus \{0\}$, let $\nu(x)$ be the integer such that $M^{\nu(x)} = xR$, which is well-defined by unique factorization. Choose $y \in M$ such that $\nu(y)$ is minimal. It follows that $xR = M^{\nu(x)} \subseteq M^{\nu(y)} = yR$. Thus, $x \in yR$, implying $M \subseteq yR$, and $M = yR$ by maximality. Because all ideals are powers of the principal ideal M, we see R is a PID.

Weak approximation: We will need one further result regarding localization and Dedekind domains. For any ideal I in a Dedekind domain R, we define its $\mathfrak{p}$-adic evaluation $\nu_{\mathfrak{p}}(I)$ to be the power of $\mathfrak{p}$ appearing in the prime decomposition of I. For any nonzero $a \in R$, we define its $\mathfrak{p}$-adic evaluation $\nu_{\mathfrak{p}}(a)$ to be $\nu_{\mathfrak{p}}(Ra)$.

Alternatively, since $R_{\mathfrak{p}}$ is a DVR, the unique maximal ideal $\mathfrak{p}R$ equals (π) for some $\pi \in R_{\mathfrak{p}}$, so any element of $R_{\mathfrak{p}}$ can be written up to units as a power of π. It is equivalent to define $\nu_{\mathfrak{p}}(a)$ to be the power of π appearing in the decomposition of a in $R_{\mathfrak{p}}$. Note that if $a \in \mathfrak{p}^n$ and $a \notin \mathfrak{p}^{n+1}$ then $\nu_{\mathfrak{p}}(a) = n$.

We can extend $\nu_{\mathfrak{p}} : R \to \mathbb{Z}$ to a map $\nu_{\mathfrak{p}} : K \to \mathbb{Z}$ by defining $\nu_{\mathfrak{p}}(x/y) = \nu_{\mathfrak{p}}(x) - \nu_{\mathfrak{p}}(y)$ for $x, y \in R$ and $y \neq 0$. With this understanding, the Weak Approximation theorem is a straightforward generalization of the Chinese Remainder theorem for commutative rings. This theorem will play an important role in understanding the structure of finitely generated torsion free modules over Dedekind domains in the next section.

Semilocal Dedekind domains: The Weak Approximation theorem gives an easy proof that semilocal Dedekind domains are principal. Suppose that I is an ideal of a Dedekind domain R. Then, we may write $I = \mathfrak{p}_1^{n_1} \cdots \mathfrak{p}_k^{n_k}$. Using the approximation result, choose $a \in R^{\times}$ such that $\nu_{\mathfrak{p}_i}(a) = n_i$ for all $1 \leq i \leq k$. Then, $I = (a)$ and R is a PID.

"To contain is to divide" in a Dedekind domain:
Let R be a Dedekind domain with ideals I and J. Then, $I \mid J$ if and only if $J \subseteq I$. Clearly, if $I \mid J$ then $J \subseteq I$. The other direction is also clearly true if our ring is a PID. Recall that the localization of a Dedekind domain at a prime ideal is a non-field PID. Factor $I = \prod \mathfrak{p}^{r_{\mathfrak{p}}}$ and $J = \prod \mathfrak{p}^{s_{\mathfrak{p}}}$. Consider $IR_{\mathfrak{p}} = \mathfrak{p}^{r_{\mathfrak{p}}}$ and $JR_{\mathfrak{p}} = \mathfrak{p}^{s_{\mathfrak{p}}}$. Because $R_{\mathfrak{p}}$ is a PID, it follows that $s_{\mathfrak{p}} \geq r_{\mathfrak{p}}$ as desired.

Ideals of Dedekind domains are generated by two elements: Again, we will use weak approximation. We follow Proposition 4.7.7 in Cohen's "A Course in Computational Algebraic Number Theory." Let $I \subseteq R$ be an ideal in Dedekind domain R. We prove the stronger statement that for any $\alpha \in I$ there exists some $\beta \in I$ such that $I = (\alpha, \beta)$.

By unique factorization, write $(\alpha) = \prod_{i=1}^{r} \mathfrak{p}_i^{a_i}$. Because "to contain is to divide," write can write $I = \prod_{i=1}^{r} \mathfrak{p}_i^{e_i}$ for $e_i \leq a_i$. Now, by weak approximation, choose

some $\beta \in R$ such that $\nu_{\mathfrak{p}_i}(\beta) = e_i$ for each i. It follows that $\beta \in I$. Then, if we set $I' = (\alpha, \beta) = \alpha R + \beta R$, we see that $\nu_{\mathfrak{p}_i}(I') = \min(\nu_{\mathfrak{p}_i}(\alpha), \nu_{\mathfrak{p}_i}(\beta)) = e_i$ and for any prime ideal $\mathfrak{q}$ not containing α, $\nu_{\mathfrak{q}}(I') = 0$. Thus, $I = I'$ and we are done.

UFD if and only if PID in a Dedekind domain: It is well know that every PID is a UFD, so we only need to prove the forward implication. Because we have a prime factorization of every ideal, it will suffice to prove that each prime is principal. Let $\mathfrak{p} \subseteq R$ be a prime ideal. Let $0 \neq \alpha \in \mathfrak{p}$ and factor $\alpha = p_1 \cdots p_k$ into irreducibles. It follows that $(p_1) \cdots (p_n) \subseteq \mathfrak{p}$, hence $(p_i) \subseteq \mathfrak{p}$ for some i. Because R is a UFD, irreducibles are prime, so (p_i) is a prime ideal, and by the Dedekind condition it is maximal. It follows that $\mathfrak{p} = (p_i)$.

There are no non-maximal Dedekind orders in a number field:
Given a number field $\mathcal{K}$, we will refer to $\mathcal{O}_\mathcal{K}$ as the maximal order in $\mathcal{K}$. Indeed, because finite extensions are integral, $\mathcal{O}_\mathcal{K}$ contains any order. We already showed that $\mathcal{O}_\mathcal{K}$ is Dedekind, and we will now show that if $\mathcal{O} \subsetneq \mathcal{O}_\mathcal{K}$ is an order, then it contains a non-invertible ideal.

Consider the quotient $\mathcal{O}_\mathcal{K}/\mathcal{O}$, which has cardinality $m \in \mathbb{Z}_{>0}$. It is easy to check that $m\mathcal{O}$ is an ideal of $\mathcal{O}$ and $\mathcal{O}_\mathcal{K}$. We know that $m\mathcal{O}$ is invertible in $\mathcal{O}_\mathcal{K}$. Suppose it is also invertible in $\mathcal{O}$. Then, $\{x \in \mathcal{K} \mid xm\mathcal{O} \subseteq m\mathcal{O}\} = \mathcal{O}$. But, we then also have $\{x \in \mathcal{K} \mid xm\mathcal{O} \subseteq m\mathcal{O}\} = \mathcal{O}_\mathcal{K}$, which is a contradiction.

1.2.3 Submodules of Free Modules over Left Hereditary Rings

We now give a full proof of a result that will be crucial to our investigation of full modules over maximal orders. The following argument may be found in the Section 1.3 of Husert's thesis, or alternatively as Theorem 2.44 in Irving Reiner's "Maximal Orders."

We say that a ring R is left hereditary if every left ideal is a projective R-module.

Theorem 1 *If R is a left hereditary ring, then every submodule of a free left R-module M of finite rank is isomorphic to an external direct sum of ideals of R, and is therefore projective.*

Proof. We proceed by induction on the rank k of the free module. The result is trivial for $k = 1$. Up to isomorphism, we know that M looks like R^k for some integer k. Define $\pi_1 : R^k \to R$ as projection onto the first component of R^k. Let $\mathcal{U}$ be a submodule of R^k and observe that $\pi_1(\mathcal{U})$ is a left ideal of R. We obtain the following short exact sequence:

$$0 \longrightarrow \ker(\pi_1|_\mathcal{U}) \longrightarrow \mathcal{U} \longrightarrow \pi_1(\mathcal{U}) \longrightarrow 0$$

Since R is left hereditary, $\pi_1(\mathcal{U})$ is projective, hence our short exact sequence splits: $\mathcal{U} \cong \ker(\pi_1|_\mathcal{U}) \oplus \pi_1(\mathcal{U})$. But, $\ker(\pi_1|_\mathcal{U})$ clearly embeds in R^{k-1}, so by

induction we may write it as the direct sum of ideals. It follows that $\mathcal{U}$ is the direct sum of ideals and the proof is complete.

■

Of course, Dedekind domains are left (and right) hereditary, so the previous theorem applies to submodules of free modules over a Dedekind domain. It is natural to consider which modules may be embedded as submodules of free modules. If R is a Dedekind domain and M is an R-lattice, then M may be viewed as a lattice in KM. We may then pick some R-free lattice N in KM and observe that we may choose $\alpha \in R^\times$ such that $\alpha M \subseteq N$. Thus, any R-lattice may be viewed up to isomorphism as living in a free R-module, and therefore is the direct sum of ideals (and projective). Since projective modules are direct summands of free modules, they are always torsion free. Therefore, a finitely generated module over a Dedekind domain is torsion-free if and only if it is projective. This is a generalization of the situation of PIDs, where the classes of finitely generated torsion-free, projective, and free modules all coincide.

When R is a Dedekind domain, we can produce a very simple decomposition of R-lattices. This is because we can "force" ideals of R to be coprime.

Lemma 1 *Let A and B be nonzero ideals of R. Then, there exists $\gamma \in K^\times$ such that $A + \gamma B = R$.*

Proof. Because R is Dedekind, we may write both ideals as the product of prime ideals: $A = \mathfrak{p}_1^{a_1} \cdots \mathfrak{p}_{k_a}^{a_k}$ and $B = \mathfrak{q}_1^{b_1} \cdots \mathfrak{q}_{k_b}^{b_k}$. By weak approximation, choose some $\gamma \in K^\times$ such that $\nu_{\mathfrak{q}_i}(\gamma) = -a_{p_i}$ for all i. Since $\gamma A = (\gamma)A$, we find that γA has nonnegative valuation at each prime ideal of R, and it follows that $\gamma A \subseteq R$. Furthermore, each prime ideal appearing the the decomposition of B necessarily has valuation zero with respect to γA, so γA is not contained in any prime containing B, and $\gamma A + B = R$.

■

Theorem 2 *If R is a Dedekind domain, then any R-lattice M of rank n is isomorphic to $A \oplus R^{n-1}$ for some ideal A in R.*

Proof. We will prove that if $\{A_i\}_{i=1}^k$ is a collection of ideals of R, then $\bigoplus_{i=1}^k A_i \cong (A_1 \cdots A_k) \oplus R^{k-1}$. By Theorem 1, $M \cong \bigoplus_{i=1}^n A_i$ for some ideals $\{A_i\}$ in R, so this will complete the proof. First, suppose that $k = 2$. By the above lemma, we may assume (up to isomorphism) that A_1 is coprime to A_2. Thus, $A_1 A_2 = A_1 \cap A_2$. We obtain the following short exact sequence:

$$0 \longrightarrow A_1 A_2 \longrightarrow A_1 \oplus A_2 \longrightarrow R \longrightarrow 0$$

by sending $a \mapsto (a, -a)$ and $(a_1, a_2) \mapsto a_1 + a_2$.

As R is projective, the sequence splits, and the result follows for $k = 2$. For the general case, we may view a direct sum of n ideals as the direct sum of the first $n - 1$ ideals with the last ideal and then apply the above case.

1.3 Ring Theory

In this section, we will recall some of the basic properties of rings, with an emphasis on the results relevant to the proof of the Jordan-Zassenhaus theorem and the structure of maximal orders in matrix algebras. The proofs in this section are short, so we will provide them in full.

1.3.1 Semisimple and Artinian Modules

We begin with the theory of semisimple rings. All modules in this section are left modules (this was implicitly assumed in previous sections, but it is best to be explicit in this section since we will encounter noncommutative rings). A nonzero R-module M is simple if it has no nontrivial submodules. If R is an algebra over a field, we often refer to R-modules as representations and simple R-modules as irreducible representations. We say that an R-module M is completely reducible (or semisimple) if every submodule N has a complementary submodule N' such that $M = N \oplus N'$. A ring is semisimple if its left regular module is semisimple. A ring is simple if it has no nontrivial two-sided ideals. Standard arguments such as those in the first chapter of Isaacs' "Character Theory of Finite Groups," establish that M is semisimple iff it is the sum of some of its simple submodules iff it is the direct sum of some of its simple submodules.

From this characterization, it is immediate that submodules and quotients of semisimple modules are semisimple. Since every module is a quotient of a free module, it follows that a ring R is semisimple if and only if every R-module is semisimple.

Next, recall that a module is called Artinian if every descending chain of submodules stabilizes. A ring is (left/right) Artinian if its (left/right) regular module is Artinian. It is not hard to show that a module is Artinian if and only if every nonempty collection of submodules contains a minimal element. An important class of Artinian rings are the finite dimensional algebras over fields.

We would like to relate Artinian rings to semisimple rings. The link between the two is the Jacobson radical. For any R-module M, we define its Jacobson radical $\mathrm{Rad}(M)$ to be the intersection of its maximal submodules. By convention, the intersection of the empty family of submodules of M is simply M. The Jacobson radical of a ring R is the intersection of the maximal left ideals of R.

Proposition 1 *Let R be a ring. Then, $\mathrm{Rad}(R)$ is precisely the elements of R that annihilate all simple R-modules.*

Proof. Suppose that m is a maximal ideal of R. Then, R/m is a simple R-module. If $x \in \mathrm{Rad}(R)$, then $x(R/m) = 0$, implying $x \in m$. Conversely, suppose x belongs to every maximal ideal of R. Let W be a simple R-module with $w \in W$, $w \neq 0$. Then, $Rw = W$ and we claim than $Ann(w)$ is a maximal

left ideal of R. Suppose that $Ann(w) \subsetneq J \subsetneq R$ for some ideal J in R. Then, there exists some $j \in J$ such that $jw \neq 0$. Thus, $Rjw = W$ and $w = rjw$ for some $r \in R$. But this implies that $1 - rj \in Ann(w) \subseteq J$, hence $1 \in J$. Therefore, $Ann(w)$ is a maximal left ideal and $x \in Ann(w)$, so x annihilates W as desired.

■

Corollary 1 *For any ring R, the ring $R/\operatorname{Rad}(R)$ is semisimple.*

Proof. Let M be a simple R-module. Then, $R \to \operatorname{End}(M)$ descends to a map $R/\operatorname{Rad}(R) \to \operatorname{End}(M)$, so M is also a simple $R/\operatorname{Rad}(R)$-module. If $[x] \in R/\operatorname{Rad}(R)$, then $[x]$ annihilates M for any simple R-module M, and it follows that $x \in \operatorname{Rad}(R) = 0$.

■

This characterization also makes clear that $\operatorname{Rad}(R)$ is a two-sided ideal.

Proposition 2 *Artinian rings are semilocal.*

Proof. Let R be a (left/right) Artinian ring. Let S be the set of all finite intersections of maximal (left/right) ideals of R, which has a minimal element $I = \bigcap_{i=1}^{k} M_i$. Let m be a maximal (left/right) ideal of R. Then, $I \cap m = I$ by minimality. Thus, $M_1 \cdots M_k \subseteq I \subseteq m$, so $M_j = m$ for some j, and it follows that $M_1, \ldots, M_k$ are the only maximal (left/right) ideals.

■

Proposition 3 *Let M be an Artinian R-module. Then, $M/\operatorname{Rad}(M)$ is semisimple. Further, M is semisimple if and only if $\operatorname{Rad}(M) = 0$.*

Proof. By definition, $\operatorname{Rad}(M) = \bigcap_{m \in I} m$ where I is the set of all maximal submodules of M. Since M is Artinian, the set of finite intersections of elements of I has a minimal element $J = \bigcap_{i=1}^{n} m_i$, which necessarily coincides with $\operatorname{Rad}(M)$. We obtain the following exact sequence:

$$0 \longrightarrow \operatorname{Rad}(M) \longrightarrow M \longrightarrow \bigoplus_{i=1}^{n} M/m_i$$

Each M/m_i is simple, so $M/\operatorname{Rad}(M)$ embeds in a semisimple module, and must itself be semisimple.

Obviously, if $\operatorname{Rad}(M) = 0$ then M is semisimple. Conversely, if $M = \bigoplus V_i$ for some simple modules $\{V_i\}$, observe that $\ker(\pi_i)$ is maximal for each i. Therefore, $\operatorname{Rad}(M) \subseteq \bigcap \ker(\pi_i) = 0$.

■

1.3.2 Matrix Rings

We will now describe the structure of arbitrary semisimple rings.

Theorem 3 *(Artin-Wedderburn) Let R be a semisimple ring. Then, R is isomorphic to the direct sum of finitely many matrix rings over division rings.*

Proof. Consider the left regular representation of R. We may write $R \cong \bigoplus_{i=1}^{k} V_i^{n_i}$ for simple pairwise nonisomorphic R-modules $\{V_i\}$. It follows that $\mathrm{End}_R(R) \cong \mathrm{End}_R(\bigoplus_{i=1}^{k} V_i^{n_i})$.

By Schur's lemma, $\mathrm{End}_R(\bigoplus_{i=1}^{k} V_i^{n_i}) \cong \bigoplus_{i=1}^{k} \mathrm{End}_R(V_i^{n_i}) \cong \bigoplus_{i=1}^{k} \mathrm{Mat}_{n_i}(\mathrm{End}_R(V_i))$.

But, $\mathrm{End}_R(R) \cong R^{op}$, so $R \cong (\bigoplus_{i=1}^{k} \mathrm{Mat}_{n_i}(\mathrm{End}_R(V_i)))^{op} \cong \bigoplus_{i=1}^{k} \mathrm{Mat}_{n_i}(\mathrm{End}_R(V_i))$ since the transpose map gives an isomorphism between matrix rings and their opposite ring.

Finally, by Schur's lemma again, $\mathrm{End}_R(V_i)$ is a division ring D_i. So, we may write $R \cong \bigoplus_{i=1}^{k} \mathrm{Mat}_{n_i}(D_i)$ as desired.

■

The following proposition leads to a converse to the Artin-Wedderburn theorem. We follow the proof in Grillet's "Abstract Algebra."

Proposition 4 *Let R be a ring. Then, for each two-sided ideal I of R, $\mathrm{Mat}_n(I)$ is a two-sided ideal of $Mat_n(R)$. All two sided ideals of $\mathrm{Mat}_n(R)$ are of this form.*

Proof. Clearly, $\mathrm{Mat}_n(I)$ is a two-sided ideal of $\mathrm{Mat}_n(R)$ for each ideal I of R. Now, suppose that J is an ideal of $\mathrm{Mat}_n(R)$. Define I as the set of all $(1,1)$-entries of elements of J, which is an ideal of R. We will show that $J = \mathrm{Mat}_n(I)$. With E_{ij} as the matrix having entry (i,j) as 1 and zeros elsewhere, observe that $E_{ij}AE_{kl} = A_{jk}E_{il}$ for any $A \in J$. Taking $i = l = 1$ and varying j and k, it follows that $J \subseteq Mat_n(I)$.

Next, if $B \in \mathrm{Mat}_n(I)$, for each i and j, we have $B_{ij} = C_{11}^{ij}$ for some $C^{ij} \in J$.

Then, $C_{11}^{ij}E_{ij} = E_{i1}C^{ij}E_{1j} \in J$. As $B = \sum_{i,j} C_{11}^{ij}$, we see $B \in J$, completing the proof.

■

An immediate consequence of the proposition is that matrix rings over simple rings are simple. Therefore, for any division ring D, the matrix ring $\mathrm{Mat}_n(D)$ is simple, and its Jacobson radical must vanish. Furthermore, ideals of $\mathrm{Mat}_n(D)$ are in bijection with submodules of D^n (from a submodule of D^n we construct a matrix whose rows are elements of D^n). Since finitely generated modules over

Artinian rings are Artinian, D^n and therefore $\mathrm{Mat}_n(D)$ is Artinian. By Proposition 3, it follows that $\mathrm{Mat}_n(D)$ is semisimple. As the product of semisimple rings is semisimple, we obtain the fundamental result that a ring is semisimple if and only if it is isomorphic to the product of finitely many matrix rings over division rings. By applying Artin-Wedderburn, it is also clear that if R is semisimple, so too is $\mathrm{Mat}_n(R)$.

The following result will be used in the proof of the Jordan-Zassenhaus theorem. If D is a division ring, then D^n is the unique irreducible representation of $\mathrm{Mat}_n(D)$ up to isomorphism. We first need to show that D^n is actually an irreducible representation. Any matrix in $\mathrm{Mat}_n(D)$ may be regarded as a map $D^n \to D^n$. Suppose that $V \subseteq D^n$ is a nontrivial subrepresentation. Choose some $w \in D^n \setminus V$. Choose a basis of V and extend it to a basis of D^n (this is possible because all modules over division rings are free). Send one of the of the basis vectors of V to w while sending the others to arbitrary elements. Define a map $\phi : D^n \to D^n$ by extending by linearity. Therefore, we obtain a matrix for which V is not an invariant submodule. Conversely, any irreducible representation F is isomorphic to some minimal left ideal I of $\mathrm{Mat}_n(D)$. We may assume without loss of generality that there exists some $A \in I$ such that the first column of A is nonzero. Then, sending any element of I to its first column defines a nonzero map of representations $I \to D^n$, which must be an isomorphism by Schur's lemma.

1.3.3 Local Properties

Proposition 5 *Let R be an integral domain with field of fractions K. Let M be a torsion-free R-module (so M embeds in the K-space KM). Then, $M = \bigcap_{\mathfrak{m}} M_{\mathfrak{m}}$, where intersection is viewed within KM, and ranges over all maximal ideals of R.*

Proof. Of course, $M \subseteq \bigcap_{\mathfrak{m}} M_{\mathfrak{m}}$. Next, suppose $x \in \bigcap_{\mathfrak{m}} M_{\mathfrak{m}}$. Define the ideal $I = \{r \in R \mid rx \in M\}$. Then, for each maximal ideal $\mathfrak{m}$, $x = m_{\mathfrak{m}}/s_{\mathfrak{m}}$ for some $m_{\mathfrak{m}} \in M$ and $s_{\mathfrak{m}} \in A \setminus \mathfrak{m}$. It follows that $s_{\mathfrak{m}} x \in M$, implying that $s_{\mathfrak{m}} \in I$. It follows that $I \subsetneq \mathfrak{m}$ for each maximal $\mathfrak{m}$. However, by Krull's theorem, this implies $I = R$, and in particular $1 \in I$, so $x \in M$ as desired.

■

We now want to prove that if R is a local ring and $\mathcal{O}$ is an R-order in an algebra A, then $\mathcal{O}$ is semilocal. In particular, a Dedekind order over a local ring is a PID. This rather obscure result will prove useful in the study of maximal matrix orders.

Lemma 2 *Let $R \subseteq F$ be an integral extension of an integral domain R by a field F. Then, R is also a field.*

Proof. By integrality, for any $x \in R$, we have $p(x^{-1}) = 0$ for some monic $p(x) \in R[x]$. Thus, $1/x^n + a_{n-1}/x^{n-1} + \cdots + a_0 = 0$. Scaling by x^{n-1} and subtracting terms, we see that $x^{-1} \in R$ as desired.

■

Lemma 3 *Let R be a local ring with residue field $R/\mathfrak{m} = K$. Let A be a finitely generated R-algebra. Then, every maximal ideal of A contains $\mathfrak{m}A$.*

Proof. Let M be a maximal ideal of A. We may view R as a subring of A and consider $R \cap M$. Since M is maximal, it is also prime, and it follows that $R \cap M$ is a prime ideal of R. Hence, $R/(R \cap M)$ is an integral domain. Further, $R/(R \cap M)$ embeds as a subring of the field A/M. By the previous lemma, noting that finite ring extensions are integral, it follows that $R/(R \cap M)$ is also a field, so $R \cap M$ is maximal and coincides with $\mathfrak{m}$. So, $\mathfrak{m} \subseteq M$ and therefore $\mathfrak{m}A \subseteq M$.

■

The following proof closely follows Lemma 2.4 of Morandi's "Maximal Orders Over Valuation Rings."

Proposition 6 *Let R be a local ring and $\mathcal{O}$ an R-order in an algebra A. Then, $\mathcal{O}$ is semilocal.*

Proof. Suppose that $M_1, \ldots, M_n$ are maximal ideals of $\mathcal{O}$. Then, by definition of the radical, $\operatorname{Rad}(O) \subseteq \bigcap_{i=1}^n M_i$. We obtain an obvious surjection $\mathcal{O}/\operatorname{Rad}(O) \twoheadrightarrow \mathcal{O}/\bigcap_{i=1}^n M_i$. Next, because R is local $\operatorname{Rad}(R) = \mathfrak{m}$ for the unique maximal ideal $\mathfrak{m} \subseteq R$. Therefore, $R/\operatorname{Rad}(R)$ is a field.

Define $\operatorname{Rad}_R(\mathcal{O}) = \operatorname{Rad}(R)\mathcal{O}$. In the standard way, $\mathcal{O}/\operatorname{Rad}_R(\mathcal{O})$ becomes an $R/\operatorname{Rad}(R)$-module (vector space). We would like to view $\mathcal{O}/\operatorname{Rad}(O)$ and $\mathcal{O}/\bigcap_{i=1}^n M_i$ as $R/\operatorname{Rad}(R)$-spaces as well. However, we need to be careful and check that the action is actually well-defined.

It suffices to check that $\operatorname{Rad}_R(\mathcal{O}) \subseteq \operatorname{Rad}(\mathcal{O})$. This follows immediately from Lemma 3, taking $\mathcal{O}$ as our finitely generated algebra over local ring R. Therefore, $\mathcal{O}/\operatorname{Rad}(O)$ and $\mathcal{O}/\bigcap_{i=1}^n M_i$ become $R/\operatorname{Rad}(R)$-spaces.

By the Chinese Remainder theorem, $\mathcal{O}/\bigcap_{i=1}^n M_i \cong \bigoplus_{i=1}^n \mathcal{O}/M_i$. Next, since $\bigcap_{i=1}^n M_i \subseteq M_j$ for any j, each $\mathcal{O}/M_j$ is an $R/\operatorname{Rad}(R)$-space with dimension greater than or equal to 1.

Observe that $n \leq [\bigoplus_{i=1}^n \mathcal{O}/M_i : R/\operatorname{Rad}(R)] \leq [\mathcal{O}/\operatorname{Rad}(\mathcal{O}) : R/\operatorname{Rad}(R)]$.

However, $[\mathcal{O}/\operatorname{Rad}_R(\mathcal{O}) : R/\operatorname{Rad}(R)] < \infty$, because $\mathcal{O}$ is a finitely generated R-module, and this serves as an upper bound on the rightmost term on the previous line. Therefore, we cannot have infinitely many maximal ideals in $\mathcal{O}$.

■

1.4 Matrices and Modules over a PID

In this section, we will review the theory of finitely generated modules over a PID. We will use the existence of a Smith Normal Form (SNF) of a matrix over a PID to give a proof of the existence of an invariant factor decomposition of a module, and give a slick proof of the uniqueness of this form. We mention some other facts about matrices that are generally useful. All of this material is classical and can be found in any good algebra textbook, so we will be brief in the proofs. However, we still include this section to have a self-contained reference.

1.4.1 Existence Results

We start by examining the submodules of finite rank free modules over a PID. We saw earlier that submodules of finite rank free modules over a left hereditary ring (such as a Dedekind domain) are projective. The following theorem shows how we can strengthen this result when we have a PID. We follow the proof given on Matt Baker's blog (included for completeness).

Theorem 4 *Let R be a PID and let M be a free module of rank m. Then, any submodule of M is free of rank less than or equal to m.*

Proof. We proceed by induction on the rank of M. Clearly, up to isomorphism, $M \cong R^m$, and this isomorphism sends free modules to free modules, so we will assume without loss of generality that $M = R^m$.

When $m = 1$, the result holds trivially because R is a PID. Suppose that the theorem holds for free modules of rank less than m. Define $\pi : R^m \to R$ as projection onto the last coordinate of R^m. Observe that $\ker \pi$ is free of rank $m-1$. Given any submodule $N \subseteq M$, define $N' = N \cap \ker \pi$. By induction, N' is free of rank less than or equal to $m-1$. Further, $\pi(N)$ is an ideal of R, so $\pi(N) = (x)$. Choose $y \in \pi^{-1}(x) \cap N$ and define $N'' = (y)$. If $\pi(N) = 0$, then $N \subseteq \ker \pi$, and $N = N'$, in which case, we are done.

If $x \neq 0$, then $N = N' \oplus N''$. Indeed, if $w \in N' \cap N''$, then $\pi(w) = 0$ and $w = ry$ for some $r \in R$. So, $r\pi(y) = rx = 0$, forcing $r = 0$ and $w = 0$. Next, let $w \in N$ and write $\pi(w) = rx$ for some $r \in R$. Then, $\pi(w - ry) = 0$, hence $w - ry \in N'$, and $w \in N' + N''$ as desired.

■

Using the SNF of a matrix over a PID (see Matt Baker's blog post linked above if SNF is unfamiliar), we can extract "aligned bases" for M and N.

Theorem 5 *Let R be a PID and let M be a free module of rank m. Let N be a rank n submodule of M. Then, there exists a basis $x_1, \ldots, x_m$ of M such that there exists $a_1, \ldots, a_n \in R$ for which $a_1 x_1, \ldots, a_n x_n$ is a basis for N and $a_1 \mid a_2 \mid \ldots \mid a_n$.*

Proof. Recall that if $y_1, \dots y_l$ is a basis of an R module, and $Y \in \mathrm{Mat}_n(R)$, then $[y_1, \dots y_l]Y$ is a basis if and only if $Y \in \mathrm{GL}_n(R)$.

Now, let $w_1, \dots, w_n$ be a basis for N and let $z_1, \dots, z_m$ be a basis for M. We have some matrix $A \in \mathrm{Mat}_{m\times n}(R)$ such that $[w_1, \dots, w_n] = [z_1, \dots z_m]A$. We can put A in Smith Form: $A = XDY$, where $D \in \mathrm{Mat}_{m\times n}(R)$ is in SNF, and X and Y are invertible of appropriate sizes. Then, $[w_1, \dots, w_n]Y^{-1} = [z_1, \dots z_m]XD$, which tells us that $[z_1, \dots z_m]XD$ is a basis of N. It is easy to see that $[z_1, \dots z_m]X$ is our desired basis $x_1, \dots, x_m$ of M and that $[z_1, \dots z_m]XD$ is our "aligned" basis of N, where the $a_1, \dots, a_n$ arise as the diagonal terms of D. These coefficients satisfy the divisibility condition because of the divisibility condition in SNF.

■

It is now easy to show the invariant factor decomposition of any finitely generated module M over a PID R.

Theorem 6 *Let R be a PID and let M be a finitely generated R-module. Then, $M \cong R^r \oplus R/(a_1) \oplus \cdots \oplus R/(a_n)$ for some integer $r \geq 0$ and nonzero, non-unit elements $a_1, \dots a_n \in R$ satisfying the divisibility condition $a_1 \mid a_2 \mid \dots \mid a_n$.*

Proof. We have already done all of the hard work for this theorem. Since M is finitely generated, choose a minimal set of generators and define $\varphi : R^n \twoheadrightarrow M$ by sending the standard basis of R^n to the chosen generators. Note that $\ker\varphi$ is a submodule of the free module R^n, so by the previous theorem we can find an aligned bases. Then, write R^n and $\ker\varphi$ in terms of the aligned bases to see that $R^r \oplus R/(a_1) \oplus \cdots \oplus R/(a_n) \cong R^n/\ker\varphi \cong M$.

■

It follows immediately that a finitely generated module over a PID is free if and only if it is torsion-free, and may be be written as the direct sum of a free and torsion-free module.

1.4.2 Uniqueness Results

Crucially, this "invaraint factor form" is unique up to units. This follows from a more general theorem, which we learned from this answer by the user Anonymous.

Theorem 7 *Let R be a commutative ring and $I_1 \subseteq I_2 \subseteq \cdots \subseteq I_n \subsetneq R$ be an increasing sequence of ideals. Let M be an R-module for which there exists an isomorphism $M \cong R/I_1 \times \cdots \times R/I_n$. Then, (1) the minimal number of generators of M is n, and (2), for $1 \leq k \leq n$, I_k is the set of all $x \in R$ such that xM is generated by fewer than k elements.*

Proof. Clearly, we have a canonical surjection $R^n \twoheadrightarrow R/I_1 \times \cdots \times R/I_n$, so M is generated by n elements. Suppose there exists a generating set of cardinality $r < n$. Let $\mathfrak{m}$ be a maximal ideal containing I_n, which exists by Krull's theorem (here, it is essential that I_n is strictly smaller than R). Because $\mathfrak{m}$ contains I_n, we obtain a well-defined R-linear map $\psi : R/I_1 \times \cdots \times R/I_n \twoheadrightarrow (R/\mathfrak{m})^n$ by sending $([r_1]_{I_1}, \ldots, [r_n]_{I_n}) \mapsto ([r_1]_{\mathfrak{m}}, \ldots, [r_n]_{\mathfrak{m}})$. Composing ψ with the isomorphism $\epsilon : M \to R/I_1 \times \cdots \times R/I_n$, we obtain an R-linear surjection onto an n-dimensional vector space. This implies $(R/\mathfrak{m})^n$ is spanned by r vectors over $R/\mathfrak{m}$, which is a contradiction.

We now prove the second statement. Define $m_x : R \to R$ as left-multiplication by $x \in R$. Define $I_k^x = m_x^{-1}(I_k)$ for $1 \leq k \leq n$.

Now, define $\epsilon' : xM \to R/I_1^x \times \cdots \times R/I_n^x$ by sending $x\alpha \mapsto ([\epsilon_1(\alpha)]_{I_1^x}, \ldots, [\epsilon_n(\alpha)]_{I_n^x})$, where ϵ_i denotes the ith component of ϵ. This map is well-defined because of the definition of I_k^x. Because the inverse image of an ideal under a ring map is an ideal, $xM \cong R/I_1^x \times \cdots \times R/I_n^x$ meets the hypotheses of part (1) of the theorem (once we remove any trivial quotients). Indeed, it is possible that $x \in I_j$ for some j and then $I_k^x = R$ for all $k > j$. Applying (1), we see that xM can be generated by fewer than k elements if and only if $R/I_k^x = 0$ if and only if $x \in I_k$.

■

Observe that this theorem implies the uniqueness of the invariant factor form of a finitely generated module over a PID. Indeed, if $R^r \oplus R/(a_1) \oplus \cdots \oplus R/(a_n) \cong R^k \oplus R/(b_1) \oplus \cdots \oplus R/(b_n)$, then $R^r \cong R^k$, which implies $r = k$ by the well-definedness of the rank of a free module over a commutative ring. We can quotient out the torsion-free parts, and the uniqueness theorem reduces to the case of finitely generated torsion modules, which is addressed by the present theorem.

Furthermore, if V is a finite-dimensional vector space over a field k and $T \in \mathrm{End}(V)$, then we obtain a $k[x]$-module structure on V by defining $p(x)v = p(T)v$. This turns V into a finitely generated torsion module over the PID $k[x]$. Then, choosing a suitable basis on the invariant factor decomposition of V, representing the k-linear multiplication by x map with respect to this basis, and using the $k[x]$-module isomorphism $V \cong k[x]/(p_1(x)) \oplus \cdots \oplus k[x]/(p_n(x))$ guaranteed by the invaraint factor form, we obtain the Rational Canonical Form (RCF) of an endomorphism. The uniqueness of the invariant factor form guarantees RCF is unique up to units. If k is algebraically closed, a slight modification of the above idea yields the Jordan Canonical Form (JCF). For details on these arguments, see Dummit and Foote's "Abstract Algebra" sections 12.2 and 12.3.

On the topic of uniqueness theorem, it seems a good time to draw attention to the uniqueness of SNF and what it can do for us. Matt Baker provides a nice proof of this uniqueness result as well. The uniquely determined diagonal entries of the SNF a matrix A are called its SNF invariant factors (the invariant factors

of a finitely generated module over a PID arise from the invariant factors of a specific presentation matrix, so this choice of name is not coincidental). The ith SNF invariant factor d_i equals $d_i(A)/d_{i-1}(A)$, where $d_j(A)$ denotes the GCD of all $j \times j$ minors of A. In particular, the determinant of a square matrix A equals the determinant of its SNF (so it is the product of the SNF invariant factors of A). This fact will be useful when discussing discriminants in the following section on field theory.

1.4.3 Matrix Similarity Over a Field

Since any matrix over a field is similar to a unique matrix in rational canonical form, the collection of distinct possible rational canonical forms is a set of representatives of the similarity classes. Before examining the similarity of two integral matrices over the integers, it is prudent to test whether they are similar over the rationals. If they are not similar over the rationals, then they cannot be similar over the integers. Therefore, when presented with two integer matrices and asked to test for similarity, our first step should be to test similarity over the rationals.

It turns out that we can use the SNF to test for similarity of matrices over a field.

Proposition 7 *Let k be a field. Then, $A, B \in \mathrm{Mat}_n(k)$ are similar over k if and only if $xI - A, xI - B \in \mathrm{Mat}_n(k[x])$ have the same SNF over $k[x]$.*

Proof. As discussed previously, any linear endomorphism $T \in \mathrm{End}(V)$ induces a $k[x]$-module structure on V. Therefore, a matrix $X \in \mathrm{Mat}_n(k)$ furnishes a $k[x]$-module M^X (this is just k^n equipped with the action induced by X). We define the k-invariant factors of a matrix X to be the invariant factors of M^X. We do not use the SNF invariant factors of X because this is always a binary matrix, and it does not give enough information in general. Using the structure theorem for finitely generated modules over a PID, we see that A is similar to B if and only if they have the same k-invariant factors. We will show that the k-invariant factors of a matrix X are precisely the invariant factors of the SNF of $xI - X$, and this will prove the proposition.

We obtain the following exact sequence of $k[x]$-modules:

$$k[x]^n \longrightarrow k[x]^n \longrightarrow M^X \longrightarrow 0$$

where the first map is given by multiplication on the left by the matrix $xI - X$ and the second map π sends the standard basis of $k[x]^n$ to the standard basis of k^n (and we extend $k[x]$-linearly). The exactness of this sequence is non-obvious. To see that $\ker(\pi) \subseteq \mathrm{im}(xI - A)$, note that any polynomial vector $p(x)$ can be divided with remainder as $p(x) = (xI - A)q(x) + p(A)$ for some $q(x) \in k[x]^n$. This follows from observing that the action of x and A is the same in the quotient $k[x]^n / \mathrm{im}(xI - A)$, i.e., $[xp(x)] = [Ap(x)]$, which extends to $w(x)$ acting as

$w(A)$ for any polynomial $w(x) \in k[x]$.

By putting $xI - X$ in SNF, noting that equivalent matrices have isomorphic cokernels, we see that $M^X \cong \bigoplus k[x]/(a_i)$ where $\{a_i\}_{i=1}^n$ are the SNF invariant factors of $xI - X$. By the uniqueness of the structure theorem, these SNF invariant factors are the invariant factors of M^X, and thus they are the k-invariant factors of X as required.

■

This is very nice because $k[x]$ is a Euclidean domain, so we can algorithmically compute Smith forms of matrices over it. MIT OCW has a nice explanation of this. Although this algorithm is not particularly efficient, it shows that we can have a way to practically test for similarity over fields without too much effort. Unfortunately, we cannot say the same for the integers.

Before turning to field theory, let us mention one particularly easy case for similarity over a field. In general, there exist non-similar matrices with the same characteristic polynomial. However, each *irreducible* characteristic polynomial only gives rise to a single similarity class of matrices over a field.

Proposition 8 *Let k be a field and $f(x) \in k[x]$. Then, all matrices in $\mathrm{Mat}_n(k)$ with characteristic polynomial $f(x)$ are similar if and only if the unique factorization of $f(x)$ into irreducibles over k is square-free.*

Proof. If the factorization of $f(x)$ is square-free, then every matrix with characteristic polynomial $f(x)$ also has minimal polynomial $f(x)$. It follows that the rational canonical form of all such matrices is exactly the companion matrix of $f(x)$, hence they are all similar.

For the other direction, observe that if $f(x)$ is not square-free, we can have matrices $A, B \in \mathrm{Mat}_n(k)$ with characteristic polynomial $f(x)$ but distinct minimal polynomials. It follows that A is not similar to B, which concludes the proof.

■

Frequently, we will seek to test whether two integral matrices with the same irreducible characteristic polynomial are similar over the integers (this is the simplest integral similarity case we can consider), so this result tells us we can skip the step of testing similarity over the rationals.

1.5 Field Theory

In this section, we address the theory of norms, traces, and discriminants of field extensions. These ideas will be essential tools for understanding the principal ideal testing algorithm. All of this theory and more can be found in Milne's "Algebraic Number Theory," but we collect the results here for completeness and to expand on some proofs.

1.5.1 Norm and Trace

Let $A \subseteq B$ be an extension of commutative rings. Let m_x denote left multiplication by $x \in B$. We will define two maps $B \to A$ by applying linear algebraic operations to m_x. We define the trace map by $\mathrm{Tr}_{B/A}(x) = \mathrm{Tr}(m_x)$ and the norm map by $\mathrm{N}_{B/A}(x) = \det(m_x)$.

We will primarily consider the trace and norm of field extensions (such as $\mathcal{K}/\mathbb{Q}$, for a number field $\mathcal{K}$). Our first goal is to understand an alternative characterization of norm and trace in terms of the roots of minimal and characteristic polynomials. We follow Keith Conrad's notes. Recall that if A is an finite-dimensional associative, unital algebra over a field k, then we can define the characteristic polynomial of any element $x \in k$ as the characteristic polynomial of the map m_x. So as not to confuse the multiplication map m_α with the minimal polynomial, given a finite field extension L/K, we will write $\pi_{\alpha,L/K}(x)$ for the minimal polynomial of $\alpha \in L$ over K. Similarly, we write $\chi_{\alpha,L/K}(x)$ for the characteristic polynomial of $\alpha \in L$ over K.

Proposition 9 *Let L/K be a degree n extension of fields. Then, we have $\chi_{\alpha,L/K}(x) = \pi_{\alpha,L/K}(x)^{n/d}$, where $d = [K(\alpha) : K]$.*

Proof. The matrix representation of the multiplication by α map on $K(\alpha) \cong K/(\pi_\alpha(x))$ is the companion matrix of the minimal polynomial of α, so $\chi_{\alpha,K(\alpha)/K}(x)$ is a monic polynomial of degree d. But, a simple application of Cayley-Hamilton tells us that any element of a field satisfies its own characteristic polynomial, so $\pi_\alpha(x) \mid \chi_{\alpha,K(\alpha)/K}(x)$, and since they are both monic of degree d they must be equal.

Let $m = [L : K(\alpha)]$ and let $\{\beta_1, \ldots, \beta_m\}$ be a $K(\alpha)$-basis for L. It follows that $\{\alpha\beta_1, \ldots, \alpha^{d-1}\beta_1, \ldots, \alpha\beta^m, \ldots, \alpha^{d-1}\beta^m\}$ is a K-basis for L. Examining this basis, we see that the matrix representation of multiplication by α on L is the direct sum of m copies of the companion matrix of $\pi_\alpha(x)$.

Therefore, $\chi_{\alpha,L/K}(x) = (\chi_{\alpha,K(\alpha)/K}(x))^m = (\pi_\alpha(x))^m$.

As degree is multiplicative, $[L : K(\alpha)][K(\alpha) : K] = [L : K]$, and we see $m = n/d$ as desired.

■

It follows that $\mathrm{N}_{L/K}(\alpha) = (-1)^n a_0$, where a_0 is the constant term of $\chi_{\alpha,L/K}(x)$.

By the proposition, this is precisely $(-1)^n b_0^{\frac{n}{d}}$, where b_0 is the constant term of $\pi_{\alpha,L/K}(x)$. If the minimal polynomial factors as $(x - \alpha_1) \cdots (x - \alpha_d)$, then the norm of α is precisely $(\alpha_1 \cdots \alpha_d)^{\frac{n}{d}}$. Similarly, by examining the second coefficient of the characteristic polynomial, we see that $\mathrm{Tr}_{L/K}(\alpha) = \frac{n}{d}(\alpha_1 + \cdots + \alpha_d)$.

We now want to relate the the norm of $x \in \mathcal{K}$ to the embeddings of $\mathcal{K} \hookrightarrow \mathbb{C}$. It is this form of the norm that will be primarily used going forward. First, recall that there are exactly $[\mathcal{K} : \mathbb{Q}]$ embeddings of $\mathcal{K}$ into $\mathbb{C}$. Indeed, by the primitive element theorem, we can write $\mathcal{K} = \mathbb{Q}(\alpha)$ for some $\alpha \in \mathcal{K}$. Then, $\pi_{\alpha,\mathcal{K}/\mathbb{Q}}(x)$ has exactly n roots in $\mathbb{C}$, say $\{\alpha_i\}_{i=1}^n$ (irreducible polynomials are separable in characteristic 0), and sending $\alpha \mapsto \alpha_i$ defines n embeddings of $\mathcal{K}$ (we use the universal property of polynomial rings and quotients to guarantee this is a well-defined ring morphism). It is easy to check that these are the only possible embeddings.

Proposition 10 *Let $\mathcal{K}$ be a degree n number field with embeddings $\sigma_i : \mathcal{K} \hookrightarrow \mathbb{C}$. Then, $\mathrm{N}_{\mathcal{K}/\mathbb{Q}}(x) = \prod_{i=1}^n \sigma_i(x)$ for all $x \in \mathcal{K}$.*

Proof. By the primitive element theorem, we may write $\mathcal{K} = \mathbb{Q}(\alpha)$. By the remarks following Proposition 9, $\mathrm{N}_{\mathcal{K}/\mathbb{Q}}(\alpha) = \prod_{i=1}^n \sigma_i(\alpha)$.

To obtain the result for any $y \in \mathcal{K}$, observe that $\mathcal{K}/\mathbb{Q}(y)$ is a finite extension, and therefore primitive, i.e., there exists $\beta \in \mathcal{K}$ such that $\mathcal{K} = \mathbb{Q}(y)(\beta)$. By the same argument as before, we know that $\mathbb{Q}(y)$ has $d = \deg(\pi_y(x))$ embeddings $\tau_1, \ldots, \tau_d$. For each such embedding τ_i, by sending β to the roots of its minimal polynomial over $\mathbb{Q}(y)$, we obtain m embeddings $\mathcal{K} \hookrightarrow \mathbb{C}$ extending τ_i. Because $m = n/d$, we see that these extended embeddings make up all of the σ_i.

Now, we have $\mathrm{N}_{\mathbb{Q}(y)/\mathbb{Q}}(y) = \prod_{i=1}^d \tau_i(y)$, so $\mathrm{N}_{\mathcal{K}/\mathbb{Q}}(y) = (\prod_{i=1}^d \tau_i(y))^{n/d}$.

Because each τ_i gives rise to m of the σ_i, we see that $\mathrm{N}_{\mathcal{K}/\mathbb{Q}}(y) = \prod_{i=1}^n \sigma_i(y)$ as desired.

■

The above argument holds in the more general setting of a finite separable field extension.

An essentially identical argument establishes that $\mathrm{Tr}_{\mathcal{K}/\mathbb{Q}}(x) = \sum_{i=1}^n \sigma_i(x)$.

We can also use the field norm to find the units of $\mathcal{O}_\mathcal{K}$. Because the minimal polynomial of any algebraic integer is integral by Gauss' lemma, we see that the characteristic polynomial of any $x \in \mathcal{O}_\mathcal{K}$ is integral, so the norm of x is an integer (arising from the constant term of the characteristic polynomial). If x is a unit, then by the multiplicativity of the norm (which comes from the multiplicativity of the determinant), we see $\mathrm{N}_{\mathcal{K}/\mathbb{Q}}(x) = \pm 1$. Further, if $\mathrm{N}_{\mathcal{K}/\mathbb{Q}}(x) = \pm 1$, then the constant term of the minimal polynomial of x is also ± 1.

Let $\pi_x(x) = x^n + a_{n-1}x^{n-1} + \cdots + a_0$ be the minimal polynomial of x. We can then write $x(x^{n-1} + a_{n-1}x^{n-2} + \cdots + a_1) = \pm 1$, and we find x is a unit. Therefore, the units of $\mathcal{O}_\mathcal{K}$ are precisely the elements with norm ± 1.

1.5.2 Numerical Norm

There is also the notion of the numerical norm of an ideal of an order $\mathcal{O} \subseteq K$. Let $I \subseteq \mathcal{O}$ be an ideal. We define $\mathbb{N}(I) = [\mathcal{O} : I]$, and we call this the numerical norm of I in $\mathcal{O}$.

We should first show that $\mathbb{N}(I)$ is finite. This is immediate from the aligned basis theorem, but it also follows from I and $\mathcal{O}$ being $\mathbb{Z}$ lattices in $\mathcal{K}$. Indeed, there exists $m \in \mathbb{Z}$ such that $m\mathcal{O} \subseteq I \subseteq \mathcal{O}$. Therefore, $\mathbb{N}(I) \leq |\mathcal{O}/m\mathcal{O}| = m^n$, where n is the degree of $\mathcal{K}$. This also shows that there are finitely many ideals in $\mathcal{O}$ with a given norm. Suppose $\mathbb{N}(I) = m$. Then, by Lagrange's theorem, $m\mathcal{O} \subseteq I \subseteq \mathcal{O}$. But, we know that ideals of the finite group $\mathcal{O}/m\mathcal{O}$ are in bijection with ideal containing $m\mathcal{O}$, so there are only finitely many possible choices for I.

There is an important relationship between the numerical norm and field norm.

Proposition 11 *Let $x \in \mathcal{O}$. Then, $\mathrm{N}_{\mathcal{K}/\mathbb{Q}}(x) = \mathbb{N}(x\mathcal{O})$.*

Proof. This is a special case of Proposition 12 below, taking $M = \mathcal{O}$ and $N = x\mathcal{O}$. Note that if $y_1, \ldots, y_n$ is a $\mathbb{Z}$-basis of $\mathcal{O}$, then $xy_1, \ldots, xy_n$ is a $\mathbb{Z}$-basis of $x\mathcal{O}$, and the change of basis matrix representing the $x\mathcal{O}$ basis in terms of the $\mathcal{O}$ basis is precisely the matrix representation for multiplication by x on $\mathcal{O}$ with respect to $y_1, \ldots, y_n$.

■

Importantly, this tells us that the numerical norm of any ideal is less than or equal to the norm of any of its elements. This fact will play a crucial role in the algorithm we use to compute minima of ideals.

One can also show that the numerical norm is multiplicative.

1.5.3 Discriminants

Milne's section on discriminants is quite clear, so we refer the reader to his notes for the basic definitions.

However, we would like to call attention to two specific points. Milne's Remark 2.25 will be especially relevant for us when discussing Minkowski theory, so we state the case we need as a proposition.

Proposition 12 *Let M be a rank n free $\mathbb{Z}$-module with basis $e_1, \ldots, e_n$ and let $N \subseteq M$ be a rank n submodule of M with basis $f_1, \ldots, f_n$. Write $f_i = \sum_{j=1}^{n} a_{ij} e_j$ and define $A = (a_{ij})_{n \times n}$. Then, $[M : N] = |\det(A)|$.*

Proof. Write $[f_1, \ldots, f_n] = [e_1, \ldots, e_n]A$. As in Theorem 6, use SNF to write $A = XDY$, and observe $[f_1, \ldots, f_n]Y^{-1} = [e_1, \ldots, e_n]XD$. We thereby obtain aligned bases for M and N with the scaling terms as the diagonal entries

$d_1, \ldots, d_n$ of D. Clearly, $[M : N] = |d_1 \cdots d_n| = |\det(D)|$, and by the uniqueness of SNF, this is precisely $|\det(A)|$.

■

This proposition tells us that the index of a full rank submodule is the product of the SNF invariant factors of the change of basis matrix.

Recall that the discriminant of an extension of fields L/K is defined up to squares of units. More precisely, if $a_1, \ldots, a_n$ and $b_1, \ldots, b_n$ are bases for L over K, then $D(a_1, \ldots, a_n) = \det(C^2)D(b_1, \ldots, b_n)$, where C is the change of basis matrix associated with our bases. It follows that the discriminant of a free $\mathbb{Z}$-module is a well-defined integer, as 1 is the only square of a unit in $\mathbb{Z}$.

Also, recall that fractional ideals of $\mathcal{O}_\mathcal{K}$ are the same as full $\mathcal{O}_\mathcal{K}$-modules in K, so they are free $\mathbb{Z}$-modules of full rank. Therefore, all ideals of $\mathcal{O}_\mathcal{K}$ have a well-defined discriminant over $\mathbb{Z}$.

Given an ideal $I \subseteq \mathcal{O}_\mathcal{K}$, we can choose integral bases for I and $\mathcal{O}_\mathcal{K}$, which are also $\mathbb{Q}$-bases for $\mathcal{K}$. By Proposition 12, $\operatorname{disc}(I/\mathbb{Z}) = [\mathcal{O}_\mathcal{K} : I]^2 \operatorname{disc}(\mathcal{O}_\mathcal{K}/\mathbb{Z})$. Using the numerical norm, $\operatorname{disc}(I/\mathbb{Z}) = \mathbb{N}(I)^2 \operatorname{disc}(\mathcal{O}_\mathcal{K}/\mathbb{Z})$.

Another result that will be important for us is the non-degeneracy of the trace pairing on a number field.

Proposition 13 *Let $\mathcal{K}$ be a degree n number field with embeddings $\sigma_i : \mathcal{K} \hookrightarrow \mathbb{C}$. Then, $D(\beta_1, \ldots, \beta_n) = \det((\sigma_i(\beta_j)_{n\times n})^2 \neq 0$, for every basis $\beta_1, \ldots, \beta_n$ of $\mathcal{K}$ over $\mathbb{Q}$.*

■

The above proposition is a special case of Proposition 2.26 in Milne.

Proposition 13 also gives us an alternative definition of the discriminant of a number field. Let us quickly clear up some terminology. Generally, the discriminant is defined with respect to a specific basis. When considering free $\mathbb{Z}$-modules, the choice of basis becomes irrelevant.

When we speak of "the discriminant of number field $\mathcal{K}$," we mean precisely $\operatorname{disc}(\mathcal{O}_\mathcal{K}/\mathbb{Z})$. For ease of notation, throughout the rest of this paper, we will denote $\mathrm{d}_\mathcal{K} = \operatorname{disc}(\mathcal{O}_\mathcal{K}/\mathbb{Z})$. If $\beta_1, \ldots \beta_n$ is an integral basis of $\mathcal{O}_\mathcal{K}$, we will call $(\sigma_i(\beta_j)_{n\times n})$ an embedding matrix of $\mathcal{O}_\mathcal{K}$.

Therefore, Proposition 13 tells us that $\mathrm{d}_\mathcal{K}$ is nonzero and equals the square of the determinant of any embedding matrix of $\mathcal{O}_\mathcal{K}$.

In other words, $|\,\mathrm{d}_\mathcal{K}\,|^{\frac{1}{2}} = |\det((\sigma_i(\beta_j)_{n\times n}))|$.

2 Notes on Husert's Dissertation

In this section, we provide supplementary notes to the first three sections of Husert's dissertation. These sections demonstrate how the integral similarity problem may be completely understood in nice cases using algebraic number theory. It is best to read this section alongside Husert's work.

2.1 Full Modules Over Orders (Section 1.1)

This section introduces the essential definitions and terminology to concisely state the generalized Latimer-MacDuffee theorem. Almost all Husert's comments here are restatements of the material presented in the Preliminary Sections 1.1 and 1.2. However, parts of Proposition (1.2) are non-obvious, so we give a slightly expanded argument below.

Proposition 1.2* *The multiplier ring is a $\mathbb{Z}$-order of $\mathcal{K} = \bigoplus_{i=1}^{s} \mathcal{K}_i$.*

Proof. Recall that the multiplier ring of a free $\mathbb{Z}$-module $\mathcal{U}$ of rank $\dim_{\mathbb{Q}} \mathcal{K}^{\mathbf{n}}$ is defined to be $\mathcal{O} = \{x \in \mathcal{K} \mid x\mathcal{U} \subseteq \mathcal{U}\}$. It is clear that $\mathcal{O}$ is a subring of $\mathcal{K}$, so it remains to show that $\mathcal{O}$ has full rank in $\mathcal{K}$.

Define $\pi : \mathcal{K}^{\mathbf{n}} \twoheadrightarrow \mathcal{K}$ by picking a copy of $\mathcal{K}$ inside of $\mathcal{K}^{\mathbf{n}}$ and projecting onto it. Clearly, π defines a module morphism, so $\pi(\mathcal{U})$ is also an $\mathcal{O}$-module. Since $\mathcal{U}$ spans $\mathcal{K}^{\mathbf{n}}$ over $\mathbb{Q}$, the piece of $\mathcal{U}$ in $\mathcal{K}$, i.e. $\pi(\mathcal{U})$, must span $\mathcal{K}$ over $\mathbb{Q}$. Since $\pi(\mathcal{U})$ is a full $\mathcal{O}$-module in $\mathcal{K}$, we may choose some $\alpha \in \mathbb{Z}$ such that $\mathcal{O}_{\mathcal{K}} \subseteq (1/\alpha)\pi(\mathcal{U})$. It follows that there is some nonzerodivisor $a \in \pi(\mathcal{U})$, since otherwise every tuple of algebraic integers would have a zero coordinate, which is clearly not the case.

It follows that $a\mathcal{O} \subseteq \pi(\mathcal{U})$, so $\mathcal{O} \subseteq (1/a)\pi(\mathcal{U})$, and therefore $\mathcal{O}$ is free of rank less than or equal to $\dim_{\mathbb{Q}} \mathcal{K} := m$ (by the fundamental theorem of finitely generated abelian groups). Next, let $\{\alpha_i\}_{i=1}^{m}$ be a $\mathbb{Z}$-basis of $\pi(\mathcal{U})$. Then, $\alpha_i\alpha_j = \sum_{k=1}^{m} x_k^{ij}\alpha_k$ for some rational numbers $\{x_k^{ij}\}$. By scaling out the denominators, we see that for some integers $\{y_k^{ij}\}$, we have $(x\alpha_i)\alpha_j = \sum_{k=1}^{m} y_k^{ij}\alpha_k \in \pi(\mathcal{U})$.

Therefore, for each i, $x\alpha_i \in \mathcal{O}$, and since these elements form a linearly independent set, the rank of $\mathcal{O}$ is at least m. It follows that $\mathcal{O}$ is an order in $\mathcal{K}$ as desired.

■

2.2 The Theorem of Latimer and MacDuffee (Section 1.2)

In this section, Husert proves what will be our main tool for studying matrix similarity over the integers.

A note on terminology is in order. We say that $X \in \mathrm{Mat}_n(k)$ is semisimple if its minimal polynomial is square-free. One can show that this is equivalent to the condition that every X-invariant subspace of k^n has an X-invariant complement. This explains the connection between semisimple matrices and semisimple rings. Indeed, X is semisimple if and only if its associated module M^X (see Section 1.4) is a semisimple $k[x]$-module. One can describe a semisimple matrix in terms of its eigenvalues. We say that an eigenvalue λ of X is semisimple if its geometric and algebraic multiplicity are equal. As explained in the notes on lemma 1.5, in characteristic 0, a matrix is semisimple if and only if all of its eigenvalues are semisimple.

We can now state the generalized Latimer-MacDuffee theorem.

Theorem (Latimer-MacDuffee-Husert)

- Let $v = (v_1, \ldots, v_s)$ be a tuple of algebraic integers with distinct minimal polynomials $\mu_1, \ldots, \mu_s$.
- Let $\mathcal{K} = \bigoplus_{i=1}^{s} \mathbb{Q}(v_i)$, and let $\mathcal{O} = \mathbb{Z}[v]$.
- Let $\mathbf{n} = (n_1, \ldots, n_s)$, $n_i \in \mathbb{Z}_{>0}$.
- Let $\mu = \mu_1 \cdots \mu_s$ and $\chi = \mu_1^{n_1} \cdots \mu_s^{n_s}$.

Then, there is a bijection between the set of similarity classes of integer matrices with minimal polynomial μ and characteristic polynomial χ and isomorphism classes of full $\mathcal{O}$-modules in $\mathcal{K}^{\mathbf{n}}$.

The proof of this powerful theorem mostly relies on clever linear algebra. For the full proof, see Husert's thesis. We collect a couple of notes on the proof below.

Notes on Lemma 1.5: A squarefree polynomial in characteristic 0 is separable. Indeed, irreducible polynomials are coprime to their derivatives, and distinct irreducibles cannot share any roots by considering minimal polynomials. Therefore, any semisimple integer matrix A is diagonalizable in some field extension (such as a splitting field of the characteristic polynomial). It follows that the algebraic and geometric multiplicities of any eigenvalues of A are the same when we compute geometric multiplicities over our splitting field. However, since eigenspaces arise as the kernel of a linear map, their dimension may be computed by row reduction (which is independent of what field we view our matrix in), and therefore the dimension of $\mathrm{Eig}(A, v_\iota)$ over $\mathcal{K}_\iota$ is n_ι, the algebraic multiplicity of v_ι. With these ideas in mind, the proof that $\mathcal{U}$ is a module over $\mathbb{Z}[\alpha]$ is straightforward.

Elementary linear algebra tells us that a basis of a subspace W of a vector space V may be extended to a basis of V. Unfortunately, general modules are not so well-behaved. Indeed, there exist free submodules of free modules whose bases may not be extended to their parent module. We can characterize free submodules of free modules whose bases may be extended. A free submodule N of a free R-module M may be extended to a basis of M if and only if M/N is free. This is useful for proving that $\mathcal{U}$ is a *full* $\mathcal{O}$-module in $\mathcal{K}^{\mathbf{n}}$.

It suffices to show that the columns of Ξ are $\mathbb{Z}$-linearly independent. Suppose this is not the case. The matrix Ξ corresponding to $A \in \mathrm{Mat}_m(\mathbb{Z})$ may be viewed as a $\mathbb{Z}$-linear map $\mathbb{Z}^m \to \mathcal{U} \subseteq \mathcal{K}^{\mathbf{n}}$. Although we do not yet know how to control the rank of $\mathcal{U}$, we can be sure that $\mathcal{U}$ is $\mathbb{Z}$-free because it a finitely generated $\mathbb{Z}$-submodule of the $\mathbb{Q}$-algebra $\mathcal{K}^{\mathbf{n}}$. Therefore, $\operatorname{im}\Xi$ is a free $\mathbb{Z}$-module. As $\mathbb{Z}/\ker\Xi \cong \operatorname{im}\Xi$, we see that $\ker\Xi$ is a nonzero free submodule (say, of rank j) with a basis that can be extended to a basis of $\mathbb{Z}^m$. We can then define a matrix $B \in \mathrm{GL}(m,\mathbb{Z})$ with columns $w_1, \dots w_m$ as the basis obtained by extending a basis of $\ker\Xi$. Without loss of generality, we can assume that the basis vectors for the kernel are given as the final j columns. Then, $\Xi B = [\Upsilon\ 0]$ where $\Upsilon = [v_1, \dots, v_k]$ has linearly independent columns. Of course, $v_i = \Xi w_i$, and the collection of such vectors is linearly independent because otherwise $\Xi(a_1 w_1 + \cdots + a_{m-j} w_{m-j}) = 0$ for nontrivial scalars $\{a_i\}$, contradicting that the first $m-j$ vectors w_i live outside of $\ker\Xi$. Taking $C = B^T$, we obtain the matrix $C \in \mathrm{Mat}_m(\mathbb{Z})$ from the thesis.

The rest of the proof is clear as written.

Notes on Lemma 1.9: Let us expand on the claim that the vectors $x_1^\iota, \dots, x_{n_\iota}^\iota$ are linearly independent. If this were not the case, we could choose some nontrivial linear combination equal to zero and form a row vector using the coefficients. Then, extend to a basis of the eigenspace. In this way, we can construct $\Gamma \in \mathrm{GL}(\mathbf{n}, \mathcal{K})$ such that $\Gamma\Xi$ has a zero row, so $\Gamma\mathcal{U}$ (and therefore $\mathcal{U}$ itself) cannot be full. Therefore, $\dim_{\mathcal{K}_\iota} \mathrm{Eig}(A, v_\iota) \geq n_\iota$.

Since the geometric multiplicity of an eigenvalue is always less than or equal to its algebraic multiplicity, $\chi = \mu_1^{n_1} \cdots \mu_s^{n_s}$ divides the characteristic polynomial of A. Since these two polynomials have the same degree, they must be equal. It follows that the algebraic multiplicity of v_ι is n_ι, so each v_ι is a semisimple eigenvalue (algebraic and geometric multiplicities coincide). Recall that a matrix is diagonalizable in a field extension (potentially diagonalizable) if and only if each of its eigenvalues are semisimple. So far, we have established that for each irreducible factor μ_ι of the characteristic polynomial of A there exists a semisimple eigenvalue, namely v_ι. By the following lemma (communicated to the author by Vanni Noferini), this implies that all eigenvalues of A are semisimple. Thus, A is potentially diagonalizable, and therefore semisimple with minimal polynomial $\mu_1 \cdots m_s$.

Lemma 4 *Let K/F be a finite extension of fields of characteristic 0 and let $A \in \mathrm{Mat}_n(F)$. Suppose that the minimal polynomial of A decomposes into irreducibles over F as $m_A(x) = m_1(x) \cdots m_s(x)$. Then, for any i, there exists a semisimple root of $m_i(x)$ if and only if all roots of $m_i(x)$ are semisimple.*

Proof. It is easy to show that the rank of a matrix is the order of its largest nonzero minor. Since a field automorphism has trivial kernel, applying a field automorphism entrywise to a matrix preserves whether a minor is zero or not.

Let v_i be a semisimple root of $m_i(x)$. Choose some basis for $\mathrm{Eig}(A, v_i)$ and let W denote the matrix whose columns consist of such basis vectors. By assumption, the geometric multiplicity of v_i equals its algebraic multiplicity n_i. Next, let $G_i = \mathrm{Gal}(m_i(x))$. Consider $(A - v_i I)W = 0$ and apply some $g \in G_i$ entrywise to obtain $(A - g(v_i)I)g(W) = 0$. Recall that g permutes the roots of $m_i(x)$. We know that $g(W)$ has full rank because g is an automorphism, so the geometric multiplicity of $g(v_i)$ is at least n_i. Since our fields are characteristic 0, the algebraic multiplicity of $g(v_i)$ is the same of that of v_i (irreducible polynomials are separable). It follows that the geometric multiplicity of $g(v_i)$ is n_i. Because $m_i(x)$ is irreducible, G_i acts transitively on its roots, and every root must be semisimple.

■

Let us discuss where we now stand with the integral matrix similarity problem. The Latimer-MacDuffee-Husert (LMH) theorem tells us that two integral matrices with minimal polynomial μ and characteristic polynomial χ are similar if and only if their corresponding modules in $\mathcal{K}^{\mathbf{n}}$ are isomorphic. Therefore, we will have made great progress if we understand how to compute isomorphism classes of full modules. Unfortunately, it appears that computing such isomorphism classes is a relatively hard problem. Indeed, in the easiest case (EC) where χ is irreducible and $\mathbb{Z}[v] = \mathcal{O}_{\mathcal{K}}$, this amount to determining whether two fractional ideals belong to the same idea class. By inverting one such fractional ideal, we see that this is equivalent to determining whether a given fractional ideal is principal. There are known algorithms to perform this task, and we cover them in depth in Section 3, but we will see that even this simple case requires a lot of work to understand.

Fortunately, as we will soon see, the principal ideal testing algorithm gets us pretty far. Without too much extra work, it allows us to test whether two full modules over the maximal order in $\mathcal{K} = \bigoplus_{i=1}^{s} \mathcal{K}_i$ are isomorphic. Therefore, for any square-free minimal polynomial μ and arbitrary characteristic polynomial χ, if we are lucky enough to have $\mathbb{Z}[v] = \mathcal{O}_{\mathcal{K}}$, then we can test for similarity.

We might also be interested in computing a full collection of representatives of the similarity classes of semisimple matrices with given characteristic and minimal polynomial. To ensure we have found all classes, we need to understand how many there are. By LMH, this amounts to understanding the number of

isomorphism classes of full $\mathcal{O}$-modules. In EC, this is precisely the class number of $\mathbb{Z}[v] = \mathcal{O}_\mathcal{K}$. In small cases, we can use Minkowski's bound to compute class numbers with relative ease, so in EC we often have a method to determine the number of similarity classes. To compute a set of representatives in EC amounts to computing a class group and mapping inequivalent ideals to their matrices under the LMH map. There are known algorithms for computing the class group of a number field. The recent paper "Computing the Ideal Class Monoid of an Order" by Stefano Marseglia claims to have an algorithm that allows for the computation of the ideal class monoid for a non-maximal order. We will consider computing class groups in greater depth later.

Already, we can use ad-hoc methods to produce a set of a representatives for similarity classes. Let us do an easy example. Consider $\mathcal{K} = \mathbb{Q}(\sqrt{-5})$ with $\mathcal{O}_\mathcal{K} = \mathbb{Z}[\sqrt{-5}]$. One can compute that $\mathfrak{p}^2 = (2)$, where $\mathfrak{p} = (2, 1 + \sqrt{-5})$. It follows that $\mathbb{N}(\mathfrak{p}) = 2$. But then there is no $(x, y) \in \mathbb{Z}^2$ such that $x^2 + 5y^2 = 2$, so $\mathfrak{p}$ is not principal. Therefore, the class group has order at least 2. It is easy to compute that Minkowski's bound $M_\mathcal{K}$ is less than 3. Let I be a non-principal ideal of norm 2. Then, by Lagrange's theorem, $2 \in I$, hence $I \mid (2) = \mathfrak{p}^2$. But then it must be that $I = \mathfrak{p}$ because I is not principal. It follows that $\mathbb{Z}[\sqrt{-5}]$ and $\mathfrak{p}$ are a full set of representatives and $\mathcal{K}$ has class number 2.

We write $\mathbb{Z}[\sqrt{-5}] = \mathbb{Z} + \sqrt{-5}\mathbb{Z}$ and $(2, 1 + \sqrt{-5}) = 2\mathbb{Z} + (1 + \sqrt{-5})\mathbb{Z}$. Recall that the LMH correspondence sends full $\mathcal{O} = \mathbb{Z}[v]$-modules to the matrix representation of the multiplication by v map. Therefore, we obtain non-similar matrices

$$\begin{pmatrix} 0 & -5 \\ 1 & 0 \end{pmatrix}_{\mathbb{Z}[v]} \qquad \begin{pmatrix} -1 & -3 \\ 2 & 1 \end{pmatrix}_{\mathfrak{p}}$$

and we know these are a full collection of representatives.

A sharp reader may notice that the matrix corresponding the class of principal ideals is the companion matrix of $\chi(x) = x^2 + 5$. This is no coincidence. The $\mathbb{Z}$-linear multiplication by v on $\mathbb{Z}[v]$ is represented by the companion of the characteristic polynomial, and $(1) = \mathbb{Z}[v]$ is principal.

Note that even if we know a full collection of representatives for similarity classes, it is not obvious how to relate a given matrix to the known classes. Over a field, there are algorithms that transform a matrix to its RCF. However, for integer matrices, the problems of similarity testing and producing a collection of representatives seem more distant because there is no known (similarity preserving) canonical form over the integers.

2.2.1 Purely Monogenic Fields

A number field $\mathcal{K}$ is called monogenic when $\mathcal{O}_\mathcal{K} = \mathbb{Z}[\alpha]$ for some $\alpha \in \mathcal{O}_\mathcal{K}$. In the EC, we are concerned with a specific type of monogenic field. Namely, we

are considering the ideal class group of a number field $\mathcal{K} = \mathbb{Q}(\alpha)$ such that $\mathcal{O}_\mathcal{K} = \mathbb{Z}[\alpha]$ for some $\alpha \in \mathcal{O}_\mathcal{K}$. There does not seem to be an accepted name in the literature for this specific case, so we will coin some terms. A number field $\mathcal{K} = \mathbb{Q}(\alpha)$ is purely monogenic if $\mathcal{O}_\mathcal{K} = \mathbb{Z}[\alpha]$. We will refer to the generator α of a purely monogenic number field as a purely monogenic integer. An integer matrix with irreducible characteristic polynomial and a purely monogenic eigenvalue will be called a purely monogenic matrix.

Some examples of purely monogenic number fields are the cyclotomic fields and quadratic fields $\mathbb{Q}(\sqrt{d})$ when $d \not\equiv 1 \bmod 4$. The cyclotomic field example is more involved, but we will compute the ring of integers of quadratic number fields.

Proposition 14 *Let d be a squarefree integer and let $\mathcal{K} = \mathbb{Q}(\sqrt{d})$. Then, $\mathcal{O}_\mathcal{K} = \mathbb{Z}[\frac{1+\sqrt{d}}{2}]$ if $d \equiv 1 \bmod 4$, and $\mathcal{O}_\mathcal{K} = \mathbb{Z}[d]$ otherwise.*

Proof. Let $\frac{1+\sqrt{d}}{2} = \omega$. Observe that $\mathbb{Z}[d] \subseteq \mathcal{O}_\mathcal{K}$ in general, and $\mathbb{Z}[\omega] \subseteq \mathcal{O}_\mathcal{K}$ when $d \equiv 1 \bmod 4$ because $\omega^2 - \omega + \frac{1-d}{4}$. Therefore, it remains to show the reverse inclusion. Let $\alpha = a + b\sqrt{d}$ be an algebraic integer. The reverse inclusion is obvious if $b = 0$, so we assume $b \neq 0$. Then, we have minimal polynomial $m_\alpha(x) = x^2 - 2ax + (a^2 - b^2 d)$, so $2a$ and $a^2 - b^2 d$ are rational integers.

Next, $4(a^2 - b^2 d) = 4a^2 - 4b^2 d \in \mathbb{Z}$, hence $4b^2 d \in \mathbb{Z}$. We will show that $2b \in \mathbb{Z}$. Write $b = \frac{m}{n}$, $(m, n) = 1$. So, $\frac{4m^2 d}{n^2} \in \mathbb{Z}$, and because d is square-free we get $n^2 \mid 4m^2$. It follows $n \mid 2m$ and writing $2m = nx$ for some $y \in \mathbb{Z}$, and we see $2b = y$.

Writing $a = \frac{x}{2}$ for some $x \in \mathbb{Z}$ and $b = \frac{y}{2}$, we see $x^2 - y^2 d \equiv 0 \bmod 4$. The only squares mod 4 are 0 and 1, corresponding to even and odd integers, respectively. So if $d \equiv 1$ then both x and y are even or both are odd. Otherwise, $d \equiv 2 \bmod 4$ or $d \equiv 3 \bmod 4$, and both a and b are even.

When $d \equiv 1 \bmod 4$, we can write $\alpha = \frac{x-y}{2} + y\omega \in \mathbb{Z}[\omega]$. Otherwise, both a and b are even, so $\alpha \in \mathbb{Z}[d]$.

■

By LMH, the easiest case of testing similarity of integer matrices is precisely equivalent to testing whether a given ideal of $\mathcal{O}_\mathcal{K}$ is principal for a purely monogenic number field $\mathcal{K}$. The number of similarity classes of a purely monogenic matrix with purely monogenic eigenvalue α is the class number of $\mathbb{Q}(\alpha)$.

We have mostly just introduced new language to talk about old ideas in this section, but it helps to make the connection between integer matrices and number theory more clear. One route that we have not yet explored but may be interesting is seeing what we can say about purely monogenic number fields by using matrix theory.

The equality of class numbers and number of similarity classes allows us to use tables of class numbers to come up with examples of matrices with a given number of similarity classes. For example, according to this table, $\mathbb{Q}(\sqrt{-74})$ has class number 10. Because $74 \equiv 2 \bmod 4$, we see $\sqrt{-74}$ is purely monogenic and there are 10 similarity classes of integer matrices with characteristic polynomial $\chi(x) = x^2 + 74$. Unlike the case of matrices over a field, the number of similarity classes with given characteristic polynomial does not seem to depend at all on the degree of the characteristic polynomial.

We may also want to identify when a field is purely monogenic. Although it does not seem to be well understood in general, there exists a simple sufficient condition.

Proposition 15 *Suppose that $\alpha \in \mathcal{O}_\mathcal{K}$ and the discriminant of $\mathbb{Z}[\alpha]$ is square-free. Then, $\mathbb{Q}(\alpha)$ is purely monogenic.*

Proof. Consider the integral basis $1, \alpha, \ldots, \alpha^{n-1}$ for $\mathbb{Z}[\alpha]$. Let $\beta_1 \ldots, \beta_n$ be an integral basis for $\mathcal{O}_\mathcal{K}$. Recall from Proposition 12 that $\operatorname{disc}(1, \alpha, \ldots, \alpha^{n-1}) = [\mathcal{O}_\mathcal{K} : \mathbb{Z}[\alpha]]^2 \, \mathrm{d}_\mathcal{K}$. Because our discriminant is assumed to be square-free, the index must be equal to 1.

■

It is easy to see that the discriminant of any $\mathbb{Z}[\alpha] = \prod_{i<j}(\alpha_i - \alpha_j)^2$ using the formula for the determinant of a Vandermonde matrix.

2.3 Equivalence Over Maximal Orders (Section 1.3)

We now want to extend the work of the previous section on the purely monogenic case (the easiest case), to a broader class of matrices. To do so, we want to understand modules over the maximal order $\mathcal{O}_\mathcal{K}$ (ring of integers) of $\mathcal{K} = \bigoplus_{i=1}^s \mathcal{K}_i$.

Suppose that $\mathcal{K} = \bigoplus_{i=1}^s \mathcal{K}_i$ and $\mathcal{U}$ is a full module in $\mathcal{K}^\mathbf{n}$ over the maximal order $\mathcal{O}_\mathcal{K}$.

Because $\mathcal{O}_\mathcal{K} = \bigoplus_{i=1}^s \mathcal{K}_i$, we have $\mathcal{U} = \bigoplus_{i=1}^s \mathcal{U}_i$ where $\{\mathcal{U}_i\}$ are full $\mathcal{O}_{\mathcal{K}_i}$-modules in $\mathcal{K}^{n_i}$.

Therefore, any two full $\mathcal{O}_\mathcal{K}$-modules in $\mathcal{K}^\mathbf{n}$ can be decomposed as the direct sum of modules, and it is not hard to see that they are isomorphic if and only if each of their corresponding summands are isomorphic. We then see that it suffices to study full $\mathcal{O}_\mathcal{K}$-modules in $\mathcal{K}^n$ where $\mathcal{K}$ is a number field and n is a positive integer and then "stitch things together block by block."

We have reduced the problem from studying full $\bigoplus_{i=1}^s \mathcal{O}_{\mathcal{K}_i}$-modules in $\mathcal{K}^\mathbf{n}$, to full $\mathcal{O}_\mathcal{K}$-modules ($\mathcal{K}$ a number field) in $\mathcal{K}^n$ (n an integer). We want to now

reduce further to studying full $\mathcal{O}_\mathcal{K}$-modules in $\mathcal{K}$.

By the results of Preliminary Section 2.2, we know that $\mathcal{O}_\mathcal{K}$ is a Dedekind domain. From Theorems 1 and 2, we know any full $\mathcal{O}_\mathcal{K}$-module $\mathcal{U}$ is isomorphic to a direct sum of ideals $\bigoplus_{i=1}^{n} A_i$ of $\mathcal{O}_\mathcal{K}$, which is in turn isomorphic to $A_1 \cdots A_n \oplus \mathcal{O}_\mathcal{K}^{n-1}$.

Define $A = A_1 \cdots A_n$ and $B = B_1 \cdots B_n$ for ideals $\{A_i\}$ and $\{B_i\}$ in $\mathcal{O}_\mathcal{K}$. Suppose that $\mathcal{U} = A \oplus \mathcal{O}_\mathcal{K}^{n-1}$ and $\mathcal{B} = B \oplus \mathcal{O}_\mathcal{K}^{n-1}$. Obviously, if $A \cong B$, then $\mathcal{U} \cong \mathcal{B}$. The other direction is also true and is Proposition 1.13 in Husert's work.

Proposition 1.13* *Suppose $\mathcal{U} = A_1 \oplus \cdots \oplus A_n$ and $\mathcal{B} = B_1 \oplus \cdots \oplus B_n$ are two full modules in $\mathcal{K}^n$. If $\Gamma \in \mathrm{GL}(n, \mathcal{K})$ satisfies $\Gamma\mathcal{U} = \mathcal{B}$, then $\det(\Gamma) A_1 \cdots A_n = B_1 \cdots B_n$.*

Proof. We write $\Gamma = (\gamma_{ij})$ and observe that $B_i = \sum_{j=1}^{n} \gamma_{ij} A_j$. Thus, $B_1 \cdots B_n = \prod_{i=1}^{n} (\sum_j \gamma_{ij} A_j)$. When we expand this product, we see that we can obtain a term involving $A_1 \cdots A_n$ in many ways. In fact, each such term corresponds to a choice of some $\sigma \in S_n$ (we get to choose from which factor we want a given A_j with no repeats).

For each $\sigma \in S_n$, we obtain a summand of the form $\prod_j \gamma_{j,\sigma(j)} A_{\sigma(j)}$. It follows that the coefficient appearing on $A_1 \cdots A_n$ in the expanded product is $\sum_{\sigma \in S_n} (\prod_j \gamma_{j,\sigma(j)})$. This is starting to look like the permutation definition of the determinant.

Indeed, $\det \Gamma = \sum_{\sigma \in A_n} (\prod_j \gamma_{j,\sigma(j)}) + \sum_{\sigma \in S_n \setminus A_n} (sgn(\sigma) \prod_j \gamma_{j,\sigma(j)})$.

By definition, $sgn(\sigma) = -1$ for $\sigma \in S_n \setminus A_n$, so we obtain the following equality:

$$\det \Gamma + \sum_{\sigma \in S_n \setminus A_n} (\prod_j \gamma_{j,\sigma(j)}) = \sum_{\sigma \in A_n} (\prod_j \gamma_{j,\sigma(j)}).$$

Therefore, $B_1 \cdots B_n = \det(\Gamma)(A_1 \cdots A_n) + \sum_{\sigma \in S_n \setminus A_n} (\prod_j \gamma_{j,\sigma(j)})(A_1 \cdots A_n)$. It follows that $\det(\Gamma)(A_1 \cdots A_n) \subseteq B_1 \cdots B_n$. By considering Γ^{-1}, we prove the other inclusion identically.

■

Therefore, testing the isomorphism class of full modules over the maximal order amounts to determining whether a series of fractional ideals are principal. Husert says that there are well known algorithms to do this, but he does not address these algorithms. In the following section, we will discuss in detail Johannes Buchmann and H.C. Williams' principal ideal testing algorithm. Minkowski's theory on the geometry of numbers and the Lenstra–Lenstra–Lovász (LLL) algorithm will be indispensable tools, and we will address them as well.

3 Principal Ideal Testing in Number Fields

In this section, we give a detailed account of how to test whether an ideal in the ring of integers of a number field is principal.

3.1 Minkowski Theory and the Finiteness of the Class Number

We mostly follow Milne's ANT Chapter 4, but skip some details that are not relevant to us in order to give a succinct account of Minkowski's bound.

Let us give a definition of a lattice in a vector space which is slightly more general than that in Section 1. Let K be a field with subring $R \subseteq K$. Then, an R-lattice Λ in a K-space V is a finitely generated R-submodule of V such that $K\Lambda = V$. When we first introduced lattices, we required R to be a Noetherian integral domain with field of fractions K. Now, we drop those additional conditions because we will want to consider $\mathbb{Z}$-lattices in $\mathbb{R}^n$. In the case that we want to differentiate between the two sorts of lattices, let us call the former "module lattices" and the latter "Euclidean lattices."

By the fundamental theorem of finitely generated abelian groups, any $\mathbb{Z}$-lattice in $\mathbb{R}^n$ is $\mathbb{Z}$-free. Note that, unlike in module lattices, the spanning condition does not fix the $\mathbb{Z}$-rank of a lattice because $\mathbb{Z}$-independence is different from $\mathbb{R}$-independence. For example, we can have a rank 3 lattice in $\mathbb{R}^2$: take the $\mathbb{Z}$-linear span of $(1,0), (0,1), (\sqrt{2}, 0)$. This is a bit weird, and we will not concern ourselves with such cases, but the given definition does not preclude their existence. We will call a Euclidean lattice $\Lambda \subseteq \mathbb{R}^n$ full if its $\mathbb{Z}$-rank is n.

Given a full lattice Λ in $\mathbb{R}^n$ with basis $\beta = \{x_1, \ldots, x_n\}$ and an anchor point x, we define the fundamental parallelepiped with respect to (x, β) as the set $F_x(\beta) = \{x + \sum_{i=1}^n \lambda_i x_i \mid \lambda_i \in [0,1)\}$. Note that for a fixed basis β, when we vary $x \in \mathbb{R}^n$ over all points, we obtain a partition of $\mathbb{R}^n$. Most of the time, we will have an implicit base point and basis, in which case we will denote a fundamental parallelepiped as just F.

Recall that there exists a canonical Haar measure μ on $\mathbb{R}^n$ such that $\mu([0,1]^n) = 1$. Of course, μ is just the Lebesgue measure on $\mathbb{R}^n$, and the above is a fancy way of stating its basic properties.

Now, we want to show that the measure of the image of the unit cube under $T \in \mathrm{End}(\mathbb{R}^n)$ is $|\det(T)|$. If T is invertible, we can write it as the product of elementary matrices. The affect of the elementary matrices on the measure of the cube precisely corresponds to its affect on the determinant of the identity matrix, so the statement holds for invertible T. If T is not invertible, then it can be written as the product of elementary matrices times a matrix that has a row of zeros. Therefore, the image of the cube under T has all zeros in

some coordinate, so we can produce a covering using cubes with a degenerate side, and the result follows. Because μ is translation invariant, it follows that for all $x \in \mathbb{R}^n$, $\mu(F_x(\beta)) = |\det(\beta)|$, where the RHS should be interpreted as magnitude of the determinant of the matrix with columns from β.

Theorem 8 *(Blichfield) Let F be a fundamental parallelepiped for a full lattice $\Lambda \subseteq \mathbb{R}^n$. Let S be a measurable subset of $\mathbb{R}^n$. If $\mu(S) > \mu(F)$, then there exist distinct $\alpha, \beta \in S$ such that $\beta - \alpha \in \Lambda$.*

Proof. Because the Λ-translates of F partition $\mathbb{R}^n$, we can write $S = \bigcup_{x \in \Lambda} S \cap F_x$. By countable additivity, it follows that $\sum_{x \in \Lambda} \mu(S \cap F_x) = \mu(S) > \mu(F)$. Next, we translate back to the origin by considering $O_x = (S \cap F_x) - x \subseteq F$, and we see that $\sum_{x \in \Lambda} \mu(O_x) > \mu(F)$. It follows that some distinct O_{x_1} and O_{x_2} overlap, so we may choose $w \in O_{x_1} \cap O_{x_2}$. But, then we have $w + x, w + y \in S$, and $(w + x) - (w + y) = x - y \in \Lambda$.

■

Theorem 9 *(Minkowski) Let $T \subseteq \mathbb{R}^n$ satisfy that $\alpha, \beta \in T \implies \frac{1}{2}(\alpha - \beta) \in T$. Let F be a fundamental parallelepiped of Λ. Then, T contains a nonzero point of Λ if $\mu(T) > 2^n \mu(F)$.*

Proof. Let $S = \frac{1}{2}T$. Then, T contains the difference of any two points of S, so it will suffice to show that $\mu(S) > \mu(F)$ by Blichfield's Theorem. But, $\mu(S) = 2^{-n}\mu(T)$, so this clearly holds.

■

The condition on T we specified may seem strange, but there is a large class of subsets which satisfy it. The class we care about are the convex, 0-symmetric subsets. Recall that $X \subseteq \mathbb{R}^n$ is 0-symmetric if $x \in X \implies -x \in X$. It is easy to verify such sets meet the condition placed on T.

Therefore, if $T \subseteq \mathbb{R}^n$ convex, 0-symmetric, and $\mu(T) > 2^n \mu(F)$, then T contains a nonzero point of lattice Λ. This is the standard formulation of Minkowski's convex body theorem. Although generally the strict inequality is required, if we also assume that T is compact, non-strict inequality suffices.

We first need a general lemma about topological groups. We follow the proof given by Tuo.

Lemma 5 *Let G be a Hausdorff topological group. Then, any discrete subgroup $H \leq G$ is closed.*

Proof. By definition of the subspace topology and H being discrete, we can find an open set $U \subseteq G$ such that $U \cap H = \{e\}$. Our first goal will be to show that we can find an open set $V \subseteq U$ such that $VV^{-1} \subseteq U$ and $e \in V$.

Let $\sigma : U \times U \to G$ be given by $(y_1, y_2) \mapsto y_1 y_2^{-1}$. This map is the restriction of the continuous multiplication map composed with the continuous inversion map, so it is continuous. Define $N = \sigma^{-1}(U)$ and observe that $N \subseteq U \times U$ is an open subset containing (e, e). By definition of the product topology, (cartesian products of open sets form a base for the topology), we know that there exist open sets $V_1, V_2 \subseteq U$ such that $e \in V_1, V_2$. Define $V = V_1 \cap V_2$ so that $V \times V \subset V_1 \times V_2$ and $VV^{-1} = \sigma(V \times V) \subseteq \sigma(V_1 \times V_2) \subset U$.

Now, let $x \in H^c$. We will construct an open neighborhood of x contained in H^c, which will show that H is closed. Define $L_x : G \to G$ as left multiplication by x. As a restriction of the continuous multiplication map $G \times G \to G$, we see L_x is continuous for any $x \in G$. It is easy to see that L_x is then a homeomorphism with inverse $L_{x^{-1}}$. In particular, L_x is an open map, so $W := L_x(V)$ is an open neighborhood of x.

Suppose that $h_1, h_2 \in W \cap H$. Then, $h_1 = xv_1$ and $h_2 = xv_2$ for some $v_1, v_2 \in V$. It follows that $h_1 h_2^{-1} = v_1 v_2^{-1} \in VV^{-1} \subseteq U$. Hence, $h_1 h_2^{-1} \in U \cap H = \{e\}$, so $h_1 = h_2$. Therefore, W intersects H at at most one point. If $W \cap H =$, then we have our desired neighborhood of x. If $W \cap H = \{h\}$, then we use the Hasdorff condition on W to find an open neighborhood of x disjoint from h.

■

Note that this lemma also implies that the notion of a shortest nonzero vector in a lattice in $\mathbb{R}^n$ is well defined. Observe that $B_r(0) \cap \Lambda$ is finite, where $B_r(0)$ is the compact n-dimensional ball of radius r centered at 0. Therefore, we can choose a sufficiently large r such that $B_r(0) \cap \Lambda$ contains nonzero lattice points, and then choose a minimum from the finite set of nonzero lengths.

Proposition 16 *Let $T \subseteq \mathbb{R}^n$ be compact, convex, and 0-symmetric. Suppose that $\mu(T) \geq 2^n \mu(F)$. Then, T contains a nonzero lattice point.*

Proof. For any $\epsilon > 0$, $\mu((1+\epsilon)T) = (1+\epsilon)^n \mu(T) > 2^n \mu(F)$, so $(1+\epsilon)T$ contains a nonzero lattice point. It is easy to see that Λ is discrete in the subspace topology, so $\Lambda \cap (1+\epsilon)T$ is a discrete, compact set (Λ is closed by the lemma and closed subsets of compact sets are compact). It follows trivially from the definition of compactness that $\Lambda \cap (1+\epsilon)T$ is finite.

Next, we show that $T = \bigcap_{\epsilon>0}(1+\epsilon)T$. One inclusion is trivial, so we only need to show that $\bigcap_{\epsilon>0}(1+\epsilon)T \subseteq T$. Suppose that $x \in \bigcap_{\epsilon>0}(1+\epsilon)T$ but $x \notin T$. Then, because T is a compact subset of a Hausdorff space, T is closed. Therefore, we can choose a ball around x entirely outside of T. By shrinking ϵ, this contradicts $x \in \bigcap_{\epsilon>0}(1+\epsilon)T$.

Let $a > 0$. Suppose that none of the finitely many points of $\Lambda \cap (1+a)T$ other than the origin are in T. Then, we can choose $\epsilon > 0$ small enough such that $\Lambda \cap (1+\epsilon)T = \{0\}$, which contradicts $(1+\epsilon)T$ containing a nonzero lattice point.

■

So, how does the geometry of lattices relate to number theory and ideal class groups? Well, for any degree n number field $\mathcal{K}$, we can associate a lattice in $\mathbb{R}^n$ to each of its nonzero ideals.

Let $\sigma_i : \mathcal{K} \hookrightarrow \mathbb{C}$ denote the embeddings of $\mathcal{K}$ in $\mathbb{C}$. We say that an embedding is *real* if $\sigma_i(\mathcal{K}) \subseteq \mathbb{R}$ and *complex* if this is not the case. Because the complex conjugate of any embedding is also an embedding, complex embeddings come in conjugate pairs. Therefore, $n = r + 2s$, where r is the number of real embeddings and s is the number of complex embeddings of $\mathcal{K}$.

Putting all of these embeddings together, we obtain a map $\sigma : \mathcal{K} \hookrightarrow \mathbb{R}^n$ given by $\sigma(x) = (\sigma_1(x), \ldots, \sigma_r(x), \Re(\sigma_{r+1}(x), \Im(\sigma_{r+1}(x)), \ldots, \Re(\sigma_{r+s}(x)), \Im(\sigma_{r+s}(x)))$, where $\Re$ and $\Im$ denote taking real and imaginary parts, respectively. The injectivity of σ will be justified shortly.

Crucially, if $\mathfrak{a}$ is any nonzero ideal of $\mathcal{O}_\mathcal{K}$, then $\sigma(\mathfrak{a})$ is actually a lattice in $\mathbb{R}^n$. Therefore, if we want to study an ideal of $\mathcal{O}_\mathcal{K}$, we can consider its associated lattice and use the tools of convex geometry to understand the ideal.

Proposition 17 *Let $\mathfrak{a}$ be a nonzero ideal of $\mathcal{O}_\mathcal{K}$. Then, $\sigma(\mathfrak{a})$ is a full lattice in $\mathbb{R}^n$. If F is a fundamental parallelpiped of $\sigma(\mathfrak{a})$, then $\mu(F) = 2^{-s}\mathbb{N}(\mathfrak{a})|\,\mathrm{d}_\mathcal{K}\,|^{\frac{1}{2}}$.*

Let $\alpha_1, \ldots, \alpha_n$ be an integral basis of $\mathfrak{a}$. To show that $\sigma(\mathfrak{a})$ is a full lattice, it suffices to show that the "lattice matrix" L with rows $\sigma(\alpha_i)$ has nonzero determinant. Recall that the embedding matrix E of $\mathfrak{a}$ has nonzero determinant by Proposition 13.

We will show how to transform E^T into the lattice matrix. We will differentiate between the columns of E^T and the L by using the subscrpits E and L, respectively.

First, note that the first r columns of these two matrices are equal. We start by examining columns $(r+1)_E$ and $(r+2)_E$. By adding column $(r+2)_E$ to $(r+1)_E$, the resulting matrix A_1 has column $r+1$ as two times column $(r+1)_L$. Multiply this column by $\frac{1}{2}$, so that the resulting matrix A_2 has column $r+1$ as $(r+1)_L$. Now, subtract $\frac{1}{2}$ times column $r+1$ of A_2 from column $r+2$ of A_2 to obtain a matrix A_3 with column $r+1$ as $(r+1)_L$ and column $r+2$ as $-i(r+2)_L$. Divide column $r+2$ of A_3 by $-i$ to obtain a matrix A_4 whose first $r+2$ columns agree with those of the lattice matrix.

Observe that all of these transformations change the determinant of E^T by a factor of $\frac{-1}{2i}$. Therefore, after performing these transformations on all s pairs of complex conjugate columns to obtain matrix L, we see that the determinant has been scaled by $\frac{-1}{(2i)^s}$. Therefore, $|\det(L)| = \frac{1}{2^s}|\det(E)|$.

Finally, $\mu(F) = |\det(L)| = 2^{-s}|\det(E)| = 2^{-s}|\operatorname{disc}(\mathfrak{a})|^{\frac{1}{2}} = 2^{-s}\mathbb{N}(\mathfrak{a})|\operatorname{d}_{\mathcal{K}}|^{\frac{1}{2}}$.

■

The above proof is a bit pedantic. Everything should be clear if one writes out the matrices on a blackboard. Note that this argument also shows σ has nontrivial kernel, implying its injectivity.

Theorem 10 *Let $\mathfrak{a}$ be a nonzero ideal of $\mathcal{O}_{\mathcal{K}}$. Then, $\mathfrak{a}$ contains a nonzero element α such that $|\operatorname{N}_{\mathcal{K}/\mathbb{Q}}(\alpha)| \leq \mathbb{N}(\mathfrak{a})M_{\mathcal{K}}$, where $M_{\mathcal{K}} = \frac{n!}{n^n}(\frac{4}{\pi})^s|\operatorname{d}_{\mathcal{K}}|^{\frac{1}{2}}$.*

Proof. First, let $V = \mathbb{R}^r \times \mathbb{C}^s$. This is an n-dimensional real vector space. We can canonically identify V with $\mathbb{R}^n$ through the map φ, which expresses the complex coordinates as their real and imaginary parts.

Then, we have a vector space norm such that $\|x\| = \sum_{i=1}^{r}|x_i| + \sum_{i=r+1}^{r+s}|z_i|$ for any $x = (x_1, \ldots, x_r, z_{r+1}, \ldots, z_{r+s})$.

Define $X(t) = \{x \in V \mid \|x\| \leq t\}$ as the norm ball of radius t. By a messy calculus calculation akin to calculating the volume of the Euclidean n-ball (done in full as lemmas 4.22 and 4.23 in Milne), $\mu(\varphi(X(t))) = 2^{r-s}\pi^s\frac{t^n}{n!}$.

Consider the lattice $\sigma(\mathfrak{a})$ and let F be one of its fundamental parallelpipeds. A general fact from convex geometry is that there is a bijection between compact, convex, 0-symmetric subsets of $\mathbb{R}^n$ with 0 as an interior point, and vector space norms on $\mathbb{R}^n$, so we immediately know $\varphi(X(t))$ satisfies the conditions from Minkowski's theorem. This is also easy to verify directly. Therefore, for t large enough such that $\mu(\varphi(X(t))) \geq 2^n\mu(F)$, the ball $\varphi(X(t))$ must contain a nonzero point $\sigma(\alpha) \in \sigma(\mathfrak{a})$.

Then,

$$\begin{aligned}|\operatorname{N}_{\mathcal{K}/\mathbb{Q}}(\alpha)| &= |\sigma_1(\alpha)|\cdots|\sigma_r(\alpha)||\sigma_{r+1}(\alpha)|^2\cdots|\sigma_{r+s}(\alpha)|^2 \\ &\leq (\sum_{i=1}^{r}|\sigma_i(\alpha)| + \sum_{i=r+1}^{r+s}2|\sigma_i(\alpha)|)^n/n^n \\ &\leq t^n/n^n\end{aligned}$$

by Proposition 10 and the AM-GM inequality.

We want to choose t such that $\mu(\varphi(X(t))) = 2^n\mu(F)$ to make our bound as tight as possible. A quick calculation tells us that the desired value for t is such that $t = (\mathbb{N}(\mathfrak{a})|\operatorname{d}_{\mathcal{K}}|^{\frac{1}{2}}n!(\frac{4}{\pi})^s)^{\frac{1}{n}}$.

Plugging this into our estimate above, we obtain the desired bound.

■

Finally, we arrive at *Minkowski's bound*.

Theorem 11 *Let $\mathfrak{c}$ be a nonzero ideal of $\mathcal{O}_\mathcal{K}$. Then, there exists an ideal J in the ideal class of $\mathfrak{c}$ such that $\mathbb{N}(J) \leq M_\mathcal{K}$.*

Proof. Let I be an integral ideal in the class of $\mathfrak{c}^{-1}$. We can choose some $\gamma \in I$ such that $\mathrm{N}_{\mathcal{K}/\mathbb{Q}}(\gamma) \leq M_\mathcal{K}\mathbb{N}(I)$. Because $\mathcal{O}_\mathcal{K}$ is Dedekind, using unique factorization we can find an integral ideal J such that $IJ = (\gamma)$. It follows that I is in the ideal class of $\mathfrak{c}$. Further, using the multiplicativity of the numerical norm, we see that $\mathbb{N}(I)\mathbb{N}(J) = |\,\mathrm{N}_{\mathcal{K}/\mathbb{Q}}(\gamma)| \leq M_\mathcal{K}\mathbb{N}(I)$. Dividing by $\mathbb{N}(I)$, we obtain the so called Minkowski bound: $\mathbb{N}(J) \leq M_\mathcal{K}$.

■

Because there are only finitely many integral ideals of a given numerical norm, it follows that the ideal class group of $\mathcal{O}_\mathcal{K}$ is finite.

We should note that the finiteness of the class group is not unique to number fields. We will prove a vastly more general finiteness theorem in the section on the Jordan-Zassenhaus Theorem. However, Minkowski's bound specifically is an inherently number theoretic result. It does not appear that there is a similarly nice computable constant that plays the role of $M_\mathcal{K}$ in the general case.

3.2 Reduced Ideals (Buchmann)

We present our understanding of Buchmann's paper "On the computation of units and class numbers by a generalization of Lagrange's algorithm."

Let us introduce notation that will be used throughout the section. Let $\mathcal{K}$ be a number field of degree n with r real embeddings $\sigma_1, \ldots, \sigma_r$ and s pairs of conjugate complex embeddings $\sigma_{r+1}, \overline{\sigma_{r+1}}, \ldots, \sigma_{r+s}, \overline{\sigma_{r+s}}$. Let $m = r + s$. For each $1 \leq i \leq m$, define $|\cdot|_i$ to be the normalized Archimedean valuation. This means that for $x \in \mathcal{K}$, $|x|_i = |\sigma_i(x)|$ when $1 \leq i \leq r$, and $|x|_i = |\sigma_i(x)|^2$ for $r+1 \leq i \leq m$ ($|\cdot|$ being the complex modulus). Let A be an integral ideal of $\mathcal{O}_\mathcal{K}$.

We now give the main definition for this section. Let J be a fractional ideal of $\mathcal{O}_\mathcal{K}$. Then, $0 \neq x \in J$ is a minimum of J if there does not exist a nonzero $\alpha \in J$ such that $|\alpha|_i < |x|_i$ for all $1 \leq i \leq m$. Because $|\,\mathrm{N}_{\mathcal{K}/\mathbb{Q}}(a)| = \prod_{i=1}^m |a|_i$, we see that any element of minimal norm must be a minimum. In particular, the units of $\mathcal{O}_\mathcal{K}$ in J are minima.

Our first goal will be to establish that A has only finitely many non-associated minima.

Proposition 18 *Let x be a minimum of A. Then, $|\,\mathrm{N}_{\mathcal{K}/\mathbb{Q}}(x)| \leq (\frac{2}{\pi})^s \mathbb{N}(A) |\,\mathrm{d}_\mathcal{K}\,|^{\frac{1}{2}}$.*

Proof. Let $C_A = (\frac{2}{\pi})^s \mathbb{N}(A) |\,\mathrm{d}_\mathcal{K}\,|^{\frac{1}{2}}$. We apply Minkowski's convex body theorem to a specific body to obtain the result. Recall that we have a canonical

identification $\varphi : \mathbb{R}^r \times \mathbb{C}^s \to \mathbb{R}^n$. Let $x \in A$ and suppose that $|\,\mathrm{N}_{\mathcal{K}/\mathbb{Q}}(x)| > (\frac{2}{\pi})^s \mathbb{N}(A)|\,\mathrm{d}_{\mathcal{K}}\,|^{\frac{1}{2}}$. We will show that x is not a minimum.

Consider the following set: $X = \{y \in \mathbb{R}^r \times \mathbb{C}^s \mid |y_i| < |\sigma_i(x)|\}$ where y_i denotes the components of y. We can then map X in $\mathbb{R}^n$ by the canonical map: $\varphi(X) = \{y \in \mathbb{R}^n \mid |y_i| < |\sigma_i(x)|,\ y_{r+i}^2 + y_{r+i+1}^2 < |\sigma_{r+i}(x)|^2,\ 1 \le i \le r\}$.

This resulting set can be thought of as a cylinder with r rectangular parts and s circular parts. By using Fubini's theorem and the change of variables formula, we have the following computation:

$$\mu(\varphi(X)) = \textstyle\int_{\varphi(X)} 1 d\mu = (\prod_{i=1}^{r} 2|\sigma_i(x)|)(\prod_{j=1}^{s} \pi|\sigma_{r+j}(x)|^2) = 2^r \pi^s |\,\mathrm{N}_{\mathcal{K}/\mathbb{Q}}(x)|.$$

Now, let σ be as in Proposition 17. Consider the lattice $\sigma(A)$ associated with A and let F be a fundamental parallelpiped of $\sigma(A)$. Then, using the results from Minkowski theory:

$$2^n \mu(F) = 2^{n-s} \mathbb{N}(A)|\,\mathrm{d}_{\mathcal{K}}\,|^{\frac{1}{2}} \le 2^{n-s} |\,\mathrm{N}_{\mathcal{K}/\mathbb{Q}}(x)||\,\mathrm{d}_{\mathcal{K}}\,|^{\frac{1}{2}} < 2^r \pi^s (\tfrac{2}{\pi})^s \mathbb{N}|\,\mathrm{d}_{\mathcal{K}}\,|^{\frac{1}{2}} < \mu(\varphi(X)).$$

Because X is compact, convex, and 0-symmetric, so is $\varphi(X)$. By Minkowski's convex body theorem, $\varphi(X)$ contains a nonzero lattice point, and this means that x is not a minimum of A.

■

This bound on the absolute norm of minima implies that A has finitely many non-associated minima. Indeed, $|\,\mathrm{N}_{\mathcal{K}/\mathbb{Q}}(x)| = \mathbb{N}((x))$, and we know that there are finitely many ideals of a given norm. Therefore, because two elements generate the same ideal if and only if they are associate, for each integer k up to the bound, there are finitely many elements of norm k. Thus, A has finitely many non-associated minima. We will refer to a set of representatives for the non-associate classes of minima of A as a *cycle* of minima of A.

It is not hard to very that if x is a minimum of A and $\alpha \in \mathcal{K}^\times$, then αx is a mimimum of αA.

Proposition 19 *Let A and A' be ideals of $\mathcal{O}_{\mathcal{K}}$. Let C be a cycle of minima in A and let μ' be a minimum of A'. Then, $A \cong A'$ if and only if there exists $\mu \in C$ such that $\frac{1}{\mu} A = \frac{1}{\mu'} A'$.*

Proof. Clearly, $\frac{1}{\mu} A = \frac{1}{\mu'} A'$ implies $A \cong A'$, so it only remains to show the other direction. If $A \cong A'$, then there exists $\alpha \in \mathcal{K}^\times$ such that $A = \alpha A'$. It follows that $\alpha\mu'$ is a minimum of A. Because C is a cycle of minima in A, we can find a unit β such that $\beta\mu = \alpha\mu'$ for some $\mu \in C$.

To conclude, observe that $\frac{1}{\mu} A = \frac{1}{\beta\mu} A = \frac{1}{\alpha\mu'} A = \frac{1}{\mu'} A'$.

■

We say that an integral ideal A is primitive if it is not divisible by any rational integer, i.e., $\frac{1}{n}A$ is not any ideal of $\mathcal{O}_\mathcal{K}$ for any $n \in \mathbb{Z}$. Define $L(A)$ to be the smallest positive rational integer in A. We say that an integral ideal A is reduced if it is primitive and $L(A)$ is a minimum of A.

There is a close connection between cycles of minima and reduced ideals as shown in the following theorem.

Theorem 12 *Let A be an ideal of $\mathcal{O}_\mathcal{K}$ and let $C = \{\mu_1, \ldots, \mu_p\}$ be a cycle of minima in A. Then, there are exactly p reduced ideals $B_1, \ldots, B_p$ in the ideal class of A. These ideals are given by the formula $\frac{1}{L(B_j)}B_j = \frac{1}{\mu_j}A$, for $1 \leq j \leq p$.*

Proof. For $1 \leq j \leq p$, let $m_j = \min\{m \in \mathbb{Z}_{>0} \mid \frac{m}{\mu_j}A \subseteq \mathcal{O}_\mathcal{K}\}$. Define the ideal $B_j := \frac{m_j}{\mu_j}A$. We will show that B_j is reduced.

First, we show that B_j is primitive. Because $B_j \subseteq \mathcal{O}_\mathcal{K}$ and $\mathcal{O}_\mathcal{K}$ is integrally closed, there does not exist any $k \in \mathbb{Z}$ not dividing m_j such that $\frac{1}{k}B_j$ is an integral ideal. Indeed, this would imply B_j contains a proper rational. Further, if $k \in \mathbb{Z}$ divides m_j and divides B_j, then this contradicts the definition of m_j because $\frac{m_j}{k}$ scales $\frac{1}{\mu_j}A$ into $\mathcal{O}_\mathcal{K}$.

Next, we note that 1 is a minimum of $\frac{1}{\mu_j}A$ because μ_j is a minimum of A. It follows that m_j is a minimum of B_j. Now, we just need to show that $m_j = L(B_j)$. Suppose that $m_j > k \in \mathbb{Z}_{>0}$ and $k \in B_j$. Then, $|\mu_j \frac{k}{m_j}|_i < |\mu_j|_i$ for $1 \leq i \leq m$, and $\mu_j \frac{k}{m_j} \in A$, which contradicts μ_j being a minimum of A.

Therefore, each B_j is a reduced ideal of A and obeys the formula from the statement of the theorem. To see that these are the only reduced ideals, observe that the previous proposition implies $\frac{1}{L(B)}B = \frac{1}{\mu_j}A$ for some j because C is a cycle.

Finally, $B_j \neq B_i$ for $j \neq i$, because μ_j is not associated to μ_i (otherwise, we could cancel the ideal by invertibility for a contradiction).

■

Because $\mathcal{O}_\mathcal{K}$ has finite class number, the above theorem implies that there finitely many reduced ideals of $\mathcal{O}_\mathcal{K}$.

We say that a minimum μ' of A is a neighbor of a minimum μ of A provided there is no $0 \neq \alpha \in A$ such that $|\alpha|_i < \max\{|\mu|_i, |\mu'|_i\}$ for $1 \leq i \leq m$, and $|\mu'| \leq |\mu|$.

It is not hard to very that if x is a neighbor of x' and $\alpha \in \mathcal{K}^\times$, then αx is a neighbor of $\alpha x'$. So, multiplication by units respects the neighbor relation.

Lemma 6 *The number of neighbors of a minimum μ of A is finite.*

Proof. Let μ' be another minimum of A. By Proposition 10, we can write $|\,\mathrm{N}_{\mathcal{K}/\mathbb{Q}}(\mu)| = \prod_{i=1}^{m} |\mu|_i$.

Copying Proposition 18, we see that $\prod_{i=1}^{m} \max\{|\mu|_i, |\mu'|_i\} \leq C_A$.

It follows that $0 \leq |\mu'|_i \leq C_A / \prod_{j=1, j\neq i} |\mu|_j$, for $1 \leq i \leq m$.

Letting $\sigma : \mathcal{K} \hookrightarrow \mathbb{R}^n$ be as in Proposition 17, we see that $\sigma(\mu')$ lies in a cylinder with r rectangular parts and s circular parts, the size of which only depends on μ. The intersection of a lattice with such a cylinder is a discrete, compact set, and therefore finite. Because σ is injective, this implies μ has finitely many neighbors.

■

A set $N \subseteq A$ of minima is called a neighbor-cycle if the elements of N are pairwise non-associated, and if for each $x \in N$, each neighbor of x is associated to some element of N. Note that Buchmann calls a neighbor-cycle a cycle of minima. The following fundamental theorem shows these two notions are equivalent.

Theorem 13 *Let $A \subseteq \mathcal{O}_\mathcal{K}$ be an ideal. Then N is a neighbor-cycle of A if and only if it is a cycle of minima of A.*

■

A crucial step in the principal ideal testing algorithm will be to compute the set of all reduced ideals in the class of principal ideals of $\mathcal{O}_\mathcal{K}$. We now give an algorithm to compute the set of reduced ideals in a given ideal class.

Cycle of Reduced Ideals Algorithm

Let $A \subseteq \mathcal{O}_\mathcal{K}$ be an ideal. We will compute the set of all reduced ideals of $\mathcal{O}_\mathcal{K}$ in the ideal class of A. We assume that we already know the $\mathbb{Z}$-basis for some reduced ideal B in the class of A. We will address how to compute such a B later.

We initialize the algorithm with the following parameters:
$k \leftarrow 1, p \leftarrow 1, B_1 \leftarrow B/L(B), \mu_1 \leftarrow 1, i \leftarrow 1$.

We then perform the following steps until $k > p$: $(*)$

1. Compute a complete system of neighbors $T_k = \{\eta_1^{(k)}, \ldots, \eta_{x_k}^{(k)}\}$ of 1 in B_k.
2. Repeat until $i > x_k$. Set $B \leftarrow (1/\eta_i^{(k)})B_k$. If $B = B_j$ for some $1 \leq j \leq p$, then we have a "repeat," so increment $i \leftarrow i + 1$. Otherwise, in order, set $p \leftarrow p + 1$, $B_p \leftarrow B$, and $\mu_p \leftarrow \eta_i^{(k)} \mu_k$, and then increment $i \leftarrow i + 1$.

3. Set $k \leftarrow k+1$.

Finally, once k exceeds p, set $B_k \leftarrow d_k B_k$ with $d_k = \min\{d \in \mathbb{Z}_{>0} \mid dB_k \subseteq \mathcal{O}_\mathcal{K}\}$, for $1 \leq k \leq p$, and terminate the algorithm.

Now, let us justify the correctness of the algorithm and explain how it works.

First, we can verify by induction that $\frac{1}{\mu_k} B_1 = B_k$ in $(*)$. The base case is trivial.

Now, suppose $(1/\mu_j)B_1 = B_1$ for positive integers up to some j. Observe that B_{j+1} is obtained during the kth iteration of $(*)$ for some integer $k \leq j$ as $B_j = (1/\eta_i^{(k)})B_k$.

By induction, we can write $B_k = (1/\mu_k)B_1$, and therefore $B_{j+1} = (1/\eta_i^{(k)})(1/\mu_k)B_1$. Because μ_{j+1} is obtained as $\mu_{j+1} = \eta_i^{(k)}\mu_k$, we see that $B_{j+1} = (1/\mu_j)B_1$.

Similarly, we can show that μ_k is a minimum of B_1 for $1 \leq k \leq p$. The base case is trivial because 1 is a unit.

Now, suppose μ_j is a minimum for positive integers up to some j. For some $k \leq j$, we can write $\mu_j = \eta_i^{(k)}\mu_k$. By induction, μ_k is a minimum of B_1. Further, $\eta_i^{(k)}$ is a neighbor of 1 in B_k, and by definition neighbors are minima. Therefore, $\eta_i^{(k)}$ is a minimum of $B_k = (1/\mu_k)B_1$, hence $\eta_i^{(k)}\mu_k = \mu_{j+1}$ is a minimum of B_1.

Because $B = B_j$ if and only if the minima $\eta\mu_k$ and μ_j are associated, $\{\mu_1, \ldots, \mu_p\}$ is a cycle of minima of B_1 by Theorem 13. By Theorem 12, it follows that the output $\{B_1, \ldots, B_p\}$ at the termination of the algorithm is a cycle of reduced ideals in the class of A.

Let us break down the algorithm in more depth. As we have shown above, to compute a cycle of reduced ideals in the class of A, it suffices to find a cycle of minima of B_1. Furthermore, by Theorem 13 we know that cycles of minima are the same as neighbor-cycles. Therefore, to obtain a cycle of minima of A, we need a set such that up to units the neighbors of its elements stay within the set, and we have no extraneous elements. Essentially, we want the smallest set that contains the minima of A up to units. This is awfully similar to the ubiquitous notion of *generation* in algebra. When dealing with "the (algebraic structure) generated by (subset of algebraic structure)," we frequently have two equivalent descriptions. For example, the span of a set of vectors E in a vector space is the smallest subspace containing E, but also can be characterized as all finite linear combination of elements of E (the set obtained by "closing" under the vector space operations).

Through the lens of generation, we can view Theorem 13 as proving the equivalence of "smallest subset containing the minima up to units" and "smallest

subset closed under the neighbor relation up to units." Because 1 is always a minimum of B_1, we see that this latter notion is the same as the "smallest subset containing 1 closed under the neighbor relation up to units." Let us call a subset closed under the neighbor relation up to units a *neighborhood*. Our goal of computing a cycle of minima of B_1 can thus be re-framed as computing the smallest neighborhood of 1 in B_1.

With this conceptual framework, we can see that all that the algorithm is doing is producing a set closed under the neighbor relation by force. Indeed, we start with $\mu_1 = 1$, and then obtain our first group of minima as all of the non-associated neighbors of 1 in B_1. Then, we obtain our second group of minima by multiplying by the neighbors of 1 in B_2, which amounts to closing under the neighbors of μ_2 because multiplication by units is compatible with the neighbor relation. Each time we close under the neighbor relation, we obtain more elements whose neighbors we must include.

A priori, it is not clear why this process must terminate. But, remember that we have proved that there are finitely many total minima and each has only finitely many neighbors. Therefore, there must exist an index m after which none of our iterations produce new minima. After m, the index p is held constant while k is incremented until it overtakes p and the algorithm terminates.

We also need to make sure that the algorithm cannot terminate before we have a full cycle of minima. To see this, observe that the kth iteration of $(*)$ produces w_k new minima for some $w_k \in \mathbb{Z}$. Therefore, after k iterations, our value for k is just $k+1$, and our value for p is kw_k+1. By definition, after k iterations, our set of minima contains all of the neighbors of $\mu_1, \ldots, \mu_k$. To have $k > p$, the next $k(w_k - 1) + 1$ iterations must be inert (in the sense that no new minima are produced; this is the case precisely when our set already contains all neighbors up to units). When the algorithm terminates, our set of minima therefore contains up to units all of the neighbors of $k + k(w_k - 1) + 1 = p$ minima. By Theorem 13, these minima form a cycle.

3.3 The LLL Algorithm

Let L be an integral lattice in $\mathbb{R}^n$. In the section on Minkowski theory, we showed how such a lattice can be associated to any ideal $\mathfrak{a}$ of the ring of integers $\mathcal{O}_\mathcal{K}$ of a number field. These will be the lattices with which we are primarily concerned, but the lattice reduction algorithm we discuss is broadly useful.

Because a lattice is a free $\mathbb{Z}$-module, any rank m lattice L can be written as $A\mathbb{Z}^m$ for some $A \in \mathrm{Mat}_{n\times m}(\mathbb{R})$. We can then show that two matrices $A, B \in \mathrm{Mat}_{n\times m}(\mathbb{R})$ with associated lattices L_A and L_B define the same lattice if and only if there exists some $C \in \mathrm{GL}_n(\mathbb{Z})$ such that $A = BC$. If such a matrix exists, it is obvious that A and B define the same lattice because $A = BC$

implies $L_B \subseteq L_A$, and inverting C we obtain the reverse inclusion. If A and B define the same lattice, then we can find a (necessarily unimodular) change of basis matrix which scales one into the other. Therefore, the determinant of any matrix defining a full lattice is the same, and we define $\det(L) := |\det(A_L)|$, where A_L is any matrix defining L.

The LLL algorithm takes as an input a basis for an integer lattice and outputs an LLL-reduced basis for the same lattice. The vectors in an LLL-reduced basis are "short" in a precise sense. We will refer the reader to Henri Cohen's "A Course in Computational Algebraic Number Theory," which addresses the algorithm in great depth. Going forward, we assume familiarity with the version of the LLL algorithm presented in Cohen's book.

A useful application of the LLL algorithm is finding a shortest vector in a full integral lattice in $\mathbb{R}^n$. We present the algorithm given by H. W. Lenstra, Jr. in "Integer Programming with a Fixed Number of Variables."

Shortest Vector Algorithm (Lenstra)

Let L be a full integer lattice in $\mathbb{R}^n$ with basis $b_1, \ldots, b_n$. By the LLL algorithm, we may assume that this basis is LLL-reduced. Let $x \in L$ and write $x = \sum_{i=1}^{n} m_i b_i$ for some $m_i \in \mathbb{Z}$.

We can define $y \in \mathbb{R}^n$ as the vector with coordinates m_i and B as the matrix with columns b_i. We can then write $By = x$. Let B_i^x be the matrix B with column i replaced with x.

By Cramer's rule, we have $|m_i| = |\det(B_i^x)|/|\det(B)|$. By Hadamard's inequality and the LLL condition:

$$|m_i| \leq \frac{|x|}{|\det(B)|} \prod_{j=1, j\neq i}^{n} |b_j| = \frac{|x|}{|\det(B)||b_i|} \prod_{j=1}^{n} |b_j| \leq c_2 \frac{|x|}{|b_i|}$$

where $c_2 = 2^{n(n-1)/4}$.

If x is a shortest vector, then $\frac{|x|}{|b_i|} \leq 1$, so $|m_i| \leq c_2$. Thus, if x is a shortest vector, it lies in the finite set $\{y \in L \mid y = \sum_{i=1}^{n} m_i b_i, m_i \in \mathbb{Z}, |m_i| \leq c_2\}$.

Since this allows for $2c_2$ possibilities for each coordinate, and there are n coordinates, the total search space is a set of cardinality $2nc_2$.

3.4 Computing Minima (Buchmann and Willams)

At long last, we reach the final step of the principal ideal testing algorithm. In this section, we describe our understanding of Buchmann and Williams' algo-

rithm for computing a minimum of an ideal $\beta \subseteq \mathcal{O}_\mathcal{K}$.

Before proceeding forward, let us describe how to test whether β is principal, provided we know how to compute minima. We will first choose some principal ideal of $\mathcal{O}_\mathcal{K}$ and compute one of its minima. Using Theorem 12, this allows us to find a reduced principal ideal. Applying the reduced cycle algorithm to this ideal yields a full cycle of the finitely many reduced principal ideals. We then compute a minimum of β and use it to find a reduced ideal R in the class of β. Finally, we check to see if R belongs to the cycle of principal ideals. If it does, then β is principal, and if it does not, then β is not principal.

3.4.1 Rational Approximation Lattices

As before, let $\mathcal{K}$ be a number field of degree n with r real embeddings $\sigma_1, \ldots, \sigma_r$ and s pairs of conjugate complex embeddings $\sigma_{r+1}, \overline{\sigma_{r+1}}, \ldots, \sigma_{r+s}, \overline{\sigma_{r+s}}$. Let $m = r + s$.

Now, let $\omega_1, \ldots, \omega_n$ be an integral basis for $\mathcal{O}_\mathcal{K}$ and let Ω_E be the corresponding embedding matrix. Adding a superscript of (i) to an element of $\mathcal{K}$ will always denote its image under σ_i. For example, $\Omega = (\omega_k^{(i)})_{n\times n}$. Let $\sigma : \mathcal{K} \hookrightarrow \mathbb{R}^n$ be the Minkowski mapping from Proposition 17. Let Ω_L be the matrix whose ith row is $\sigma(\omega_i)$. By the LLL algorithm, we can replace $\sigma(\omega_1), \ldots, \sigma(\omega_n)$ with some LLL reduced basis. Pulling back under the injective map σ, we obtain an integral basis for $\mathcal{O}_\mathcal{K}$ whose lattice comes with an LLL-reduced basis. Therefore, without loss of generality, we can assume $\sigma(\omega_1), \ldots, \sigma(\omega_n)$ is LLL-reduced.

By Proposition 13, we know that $|\det(\Omega_E)| = |\operatorname{d}_\mathcal{K}|^{\frac{1}{2}}$. Furthermore, by Proposition 17, we know that $|\det(\Omega_L)| = 2^{-s}|\det(\Omega_E)|$.

Lemma 7 *For all* $1 \le i, j \le n$, $|\omega_j^{(i)}| < 2^{n^2/4}|\det(\Omega_E)|$.

Proof. We will first prove that $|\sigma(\omega_j)| \ge 1$ for all j. Because $\omega_j \in \mathcal{O}_\mathcal{K}$, we know that $\mathrm{N}_{\mathcal{K}/\mathbb{Q}}(\omega_j) \in \mathbb{Z}$, so its absolute norm is greater than or equal to 1.

It follows that $\prod_{i=1}^n |\omega_j^{(i)}| \ge 1$. Therefore, using the AM-GM inequality, we obtain:

$$1 \le (\textstyle\prod_{i=1}^n |\omega_j^{(i)}|)^{\frac{1}{n}} \le \frac{1}{n}\sum_{i=1}^n |\omega_j^{(i)}| \le \max_i |\omega_j^{(i)}| \le |\sigma(\omega_j)|.$$

Because $\omega_1, \ldots, \omega_n$ is LLL-reduced, we know that $\prod_{i=1}^n |\sigma(\omega_j)| \le 2^{n(n-1)/4}|\det(\Omega_E)|$.

It follows that $|\omega_j^{(i)}| \le |\sigma(\omega_j)| \le 2^{n(n-1)/4}|\det(\Omega_E)| < 2^{n^2/4}|\det(\Omega_E)|$ as claimed.

■

This bound will be used later in the justification of the algorithm to come.

Recall that to any ideal β of $\mathcal{O}_\mathcal{K}$, we can associate a lattice $\sigma(\beta)$. We want to approximate this lattice by a lattice $\hat{\beta}$ with rational coordinates, i.e., $\hat{\beta} \subseteq \mathbb{Q}^n$. Because β is generated over $\mathbb{Z}$ by the $\omega_j^{(i)}$, it will suffice to define a rational approximation to such elements.

For any entry $z = a + bi$ of Ω_E, define its level q rational approximation as $\hat{z}(q) = [2^q a]/2^q + i[2^q b]/2^q$, where the brackets denote the floor function.

For any $\alpha \in \mathcal{O}_\mathcal{K}$, we can write $\alpha = \sum_{j=1}^n x_j \omega_j$ uniquely.

Therefore, $\alpha^{(i)} = \sum_{i=j}^n x_j \omega_j^{(i)}$ and we define $\hat{\alpha}^{(i)} = \sum_{i=j}^n x_j \hat{\omega}_j^{(i)}$.

We now define a rational Minkowski mapping $\hat{\sigma} : \mathcal{O}_\mathcal{K} \to \mathbb{Q}^n \subseteq \mathbb{R}^n$ by sending $\alpha \in \mathcal{O}_\mathcal{K}$ to $(\hat{\alpha}^{(1)}, \ldots, \hat{\alpha}^{(r)}, \Re(\hat{\alpha}^{(r+1)}), \Im(\hat{\alpha}^{(r+1)}), \ldots, \Re(\hat{\alpha}^{(r+s)}), \Im(\hat{\alpha}^{(r+s)})$, where $\Re$ and $\Im$ denote real and imaginary parts, respectively. Let $\hat{\beta} = \hat{\sigma}(\beta)$.

We now define $\hat{\Omega}_L$ to be the matrix whose ith column is $\hat{\sigma}(\omega_i)$.

A simple calculation then yields $|\Omega_L - \hat{\Omega}_L|_{max} < 2^{-q}$, where $|X|_{max}$ denotes the maximum modulus of the entries of matrix X. Also, $|\omega_k^{(i)} - \hat{\omega}_k^{(i)}|_{max} < 2^{-q+1/2}$.

Note the difference in the bounds. The matrices Ω_L and $\hat{\Omega}_L$ have real entries while $\omega_k^{(i)}$ and $\hat{\omega}_k^{(i)}$ may have nontrivial imaginary parts. Although some of our embeddings are complex, the rational Minkowski map $\hat{\sigma}$ converts complex coordiates to real and imaginary parts. The magnitude of the difference between a real number and its floor is always less than or equal to 1, but when dealing with the complex modulus and nontrivial imaginary parts, we end up adding together two such differences, resulting in an extra $\sqrt{2}$ factor.

We also define $\hat{\Omega}_E$ to be the matrix obtained by taking rational approximations component-wise. Note that the entries of this matrix may have nontrivial imaginary parts.

We now want to show that like the Minkowski map σ, the rational Minkowski map sends an ideal $\beta \subseteq \mathcal{O}_\mathcal{K}$ to a lattice in $\mathbb{R}^n$ (for sufficiently large q). Essentially, because we know the determinant of the lattice $\sigma(\beta)$ is nonzero, by choosing a sufficiently precise rational approximation, we can ensure that $\hat{\Omega}_L$ has nonzero determinant as well. Because the columns of $\hat{\Omega}_L$ span $\hat{\sigma}(\mathcal{O}_\mathcal{K})$, its determinant determines whether $\hat{\sigma}(\mathcal{O}_\mathcal{K})$ is a lattice. Further, if $\hat{\sigma}(\mathcal{O}_\mathcal{K})$ is a lattice, the map $\hat{\sigma}$ is injective, so $\hat{\sigma}(\beta) = \hat{\beta}$ is a lattice as well.

Theorem 14 *There is a constant $D = D(n)$ such that if $q > \log_2(2^s D |\, \mathrm{d}_\mathcal{K} \,|^{(n-2)/2})$, then $\hat{\Omega}_L(q)$ is invertible.*

Proof. By the reverse triangle inequality we obtain:

$$|\det(\hat{\Omega}_L)| \geq |\det(\Omega_L)| - |\det(\hat{\Omega}_L) - \det(\Omega_L)|.$$

Recall that $|\det(\Omega_L)| = 2^{-s}|\det(\Omega_E)| > 0$ by Proposition 17. Thus, if we can bound the difference $|\det(\hat{\Omega}_L) - \det(\Omega_L)|$ above by a function of q whose limit $q \to \infty$ is 0, then for large enough q we can force $|\det(\hat{\Omega}_L)| > 0$.

By the multi-linearity of the determinant, we obtain that

$$|\det(\hat{\Omega}_L) - \det(\Omega_L)| \leq \sum_{i=1}^{n} |\det(\Omega_j)|,$$

where Ω_j has columns $[\sigma(\omega_1), \ldots, \sigma(\omega_{j-1}), \hat{\sigma}(\omega_j) - \sigma(\omega_j), \hat{\sigma}(\omega_{j+1}), \ldots, \hat{\sigma}(\omega_n)]$.

By Hadamard's inequality, $|\det(\Omega_j)| \leq \prod_{i=1}^{n} |v_i|$, where the v_i are the columns of Ω_j. We see that we have $\frac{n-1}{2}$ columns coming from Ω_L and $\hat{\Omega}_L$, respectively, and a single column coming from $\hat{\Omega}_L - \Omega_L$.

By Lemma 7 and our previous bound on $|\hat{\Omega}_L - \Omega_L|_{max}$, we have the following inequalities:

$$\begin{aligned}\prod_{i=1}^{n} |v_i| &\leq (\sqrt{n}|v_j|_{max}) \prod_{i=1, i\neq j}^{n} \sqrt{n} 2^{n^2/4} |\operatorname{d}_{\mathcal{K}}|^{\frac{1}{2}} \\ &\leq (2^{n^2(n-1)/4} n^{\frac{n}{2}}) 2^{-q} |\operatorname{d}_{\mathcal{K}}|^{(n-1)/2}.\end{aligned}$$

Let $D_1 = 2^{n^2(n-1)/4} n^{\frac{n}{2}}$. We then see that $|\det(\hat{\Omega}_L) - \det(\Omega_L)| \leq n D_1 2^{-q} |\operatorname{d}_{\mathcal{K}}|^{(n-1)/2}$.

We want to guarantee that $n D_1 2^{-q} |\operatorname{d}_{\mathcal{K}}|^{(n-1)/2} < |\det(\Omega_L)| = 2^{-s} |\operatorname{d}_{\mathcal{K}}|^{\frac{1}{2}}$ to force $\hat{\Omega}_L$ to be invertible.

This amounts to $q > \log_2(2^s E |\operatorname{d}_{\mathcal{K}}|^{\frac{n-2}{2}})$, where $D = n D_1 = 2^{n^2(n-1)/4} n^{\frac{n+1}{2}}$.

■

Now that we know $\hat{\beta}(q)$ is a lattice for sufficiently large q, we can use Lenstra's algorithm to compute a shortest vector $\hat{\mu} \in \hat{\beta}(q)$.

3.4.2 Definitions and Bounds

Let $\hat{\mu} = \sum_{j=1}^{n} y_j \hat{\sigma}(\beta_j) = \sum_{j=1}^{n} x_j \hat{\sigma}(\omega_j)$, be a shortest vector in $\hat{\beta}$, where $\beta_1, \ldots, \beta_n$ is an integral basis of β. Similarly, define $\mu = \sum_{j=1}^{n} y_j \beta_j = \sum_{j=1}^{n} x_j \omega_j$. We will call μ the element associated with $\hat{\mu}$.

Define $T = (\operatorname{Tr}(\omega_{ij}))_{n\times n} = \Omega_E^T \Omega_E$ and $W = |\hat{\Omega}_E T^{-1}|_{max} + 2^{-q+1/2} n |T^{-1}|_{max}$.

Observe that $\Omega_E^{-T} = \hat{\Omega}_E T^{-1} + (\Omega_E - \hat{\Omega}_E)T^{-1}$. Although the norm $|\cdot|_{max}$ is not sub-multiplicative, one can check that $|AB|_{max} \leq n|A|_{max}|B|_{max}$ for $A, B \in \mathrm{Mat}_n(\mathbb{C})$.

It follows that $|\Omega_E^{-1}|_{max} \leq |\hat{\Omega}_E T^{-1}|_{max} + 2^{-q+1/2} n |T^{-1}|_{max} = W$.

Next, define $\delta(q) = 2^{-q+1/2} n^2 W$ and $\hat{M} = |x|_{max}/(nWN)$, where $N = \mathbb{N}(\beta)^{\frac{1}{n}}$, and $\mathbf{x} = (x_1, \ldots, x_n)$ is such that $\hat{\mu} = \sum_{j=1}^n x_j \hat{\omega}_j$.

Choose $\hat{M}_i \in \mathbb{Q}$ such that $|\hat{\mu}^{(i)}| \leq \hat{M}_i N$ for $1 \leq i \leq n$. Define $M_i = \hat{M}_i + \delta\hat{M}$ and $M = \max_i M_i$.

3.4.3 Shortest Vectors and Minima

We will now show that μ is a minimum of β for sufficiently large q. We will proceed by contradiction.

Suppose there exists $0 \neq \mu' \in \beta$ such that $|\mu'^{(i)}| < |\mu^{(i)}|$ for all $1 \leq i \leq n$. $(*)$

Proposition 20 *If $(*)$ holds, then there is some $1 \leq i \leq n$ such that the following bounds hold: $|\mu^{(i)}| - |\mu'^{(i)}| < \delta c_1 N$ and $|\mu^{(i)}|^2 - |\mu'^{(i)}|^2 < \delta c_2 N$, where $c_1 = \hat{M} + M$ and $c_2 = 2M_i c_1$.*

Proof. First, observe that

$$|\mu^{(i)} - \hat{\mu}^{(i)}| = |\textstyle\sum_{j=1}^n x_k \omega_k^{(i)} - \sum_{j=1}^n x_k \hat{\omega}_k^{(i)}| \leq n|x|_{max} 2^{-q+1/2} = \delta \hat{M} N^2,$$

by using $|\omega_k^{(i)} - \hat{\omega}_k^{(i)}| < 2^{-q+1/2}$ and the triangle inequality.

Then, by the reverse triangle inequality and our bound on $|\hat{\mu}^{(i)}|$, we obtain $|\mu^{(i)}| \leq (\hat{M}_i + \delta\hat{M})N = M_i N$.

Next, let $x' \in \mathbb{Z}^n$ be such that $\sigma(\mu') = \Omega_E x'$, which exists because the columns of Ω_E span $\sigma(\mathcal{O}_\mathcal{K})$. It follows that $|x'|_{max} \leq n|\Omega_E^{-1}|_{max}\ |\sigma(\mu')|_{max} \leq nWMN$. Further, we see that $|\mu'^{(i)} - \hat{\mu}'^{(i)}| \leq \delta MN$ by the same argument used in the first string of inequalities in this proof.

Because $\hat{\mu} = \hat{\sigma}(\mu)$ is a shortest vector in $\hat{\beta}$, there exists some $1 \leq i \leq n$ such that $|\hat{\mu}'^{(i)}| - |\hat{\mu}^{(i)}| \geq 0$.

For this i, we have the following inequalities:

$$|\mu'^{(i)}| - |\mu^{(i)}| \leq |\mu^{(i)}| - |\mu'^{(i)}| + |\hat{\mu}'^{(i)}| - |\hat{\mu}^{(i)}| \leq |\mu^{(i)} - \hat{\mu}^{(i)}| + |\mu'^{(i)} - \hat{\mu}'^{(i)}| \leq c_1 \delta N.$$

This establishes our first claim. To see the second claim, set $a = |\mu^{(i)}|$ and $b = |\mu'^{(i)}|$ and observe that $a^2 - b^2 = (a-b)(a+b) \leq c_1 \delta N (2NM_i) = c_2 \delta N^2$.

■

Proposition 21 *If* $(*)$ *holds, then* $|\mu^{(i)}| - |\mu'^{(i)}| > N/c_3$ *for* $1 \leq i \leq r$ *and* $|\mu^{(i)}|^2 - |\mu'^{(i)}|^2 > N/c_4$ *for* $r+1 \leq i \leq n$, *where* $c_3 = 2^{n-1}M^{n-1}$ *and* $c_4 = 2^{n^2-2}M^{n^2-2}$.

Proof. Define $\gamma_i = |\mu^{(i)}| - |\mu'^{(i)}| > 0$ and let $1 \leq i \leq r$. Recall that the numerical norm of an ideal is always less than or equal to the modulus of the norm of any of its elements. Further, one can show $\mathbb{N}(\beta) = \mathbb{N}(\beta^{(i)})$, where the norm of $\beta^{(i)} := \sigma_i(\beta)$ is taken as an ideal of $\mathcal{K}^{(i)} := \sigma_i(\mathcal{K})$.

Therefore, $N^n \leq |\,\mathrm{N}_{\mathcal{K}^{(i)}/\mathbb{Q}}(\gamma_i)| = \prod_{j=1}^n |\tau_j(\gamma_i)|$, where $\tau_j \in \mathrm{Hom}(\mathcal{K}^{(i)}, \mathbb{C})$. Because the embeddings σ_i are real-valued for $1 \leq i \leq r$, our complex modulus reduces to the real absolute value. Therefore, because the τ_j are ring maps, we can commute them with the absolute value to obtain $\prod_{j=1}^n |\tau_j(\gamma_i)| = \prod_{j=1}^n |\tau_j(\mu^{(i)})| - |\tau_j(\mu'^{(i)})|$. But because τ_j is an embedding, it permutes roots of minimal polynomials, so $|\tau_j(\mu^{(i)})|$ and $|\tau_j(\mu'^{(i)})|$ are conjugates of $\mu^{(i)}$ and $\mu'^{(i)}$, respectively.

Using our bound on the magnitude of the conjugates and the triangle inequality, we see that $\tau_j(\gamma_i) \leq 2NM$. Noting that one of the τ_j is the identity, we obtain bound $N^n \leq 2^{n-1}N^{n-1}M^{n-1}\gamma_i$. Thus, $\gamma_i \geq N/(2^{n-1}M^{n-1})$.

Now, let $\xi_i = |\mu^{(i)}|^2 - |\mu'^{(i)}|^2$ and $r+1 \leq i \leq n$. Let $L := \mathcal{K}^{(i)}\mathcal{K}^{(i+1)}$ be the composite field of $\mathcal{K}^{(i)}$ and its conjugate field. Let $l = [L : \mathcal{K}^{(i)}] = [L : \mathcal{K}^{(i+1)}$. Let $B = \beta^{(i)}\beta^{(i+1)}$. It is easy to verify that $\xi_i \in B$.

Using the transitivity and multiplicativity of the relative ideal norm, and its value on extended ideals, we see that $\mathbb{N}(B\mathcal{O}_L) = N^{2nl}$, where the norm is taken viewing $B\mathcal{O}_L$ as an ideal of $\mathcal{O}_L$.

As we saw in Proposition 10, each embedding of $\mathcal{K}^{(i)}$ or $\mathcal{K}^{(i+1)}$ are precisely l extensions to an embedding $L \hookrightarrow \mathbb{C}$. Let $\eta_1, \ldots, \eta_{nl}$ denote the collection of embeddings $L \hookrightarrow \mathbb{C}$. Without loss of generality, assume η_1 is the identity and η_2 is complex conjugation. Note that because $r+1 \leq i \leq n$, $\mathcal{K}^{(i)} \not\subseteq \mathbb{R}$, so these maps are distinct (in the previous case they were the same).

Proceeding as before, we see that

$$N^{2nl} \leq \xi_i^2 \prod_{j=3}^{nl} |\eta_j(\xi_i)| \leq \prod_{j=3}^{nl} |\eta_j(\mu^{(i)}\mu^{(i+1)})| + |\eta_j(\mu'^{(i)}\mu'^{(i+1)})|$$

by the triangle inequality twice.

We then see that $N^{2nl} \leq \xi_i^2 \prod_{j=3}^{nl} 2N^2M^2 = \xi_i^2 2^{nl-2}N^{2(nl-2)}M^{2(nl-2)}$ by using our bounds on conjugates.

It follows that $N^4/(2^{2(nl-2)}M^{2(nl-2)}) \leq \xi_i^2$. And therefore, $\xi_i \geq N/(2^{nl-2}M^{nl-2})$.

Finally, because the index of a compositite of field extensions is less than or equal to the product of their indexes, we see that $l \leq n$, and we obtain $\xi_i \geq N/(2^{n^2-2}M^{n^2-2})$.

■

Corollary 2 *If $\delta(q) < c_5 := \min\{c_1^{-1}, c_2(i)^{-1}, c_3^{-1}, c_4^{-1} \mid 1 \leq i \leq m\}$, then μ is a minimum of β.*

Proof. Suppose that μ is not a minimum, i.e., $(*)$ is true. By Proposition 20, there exists $1 \leq i \leq r$ such that $|\mu^{(i)}| - \mu'^{(i)}| < N/c_3$ and $|\mu^{(i)}|^2 - \mu'^{(i)}|^2 < N/c_4$. By Proposition 21, either $|\mu^{(i)}| - \mu'^{(i)}| \geq N/c_3$ or $|\mu^{(i)}|^2 - \mu'^{(i)}|^2 \geq N/c_4$, so we reach a contradiction.

■

Lemma 8 *If Λ is a full lattice in $\mathbb{R}^n$, then there exists a nonzero lattice point x such that $|x|_{max} \leq \det(\Lambda)^{\frac{1}{n}}$.*

Proof. This lemma is a simple application of Minkowksi's convex body theorem. Indeed, let $m = \min\{|x|_{max} \mid 0 \neq x \in \Lambda\}$ and assume for a contradiction that $m > \det(\Lambda)^{\frac{1}{n}}$. Consider the cube $C = \{x \in \mathbb{R}^n \mid |x|_{max} < m\}$, and observe that $\mu(C) = (2m)^n > 2^n \det(\Lambda)$. Therefore, C contains a nonzero lattice point, which is a contradiction.

■

Theorem 15 *For sufficiently large q, $\delta(q) < c_5(q)$.*

Proof. Step 1:

Recall that $W = |\hat{\Omega}_E T^{-1}|_{max} + 2^{-q+1/2} n |T^{-1}|_{max}$ and $\delta(q) = 2^{-q+1/2} W n^2$.

Our first goal will be to show that $\delta(q) \to 0$ as $q \to \infty$. Clearly, it suffices to give an upper bound on W that does not depend on q. In what follows, a_i, $i \geq 1$, will denote constants only depending on n.

First, $|\hat{\Omega}_E|_{max} \leq |\Omega_E|_{max} \leq a_1 |\operatorname{d}_{\mathcal{K}}|^{\frac{1}{2}}$ by Lemma 7.

By Cramer's rule, $\Omega_E^{-1} = \frac{1}{\det(\Omega_E)} \operatorname{adj}(\Omega_E)$. Recall that the ith column of $\operatorname{adj}(\Omega_E)$ is $\det(A_1^{e_i}, \ldots, A_n^{e_i})$, where $A_j^{e_i}$ denotes the matrix where the jth column of Ω_E is replaced by e_i.

By Lemma 7 again, we know that each of the columns of Ω_E have magnitude less than $a_1 |\det(\Omega_E)|$, hence by Hadamard's inequality, $|\Omega_E^{-1}|_{max} \leq a_2 |\operatorname{d}_{\mathcal{K}}|^{\frac{n-2}{2}}$.

Because $T^{-1} = \Omega_E^{-1} \Omega_E^{-T}$, we see that $|T^{-1}|_{max} \leq a_3 |\operatorname{d}_{\mathcal{K}}|^{n-2}$. Thus, for sufficiently large q,

$$W \leq n|\hat{\Omega}_E|_{max}|T^{-1}|_{max} + 2^{-q+1/2}n|T^{-1}|_{max} \leq a_4|\operatorname{d}_{\mathcal{K}}|^{n-1}.$$

It follows that $\delta(q) \to 0$ as $q \to \infty$. This completes the first step of the proof.

Step 2:

We next need to give upper bounds on $\hat{M}_i(q)$ and $\hat{M}(q)$ to ensure that $c_5(q)$ cannot also get arbitrarily small as $q \to \infty$.

By Hadamard's inequality, $|\det(X)| \leq n^n|X|_{max}^n$, for any matrix complex matrix X, so $|\det(\hat{\Omega}_E)| \leq a_5|\operatorname{d}_{\mathcal{K}}|^{\frac{n}{2}}$. It follows by the argument in Proposition 17 relating the determinant of embedding matrices and lattice matrices that $|\det(\hat{\Omega}_L)| \leq a_5|\operatorname{d}_{\mathcal{K}}|^{\frac{n}{2}}$.

Further, $|\det(\hat{\Omega}_L)|$ is precisely $\det(\hat{\sigma}(\mathcal{O}_{\mathcal{K}}))$, and $\det(\hat{\sigma}(\beta))$ is $\det(\hat{\sigma}(\mathcal{O}_{\mathcal{K}}))$ times the determinant of the change of basis matrix corresponding to the respective lattice bases. By Proposition 12, this scale factor is precisely $\mathbb{N}(\beta)$, as the change of basis matrix from $\{\beta_1, \ldots, \beta_n\}$ to $\{\omega_1, \ldots, \omega_n\}$ is the same as the change of basis matrix for the image of these bases under $\hat{\sigma}$.

It follows that $\det(\hat{\sigma}(\beta)) \leq a_6|\operatorname{d}_{\mathcal{K}}|^{\frac{n}{2}}N^n$.

By Lemma 8, $|\hat{\mu}|_{max} \leq |\hat{\mu}| \leq \sqrt{n}\det(\hat{\beta})^{\frac{1}{n}} = a_7|\operatorname{d}_{\mathcal{K}}|^{\frac{1}{2}}N$.

Therefore, we may choose $\hat{M}_i(q) = a_7|\operatorname{d}_{\mathcal{K}}|^{\frac{1}{2}}$, which gives an upper bound on $\hat{M}_i(q)$ independent of q.

Next, we have $\hat{\mu} = \hat{\Omega}_L\mathbf{x}$. Our goal will be to bound the components of $\mathbf{x}$. This will be somewhat involved.

Note that the columns of $\hat{\Omega}_L$ are the $\{\hat{\sigma}(\omega_j)\}_{j=1}^n$ and these form a basis for the lattice $\hat{\beta}$. Further, $|\hat{\sigma}(\omega_j)| \leq |\sigma(\omega_j)|$, so $\prod_{j=1}^n |\hat{\sigma}(\omega_j)| \leq \prod_{j=1}^n |\sigma(\omega_j)|$.

Proceeding as in the proof of Lenstra's shortest vector algorithm, we see that

$$|x_i| \leq \frac{|\hat{\mu}|}{|\det(\hat{\Omega}_L)||\hat{\sigma}(\beta_i)|}\prod_{j=1}^n |\sigma(\omega_j)| \leq \frac{a_1|\hat{\mu}|}{|\hat{\sigma}(\beta_i)|}\frac{|\det(\Omega_L)|}{|\det(\hat{\Omega}_L)|}.$$

By Theorem 14, $|\det(\hat{\Omega}_L)| \geq |\det(\Omega_L)| - 2^{-q}D|\operatorname{d}_{\mathcal{K}}|^{\frac{n-1}{2}} \geq |\det(\Omega_L)| - 1$ for sufficiently large q. Similarly, $|\hat{\sigma}(\beta_i)|$ approaches $|\sigma(\beta_i)|$ as $q \to \infty$. So, for large q, $|\hat{\sigma}(\beta_i)| \geq |\sigma(\beta_i)| - 1$.

Putting everything together, we see that $|x_i| \leq a_1a_7N|\operatorname{d}_{\mathcal{K}}|^{\frac{1}{2}}\frac{|\det(\Omega_L)|}{(|\sigma(\beta_i)|-1)(|\det(\Omega_L)|-1)}$.

Therefore, we have some constant E not depending on q such that $|\mathbf{x}|_{max} \leq E$.

Observe that $W \geq |\Omega_E|_{max}^{-1} \geq \frac{1}{n}|\Omega_E|_{max}^{-1}$.

By definition, $\hat{M}(q) = \frac{|\mathbf{x}|_{max}}{nWN}$. Using our bounds on $|\Omega_E|_{max}$ and $|\mathbf{x}|_{max}$, we see that $\hat{M}(q) \leq \frac{a_1 E |\, \mathrm{d}_{\mathcal{K}} \,|^{\frac{1}{2}}}{N}$ for sufficiently large q, and our bound is independent of q.

Therefore, as $q \to \infty$, $\delta(q) \to 0$, while $c_5(q) \not\to 0$ because each of c_j, $1 \leq j \leq 4$ cannot grow arbitrarily large. It follows that $\delta(q)$ eventually falls below $c_5(q)$.

■

Computation of minima algorithm

We now have enough tools to compute minima of an ideal $\beta \subseteq \mathcal{O}_{\mathcal{K}}$. The process is relatively straightforward but rather computationally expensive. Remember that we have assumed that $\omega_1, \ldots, \omega_n$ is an LLL-reduced basis of $\mathcal{O}_{\mathcal{K}}$ throughout this section. Therefore, before proceeding forward, we must compute such a reduced basis (1). Then, we use Theorem 14 to see how large q must be for $\hat{\Omega}_L(q)$ to be invertible (2). After choosing q sufficiently large, we compute a shortest vector $2^q\hat{\mu}$ in $2^q\hat{\beta}(q) \subseteq \mathbb{Z}^n$. We then scale $2^q\hat{\mu}$ by 2^{-q} to obtain a shortest vector $\hat{\mu}$ in $\hat{\beta}$ with associated element μ (3). We then compute the c_j, $1 \leq j \leq 4$, corresponding to μ and check if $\delta(q) < c_5$ (4). If so, then μ is a minimum, and we terminate the algorithm. Otherwise, increment q and return to step (3). By the above theorem, $\delta(q)$ is guaranteed to eventually drop below $c_5(q)$, so the algorithm will terminate.

---

4 On the Number of Similarity Classes

For matrices over a field F, it is easy to determine the number of similarity classes of matrices with characteristic polynomial $\chi = \mu_1^{n_1} \cdots \mu_k^{n_k}$. Indeed, this number is simply the product $\prod_{i=1}^{k} p(n_i)$, where p denotes the number-theoretic partition function by the structure of the rational canonical form.

We can also write down the number of similarity classes of matrices with minimal polynomial $\mu_1^{j_1} \cdots \mu_k^{j_k}$ and characteristic polynomial $\mu_1^{n_1} \cdots \mu_k^{n_k}$ where $j_i \leq n_i$ for $1 \leq i \leq k$. Recall that the characteristic polynomial of a matrix is the product of its invariant factors (the largest of which is its minimal polynomial), and two matrices are similar if and only if they have the same invariant factors. Once we have fixed the minimal polynomial, the exponents on the ith irreducible factors of all smaller invariant factors must be $\leq j_i$. Further, these exponents are weakly decreasing (as we move from the largest to the smallest invariant factors) and their sum must be $n_i - j_i$. It follows that the number of similarity classes with given minimal and characteristic polynomial is $\prod_{i=1}^{k} p_{j_i}(n_i - j_i)$, where $p_{j_i}(h)$ is the number of partitions of h with terms less than or equal to j_i.

In particular, the characteristic polynomial is a similarity invariant for semisimple matrices over a field.

Unfortunately, the story for integer matrices is much more complicated. By the generalized Latimer-MacDuffee theorem, we know that similarity classes of semisimple integer matrices with given characteristic polynomial are in bijection with the isomorphism classes of certain modules over an order of a direct sum of number fields. In the simplest case of integer matrices with irreducible characteristic polynomial and purely monogenic eigenvalue α, this amounts to the existence of a bijection between similarity classes and the ideal class group of $\mathbb{Q}(\alpha)$. If we drop the purely monogenic assumption, we get a bijection between similarity classes and the ideal class monoid of $\mathbb{Z}[\alpha]$. Obviously, the characteristic polynomial is not a similarity invariant for semisimple integer matrices because there exist number fields of class number greater than 1.

The general theory of finitely generated torsion-free modules over hereditary rings tells us that the similarity classes of semisimple integer matrices with maximal associated order correspond to the similarity classes of matrices corresponding to each direct summand of the maximal order.

Explicitly, the number of similarity classes of integer matrices with minimal polynomial $\mu = \mu_{\alpha_1} \cdots \mu_{\alpha_k}$, characteristic polynomial $\chi = \mu_{\alpha_1}^{n_1} \cdots \mu_{\alpha_k}^{n_k}$, and $\mathbb{Z}[\alpha_1, \ldots, \alpha_k] = \bigoplus_{i=1}^{k} \mathcal{O}_{\mathbb{Q}(\alpha_i)}$ (associated order is maximal), is $\prod_{i=1}^{k} h_{\mathbb{Q}(\alpha_i)}$. Therefore, when dealing with the case of the maximal order, the number of similarity classes falls entirely within the field of algebraic number theory. In the case of general semisimple integer matrices, it seems that very little is known about the number of similarity classes. However, we do know that this number is guaranteed to be finite. Most of the rest of this section will be dedicated to establishing this result.

The proofs given here mostly follow material presented in sections 20 and 79 in Curtis and Reiner's classic representation theory text.

Let L be a skewfield of dimension n over $\mathbb{Q}$. Let Λ be a $\mathbb{Z}$-order in L. Observe that any map $\varphi : M_1 \to M_2$ of L_0-lattices in L extends uniquely to a map $\phi \in \mathrm{End}_L(L) \cong L^{\mathrm{op}}$, hence is given by multiplication by some $\alpha \in L$. Therefore, the isomorphism classes of L_0-lattices are the same as the equivalence classes of lattices under the relation $M_1 \sim M_2 \iff M_1 = M_2\alpha$ for some $0 \neq \alpha \in L$.

Theorem 16 *The number of isomorphism classes of L_0-lattices in L is finite.*

Before proceeding with the proof, notice that the finiteness of the class number immediately follows from the theorem by taking L as a number field and $L_0 = \mathcal{O}_L$. The cost of this greater generality is the loss of a bound such as Minkowski's bound for number fields. The structure of the proof is similar to

the proof of Minkowski's bound, except the bound on the norm of certain elements of isomorphism classes is obtained without considering the Minkowski lattice in $\mathbb{R}^n$.

Proof. For each element $\alpha \in L$, we can consider the "multiplication by α" map. This map is $\mathbb{Q}$-linear, and we define the norm $N(\alpha)$ to be its determinant. For each ideal $I \subseteq L_0$, we define its numerical norm $\mathbb{N}(I) = [L_0 : I]$. This is guaranteed to be finite because L_0 is Noetherian, hence I is an L_0-lattice and there exists some positive integer m such that $mL_0 \subseteq I$. It follows that $[L_0 : I] \leq [L_0 : mL_0] = m^n$. We also have that $\mathbb{N}(xL_0) = |N(x)|$ by an argument identical to the one used in the case of number fields.

Next, we want to show that there exists a positive constant c such that if X is an L_0-lattice, then there exists some $0 \neq x \in X$ such that $|N(x)| \leq c\mathbb{N}(X)$. In the case of number fields, this constant c is precisely Minkowski's bound $M_{\mathcal{K}}$.

Define the homogenous degree n polynomial $f(\xi_1, \ldots, \xi_n) = \det(\xi_1 M(\alpha_1) + \cdots + \xi_n M(\alpha_n))$, where $\{\alpha_i\}$ is a $\mathbb{Z}$-basis for L_0, and $M(\alpha_i)$ denotes the matrix representation for multiplication by α_i with respect to this basis. There exists $c > 0$ such that $|f(\xi_1, \ldots, \xi_n)| \leq ca^n$, where $|\xi_i| \leq a$. Let $X \subseteq L_0$ be an ideal and consider the set of all $\mathbb{Z}$-linear combinations of the integral basis $\alpha_1, \ldots, \alpha_n$ with scalars b_i between 0 and $\mathbb{N}(X)^{\frac{1}{n}}$.

Because this set has more than $\mathbb{N}(X)$ elements, there must be at least two distinct elements with equal classes mod X. Their difference is some element $x = \sum_{i=1}^{n} a_i \alpha_i \in X$ where $|a_i| \leq \mathbb{N}(X)^{\frac{1}{n}}$. Hence, $|N(x)| = |f(a_1, \ldots, a_n)| \leq c\mathbb{N}(X)$.

In the number field case, we could just conclude by Theorem 11, but we need a slightly more subtle argument for the general case because we are not assuming that L_0 is Dedekind (so some ideals may not be invertible). We will show that for each L_0-lattice X there exists an isomorphic L_0-lattice X' containing L_0 such that $[X' : L_0] \leq c$. By scaling, we can assume that $X \subseteq L_0$. Then, choose $x \in X$ such that $|N(x)| \leq c\mathbb{N}(X)$ and define $X' = Xx^{-1}$, so $X' \cong X$.

Observe that $L_0 x \subseteq X$, so $L_0 \subseteq X'$ and we have the following string of equalities:

$$[X' : L_0] = [Xx^{-1} : L_0] = [X : L_0 x] = [L_0 : L_0 x]/[L_0 : X] \leq |N(x)|/\mathbb{N}(X) \leq c.$$

To conclude, it will suffice to show that there are only finitely many L_0-lattices X' which contain L_0 such that $[X' : L_0] \leq c$. This is very similar to the statement that there are finitely many ideals of a given norm which we use in the number field case. Let $\ell = [X' : L_0]$ and observe that $\ell L_0 \subseteq \ell X' \subseteq L_0$, so $\ell X'$ is a subgroup of L_0 containing ℓL_0. Because these subgroups are in bijection with subgroups of the finite group $L/\ell L_0$, we see that there are finitely many choices for X', which concludes the proof.

■

The above theorem will serve as a stepping stone for the more general Jordan-Zassenhaus theorem. Let A be a finite-dimensional semisimple algebra over $\mathbb{Q}$. Let G be a $\mathbb{Z}$-order in A. For brevity, we will refer to finite rank $\mathbb{Z}$-free G-modules as simply G-modules. For example, when H is a finite group, and $A = \mathbb{Q}[H]$ and $G = \mathbb{Z}[H]$, then G-modules are precisely integral representations of G. Let L^* be an A-module. We define a *full* G-module in L^* to be a G-module that is a subset of L^* whose $\mathbb{Q}$-linear span is all of L^*. Let Λ denote the set of G-module isomorphism classes of full G-modules in L^*

Theorem 17 *(Jordan-Zassenhaus) The number of G-module isomorphism classes of full G-modules in L^* is finite, i.e., Λ is a finite set.*

Before proceeding with the proof, let us observe that this theorem implies that the number of similarity classes of semisimple integer matrices with given characteristic polynomial is finite. Indeed, following the notation is Section 2, if we take $A = \mathcal{K}$, $L^* = \mathcal{K}^{\mathbf{n}}$, and $G = \mathbb{Z}[v]$, the Jordan-Zassenhaus theorem tells us that the number of isomorphism classes of full $\mathbb{Z}[v]$-modules in $\mathcal{K}^{\mathbf{n}}$ is finite. By the generalized Latimer-MacDuffee theorem, these isomorphism classes are in bijection with the similarity classes of semisimple integer matrices with corresponding characteristic polynomial.

Proof. We will first address the case where L^* is an irreducible A-module. Because A is a semisimple algebra, up to isomorphism, we can view L^* as a subset of A. In particular, the irreducibility assumption tells us that L^* is a minimal left ideal of A. Further, by the Artin-Wedderburn theorem, we know that A is isomorphic to a finite product of matrix rings over division rings. Because ideals in a product of rings arise as the product of ideals in each component ring, the projection of L^* must be minimal in some matrix ring $B = \mathrm{Mat}_f(D)$, where D is an n-dimensional division ring over $\mathbb{Q}$ (this naturally arises as $\mathrm{End}_B(L^*)$ as in the proof of Artin-Wedderburn). By Schur's lemma, our surjective projection map is actually an isomorphism of representations. Therefore, up to isomorphism, L^* is a minimal left ideal of B.

We see immediately that B has dimension f^2n over $\mathbb{Q}$. We want to compute the dimension of L^* over $\mathbb{Q}$. By the discussion following Proposition 4, we know that D^f is the unique irreducible representation of B up to isomorphism. Because D^f has dimension fn over $\mathbb{Q}$, and $L^* \cong D^f$ as A-representations, we see that L^* has dimension fn over $\mathbb{Q}$ as well.

We can also view $D \subseteq B \subseteq A$, which allows us to define $D_0 = D \cap G$. As a submodule of G, we know that D_0 is a free $\mathbb{Z}$-module. We want to show that D_0 has full rank in D. Observe that $\mathbb{Q}D_0 = \mathbb{Q}D \cap \mathbb{Q}G = D \cap A = D$, hence D_0 has $\mathbb{Z}$-rank n. [*in progress*]

■

It turns out that semisimple matrices are the only class of matrices that admit finitely many $\mathbb{Z}$-similarity classes.

Theorem 18 *The number of $\mathbb{Z}$-similarity classes of integer matrices with minimal polynomial μ and characteristic polynomial χ is finite if and only if μ is square-free.*

If μ is semisimple, then by the Jordan-Zassenhaus theorem and the generalized Latimer-MacDuffee theorem, the number of similarity classes if finite. So, it remains to show that if μ is not square-free then there are infinitely many similarity classes of integer matrices with minimal polynomial μ and characteristic polynomial χ.

We will first show the result for nilpotent matrices, and then use the Jordan-Chevalley decomposition to obtain the general case.

Lemma 9 *Suppose A is a nonzero nilpotent integer matrix with minimal polynomial μ and characteristic polynomial χ. Then, $\{nA\}_{n=1}^{\infty}$ is an infinite collection of non-conjugate integer matrices with minimal polynomial μ and characteristic polynomial χ.*

Proof. We will show that nA is not similar to A over $\mathbb{Z}$ for any nonzero integer n. This follows from the fact that the one of the SNF invariant factors of A is $\gcd(A)$ and similar matrices have the same Smith Normal Form. Clearly, $\gcd(nA) = n\gcd(A) \neq \gcd(A)$, so nA and A have different invariant factors and cannot be similar. However, scaling by n does not change the minimal or characteristic polynomial of a nilpotent matrix, so we obtain infinitely many distinct similarity classes.

Further, all such scaling fall into the same similarity class over $\mathbb{Q}$. To see this, observe that N has a Jordan form over $\mathbb{Q}$ because its characteristic polynomial splits. For a block of the Jordan form of size r, we can conjugate by $\mathrm{diag}(j, j^2, \ldots, j^r)$ to scale by r. By stitching these diagonal matrices together, we obtain a diagonal matrix D_j that conjugates N to rN over $\mathbb{Q}$.

■

Recall that the Jordan-Chevalley decomposition for matrices over an algebraically closed field F asserts that if $A \in \mathrm{Mat}_n(F)$, then $A = S+N$, where S is semisimple, N is nilpotent, and $SN = NS$. By Galois theory, we are guaranteed more: the Jordan-Chevalley decomposition exists over any perfect field.

Theorem 19 *Suppose A is an integer matrix with minimal polynomial μ. If μ is not square-free, then the $\mathbb{Q}$-similarity class of A splits into infinitely many classes over $\mathbb{Z}$. In particular, the number of $\mathbb{Z}$-similarity classes of integer matrices with characteristic polynomial χ and minimal polynomial μ is infinite.*

Proof. Suppose $A \in \mathrm{Mat}_n(\mathbb{Z})$ is not semisimple and has minimal polynomial μ. Matrices similar over the integers have the same Smith form and thus the same GCD. It follows that if $kA \sim_{\mathbb{Z}} B$, then B is divisible by k and $A \sim_{\mathbb{Z}} \frac{1}{k}B$. Therefore, if the $\mathbb{Q}$-similarity class of kA splits into j integer similarity classes, then so does the $\mathbb{Q}$-similarity class of A.

By the Jordan-Chevalley decomposition, we can uniquely write $A = S + N$, where S is a semisimple rational matrix and N is a nilpotent rational matrix, and $SN = NS$. Because μ is not square-free, N is guaranteed to be nonzero. By choosing sufficiently large $k_1 \in \mathbb{N}$, we can write $k_1 A = k_1 S + k_1 N$, where S and N are integer matrices. Because $k_1 N$ is nilpotent, its minimal polynomial splits over $\mathbb{Q}$, so we can choose a rational matrix Q that conjugates it to its Jordan form: $k_1 QAQ^{-1} = k_1 QSQ^{-1} + k_1 QNQ^{-1}$.

It is possible that $k_1 QSQ^{-1}$ is not an integer matrix, which is problematic because we want to construct non-conjugate *integer* matrices with minimal polynomial μ. Therefore, choose $k_2 \in \mathbb{N}$ large enough so that $k_2 k_1 QSQ^{-1} \in \mathrm{Mat}_n(\mathbb{Z})$ and $k_2 k_1 QAQ^{-1} = k_2 k_1 QSQ^{-1} + k_2 J k_1 QNQ^{-1}$.

Let $B(i)$ denote the ith Jordan block of $k_2 k_1 QNQ^{-1}$ and suppose that $B(i)$ has size n_i. Observe that if $X \in \mathrm{Mat}_r(\mathbb{Q})$, conjugation by $\mathrm{diag}(j, j^2, \ldots, j^r)$ acts by scaling the mth superdiagonal of X by j^m and by scaling by j^{-m} on the mth subdiagonal. Therefore, conjugation by $\mathrm{diag}(j, j^2, \ldots, j^{n_i})$ acts on $B(i)$ by scaling the matrix by j.

Let D_j be the block diagonal matrix that scales $k_2 k_1 QNQ^{-1}$ by j through conjugation. We want to conjugate by D_j, but we need to ensure that $k_2 k_1 QSQ^{-1}$ remains an integer matrix after conjugation. Therefore, let $k_3 = j^n$ and scale to obtain $k_3 k_2 k_1 QAQ^{-1} = k_3 k_2 k_1 QSQ^{-1} + k_3 k_2 k_1 QNQ^{-1}$. Because conjugation by D_j divides entries by at most j^n, our choice of k_3 ensures that we can safely conjugate $k_3 k_2 k_1 QSQ^{-1}$.

Let $k = k_3 k_2 k_1$ and $\Lambda_j = D_j Q$. Conjugating by D_j, we get $k\Lambda_j A \Lambda_j^{-1} = k\Lambda_j S \Lambda_j^{-1} + k\Lambda_j N \Lambda_j^{-1}$, where both of the RHS summands are integer matrices. For any $X \in \mathrm{Mat}_n(\mathbb{Q})$, define $X_j = k\Lambda_j X \Lambda_j^{-1}$, so $A_j = S_j + N_j$ and $S_j N_j = N_j S_j$. Furthermore, S_j is semisimple and N_j is nilpotent. It is obvious that N_j is nilpotent. To see S_j is semisimple, recall that semisimplicity and potential diagonalizablility are equivalent over a perfect field like $\mathbb{Q}$. Therefore, conjugation and scaling do not change the semisimplicity of a matrix.

Now, $D_\ell A_j D_\ell^{-1}$ is an integer matrix similar to A_j over $\mathbb{Q}$, for $1 \leq \ell \leq j$. However, $D_\ell A_j D_\ell^{-1}$ is not similar to A_j over $\mathbb{Z}$. To see this, observe that if $D_\ell A_j D_\ell^{-1} \sim_{\mathbb{Z}} A_j$, then $A_j = P D_\ell S_j D_\ell^{-1} P^{-1} + P D_\ell N_j D_\ell^{-1} P^{-1}$ for some unimodular matrix P. By the uniqueness of the Jordan-Chevalley decomposition, we find that $P D_\ell N_j D_\ell^{-1} P^{-1} = N_j$. However, $D_\ell N_j D_\ell^{-1} = \ell N_j$, which cannot

be similar to N_j by Lemma 9. Therefore, the rational similarity class of A_j splits into at least j integral similarity classes. By the discussion at the start of the proof, this implies that the rational similarity class of A splits into at least j integral similarity classes. Since j was arbitrary, this implies the rational similarity class splits into infinitely many integer similarity classes.

■